An extended abstract of this paper appears in Christian Cachin, editor, *International Colloquium on Automata, Languages and Programming – ICALP 2007*, Volume ?? of Lecture Notes in Computer Science, pages ??–??, Wroclaw, Poland, July 9–13, 2007. Springer-Verlag, Berlin, Germany. This is the full version.

Unrestricted Aggregate Signatures

Mihir Bellare[1] University of California San Diego
mihir@cs.ucsd.edu, http://www-cse.ucsd.edu/users/mihir

Chanathip Namprempre[2] Thammasat University
nchanath@engr.tu.ac.th, http://chanathip.ee.engr.tu.ac.th

Gregory Neven[3] Katholieke Universiteit Leuven
Gregory.Neven@esat.kuleuven.be, http://www.neven.org

June 2006

Abstract

Secure use of the BGLS [7] aggregate signature schemes is restricted to the aggregation of distinct messages (for the basic scheme) or per-signer distinct messages (for the enhanced, prepend-public-key version of the scheme). We argue that these restrictions preclude interesting applications, make usage of the schemes error-prone and are generally undesirable in practice. Via a new analysis and proof, we show how the restrictions can be lifted, yielding the first truly unrestricted aggregate signature scheme. Via another new analysis and proof, we show that the distinct signer restriction on the sequential aggregate signature schemes of [18] can also be dropped, yielding an unrestricted sequential aggregate signature scheme. Finally, we present variants of these schemes with tight security reductions.

[1] Supported by NSF grants CNS-0524765 and CNS-0627779, and a gift from Intel Corporation.
[2] Supported by the Thailand Research Fund.
[3] Postdoctoral Fellow of the Research Foundation Flanders, supported in part by the Concerted Research Action Ambiorics 2005/11 of the Flemish Government and the European Commission through the IST Programme under Contract IST-2002-507932 ECRYPT.

1 Introduction

Aggregate signatures. An aggregate signature (AS) scheme [7] is a digital signature scheme with the additional property that a sequence $\sigma_1, \ldots, \sigma_n$ of individual signatures —here σ_i is the signature, under the underlying base signature scheme, of some message m_i under some public key pk_i— can be condensed into a single, compact aggregate signature σ that simultaneously validates the fact that m_i has been signed under pk_i for all $i = 1, \ldots, n$. There is a corresponding aggregate verification process that takes input $(pk_1, m_1), \ldots, (pk_n, m_n), \sigma$ and accepts or rejects. Aggregation is useful to reduce bandwidth and storage, and is especially attractive for mobile devices like sensors, cell phones, and PDAs where communication is more power-expensive than computation and contributes significantly to reducing battery life.

Schemes. Boneh, Gentry, Lynn, and Shacham [7] present an aggregate signature scheme based on the BLS signature scheme [9]. We call it $\mathcal{AS}$*-1* and represent it succinctly in the first row of Table 1. $\mathcal{AS}$*-1*, however, has some limitations. As the table shows, the aggregate verification process, on inputs (pk_1, m_1), $\ldots, (pk_n, m_n), \sigma$, rejects if the messages $m_1, \ldots, m_n$ are not distinct. The restriction is crucial because, without it, as shown in [7], the scheme is subject to a forgery attack. The consequence, however, is to restrict the ability to aggregate to settings where the messages signed by the signers are all different. BGLS recognize this limitation and suggest a workaround. Specifically, they say: "It is easy to ensure the messages are distinct: The signer simply prepends her public key to every message she signs ..." [7, Section 3.2]. They stop short of specifying a scheme in full, but since it is clearly their intention to "reduce to the previous case," our understanding is that they are advocating the modified scheme in which the signature of a message m under pk is the BLS signature of the *enhanced message* $M = pk\|m$ under pk while aggregate verification is done by applying the aggregate verification procedure of $\mathcal{AS}$*-1* to $(pk_1, pk_1\|m_1), \ldots, (pk_n, pk_n\|m_n), \sigma$. However, if so, in this scheme, which we call $\mathcal{AS}$*-2*, the aggregate verification process will reject unless the enhanced messages $pk_1\|m_1, \ldots, pk_n\|m_n$ are distinct. (Why? Because the aggregate verification process of $\mathcal{AS}$*-1* rejects unless the messages are distinct, and the role of the messages is now played by the enhanced messages.) The consequence is that the ability to aggregate is restricted to settings where the enhanced messages signed by the signers are all different. That is, the limitations have been reduced, but not eliminated.

Our result. We ask whether there exists a truly unrestricted proven-secure aggregate signature scheme. Namely, there should be no distinctness-based restriction of any kind, whether on messages or enhanced messages. We show that the answer is yes. Our result is a new, direct analysis of the security of enhanced-message signature aggregation which shows that the distinctness condition in the aggregate verification process of $\mathcal{AS}$*-2* —namely that this process rejects if any two enhanced messages are the same— can be dropped without compromising security. In other words, an unrestricted scheme can be obtained by the natural adaptation of $\mathcal{AS}$*-2* in which the distinctness condition in the verification is simply removed and all else is the same. This scheme, which we denote $\mathcal{AS}$*-3*, is summarized in the last row of Table 1. The fact that $\mathcal{AS}$*-3* is very close to $\mathcal{AS}$*-2* is a plus because it means existing implementations can be easily patched.

We clarify that the security of $\mathcal{AS}$*-3* is not proved in [7]. They prove secure only $\mathcal{AS}$*-1*. The security of $\mathcal{AS}$*-2* is a consequence, but the security of $\mathcal{AS}$*-3* is not. What we do instead is to *directly* analyze security in the case that signatures are on enhanced messages. Our analysis explicitly uses and exploits the presence of the prepended public keys to obtain the stronger conclusion that $\mathcal{AS}$*-3* (not just $\mathcal{AS}$*-2*) is secure.

Motivation. The limitation of $\mathcal{AS}$*-2* —namely that aggregation is restricted to settings where no two enhanced messages are the same— may seem minor, because all it says is that a set of signatures to be aggregated should not contain duplicates, meaning multiple signatures by a particular signer of a particular message. However, as we now explain, there are several motivations for schemes like $\mathcal{AS}$*-3* where aggregation is possible even in the presence of duplicates.

Consider a sensor network deployed in a remote area, for example, an environment monitoring network such as the tsunami early warning system already in operation in the Indian Ocean [15]. The sensors periodically record measurements from the environment and send them to a monitoring center. In applications where integrity is important (say, to prevent attackers from raising a false alarm that a tsunami is coming), each sensor node must sign its data. The center aggregates these data and the signatures

Scheme	Sign	Aggregate verification process accepts iff
$\mathcal{AS}$-*1* [7]	$\mathrm{H}(m)^x$	$\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(\mathrm{H}(m_i), g^{x_i})$ and $m_1, \ldots, m_n$ all distinct
$\mathcal{AS}$-*2* [7]	$\mathrm{H}(g^x \| m)^x$	$\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(\mathrm{H}(g^{x_i} \| m_i), g^{x_i})$ and $g^{x_1} \| m_1, \ldots, g^{x_n} \| m_n$ all distinct
$\mathcal{AS}$-*3*	$\mathrm{H}(g^x \| m)^x$	$\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(\mathrm{H}(g^{x_i} \| m_i), g^{x_i})$

Table 1: The aggregate signature schemes we discuss. Here $\mathbf{e}\colon \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is a bilinear map, g is a generator of $\mathbb{G}_2$ known to all parties, and $\mathrm{H}\colon \{0,1\}^* \to \mathbb{G}_1$ is a hash function. The second column shows the signature of a message m under public key g^x, generated using secret key x. In all cases, a sequence of signatures is aggregated by simply multiplying them in $\mathbb{G}_1$. The third column shows under what conditions the aggregate verification algorithm accepts σ as a valid aggregate signature of messages $m_1, \ldots, m_n$ under public keys $g^{x_1}, \ldots, g^{x_n}$ respectively.

---

to save storage. (The schemes discussed here permit on-line aggregation in which one can maintain a current aggregate and aggregate into it a received signature.) However, environmental measurements can certainly repeat over time! Indeed, especially in stable conditions, we would expect frequent repeats. Thus, a single signer (sensor) may sign the same data (measurement) many times. In general, whenever the data being signed is drawn from a small space, not just messages, but even enhanced messages can repeat and an unrestricted aggregate signature scheme is necessary.

Perhaps an even more important reason to prefer unrestricted schemes is that they are less likely to be misused or to result in unexpected errors. An application designer contemplating using $\mathcal{AS}$-*2* must ask herself whether, in her application, enhanced messages might repeat. This may not only be hard to know in advance, but might also change with time. (Experience has repeatedly shown that once a piece of cryptography is deployed, it is used for purposes or applications beyond those originally envisaged.) With an unrestricted scheme, the application designer is freed from the need to worry about whether the setting of her application meets the restrictions, reducing the chance of error. In general, application designers and programmers have a hard enough time as it is to make error-free use of cryptography. Asking them to understand message distinctness restrictions and anticipate whether their application meets them is an added burden and one more place where things can go wrong.

Possible workarounds. Various ways to get around message distinctness restrictions may come to mind, but these workarounds either do not work or tend to be more complex or restrictive than direct use of an unrestricted scheme. For example, one could have the verifier drop from its input list (pk_1, m_1), $\ldots, (pk_n, m_n), \sigma$ any pair pk_i, m_i that occured previously in the list, but then, assuming σ was a correct aggregate signature of the given list, it will not be a correct aggregate signature for the truncated list, and verification will fail. Another possibility is that the aggregator refrain from including in the aggregate any signature corresponding to a public key and message whose signature is already in the aggregate. But aggregation may be done on-line, and the aggregator may know only the current aggregate and the incoming signature, and have no way to tell whether or not it is a duplicate. Nonces (timestamps or sequence numbers) could be added to messages to render them unique, but this would complicate the application and increase cost.[1] Being able to use $\mathcal{AS}$-*3* without any worries about signature or message replication is simpler, easier, and more practical.

Results for SASs. A sequential aggregate signature (SAS) scheme [18] permits a more restrictive kind of aggregation in which the signers must participate in the process and use their secret keys. Imagine the signers forming a chain. In step i, the i-th signer receives from the previous one a current aggregate, and, using its secret key, it aggregates into this its own signature, passing the new aggregate on to the

[1] Nonces might be added anyway to prevent replay attacks. If so, well and good. But replay attacks are not always a concern. In some settings, an adversary might be able to inject bogus data, yet be unable to eavesdrop. Without eavesdropping, the adversary clearly cannot replay data. This could be true for the ocean sensor scenario we discussed above.

next signer in the chain. The output of the final signer is the aggregate signature. Although clearly less powerful than general aggregate signature (GAS) schemes —following [8] we now use this term for the BGLS-type aggregate signatures discussed above in order to distinguish them from SASs— the argument of [18] is that sequential aggregation is possible for base signature schemes other than BLS or may be done more cheaply. Specifically, Lysyanskaya, Micali, Reyzin, and Shacham [18] present a SAS scheme based on certified [5] claw-free trapdoor permutations.

However, the model and schemes of [18] also have some limitations. They require that no public key can appear more than once in a chain. That is, the signers who are part of a signing chain must be distinct. But in practice there are many reasons to prefer to allow aggregation even when a particular signer signs multiple messages, meaning when there are loops in the chain and public keys can repeat. Certainly, the previously discussed motivations are still true. Namely, in applications like signing in sensor nets or ad hoc networks, a particular signer (sensor) will be producing many signatures and it would be convenient to allow these to be aggregated. More importantly, an unrestricted SAS scheme is more misuse resistant because it does not require the application designer to try to predict in advance whether or not there will be repeated signers in a prospective chain. But in fact the restrictions in [18] are even greater than the ones for $\mathcal{AS}$-*2*, for they say not only that one cannot aggregate when a particular signer signs a particular message twice, but that one cannot aggregate even when a particular signer signs two or more messages that are distinct. This makes the number of excluded applications even larger, for it is common that a particular signer needs to sign multiple messages.

Our result —analogous to the GAS case— is a new analysis and proof that shows that the restrictions imposed by [18] on their schemes can be lifted without impacting security (that is, verification can just drop the condition that one reject if there are repeated public keys, and security is preserved), yielding unrestricted schemes. Again, that only a minor modification to the scheme is needed is convenient if existing implementations need to be updated.

Tight reduction results. The security of $\mathcal{AS}$-*1*, $\mathcal{AS}$-*2*, and $\mathcal{AS}$-*3* is based on the coCDH problem. However, none of the security reductions —namely, those of [7] in the first two cases and ours in the last— are tight. By applying the technique of Katz and Wang [16], we obtain an alternative scheme $\mathcal{AS}$-*4* with a tight security reduction to coCDH. When implemented over the same group, $\mathcal{AS}$-*1*, $\mathcal{AS}$-*2*, $\mathcal{AS}$-*3*, and $\mathcal{AS}$-*4* all have about the same computational cost, the price for $\mathcal{AS}$-*4* being a slightly larger aggregate signature. (Specifically, one needs one extra bit per constituent signature.) However, due to the tight reduction, $\mathcal{AS}$-*4* is actually *more* efficient than the other schemes when one compares them at the same (provable) security level. (Because to achieve a given level of security, $\mathcal{AS}$-*4* needs a smaller group than the other schemes.) We also obtain an analogous result for the SAS case. The statements and the proofs of our tight reduction results are in Section 5.

Related work. Lu, Ostrovsky, Sahai, Shacham, and Waters [17] present a SAS scheme for which they can lift the distinct signer restriction, as follows. If a signer wishes to add its signature of a message M_{new} to an aggregate S which already contains its signature S_{old} on some message M_{old}, then it removes S_{old} from S and then adds back in a signature on the message $M_{\mathrm{new}} \| M_{\mathrm{old}}$. However, their scheme is weaker than the others we have discussed —ours or those of [7, 18]— with regard to some security properties and also with regard to efficiency. Specifically, [17] uses the certified public key model [1, 6], which reflects the assumption that signers provide strong ZK proofs of knowledge of secret keys to the CA at key-registration, an assumption that we, following [7, 18], do not make. Also, in the scheme of [17], public keys are very large, namely 162 group elements, which is particularly expensive if public keys have to be transmitted along with signatures. In contrast, other schemes have short public keys. On the other hand, the proofs of [17] are in the standard model, while ours, following [7, 18], are in the random oracle model of [3].

Interestingly, in a survey paper, Boneh, Gentry, Lynn, and Shacham [8] present $\mathcal{AS}$-*3*, claiming that the results of [7] prove it secure. However, this appears to be an oversight because, as we have seen, the results of [7] prove $\mathcal{AS}$-*2* secure, not $\mathcal{AS}$-*3*. By proving $\mathcal{AS}$-*3* secure via our direct analysis, we are filling the gap and validating the claim of [8]. Shacham's Ph.D thesis [19] notes that the concrete security of the reduction of [7] can be slightly improved by replacing messages with enhanced ones, but he does not claim security of $\mathcal{AS}$-*3*.

2 Notation and Basic Definitions

Notation and conventions. If x is a string, then $|x|$ is the length of x. We denote by $x_1\|\cdots\|x_n$ an encoding of objects $x_1,\ldots,x_n$ as a binary string from which the constituent objects are uniquely recoverable. When the objects can be encoded as strings whose length is known from the context, simple concatenation will serve the purpose. If S is a finite set, then $|S|$ is its size, and $s \xleftarrow{\$} S$ means that s is chosen at random from S. We let e denote the base of the natural logarithm. An algorithm may be randomized unless otherwise indicated. An adversary is an algorithm. If A is an algorithm then $y \xleftarrow{\$} A(x_1,x_2,\ldots)$ means that y is the result of executing A on fresh random coins and inputs $x_1,x_2,\ldots$. We denote by $[A(x_1,x_2,\ldots)]$ the set of all possible outputs of A on the indicated inputs, meaning the set of all strings that have a positive probability of being output by A on inputs $x_1,x_2,\ldots$. We let $\mathsf{Maps}(D)$ denote the set of all functions with domain $\{0,1\}^*$ and range D.

Digital signature schemes. We recall definitions for (standard) signature schemes [13] in the random-oracle (RO) model [3]. A signature scheme $\mathcal{DS} = (\mathsf{Kg},\mathsf{Sign},\mathsf{Vf})$ is specified by three algorithms, the last deterministic. Via $(pk,sk) \xleftarrow{\$} \mathsf{Kg}$, a signer generates its public key pk and matching secret key sk, where $\mathrm{H}\colon \{0,1\}^* \to D$ is a random oracle whose range D is a parameter of the scheme. Via $\sigma \xleftarrow{\$} \mathsf{Sign}^{\mathrm{H}}(sk,m)$ the signer can generate a signature σ on a message m. A verifier can run $\mathsf{Vf}^{\mathrm{H}}(pk,m,\sigma)$ to obtain a bit, with 1 indicating accept and 0 reject. The consistency (or correctness) condition is that

$$\Pr\left[\mathsf{Vf}^{\mathrm{H}}(pk,m,\sigma)=1 \;:\; (pk,sk)\xleftarrow{\$}\mathsf{Kg}\;;\;\mathrm{H}\xleftarrow{\$}\mathsf{Maps}(D)\;;\;\sigma\xleftarrow{\$}\mathsf{Sign}^{\mathrm{H}}(sk,m)\right] \;=\; 1$$

for all messages $m \in \{0,1\}^*$. To capture security (unforgeability under chosen-message attack) we define the advantage of an adversary B as

$$\mathbf{Adv}^{\mathrm{uf\text{-}cma}}_{\mathcal{DS}}(\mathsf{B}) \;=\; \Pr\left[\mathsf{Vf}^{\mathrm{H}}(pk,m,\sigma)=1 \;:\; (pk,sk)\xleftarrow{\$}\mathsf{Kg}\;;\;\mathrm{H}\xleftarrow{\$}\mathsf{Maps}(D)\;;\;(m,\sigma)\xleftarrow{\$}\mathsf{B}^{\mathsf{Sign}^{\mathrm{H}}(sk,\cdot),\mathrm{H}}\right]\;.$$

To make this meaningful, we only consider adversaries that are *legitimate* in the sense that they never queried the message in their output to their sign oracle. We say that $\mathcal{DS}$ is $(t,q_{\mathrm{S}},q_{\mathrm{H}},\epsilon)$-secure if no adversary B running in time at most t, invoking the signature oracle at most q_{S} times and the random oracle at most q_{H} times, has advantage more than ϵ.

3 Unrestricted General Aggregate Signatures

GAS schemes. A general aggregate signature (GAS) scheme [7] $\mathcal{AS} = (\mathsf{Kg},\mathsf{Sign},\mathsf{Agg},\mathsf{AVf})$ consists of four algorithms, the last deterministic. The key generation and signing algorithms are exactly as for standard digital signatures. Via $\sigma \xleftarrow{\$} \mathsf{Agg}^{\mathrm{H}}((pk_1,m_1,\sigma_1),\ldots,(pk_n,m_n,\sigma_n))$, anyone can aggregate a sequence of public key, message, and signature triples to yield an aggregate signature σ. A verifier can run $\mathsf{AVf}^{\mathrm{H}}((pk_1,m_1),\ldots,(pk_n,m_n),\sigma)$ to obtain a bit, with 1 indicating accept and 0 reject.

Security. The security requirement of [7] is strong, namely that an adversary find it computationally infeasible to produce an aggregate forgery involving an honest signer, even when it can play the role of all other signers, in particular choosing their public keys, and can mount a chosen-message attack on the honest signer. To capture this, we define the advantage of an adversary A as

$$\mathbf{Adv}^{\mathrm{agg\text{-}uf}}_{\mathcal{AS}}(\mathsf{A}) \;=\; \Pr\left[\mathsf{AVf}^{\mathrm{H}}((pk_1,m_1),\ldots,(pk_n,m_n),\sigma)=1\right]$$

where the probability is over the experiment

$$(pk,sk)\xleftarrow{\$}\mathsf{Kg}\;;\;\mathrm{H}\xleftarrow{\$}\mathsf{Maps}(D)\;;\;((pk_1,m_1),\ldots,(pk_n,m_n),\sigma)\xleftarrow{\$}\mathsf{A}^{\mathsf{Sign}^{\mathrm{H}}(sk,\cdot),\mathrm{H}}(pk)\;.$$

To make this meaningful, we only consider adversaries that are *legitimate* in the sense that there must exist $i \in \{1,\ldots,n\}$ such that $pk_i = pk$ but A never queried m_i to its signing oracle. Thus, the honest signer here is the one whose keys are pk, sk, and we are asking that the aggregate forgery include some message and signature corresponding to this honest signer, but the adversary never legitimately obtained a signature of this message. We say that A $(t,q_{\mathrm{S}},n_{\max},q_{\mathrm{H}},\epsilon)$-breaks $\mathcal{AS}$ if it runs in time at most t, invokes the signature oracle at most q_{S} times, invokes the hash oracle at most q_{H} times, outputs a forgery containing at most $n_{\max}$ public-key-message pairs, and has advantage strictly greater than ϵ. We say that $\mathcal{AS}$ is $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-secure if there is no adversary that $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-breaks $\mathcal{AS}$.

A significant feature of this definition, highlighted in [7], is that A can choose $pk_1, \ldots, pk_n$ as it wishes, in particular as a function of pk. Unlike [1, 6, 17], there is no requirement that the adversary "know" the secret key corresponding to a public key it produces, and this makes the system more practical since it avoids the need for strong zero-knowledge proofs of knowledge [2] of secret keys done to the CA during key-registration. Our results continue to achieve this strong notion of security.

Bilinear maps and coCDH. Let $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ be groups, all of the same prime order p. Let $\mathbf{e}$: $\mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ be a non-degenerate, efficiently computable bilinear map, also called a pairing. Let g be a generator of $\mathbb{G}_2$. Note that following [9, 7] we use the asymmetric setting ($\mathbb{G}_1, \mathbb{G}_2$ are not necessarily equal) and must also assume there exists an isomorphism ψ: $\mathbb{G}_2 \to \mathbb{G}_1$. (The first is in order to make signatures as short as possible, and the second is required for the security proofs.) For the rest of the paper, we regard $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e}, g, \psi$ as fixed, globally known parameters, and also let $t_{\exp}$ denote the time to perform an exponentiation in $\mathbb{G}_1$. We define the advantage of an adversary A in solving the coCDH problem as

$$\mathbf{Adv}^{\text{co-cdh}}(\mathsf{A}) \;=\; \Pr\left[\, \mathsf{A}(g, g^a, h) = h^a \;:\; h \xleftarrow{\$} \mathbb{G}_1 \,;\, a \xleftarrow{\$} \mathbb{Z}_p \,\right] \;.$$

We say that the coCDH problem is (t', ϵ')-hard if no algorithm A running in time at most t' has advantage strictly more than ϵ' in solving it. Note that when $\mathbb{G}_1 = \mathbb{G}_2$, the coCDH problem becomes the standard CDH problem in $\mathbb{G}_1$, whence the name.

The $\mathcal{BLS}$ scheme. We recall the $\mathcal{BLS}$ standard signature scheme of [9]. The signer chooses a secret key $x \xleftarrow{\$} \mathbb{Z}_p$ and computes the corresponding public key $X \leftarrow g^x$. Let $\mathrm{H} : \{0,1\}^* \to \mathbb{G}_1$ be a random oracle. The signature of message m is $\sigma = \mathrm{H}(m)^x$, which can be verified by checking that $\mathbf{e}(\sigma, g) = \mathbf{e}(\mathrm{H}(m), X)$. Regarding security, we have the following:

Lemma 3.1 **[9]** *If the coCDH problem is (t', ϵ')-hard, then the $\mathcal{BLS}$ standard signature scheme is $(t, q_{\mathrm{S}}, q_{\mathrm{H}}, \epsilon)$-secure for any $t, q_{\mathrm{S}}, q_{\mathrm{H}}, \epsilon$ satisfying*

$$\epsilon \;\geq\; e(q_{\mathrm{S}} + 1) \cdot \epsilon' \quad \textit{and} \quad t \;\leq\; t' - t_{\exp}(q_{\mathrm{H}} + 2q_{\mathrm{S}}) \;. \;\blacksquare \tag{1}$$

The GAS schemes we consider. We consider four closely related aggregate signature schemes that we denote $\mathcal{AS}\text{-}0, \mathcal{AS}\text{-}1, \mathcal{AS}\text{-}2, \mathcal{AS}\text{-}3$. These schemes all use a random oracle $\mathrm{H} : \{0,1\}^* \to \mathbb{G}_1$. The key generation algorithm of $\mathcal{AS}\text{-}0$ and $\mathcal{AS}\text{-}1$ picks x at random from $\mathbb{Z}_p$, returns x as the secret key, and $X = g^x$ as the public key. The key generation algorithm of $\mathcal{AS}\text{-}2$ and $\mathcal{AS}\text{-}3$ picks x at random from $\mathbb{Z}_p$, sets $X = g^x$, returns (x, X) as the secret key, and X as the public key. For $\mathcal{AS}\text{-}0$ and $\mathcal{AS}\text{-}1$, the signing algorithm is the $\mathcal{BLS}$ one, namely the signature on message m is $\sigma = \mathrm{H}(m)^x$. For $\mathcal{AS}\text{-}2$ and $\mathcal{AS}\text{-}3$, a signature on m under public key X is $\sigma = \mathrm{H}(X \| m)^x$. For all schemes, aggregation is done by simply multiplying the signatures, i.e. $\sigma = \prod_{i=1}^{n} \sigma_i$ in $\mathbb{G}_1$. Verification is different for each scheme. On inputs $(X_1, m_1), \ldots, (X_n, m_n), \sigma$, the verification algorithm of $\mathcal{AS}\text{-}0$ accepts iff $\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(\mathrm{H}(m_i), X_i)$. The verification algorithms of the other schemes are depicted in Table 1. In particular, $\mathcal{AS}\text{-}1$ rejects if $m_1, \ldots, m_n$ are not all distinct, $\mathcal{AS}\text{-}2$ rejects if $X_1 \| m_1, \ldots, X_n \| m_n$ are not all distinct, while $\mathcal{AS}\text{-}3$ performs no such checks.

Consistency conditions. The consistency condition (under what conditions correctly generated aggregates are accepted by the verifier) differs from scheme to scheme, and is in fact the place where the restrictions they make surface in a formal sense. $\mathcal{AS}\text{-}0$ and $\mathcal{AS}\text{-}3$ meet the natural, strongest possible consistency requirement, namely that

$$\Pr\left[\, \mathsf{AVf}^{\mathrm{H}}((pk_1, m_1), \ldots, (pk_n, m_n), \sigma) = 1 \,\right] \;=\; 1$$

for all positive integers n, all messages $m_1, \ldots, m_n \in \{0,1\}^*$ and all $(pk_1, sk_1), \ldots, (pk_n, sk_n) \in [\mathsf{Kg}]$, where the probability is over the experiment $\mathrm{H} \xleftarrow{\$} \mathsf{Maps}(D)\,;\, \sigma_1 \xleftarrow{\$} \mathsf{Sign}^{\mathrm{H}}(sk_1, m_1)\,;\, \cdots\,;\, \sigma_n \xleftarrow{\$} \mathsf{Sign}^{\mathrm{H}}(sk_n, m_n)\,;\, \sigma \xleftarrow{\$} \mathsf{Agg}^{\mathrm{H}}((pk_1, m_1, \sigma_1), \ldots, (pk_n, m_n, \sigma_n))$. However, $\mathcal{AS}\text{-}1$ meets this condition only when $m_1, \ldots, m_n$ are distinct and $\mathcal{AS}\text{-}2$ when $pk_1 \| m_1, \ldots, pk_n \| m_n$ are distinct.

Discussion of security. An attack provided in [7] shows that $\mathcal{AS}\text{-}0$ is insecure. In this attack, however, the forgery output by the adversary contains repeated messages. To exclude the attack, [7] defines $\mathcal{AS}\text{-}1$,

Subroutine H-SIM(M)
If $(\exists m : M = X^*\|m)$ then return $\mathrm{H}_{\mathcal{BLS}}(M)$
Else If $\mathrm{HT}[M] = \perp$ then $\mathrm{y}[M] \xleftarrow{\$} \mathbb{Z}_p$; $\mathrm{HT}[M] \leftarrow \psi(g)^{\mathrm{y}[M]}$
Return $\mathrm{HT}[M]$

Subroutine SIGN-SIM(m)
Return $\mathsf{Sign}_{\mathcal{BLS}}(x, X^*\|m)$

Figure 1: The subroutines for B to simulate the random oracle $\mathrm{H}_{\mathcal{AS}\text{-}3}(\cdot)$ and the sign oracle $\mathsf{Sign}_{\mathcal{AS}\text{-}3}((x, X^*), \cdot)$. Above, HT and y are associative arrays assumed initially to have value $\perp$ everywhere.

---

where the aggregate verification process rejects when messages repeat. They are able to show that this suffices to guarantee security, meaning that they prove *AS-1* is secure if the coCDH problem is hard. This is their main result. Then they suggest to alleviate the message-distinctness restriction of *AS-1* by having each signer prepend its public key to the message before signing. However, they appear to want to argue the security of the resulting aggregate signature scheme as a corollary of their main result on the security of *AS-1*. If so, verification still needs to check that $X_1\|m_1, \ldots, X_n\|m_n$ are all distinct (otherwise, the result about *AS-1* does not apply), leading to the *AS-2* scheme.

As we have discussed, however, for practical reasons, *AS-3* is a preferable scheme. But the results of [7] do not prove it secure. Here is an example that helps to see what the problem is. Suppose there was an adversary A that, on input $pk = X$ and without making oracle query m, produced a forgery of the form $(X, m), (X', m'), (X', m'), \sigma$, for some $m' \neq m$ and $X' \neq X$, that was accepted by the verification procedure of *AS-3*. Since the output of A contains repeated enhanced messages, the results of [7] do not allow us to rule out the existence of A. Yet, showing that *AS-3* meets the notion of security that we have defined does require ruling out the existence of such an A.

Theorem 3.2 *If the coCDH problem is (t', ϵ')-hard, then the AS-3 aggregate signature scheme is $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-secure for any $t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon$ satisfying*

$$\epsilon \;\geq\; e(q_{\mathrm{S}}+1)\cdot\epsilon' \;\text{ and }\; t \;\leq\; t' - t_{\exp}(2q_{\mathrm{H}} + 2q_{\mathrm{S}} + 3n_{\max} + 1)\;. \blacksquare \tag{2}$$

Our approach to the proof is different from the one used by [7] to prove that *AS-1* is secure if coCDH is hard. They gave a direct reduction to coCDH, meaning, given an adversary attacking *AS-1* they construct and analyze an adversary attacking coCDH. But, in so doing, they end up duplicating a lot of the proof of the security of the *BLS* scheme as given in [9]. Instead, we reduce the security of *AS-3* to the security of *BLS*. That is, we prove the following:

Lemma 3.3 *If the BLS standard signature scheme is $(t', q'_{\mathrm{S}}, q'_{\mathrm{H}}, \epsilon')$-secure then the AS-3 aggregate signature scheme is $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-secure for any $t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon$ satisfying*

$$\epsilon \geq \epsilon',\; q_{\mathrm{S}} \leq q'_{\mathrm{S}} - n_{\max},\; q_{\mathrm{H}} \leq q'_{\mathrm{H}} \;\text{ and }\; t \leq t' - t_{\exp}\cdot(q_{\mathrm{H}} + n_{\max} + 1)\;. \blacksquare \tag{3}$$

Theorem 3.2 follows easily from Lemma 3.3 and Lemma 3.1. Our modular approach yields a simple proof even though we obtain a somewhat stronger result.

An interesting element of the proof of Lemma 3.3 is that it involves reducing the security of one random oracle model scheme to another one. Given a forger A against *AS-3* that queries a random oracle, we must build a forger B against *BLS*. But B is itself given a random oracle. The idea is that B will answer some of A's queries via its own random oracle and directly simulate the others.

Proof of Lemma 3.3: Given a forger A that $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-breaks *AS-3*, consider the following forger B against the *BLS* standard signature scheme. B is given public key $X^* = g^x$ as input, and has access to a random oracle $\mathrm{H}_{\mathcal{BLS}}(\cdot)$ and a signing oracle $\mathsf{Sign}_{\mathcal{BLS}}(x, \cdot) = \mathrm{H}_{\mathcal{BLS}}(\cdot)^x$. It runs A on input X^* and responds to its $\mathrm{H}_{\mathcal{AS}\text{-}3}(\cdot)$ and $\mathsf{Sign}_{\mathcal{AS}\text{-}3}((x, X^*), \cdot)$ oracle queries using the subroutines in Figure 1.

When A submits a query M to its random oracle $\mathrm{H}_{\mathcal{AS}\text{-}3}(\cdot)$, the forger B executes H-SIM(M). We note here that in some cases the subroutine H-SIM can in turn submit queries to B's random oracle $\mathrm{H}_{\mathcal{BLS}}(\cdot)$. When A submits a query m to its sign oracle $\mathsf{Sign}_{\mathcal{AS}\text{-}3}((x, X^*), \cdot)$, the forger B executes SIGN-SIM(m). Eventually, A halts and outputs a forgery $(X_1, m_1), \ldots, (X_n, m_n), \sigma$. Since A is legitimate, we know that there exists $i \in \{1, \ldots, n\}$ such that $X_i = X^*$ and m_i has never been queried to $\mathsf{Sign}_{\mathcal{AS}\text{-}3}((x, X^*), \cdot)$. We let i^* denote the smallest integer i for which this happens. Now B defines the sets:

$$
\begin{array}{rcl}
I & = & \{\, i \mid X_i = X^* \text{ and } m_i = m_{i^*} \,\} \\
J & = & \{\, i \mid X_i = X^* \text{ and } m_i \neq m_{i^*} \,\} \\
K & = & \{\, i \mid X_i \neq X^* \,\} \,.
\end{array}
$$

Note that I is non-empty since $i^* \in I$. Clearly, we have that $I \cup J \cup K = \{1, \ldots, n\}$ and that I, J, K are disjoint. Now, we can assume without loss of generality that $n < p$, because otherwise B can trivially forge and output a $\mathcal{BLS}$ signature under X^* in time $O(nt_{\exp})$ via exhaustive search for x. This means that $|I| \in \mathbb{Z}_p^*$, and hence has an inverse modulo p that we denote by l. Now, for each $i \in J$, our adversary B executes $\textsc{Sign-sim}(m_i)$ to obtain $\sigma_i \leftarrow \mathsf{Sign}_{\mathcal{BLS}}(x, X^* \| m_i)$. For each $i \in K$, it calls its subroutine $\text{H-}\textsc{sim}(X_i \| m_i)$, thereby ensuring that $\mathrm{y}[X_i \| m_i]$ is defined, lets $y_i \leftarrow \mathrm{y}[X_i \| m_i]$, and then lets $\sigma_i \leftarrow \psi(X_i)^{y_i}$, which we note is the $\mathcal{BLS}$ signature of $X_i \| m_i$ under public key X_i. Finally, B lets

$$
M^* \leftarrow X^* \| m_{i^*} \qquad \text{and} \qquad \sigma^* \leftarrow \left(\sigma \cdot \textstyle\prod_{i \,\in\, J \cup K} \sigma_i^{-1}\right)^l \,, \tag{4}
$$

and outputs (M^*, σ^*) as its forgery.

For the analysis, we first argue that if A's forgery is valid then B's forgery is valid too. Assuming the former, the verification equation of $\mathcal{AS}\text{-}3$ tells us that

$$
\begin{aligned}
\mathbf{e}(\sigma, g) &= \prod_{i=1}^{n} \mathbf{e}\big(\mathrm{H}_{\mathcal{AS}\text{-}3}(X_i \| m_i)\ ,\ X_i\big) \\
&= \prod_{i \in I} \mathbf{e}\big(\mathrm{H}_{\mathcal{BLS}}(X^* \| m_{i^*})\ ,\ X^*\big) \cdot \prod_{i \in J} \mathbf{e}\big(\mathrm{H}_{\mathcal{AS}\text{-}3}(X^* \| m_i)\ ,\ X^*\big) \cdot \prod_{i \in K} \mathbf{e}\big(\mathrm{H}_{\mathcal{AS}\text{-}3}(X_i \| m_i)\ ,\ X_i\big) \\
&= \mathbf{e}\big(\mathrm{H}_{\mathcal{BLS}}(X^* \| m_{i^*})\ ,\ X^*\big)^{|I|} \cdot \prod_{i \in J \cup K} \mathbf{e}(\sigma_i\ ,\ g) \ .
\end{aligned} \tag{5}
$$

Above, (5) is true because σ_i, as computed above by B, is the $\mathcal{BLS}$ signature of $X_i \| m_i$ under public key X_i, for all $i \in J \cup K$. We then applied the verification equation of the $\mathcal{BLS}$ scheme. Now we see that if σ^* is defined by (4) then, from the above and the fact that $|I| \cdot l \equiv 1 \pmod{p}$ we have

$$
\begin{aligned}
\mathbf{e}(\sigma^*, g) &= \mathbf{e}(\sigma, g)^l \cdot \prod_{i \in J \cup K} \mathbf{e}(\sigma_i, g)^{-l} \\
&= \mathbf{e}\big(\mathrm{H}_{\mathcal{BLS}}(X^* \| m_{i^*})\ ,\ X^*\big)^{|I| \cdot l \bmod p} \cdot \prod_{i \in J \cup K} \mathbf{e}(\sigma_i\ ,\ g)^l \cdot \prod_{i \in J \cup K} \mathbf{e}(\sigma_i, g)^{-l} \\
&= \mathbf{e}\big(\mathrm{H}_{\mathcal{BLS}}(M^*)\ ,\ X^*\big)\ ,
\end{aligned}
$$

which means that σ^* is a valid $\mathcal{BLS}$ signature of M^* under public key X^*.

Furthermore, it is easy to see that the answers that B provided to the oracle queries of A are distributed identically to the ones that A would have obtained from its oracles in the game defining its advantage. The last thing we need to check is that B is legitimate, meaning did not query $M^* = X^* \| m_{i^*}$ to its $\mathsf{Sign}_{\mathcal{BLS}}(x, \cdot)$ oracle. But it did not do so while answering $\mathsf{Sign}_{\mathcal{AS}\text{-}3}((x, X^*), \cdot)$ oracle queries of A because A, being legitimate itself, did not query m_{i^*} to its $\mathsf{Sign}_{\mathcal{AS}\text{-}3}((x, X^*), \cdot)$ oracle. Now B also called $\mathsf{Sign}_{\mathcal{BLS}}(x, \cdot)$ on $X^* \| m_i$ for all $i \in J$, but by definition of J, we know that $m_i \neq m_{i^*}$. Putting everything together, we get $\mathbf{Adv}^{\text{uf-cma}}_{\mathcal{BLS}}(\mathsf{B}) \geq \mathbf{Adv}^{\text{agg-uf}}_{\mathcal{AS}\text{-}3}(\mathsf{A})$.

Finally, we analyze the resource usage of B. For q_{S}, we note that B makes sign queries in only two situations: (1) whenever A makes a sign query, so does B, and (2) once A outputs a forgery, B makes $|J| \leq n_{\max}$ additional sign queries. For q_{H}, it is easy to see that B makes at most the same number of random oracle queries to $\mathrm{H}_{\mathcal{BLS}}$ as A makes to $\mathrm{H}_{\mathcal{AS}\text{-}3}$. For t, notice that B (1) answers q_{H} random oracle queries, each of which results in a call to $\text{H-}\textsc{sim}$, (2) possibly makes $n_{\max}$ more calls to $\text{H-}\textsc{sim}$ after A outputs its forgery, and (3) computes one exponentiation in the last step to convert A's forgery into its own. Thus, the claimed running time bound follows, and the proof is complete. ▮

4 Unrestricted Sequential Aggregate Signatures

SAS SCHEMES. A sequential aggregate signature (SAS) scheme [18] $\mathcal{SAS} = (\mathsf{Kg}, \mathsf{SASign}, \mathsf{SAVf})$ consists of three algorithms, the last deterministic. The key generation algorithm is exactly as for standard digital signatures. The first signer computes a signature on message m by calling $\sigma \leftarrow \mathsf{SASign}^{\mathrm{H}}(sk, m)$. Subsequent signers run $\sigma \xleftarrow{\$} \mathsf{SASign}^{\mathrm{H}}(sk, m, \sigma', (pk_1, m_1), \ldots, (pk_n, m_n))$, where $n > 0$, to aggregate their signature on a message m into a given sequential aggregate signature σ' corresponding to a sequence of public-key-message pairs. A verifier can run $\mathsf{SAVf}^{\mathrm{H}}((pk_1, m_1), \ldots, (pk_n, m_n), \sigma)$, where $n \geq 0$, to obtain a bit, with 1 indicating accept and 0 reject.

SECURITY. The security of a SAS scheme in the random oracle model requires that it be computationally infeasible for an adversary to produce a sequential aggregate forgery involving an honest signer, even when it can play the role of all other signers and can mount a chosen-message attack on the honest signer. Formally, the advantage of an adversary A is

$$\mathbf{Adv}^{\text{seq-agg-uf}}_{\mathcal{SAS}}(\mathsf{A}) \;=\; \Pr\left[\mathsf{SAVf}^{\mathrm{H}}((pk_1, m_1), \ldots, (pk_n, m_n), \sigma) = 1\right]$$

where the probability is over the experiment

$$(pk, sk) \xleftarrow{\$} \mathsf{Kg}\;;\; \mathrm{H} \xleftarrow{\$} \mathsf{Maps}(D)\;;\; ((pk_1, m_1), \ldots, (pk_n, m_n), \sigma) \xleftarrow{\$} \mathsf{A}^{\mathsf{SASign}^{\mathrm{H}}(sk, \cdots), \mathrm{H}}(pk)\;.$$

To make this meaningful, we impose that A be *legitimate* in the sense that there exists $i \in \{1, \ldots, n\}$ such that $pk_i = pk$ and there exists no $\sigma_{i-1} \in \{0,1\}^*$ such that A submitted query $m_i, \sigma_{i-1}, (pk_1, m_1), \ldots, (pk_{i-1}, m_{i-1})$ to the signing oracle. Thus, the honest signer here is the one whose keys are pk, sk, and we are asking that the sequential aggregate forgery include some message and signature corresponding to this honest signer, but the adversary never legitimately obtained a sequential aggregate signature of this message. We say that A $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-breaks $\mathcal{SAS}$ if it runs in time at most t, invokes its signing oracle at most q_{S} times, invokes its random oracle at most q_{H} times, has advantage strictly greater than ϵ, and there are at most $n_{\max}$ public keys involved in either its forgery or any of its queries to its signing or random oracles. We say that $\mathcal{SAS}$ is $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-secure if there is no adversary that $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-breaks $\mathcal{SAS}$.

The security requirement we are making is stronger than the one of [18]. One can view their definition as asking for the same condition but for a more restricted (smaller) class of adversaries, namely ones whose output must not contain repeated public keys and who may not submit any queries containing repeating public keys to the signing oracle. We will show that a construction proposed in [18] remains secure even under our more stringent definition.

CERTIFIED CLAW-FREE TRAPDOOR PERMUTATIONS. The schemes we consider use a family of certified claw-free trapdoor permutations over a group [18]. Let us recall the definitions. Let $\mathbb{G}$ be a multiplicative group. A *family of certified claw-free trapdoor permutations* Π over $\mathbb{G}$ is a 4-tuple $(\mathsf{Gen}, \mathsf{Eval}, \mathsf{Inv}, \mathsf{Test})$ of algorithms, all but the first being deterministic. Via $(\pi, \overline{\pi}, \pi^{-1}) \xleftarrow{\$} \mathsf{Gen}$, one can generate (the descriptions of) a pair of permutations $\pi, \overline{\pi}$ on $\mathbb{G}$ and (the trapdoor information describing) the inverse permutation π^{-1}. For all $(\pi, \overline{\pi}, \pi^{-1}) \in [\mathsf{Gen}]$ and all $x \in \mathbb{G}$, the evaluation algorithm Eval, on inputs π, x, returns $\pi(x)$ in time at most T_Π, and on inputs $\overline{\pi}, x$ returns $\overline{\pi}(x)$ in time at most T_Π, where T_Π is a number associated to Π. The inversion algorithm Inv takes input π^{-1} and $y \in \mathbb{G}$ and returns $\pi^{-1}(y)$ in time at most T_Π. The map Test takes input any string π' and, in time at most T_Π, returns a bit, this being 1 if and only if $\mathsf{Eval}(\pi', \cdot)$ is a permutation on $\mathbb{G}$ whose computation takes time at most T_Π. We assume the time for multiplication, inversion, and sampling random elements in $\mathbb{G}$ is also at most T_Π. We define the advantage of an algorithm B in finding a claw in Π as

$$\mathbf{Adv}^{\text{claw}}_{\Pi}(\mathsf{B}) \;=\; \Pr\left[\pi(x) = \overline{\pi}(y) \;\middle|\; (\pi, \overline{\pi}, \pi^{-1}) \xleftarrow{\$} \mathsf{Gen}\;;\; (x, y) \xleftarrow{\$} \mathsf{B}(\pi, \overline{\pi})\right]\;.$$

We say that Π is (t', ϵ')-claw-free if there is no adversary B that runs in time at most t' yet has advantage strictly more than ϵ'. From now on we will confound the permutations and their descriptions, writing $\pi(x)$ for $\mathsf{Eval}(\pi, x)$, and so on.

THE SAS SCHEMES WE CONSIDER. The $\mathcal{SAS}$*-0* scheme of [18] associated to a family of certified claw-free trapdoor permutations Π works as follows. Each signer generates a key pair (pk, sk) by computing

Alg $\mathsf{SASign}^{\mathrm{H}}((\pi,\pi^{-1}),m,\sigma_n,(\pi_1,m_1),\ldots,(\pi_n,m_n))$
If $n=0$ then $\sigma_0 \leftarrow 1$
If $\mathsf{SAVf}^{\mathrm{H}}((\pi_1,m_1),\ldots,(\pi_n,m_n),\sigma_n) \neq 1$
then Return $\perp$
$h \leftarrow \mathrm{H}(\pi_1\|m_1\|\ldots\|\pi_n\|m_n\|\pi\|m)$
$\sigma \leftarrow \pi^{-1}(h\cdot\sigma_n)$
Return σ

Alg $\mathsf{SAVf}^{\mathrm{H}}((\pi_1,m_1),\ldots,(\pi_n,m_n),\sigma)$
$\boxed{\text{If } \pi_1,\ldots,\pi_n \text{ not all different then Return } 0}$
$\sigma_n \leftarrow \sigma$
For $i=n,\ldots,1$ do
If $\mathsf{Test}(\pi_i)=0$ then Return 0
$h_i \leftarrow H(\pi_1\|m_1\|\ldots\|\pi_i\|m_i)$
$\sigma_{i-1} \leftarrow \pi_i(\sigma_i)\cdot h_i^{-1}$
If $\sigma_0=1$ then Return 1 else Return 0

Figure 2: The sequential aggregation and verification algorithms of $\mathcal{SAS}$-*0*. Removing the framed text yields $\mathcal{SAS}$-*1*.

$(\pi,\overline{\pi},\pi^{-1}) \xleftarrow{\$} \mathsf{Gen}$; $pk \leftarrow \pi$; $sk \leftarrow (\pi,\pi^{-1})$. Let $\mathrm{H}:\{0,1\}^* \to \mathbb{G}$ be a random oracle. The sequential aggregation and verification algorithms are described in Figure 2.

The $\mathcal{SAS}$-*0* scheme was proposed and proven secure in [18]. We consider here the $\mathcal{SAS}$-*1* scheme that is identical to the $\mathcal{SAS}$-*0* scheme, except that the boxed distinctness test in the verification algorithm in Figure 2 is dropped. (Note that this also affects the signing algorithm because it invokes the verification algorithm as a subroutine.)

Consistency conditions. Similar to general aggregate signatures, the consistency condition for sequential aggregate signatures differs from scheme to scheme. $\mathcal{SAS}$-*1* meets the natural, strongest possible consistency requirement, namely that

$$\Pr\Big[\mathsf{SAVf}^{\mathrm{H}}((pk_1,m_1),\ldots,(pk_n,m_n),\sigma_n)=1\Big] \;=\; 1$$

for all positive integers n and all messages $m_1,\ldots,m_n \in \{0,1\}^*$, the probability being in the experiment where we select $\mathrm{H} \xleftarrow{\$} \mathsf{Maps}(D)$; $(pk_1,sk_1),\ldots,(pk_n,sk_n) \in [\mathsf{Kg}]$, then select $\sigma_1 \xleftarrow{\$} \mathsf{SASign}^{\mathrm{H}}(sk_1, m_1)$; $\sigma_2 \xleftarrow{\$} \mathsf{SASign}^{\mathrm{H}}(sk_2,m_2,\sigma_1,(pk_1,m_1))$; $\cdots$; $\sigma_n \xleftarrow{\$} \mathsf{SASign}^{\mathrm{H}}(sk_n,m_n,\sigma_{n-1},(pk_1,m_1),\ldots,(pk_{n-1}, m_{n-1}))$. However, $\mathcal{SAS}$-*0* meets this condition only when $pk_1,\ldots,pk_n$ are distinct.

Discussion of security. No claims were made in [18] regarding the security of the $\mathcal{SAS}$-*1* scheme. Unlike the case of the $\mathcal{AS}$-*1* general aggregate signature scheme, there does not seem to be an easy attack when the distinctness condition is dropped. At the same time, the security proof of [18] clearly ceases to go through for the $\mathcal{SAS}$-*1* scheme, because the simulation of signatures and the way the forgery is exploited explicitly rely on all public keys being distinct. One may therefore rightfully wonder what the reason for this restriction is, and whether it is strictly necessary.

Our result. The following implies that the distinctness restriction in the $\mathcal{SAS}$-*0* scheme can be dropped:

Theorem 4.1 *If the family of certified claw-free trapdoor permutations Π is (t',ϵ')-claw-free, then the $\mathcal{SAS}$-1 sequential aggregate signature scheme is $(t,q_{\mathrm{S}},n_{\max},q_{\mathrm{H}},\epsilon)$-secure for any $t,q_{\mathrm{S}},n_{\max},q_{\mathrm{H}},\epsilon$ satisfying*

$$\epsilon \;\geq\; e(2q_{\mathrm{S}}+1)\cdot\epsilon' \quad and \quad t \leq t' - n_{\max}\cdot(q_{\mathrm{H}}+2q_{\mathrm{S}}+1)\cdot O(T_\Pi)\ . \ \blacksquare$$

We achieve this result through a number of refinements to the security proof of [18]. As in their proof, our simulator responds to random oracle queries in such a way that it either knows the corresponding sequential aggregate signature, or embeds its own challenge into the response. It then hopes that all the forger's signature queries correspond to hash queries of the first type, and that its forgery corresponds to one of the second type. Our refinements allow us to handle signature queries with multiple occurrences of the target public key, either by answering them correctly, or by directly exploiting the aggregate signature provided by the adversary to find a claw for the underlying permutation. The proof, written using the code-based game-playing technique [4], follows.

Games. Our proof of Theorem 4.1 will use code-based game-playing [4], and we begin by recalling some background from [4]. A game —look at Figure 4 for examples— has an **Initialize** procedure, procedures to respond to adversary oracle queries, and a **Finalize** procedure. A game G is executed with an adversary A, as follows. First, **Initialize** executes, and its outputs are the inputs to A. Then the latter executes, its oracle queries being answered by the corresponding procedures of G. When A terminates, its output becomes the input to the **Finalize** procedure. The output of the latter, denoted G^{A}, is called the output of the game, and we let "$G^{\mathsf{A}} \Rightarrow y$" denote the event that this game output takes value y. The boolean flag bad is assumed initialized to `false`, and GOOD will always denote the event that it is never set to `true`. (This event is defined in all games.) Games G_i, G_j are *identical until* bad if their code differs only in statements that follow the setting of bad to `true`. For examples, games G_0, G_1 of Figure 4 are identical until bad. The following is one version of the Fundamental Lemma of [4].

Lemma 4.2 [4] *Let G_i, G_j be identical until* bad *games, and* A *an adversary. Then*

$$\Pr\left[G_i^{\mathsf{A}} \Rightarrow 1 \ \wedge \ \text{GOOD} \right] \ = \ \Pr\left[G_j^{\mathsf{A}} \Rightarrow 1 \ \wedge \ \text{GOOD} \right] . \ \blacksquare$$

PROOF OF THEOREM 4.1. We say that an adversary A against *SAS-1* is *simplified* if it has the following properties, where π denotes the public key input to A and π^{-1} its inverse:

1. A never repeats a query to its H-oracle.
2. Any H-query of A has the form $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$ for some $n \geq 1$, some $\pi_1, \ldots, \pi_n$ such that $\mathsf{Test}(\pi_i) = 1$ for all $1 \leq i \leq n$, and some $m_1, \ldots, m_n \in \{0,1\}^*$.
3. If A makes a query $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ then it previously made H-query $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n \| \pi \| m_{n+1}$.
4. If A outputs $((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma)$ then it previously made H-query $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$.
5. If A makes H-query $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$ then it previously made H-query $\pi_1 \| m_1 \| \cdots \| \pi_{n-1} \| m_{n-1}$. (And hence, inductively, has already queried $\pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$ for all $1 \leq i \leq n-1$, in this order.)
6. If A makes a query $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ then $\mathsf{SAVf}^{\mathrm{H}}((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma_n) = 1$.
7. If A makes a query $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ then for every $i \in \{1, \ldots, n\}$ such that $\pi_i = \pi$, it is the case that A previously queried $m_i, \sigma_{i-1}, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$.
8. Let T denote the set of all $\pi_1 \| m_1 \| \ldots \| \pi_l \| m_l \| \pi \| m_{l+1}$ such that for some σ_l, adversary A queried $m_{l+1}, \sigma_l, (\pi_1, m_1), \ldots, (\pi_l, m_l)$ to its $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ oracle. Then, $|T|$ is always *exactly* q_{S} (rather than at most q_{S}, which must be true because A makes at most q_{S} queries to the oracle). By "always," we mean regardless of A's input, coin tosses, and the answers it receives to its oracle queries.

There are two stages in our proof. First, we transform a given adversary A′ attacking *SAS-1* into a simplified adversary A attacking *SAS-1* without loss in advantage although at some cost in resources:

Lemma 4.3 *Let* A′ *be an adversary that $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-breaks SAS-1. Then, we can construct a simplified adversary* A *that $(t^*, 2q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}^*, \epsilon)$-breaks SAS-1 where*

$$t^* \ = \ t + n_{\max}(q_{\mathrm{H}} + 2q_{\mathrm{S}} + 1) \cdot O(T_{\Pi}) \ \ , \ \ q_{\mathrm{H}}^* \ = \ n_{\max}(q_{\mathrm{H}} + 2q_{\mathrm{S}} + 1) \ . \ \blacksquare$$

Next, we construct an adversary B which enlists A's help in breaking the family Π underlying *SAS-1*:

Lemma 4.4 *Let* A *be a simplified adversary that $(t^*, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}^*, \epsilon)$-breaks SAS-1. Then, we can construct an adversary* B *attacking Π such that*

$$\mathbf{Adv}_{\Pi}^{\mathrm{claw}}(\mathsf{B}) \ \geq \ \frac{\mathbf{Adv}_{\mathcal{SAS}\text{-}1}^{\text{seq-agg-uf}}(\mathsf{A})}{e(q_{\mathrm{S}} + 1)} \ , \qquad (6)$$

and the running time of B *is at most that of* A *plus $(q_{\mathrm{H}}^* + 1) \cdot O(T_{\Pi})$.* $\blacksquare$

Theorem 4.1 follows directly from these two lemmas. We now proceed to prove the lemmas. A convention made in writing code is that all array elements are assumed to initially be $\perp$.

Proof of Lemma 4.3: The adversary A has access to oracles H and $\mathsf{SASign}^{\mathrm{H}}$. On input a public key π, it begins by initializing S and T to $\emptyset$. It then runs $\mathsf{A}'(\pi)$, answering hash and sign oracle queries of A′

Subroutine H-SIM(x)
If $\exists\, n, \pi_1, \ldots, \pi_n,\, m_1, \ldots, m_n$ such that
  – $n \geq 1$ and $x = \pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$
  – $\forall i \,:\, 1 \leq i \leq n \,:\, m_i \in \{0,1\}^*$ and $\mathsf{Test}(\pi_i) = 1$
Then
  For $i = 1, \ldots, n$ do $Q_i \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$; If $\mathrm{HT}[Q_i] = \bot$ then $\mathrm{HT}[Q_i] \leftarrow \mathrm{H}(Q_i)$
Else If $\mathrm{HT}[x] = \bot$ then $\mathrm{HT}[x] \stackrel{\$}{\leftarrow} \mathbb{G}$
Return $\mathrm{HT}[x]$

Subroutine SIGN-SIM$(m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))$
$\pi_{n+1} \leftarrow \pi$; $\sigma_0 \leftarrow 1$; $T \leftarrow T \cup \{\pi_1 \| m_1 \| \ldots \| \pi_l \| m_l \| \pi \| m_{l+1}\}$
For $i = 1, \ldots, n+1$ do If $\mathsf{Test}(\pi_i) = 0$ then return $\bot$; $Q_i \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$
$y \leftarrow$ H-SIM(Q_{n+1})
For $i = n, \ldots, 1$ do $\sigma_{i-1} \leftarrow \pi_i(\sigma_i) \cdot \mathrm{HT}[Q_i]^{-1}$
If $\sigma_0 \neq 1$ then return $\bot$
$S \leftarrow S \cup \{(m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))\}$
For $i = 1, \ldots, n$ do
  If $\pi_i = \pi$ and $(m_i, \sigma_{i-1}, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})) \notin S$ then Output $((\pi_1, m_1), \ldots, (\pi_i, m_i), \sigma_i)$
$\sigma_{n+1} \leftarrow \mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))$
Return σ_{n+1}

Figure 3: Subroutines for adversary A.

---

via the subroutines H-SIM and SIGN-SIM, respectively, of Figure 3. Note that these subroutines call A's oracles. When A' halts with some output $((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma_n)$, adversary A

(1) makes an additional $q = |T| - q_{\mathrm{S}}$ queries $(m'_1, 1), \ldots, (m'_q, 1)$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ oracle, where $m'_1, \ldots, m'_q$ are distinct and $\pi \| m'_i \notin T$ for all $1 \leq i \leq q$,
(2) calls H-SIM$(\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n)$,
(3) outputs $((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma_n)$, and halts.

Let us now explain how this ensures that A has the properties listed above without loss in advantage as compared to A'.

Property **1** is achieved because A stores as $\mathrm{HT}[x]$ the answer to H-query x of A', and, if this query is repeated, returns $\mathrm{HT}[x]$ without re-querying the oracle. A answers H-query x of A' via its own H oracle only when x has the form $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$ with $\mathsf{Test}(\pi_i) = 1$ for all $1 \leq i \leq n$, and, otherwise, itself picks a random value to play the role of the answer. This ensures property **2**, yet will not decrease the advantage of A because the algorithms of $\mathcal{SAS}$*-1* never invoke the H oracle on inputs not of the above form. Properties **3** and **4** are obtained by having A make the extra H-query if necessary. Property **5** is provided by having A query all appropriate un-queried prefixes of a H-query $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$ before querying the latter. Algorithm $\mathsf{SASign}^{\mathrm{H}}(\pi^{-1}, \cdots)$ returns $\bot$ on input $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$ unless $\mathsf{SAVf}^{\mathrm{H}}((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma_n) = 1$, so we have A do this test and refrain from making the query unless the answer is one, providing property **6**. Property **7** is the most interesting, and an important element in dealing with loops in signing chains. To explain how A provides it, suppose A' made a query $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ such that for some $1 \leq i \leq n$ it was the case that $\pi_i = \pi$ but A did not previously query $m_i, \sigma_{i-1}, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$. Then A, rather than making query $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$, outputs $((\pi_1, m_1), \ldots, (\pi_i, m_i), \sigma_i)$ as its forgery and halts. Property **6** tells us that $\mathsf{SAVf}^{\mathrm{H}}((\pi_1, m_1), \ldots, (\pi_i, m_i), \sigma_i) = 1$, and the fact that A did not previously query $m_i, \sigma_{i-1}, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})$ to $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ means that A remains legitimate. So this can only increase the advantage of A compared to that of A'. Property **8** holds because, once A' terminates, A makes additional sign-oracle queries that ensures that $|T|$ becomes

Initialize Game G_0 / $\boxed{G_1}$

100 $(\pi, \overline{\pi}, \pi^{-1}) \xleftarrow{\$} \mathsf{Gen}$; $\sigma[\varepsilon] \leftarrow 1$; $S \leftarrow \emptyset$
101 Return π

On H**-query** $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$
110 $Q_{n-1} \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_{n-1} \| m_{n-1}$; $Q_n \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$; $\sigma[Q_n] \xleftarrow{\$} \mathbb{G}$
111 If $\pi_n = \pi$ then $\mathrm{c}[Q_n] \xleftarrow{\delta} \{0,1\}$ else $\mathrm{c}[Q_n] \leftarrow 0$
112 If $\mathrm{c}[Q_n] = 0$ then $\mathrm{HT}[Q_n] \leftarrow \pi_n(\sigma[Q_n]) \cdot \sigma[Q_{n-1}]^{-1}$ else $\mathrm{HT}[Q_n] \leftarrow \overline{\pi}(\sigma[Q_n]) \cdot \sigma[Q_{n-1}]^{-1}$
113 Return $\mathrm{HT}[Q_n]$

On $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$**-query** $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$
120 $S \leftarrow S \cup \{(m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))\}$; $\pi_{n+1} \leftarrow \pi$; $\sigma_0 \leftarrow 1$
121 For $i = 1, \ldots, n+1$ do $Q_i \leftarrow \pi_1 \| m_1 \| \ldots \| \pi_i \| m_i$
122 $\sigma_{n+1} \leftarrow \sigma[Q_{n+1}]$
123 If $(\exists\, i : 1 \le i \le n+1 : \mathrm{c}[Q_i] = 1)$ then $\mathtt{bad} \leftarrow \mathtt{true}$; $\boxed{\sigma_{n+1} \leftarrow \pi^{-1}(\sigma_n \cdot \mathrm{HT}[Q_{n+1}])}$
124 Return σ_{n+1}

Finalize$((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma)$
130 $s \leftarrow \min\{\, i \mid 1 \le i \le n,\ \pi_i = \pi \text{ and } \forall \tau : (m_i, \tau, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})) \notin S \,\}$
131 $\sigma_n \leftarrow \sigma$
132 For $i = n, \ldots, 1$ do $Q_i \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$; $\sigma_{i-1} \leftarrow \pi_i(\sigma_i) \cdot \mathrm{HT}[Q_i]^{-1}$
133 If $\sigma_0 = 1$ then $d \leftarrow 1$ else $d \leftarrow 0$
134 If $(\exists\, i : 1 \le i \le s-1 : \mathrm{c}[Q_i] = 1)$ then $\mathtt{bad} \leftarrow \mathtt{true}$
135 If $\mathrm{c}[Q_s] = 0$ then $\mathtt{bad} \leftarrow \mathtt{true}$
136 $\omega \leftarrow (\sigma_s, \sigma[Q_s])$
137 Return d

Figure 4: Game G_1 includes the boxed statement, while game G_0 does not. Also, notice that G_1 uses π^{-1}, but G_0 does not.

---

exactly q_{S}. This does not affect the advantage of A but does increase the number of sign-oracle queries by at most a factor of 2. ▮

Proof of Lemma 4.4: We consider the games of Figure 4. The array entry $\mathrm{HT}[\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n]$ plays the role of $\mathrm{H}(\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n)$. The notation $c \xleftarrow{\delta} \{0,1\}$ means that c is assigned 0 with probability δ and 1 with probability $1-\delta$, where $\delta \in [0,1]$ is a parameter whose value will be chosen later [11]. These games rely on some of the properties of A listed above. For example, when answering H-query $Q_n = \pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$, property **5** allows us to assume $\mathrm{HT}[Q_{n-1}]$, and thus also $\sigma[Q_{n-1}]$, are already defined, where $Q_{n-1} = \pi_1 \| m_1 \| \cdots \| \pi_{n-1} \| m_{n-1}$. Similarly, property **3** tells us that the relevant $\mathrm{HT}[\cdot]$ entries are defined at the time a $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$ query is made. Property **4** tell us that the relevant $\mathrm{HT}[\cdot]$ entries are defined at the time **Finalize** is run, and the legitimacy of A tells us that s computed at line 130 is well-defined in the sense that the set over which the minimum is taken is not empty. Notice that G_1 uses π^{-1} but G_0 does not.

Let us say that a H-query $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$ is *simulation signed* if $\mathrm{c}[Q_i] = 0$ for all $1 \le i \le n$, where $Q_i = \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$. Then line 112 tells us that, in both G_0 and G_1, we have:

Claim 1. Let $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$ be a simulation signed hash query, and let $Q_i = \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$ for $0 \le i \le n$. Then for all $1 \le i \le n$ we have $\sigma[Q_i] = \pi_n^{-1}(\sigma[Q_{i-1}] \cdot \mathrm{HT}[Q_i])$. □

Initialize Game $\boxed{G_2}$ / G_3

$(\pi, \overline{\pi}, \pi^{-1}) \xleftarrow{\$} \mathsf{Gen}$; $S \leftarrow \emptyset$
Return π

On H**-query** $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$

$Q_n \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$
If $\pi_n = \pi$ then $\mathrm{c}[Q_n] \xleftarrow{\delta} \{0,1\}$ else $\mathrm{c}[Q_n] \leftarrow 0$
$\mathrm{HT}[Q_n] \xleftarrow{\$} \mathbb{G}$
Return $\mathrm{HT}[Q_n]$

On $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$**-query** $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$

$S \leftarrow S \cup \{(m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))\}$; $\pi_{n+1} \leftarrow \pi$; $\sigma_0 \leftarrow 1$
For $i = 1, \ldots, n+1$ do $Q_i \leftarrow \pi_1 \| m_1 \| \ldots \| \pi_i \| m_i$
If $\mathrm{c}[Q_{n+1}] = 1$ then $\mathsf{bad} \leftarrow \mathtt{true}$
$\boxed{\text{If } (\exists\, i : 1 \leq i \leq n : \mathrm{c}[Q_i] = 1) \text{ then } \mathsf{bad} \leftarrow \mathtt{true}}$
$\sigma_{n+1} \leftarrow \pi^{-1}(\sigma_n \cdot \mathrm{HT}[Q_{n+1}])$
Return σ_{n+1}

Finalize$((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma)$

$s \leftarrow \min \{\, i \mid 1 \leq i \leq n,\ \pi_i = \pi \text{ and } \forall \tau : (m_i, \tau, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})) \notin S \,\}$
$\sigma_n \leftarrow \sigma$
For $i = n, \ldots, 1$ do $Q_i \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$; $\sigma_{i-1} \leftarrow \pi_i(\sigma_i) \cdot \mathrm{HT}[Q_i]^{-1}$
If $\sigma_0 = 1$ then $d \leftarrow 1$ else $d \leftarrow 0$
$\boxed{\text{If } (\exists\, i : 1 \leq i \leq s-1 : \mathrm{c}[Q_i] = 1) \text{ then } \mathsf{bad} \leftarrow \mathtt{true}}$
If $\mathrm{c}[Q_s] = 0$ then $\mathsf{bad} \leftarrow \mathtt{true}$
Return d

Figure 5: Game G_2 includes the boxed statements while G_3 does not.

Claim 2. Let $\omega = (\sigma_s, \sigma[Q_s])$ be as per line 136 and d as per line 133. Then

$$\Pr\left[\, \pi(\sigma_s) = \overline{\pi}(\sigma[Q_s]) \,\right] \ \geq\ \Pr\left[\, d = 1 \ \wedge\ \textsc{good} \,\right],$$

where both probabilties are over the execution of G_0 with A.

Proof. If $\textsc{good}$ holds then line 134 tells us that $\pi_1 \| m_1 \| \cdots \| \pi_{s-1} \| m_{s-1}$ is simulation signed. If $d = 1$ then from Claim 1 and line 132 we get $\sigma_i = \sigma[Q_i]$ for all $1 \leq i \leq s-1$, and, in particular, $\sigma_{s-1} = \sigma[Q_{s-1}]$. If $\textsc{good}$ holds then line 135 implies $\mathrm{c}[Q_s] = 1$, and then line 112 implies $\mathrm{HT}[Q_s] = \overline{\pi}(\sigma[Q_s]) \cdot \sigma[Q_{s-1}]^{-1}$. Thus we have

$$\pi(\sigma_s) \ =\ \pi_s(\sigma_s) \ =\ \mathrm{HT}[Q_s] \cdot \sigma_{s-1} \ =\ \overline{\pi}(\sigma[Q_s]) \cdot \sigma[Q_{s-1}]^{-1} \cdot \sigma_{s-1} \ =\ \overline{\pi}(\sigma[Q_s])\,. \quad \Box$$

Now define adversary B against Π as follows. On inputs $\pi, \overline{\pi}$, it initializes $\sigma[\varepsilon] \leftarrow 1$ and $S \leftarrow \emptyset$. It then runs A on input π, answering its oracle queries as per the code of game G_0. When A halts, B runs the **Finalize** code of G_0 on input the output of A. It outputs ω and halts. Note that B is based on G_0 rather than G_1 and thus does not need to know π^{-1}. Then by Claim 2 we have

$$\mathbf{Adv}^{\mathrm{claw}}_{\Pi}(\mathsf{B}) \ \geq\ \Pr\left[\, G_0^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \textsc{good} \,\right]. \tag{7}$$

However, G_0 and G_1 are identical-until-bad games, and so Lemma 4.2 implies that

$$\Pr\left[\, G_0^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \textsc{good} \,\right] \ =\ \Pr\left[\, G_1^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \textsc{good} \,\right]. \tag{8}$$

Claim 3. In the execution of G_1 with A, all oracle queries of the latter are answered correctly.

Initialize Game G_4

$(\pi, \overline{\pi}, \pi^{-1}) \xleftarrow{\$} \mathsf{Gen}$; $S \leftarrow \emptyset$
Return π

On H**-query** $\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$

$\mathrm{HT}[\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n] \xleftarrow{\$} \mathbb{G}$
Return $\mathrm{HT}[\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n]$

On $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$**-query** $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$

$S \leftarrow S \cup \{(m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))\}$; $\sigma_0 \leftarrow 1$
$\sigma_{n+1} \leftarrow \pi^{-1}(\sigma_n \cdot \mathrm{HT}[\pi_1 \| m_1 \| \cdots \| \pi_n \| m_n \| \pi \| m_{n+1}])$
Return σ_{n+1}

Finalize$((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma)$

$s \leftarrow \min \{ i \mid 1 \le i \le n,\ \pi_i = \pi \text{ and } \forall \tau : (m_i, \tau, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})) \notin S \}$
$\sigma_n \leftarrow \sigma$
For $i = n, \ldots, 1$ do $Q_i \leftarrow \pi_1 \| m_1 \| \cdots \| \pi_i \| m_i$; $\sigma_{i-1} \leftarrow \pi_i(\sigma_i) \cdot \mathrm{HT}[Q_i]^{-1}$
If $\sigma_0 = 1$ then $d \leftarrow 1$ else $d \leftarrow 0$
For all $\pi'_1 \| m'_1 \| \cdots \| \pi'_l \| m'_l$ such that $\mathrm{HT}[\pi'_1 \| m'_1 \| \cdots \| \pi'_l \| m'_l] \neq \bot$ do
  If $\pi'_l = \pi$ then $\mathrm{c}[\pi'_1 \| m'_1 \| \cdots \| \pi'_l \| m'_l] \xleftarrow{\delta} \{0, 1\}$ else $\mathrm{c}[\pi'_1 \| m'_1 \| \cdots \| \pi'_l \| m'_l] \leftarrow 0$
For all $(m'_{l+1}, \sigma'_l, (\pi'_1, m'_1), \ldots, (\pi'_l, m'_l)) \in S$ do
  If $\mathrm{c}[\pi'_1 \| m'_1 \| \ldots \| \pi'_l \| m'_l \| \pi \| m'_{l+1}] = 1$ then $\mathsf{bad} \leftarrow \mathtt{true}$
If $\mathrm{c}[Q_s] = 0$ then $\mathsf{bad} \leftarrow \mathtt{true}$
Return d

Figure 6: Game G_4.

Proof. First consider a H-query $Q_n = \pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$. The random choice of $\sigma[Q_n]$ at line 110, together with the fact that $\pi_n, \overline{\pi}$ are permutations, then implies that $\mathrm{HT}[Q_n]$ is uniformly distributed, meaning the answer to this query is exactly as would be given by a random oracle. Next consider a $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$-query $m_{n+1}, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$. We consider two cases. If $Q_{n+1} = \pi_1 \| m_1 \| \ldots \| \pi_n \| m_n \| \pi \| m_{n+1}$ is simulation signed then Claim 1 tells us that the value $\sigma[Q_{n+1}]$ returned is correct. Otherwise, the value σ_{n+1} returned is computed by the boxed statement at line 123 and is correct because it is computed just as in $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$. Here we use property **6**, which tells us that σ_{n-1} was correct. $\square$

Claim 3 implies that

$$\Pr\left[\, G_1^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \text{GOOD} \,\right] \;=\; \Pr\left[\, G_2^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \text{GOOD} \,\right], \tag{9}$$

where game G_2 is in Figure 5. Game G_2 directly answers all oracle queries correctly, meaning just as in the game defining the advantage of A. Additionally it splits up the setting of bad as done by line 123 of G_1 into lines 222, 223. Next we claim that

$$\Pr\left[\, G_2^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \text{GOOD} \,\right] \;=\; \Pr\left[\, G_3^{\mathsf{A}} \Rightarrow 1 \ \wedge\ \text{GOOD} \,\right]. \tag{10}$$

To justify this, we explain why lines 223, 234 of G_2 are redundant and can simply be dropped to arrive at G_3. First consider line 223. Suppose $\mathrm{c}[Q_i] = 1$ for some $1 \le i \le n$. Then it must be that $\pi_i = \pi$, since, otherwise, due to line 211, $\mathrm{c}[Q_i]$ can only be 0. But then property **7** says that A must have previously made $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$-query $m_i, \sigma_{i-1}, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})$. If so, line 222 would have set bad at the time this query was made. Now consider line 234. The definition of s and line 211 tell us that if there is a $1 \le i \le s-1$ such that $\mathrm{c}[Q_i] = 1$ then it must be that $\pi_i = \pi$ and $(m_i, \tau, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})) \in S$ for some τ. But then, again, bad would have been set by line 222 when $\mathsf{SASign}^{\mathrm{H}}((\pi, \pi^{-1}), \cdots)$-query

$m_i, \tau, (\pi_1, m_1), \ldots, (\pi_{i-1}, m_{i-1})$ was made.

In Game G_3, the responses to oracle queries do not depend on the value of the flag bad. Thus, the choices of c[·] and the setting of bad can be postponed, meaning G_3 is equivalent to game G_4 of Figure 6. Now we claim that, in the execution of A with Game G_4, the events "$G_4^{\mathsf{A}} \Rightarrow 1$" and GOOD are independent, so that

$$\begin{aligned}
\Pr\left[\, G_3^{\mathsf{A}} \Rightarrow 1 \ \wedge \ \textsc{good} \,\right] &= \Pr\left[\, G_4^{\mathsf{A}} \Rightarrow 1 \ \wedge \ \textsc{good} \,\right] \\
&= \Pr\left[\, G_4^{\mathsf{A}} \Rightarrow 1 \,\right] \cdot \Pr\left[\, \textsc{good} \,\right] \\
&= \Pr\left[\, G_4^{\mathsf{A}} \Rightarrow 1 \,\right] \cdot \Pr\left[\, G_4^{\mathsf{A}} \text{ doesn't set } \mathsf{bad} \,\right] . \qquad (11)
\end{aligned}$$

Let us justify the independence claim. We observe that the random choices c[·] are made after the game output is determined. Since GOOD depends only on these choices, we would like to conclude the independence claim. But there is a subtle point. We also need that the number of (biased) coins involved in determining GOOD is the same in all runs. This is ensured by Property **8** of the simplified adversary.

Now, the output d of G_4^{A} is 1 exactly when A succeeds in forgery, meaning

$$\Pr\left[\, G_4^{\mathsf{A}} \Rightarrow 1 \,\right] \;=\; \mathbf{Adv}^{\text{seq-agg-uf}}_{\mathcal{SAS}\text{-}1}(\mathsf{A}) \,. \qquad (12)$$

On the other hand

$$\Pr\left[\, G_4^{\mathsf{A}} \text{ doesn't set } \mathsf{bad} \,\right] \;=\; \delta^{q_{\mathrm{S}}} \cdot (1-\delta) \,. \qquad (13)$$

We now select $\delta \in [0,1]$ to maximize the function $f(\delta) = \delta^{q_{\mathrm{S}}}(1-\delta)$, which yields $\delta = 1 - 1/(q_{\mathrm{S}}+1)$ and we have

$$\delta^{q_{\mathrm{S}}} \cdot (1-\delta) \;=\; \left(1 - \frac{1}{q_{\mathrm{S}}+1}\right)^{q_{\mathrm{S}}} \cdot \frac{1}{q_{\mathrm{S}}+1} \;>\; \frac{1}{e(q_{\mathrm{S}}+1)} \,. \qquad (14)$$

Putting together (7), (8), (9), (10), (11), (12), (13), and (14), we get (6). The running time of B is that of A plus an overhead of $(q_{\mathrm{H}}^* + 1) \cdot O(T_\Pi)$ due to line 112. ▮

INSTANTIATING $\mathcal{SAS}$-*1* WITH RSA. RSA does not directly give rise to a family of certified trapdoor permutations. (Each instance comes with its own modulus, so that different instances are over different groups. Also it is not clear in general how to test that a given number is a valid RSA exponent relative to a given modulus.) An RSA-based instantiation of $\mathcal{SAS}$-*1* can however be obtained by using the family of certified claw-free trapdoor permutations from [14]. This comes at the cost of efficiency, however, as a permutation evaluation now takes two RSA exponentiations. Another option is to use techniques in [18], which either require the signers to be ordered by increasing moduli, or add one bit to the signature for each signer. (However, we note that requiring that the signers be ordered according to their moduli excludes some pattern in which the repeats can occur, for example, non-trivial loops, namely a loop containing at least two distinct signers.) To ensure the certification, they suggest that the encryption exponent e can be chosen to be a prime larger than N. Alternatively, one can use an arbitrary encryption exponent e such that $\gcd(e, \varphi(N)) = 1$ at the cost of longer public keys as proposed in [10]. As the group operation, one must use addition (modulo $2^{|N|}$ when using [14] or modulo N when using the techniques of [18]) since multiplication is only a group operation over $\mathbb{Z}_N^*$.

5 Tightening the security reductions

We present variants of $\mathcal{AS}$-*3* and $\mathcal{SAS}$-*1* where the reductions in the security proofs are tight. These schemes can provide the same security as the previous ones in smaller groups, and thus end up being more efficient.

5.1 Tight security for GAS

The security reduction for the $\mathcal{BLS}$ standard signature scheme given in [9] (see our Lemma 3.1) is not tight. Since our security result for $\mathcal{AS}$-*3* is based on Lemma 3.1, we inherit this loss in security. The

Alg $\mathsf{Sign}((x, X), M)$
If $\mathrm{ST}[M] = \perp$ then $b \stackrel{\$}{\leftarrow} \{0,1\}$; $\sigma \leftarrow H(b\|X\|m)^x$; $\mathrm{ST}[M] \leftarrow b\|\sigma$
Return $\mathrm{ST}[M]$

Alg $\mathsf{Agg}((X_1, m_1, \sigma_1), \ldots, (X_n, m_n, \sigma_n))$
For $i = 1$ to n do Parse σ_i as $b_i\|\sigma'_i$
$\sigma \leftarrow \prod_{i=1}^{n} \sigma'_i$
Return $(b_1, \ldots, b_n, \sigma)$

Alg $\mathsf{AVf}((X_1, m_1), \ldots, (X_n, m_n), \sigma')$
Parse σ' as $(b_1, \ldots, b_n, \sigma)$
If $\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(H(b_i\|X_i\|m_i), X_i)$ then return 1 else return 0

Figure 7: The signing, aggregation, and verification algorithms for the $\mathcal{AS}$-*4* GAS scheme.

security reductions for $\mathcal{AS}$-*1* and $\mathcal{AS}$-*2* from [7] are not tight either. Here, we will present $\mathcal{AS}$-*4*, which has about the same efficiency as $\mathcal{AS}$-*3*, but is proven secure via a tight reduction to coCDH. $\mathcal{AS}$-*4* is based on a variant of the Katz-Wang variant of the $\mathcal{BLS}$ scheme. Below, we use the notation and definitions of Section 3. In particular, we fix pairing parameters $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e}, g$, and ψ.

Ideas. Let us first recall the Katz-Wang variant [16] of the $\mathcal{BLS}$ scheme, which has a tight reduction to coCDH. The signer has the same keys as in $\mathcal{AS}$-*2* and $\mathcal{AS}$-*3*, namely secret key (x, X), where $x \in \mathbb{Z}_p$ and $X = g^x \in \mathbb{G}_2$, and public key X. To sign a message m, the signer picks a random bit b and returns $\sigma = \mathrm{H}(b\|m)^x$, where H: $\{0,1\}^* \to \mathbb{G}_1$ is a random oracle. (It is important that if m is re-signed, the signer not generate a fresh bit b, but reuse the one generated the first time this message was signed. One way to do this is to maintain state in the form of a table of message-signature pairs. Another way is not to pick b at random, but instead specify it as a PRF applied to the message, where the key for the PRF is part of the signer's secret key.) On input public key X, message m, and candidate signature σ, the verifier accepts if $\mathbf{e}(\sigma, g) = \mathbf{e}(\mathrm{H}(0\|m), X)$ or $\mathbf{e}(\sigma, g) = \mathbf{e}(\mathrm{H}(1\|m), X)$, and rejects otherwise. However, suppose now that we aggregate these signatures. Given an aggregate of n signatures, the verifier will need to try each value of the bit for each constituent signature, leading to a verification algorithm that takes 2^n steps. To solve this problem, we slightly modify the base scheme. Rather than having the verifier perform the two checks for $0\|m$ and $1\|m$, we instead output the bit b as part of the signature and make the verifier accept iff $\mathbf{e}(\sigma, g) = \mathbf{e}(\mathrm{H}(b\|m), X)$. We now present an unrestricted GAS scheme with this as the base scheme.

The $\mathcal{AS}$-*4* scheme. Let $\mathcal{AS}\text{-}4 = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Agg}, \mathsf{AVf})$ be the GAS scheme defined as follows. The key generation algorithm Kg picks x at random from $\mathbb{Z}_p$, sets $X = g^x$, returns (x, X) as the secret key and X as the public key. Let H: $\{0,1\}^* \to \mathbb{G}_1$ be a random oracle. The other three algorithms are shown in Figure 7. We have presented the signing algorithm as stateful, but it can be rendered stateless as discussed above. Namely, rather than having it pick the bit b at random, it can compute it by applying a PRF to the message, where the key for the PRF is part of the secret signing key.

Notice that the length of an aggregate signature in this scheme is not equal to that of a regular signature, but instead grows by one bit for each signature added to the aggregate. From an asymptotic point of view, this might be considered terrible since the aggregate grows linearly in the number of signatures. But in fact, only a single bit is being added per signature, and in practice, the effect of this may be negligible. For example, a typical certificate chain is unlikely to contain more than 20 certificates. In this case, $\mathcal{AS}$-*4* adds only 20 bits compared to $\mathcal{AS}$-*3*. Meanwhile, the costs for signing, aggregation, and verification are the same as for $\mathcal{AS}$-*3*. But, for a given desired level of provable security, the tight reduction for $\mathcal{AS}$-*4* means it can use smaller groups than $\mathcal{AS}$-*3*. This means that it is actually more efficient. Furthermore, a smaller group means a smaller base signature, so the extra bits may even be for free. Precise estimates are difficult because we lack precise information about the size of groups needed for different levels of security [12].

Security of $\mathcal{AS}$-*4*. Our proof for $\mathcal{AS}$-*3* took a modular approach, reducing security to that of the underlying base signature scheme. Unfortunately, that does not seem to work here. Instead, we reduce directly to the coCDH problem. The result is as follows.

Theorem 5.1 *If the coCDH problem is (t', ϵ')-hard, then the $\mathcal{AS}$-4 scheme is $(t, q_S, n_{max}, q_H, \epsilon)$-secure for any $t, q_S, n_{max}, q_H, \epsilon$ satisfying*

$$\epsilon \ \geq\ 2 \cdot \epsilon' \ \text{ and } \ t \ \leq\ t' - t_{exp}(2q_H + 3q_S + 3n_{max} + 2) - O(q_S + n_{max}) \ . \ \blacksquare$$

Proof of Theorem 5.1 We say that an adversary A against $\mathcal{AS}$-4 is *simplified* if it has the following properties. Below, $X^* = g^x$ so that the secret key is (x, X^*) and X^* is the public key input to A:

1. If $((X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma))$ is the forgery output by A, then it is always the case that $n \geq 1$ and $X_1 = X^*$ and m_1 was not queried to $\mathsf{Sign}_{\mathcal{AS}\text{-}4}((x, X^*), \cdot)$.

2. If A makes a sign query m, then there is a bit $b \in \{0, 1\}$ such that it previously made a hash query $b \| X^* \| m$.

3. If A outputs forgery $((X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma))$, then it has previously made hash queries $b_i \| X_i \| m_i$ for all $1 \leq i \leq n$.

There are two stages in our proof. First, we transform a given adversary A′ attacking $\mathcal{AS}$-4 into a simplified adversary A attacking $\mathcal{AS}$-4 without loss in advantage although at some cost in resources:

Lemma 5.2 *Let A′ be an adversary that $(t, q_S, n_{max}, q_H, \epsilon)$-breaks $\mathcal{AS}$-4. Then, we can construct a simplified adversary A that $(t^*, q_S, n_{max} + 1, q_H^*, \epsilon)$-breaks $\mathcal{AS}$-4 where*

$$t^* = t + O(q_S + n_{max}) \ , \ \ q_H^* = q_H + q_S + n_{max} + 1 \ . \ \blacksquare$$

Next, we construct an adversary B which enlists A's help in solving the coCDH problem:

Lemma 5.3 *Let A be a simplified adversary that $(t^*, q_S, n^*_{max}, q_H^*, \epsilon)$-breaks $\mathcal{AS}$-4. Then, we can construct an adversary B solving the coCDH problem such that*

$$\mathbf{Adv}^{\text{co-cdh}}(\mathsf{B}) \ \geq\ \frac{\mathbf{Adv}^{\text{agg-uf}}_{\mathcal{AS}\text{-}4}(\mathsf{A})}{2} \ , \tag{15}$$

and the running time of B is at most that of A plus $t_{exp}(2q_H^* + q_S + n^*_{max} + 1)$. $\blacksquare$

Theorem 5.1 follows directly from these two lemmas. We now proceed to prove the lemmas. A convention made in writing code is that all array elements are assumed to initially be $\perp$.

Proof of Lemma 5.2: The adversary A is depicted in Figure 8. In the case that $F = \emptyset$, adversary A′ will not win anyway, so the addition made by A to the forgery does not decrease the advantage but ensures that the forgery contains b_j, X_j, m_j such that $X_j = X^*$ but m_j was not a sign query. The components of the forgery are then reordered to bring the j-th component into the first place. Since the groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are cyclic, they are also commutative, so this does not decrease the probability that the forgery is valid. $\blacksquare$

Proof of Lemma 5.3: We consider the games of Figure 9. (See Section 4 for background on game playing.) Adversary B works as follows. On inputs g, X^*, h, it runs A on input X^*. It answers A's oracle queries via subroutines H-sim and Sign-sim, the code for these being exactly that of the corresponding procedures of game G_0. When A halts with forgery $((X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma))$, adversary B executes lines 030–036 of the **Finalize** procedure of G_0 and outputs γ.

Before proceeding to the analysis, let us provide some intuition.

If A makes a hash query $b \| X \| m$ with $X \neq X^*$ then B answers in such a way that it knows the signature corresponding to this query. Specifically, it picks a random y and answers $\psi(g)^y$. The signature is $\psi(X)^y$.

If A makes a hash query $b \| X \| m$ with $X = X^*$, then B picks a random bit $\mathrm{B}[X^* \| m]$. If $\mathrm{B}[X^* \| m] = b$, then it answer as in the previous case. Otherwise, it embeds h into the answer. Specifically, it picks a

Subroutine $\textsc{Sign-sim}(m)$
$\mathrm{H}_{\mathcal{AS}\text{-}4}(0\|X^*\|m)$; $S \leftarrow S \cup \{m\}$; Return $\mathsf{Sign}_{\mathcal{AS}\text{-}4}((x, X^*), m)$

Adversary $\mathsf{A}(X^*)$
Run A' on X^* answering its hash queries using $\mathrm{H}_{\mathcal{AS}\text{-}4}$ and sign queries using $\textsc{Sign-sim}$
Until A' halts and outputs a forgery $(X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma)$
$F \leftarrow \{\, i \mid X_i = X^* \text{ and } m_i \notin S \,\}$
If $F = \emptyset$ then
  $b_{n+1} \leftarrow 0$; $X_{n+1} \leftarrow X^*$; $s \leftarrow n+1$; $j \leftarrow n+1$
  Let m_{n+1} be a message not in S
Else $s \leftarrow n$; $j \stackrel{\$}{\leftarrow} F$
For $i = 1, \ldots, s$ do $\mathrm{H}_{\mathcal{AS}\text{-}4}(b_i\|X_i\|m_i)$
$\sigma' \leftarrow (b_j, b_1, \ldots, b_{j-1}, b_{j+1}, \ldots, b_s, \sigma)$
Return $(X_j, m_j), (X_1, m_1), \ldots, (X_{j-1}, m_{j-1}), (X_{j+1}, m_{j+1}), \ldots, (X_s, m_s), \sigma'$

Figure 8: Adversary A for the proof of Lemma 5.2. The adversary has oracles $\mathrm{H}_{\mathcal{AS}\text{-}4}$ and $\mathsf{Sign}_{\mathcal{AS}\text{-}4}((x, X^*), \cdot)$. Above, S is initialized to empty.

---

random y and answers $\psi(g)^y h$. The first time there is a hash query of the form $b\|X^*\|m$, the bit $\mathrm{B}[X^*\|m]$ and the values $\mathrm{H}(0\|X^*\|m), \mathrm{H}(1\|X^*\|m)$ are defined so that subsequent queries do not redefine $\mathrm{B}[X^*\|m]$.

If A makes a sign query m, then B looks up the value y corresponding to the hash query $0\|X^*\|m$ and replies with $0\|\psi(X^*)^y$.

Once A outputs a forgery $((X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma))$, adversary B divides out all the signatures that it knows, i.e. the ones for which $X_i \neq X^*$ and for which $X_i = X^*$ and $b_i = \mathrm{B}[X^*\|m_i]$. For each of the rest, i.e. the ones for which $X_i = X^*$ but $b_i \neq \mathrm{B}[X^*\|m_i]$, it observes that the signature is of the form $\psi(g)^{xy_i}h^x = \psi(X^*)^{y_i}h^x$ where y_i is the random value corresponding to the hash query $b_i\|X^*\|m_i$ and x is the discrete logarithm of X^*, which it does not know. To compute h^x, the adversary B collects into τ the signatures it knows along with the blinding factors $\psi(X^*)^{y_i}$ in the hash values $\mathrm{H}(\mathrm{B}[X^*\|m_i]\|X^*\|m_i)$. Then, $\sigma/\tau = h^{xs}$, so h^x can be obtained as long as $s \bmod p \neq 0$.

The analysis needs to show that $s \bmod p \neq 0$ with probability at least $1/2$, and moreover, that this is independent of the success of A. It is to argue this clearly and rigorously that we use games.

For the analysis, let $\textsc{good}_i$ be the event that the flag bad is not set in game G_i, for $i = 0, 1$. We claim that

$$\mathbf{Adv}^{\text{co-cdh}}(\mathsf{B}) \geq \Pr\left[\, G_0^{\mathsf{A}} \Rightarrow \mathtt{true} \,\wedge\, \textsc{good} \,\right] . \tag{16}$$

We justify this as follows. The advantage of B is exactly the probability that $\gamma = h^x$ in game G_0. We now argue that, if A's forgery is valid and bad is not set then $\gamma = h^x$. The following chain of equalities is justified below:

$$\begin{aligned}
\mathbf{e}(\sigma, g) &= \prod_{i=1}^{n} \mathbf{e}(\mathrm{HT}[b_i\|X_i\|m_i], X_i) && (17)\\
&= \prod_{i\in I} \mathbf{e}(\mathrm{HT}[b_i\|X^*\|m_i], X^*) \cdot \prod_{i\in J} \mathbf{e}(\mathrm{HT}[b_i\|X_i\|m_i], X_i) && (18)\\
&= \prod_{i\in I} \mathbf{e}(\psi(g)^{y_i}h, X^*) \cdot \prod_{i\in J} \mathbf{e}(\psi(g)^{y_i}, X_i) && (19)\\
&= \prod_{i\in I} \mathbf{e}(h, X^*) \cdot \prod_{i=1}^{n} \mathbf{e}(\psi(g)^{y_i}, X_i) && (20)
\end{aligned}$$

Initialize Game G_0

000 $x \xleftarrow{\$} \mathbb{Z}_p^*$; $X^* \leftarrow g^x$; $h \xleftarrow{\$} \mathbb{G}_1$
001 Return g, X^*, h

On H**-query** $b\|X\|m$

010 If $\mathrm{HT}[0\|X\|m] = \perp$ then
011 $\quad c \xleftarrow{\$} \{0,1\}$; $\mathrm{B}[X\|m] \leftarrow c$
012 $\quad \mathrm{y}[0\|X\|m] \xleftarrow{\$} \mathbb{Z}_p$; $\mathrm{y}[1\|X\|m] \xleftarrow{\$} \mathbb{Z}_p$
013 If $X = X^*$ then
014 $\quad \mathrm{HT}[c\|X^*\|m] \leftarrow \psi(g)^{\mathrm{y}[c\|X\|m]}$
015 $\quad \mathrm{HT}[1-c\|X^*\|m] \leftarrow \psi(g)^{\mathrm{y}[1-c\|X\|m]}h$
016 Else
017 $\quad \mathrm{HT}[0\|X\|m] \leftarrow \psi(g)^{\mathrm{y}[0\|X\|m]}$
018 $\quad \mathrm{HT}[1\|X\|m] \leftarrow \psi(g)^{\mathrm{y}[1\|X\|m]}$
019 Return $\mathrm{HT}[b\|X\|m]$

On $\mathsf{SASign}^{\mathrm{H}}$**-query** m

020 $b \leftarrow \mathrm{B}[X^*\|m]$
021 Return $b\|\psi(X^*)^{\mathrm{y}[b\|X^*\|m]}$

Finalize$((X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma))$

030 $I \leftarrow \{\, i \mid X_i = X^* \text{ and } b_i \neq \mathrm{B}[X_i\|m_i] \,\}$
031 $J \leftarrow \{1, \ldots, n\} \setminus I$
032 $\tau \leftarrow 1$; $s \leftarrow |I|$
033 For $i = 1, \ldots, n$ do
034 $\quad y_i \leftarrow \mathrm{y}[M_i]$; $\sigma_i \leftarrow \psi(X_i)^{y_i}$; $\tau \leftarrow \tau\sigma_i$
035 If $\mathrm{B}[X^*\|m_1] = b_1$ then $\mathtt{bad} \leftarrow \mathtt{true}$; $\gamma \leftarrow \perp$
036 Else $s' \leftarrow s^{-1} \bmod p$; $\gamma \leftarrow (\sigma\tau^{-1})^{s'}$
037 Return $\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(\mathrm{HT}[b_i\|X_i\|m_i], X_i)$

Initialize Game G_1

100 $x \xleftarrow{\$} \mathbb{Z}_p^*$; $X^* \leftarrow g^x$; $h \xleftarrow{\$} \mathbb{G}_1$
101 Return g, X^*, h

On H**-query** $b\|X\|m$

110 If $\mathrm{HT}[0\|X\|m] = \perp$ then
111 $\quad \mathrm{HT}[0\|X\|m] \xleftarrow{\$} \mathbb{G}_1$; $\mathrm{HT}[1\|X\|m] \xleftarrow{\$} \mathbb{G}_1$
112 Return $\mathrm{HT}[b\|X\|m]$

On $\mathsf{Sign}^{\mathrm{H}}$**-query** m

120 If $\mathrm{B}[X^*\|m] = \perp$ then
121 $\quad \mathrm{B}[X^*\|m] \xleftarrow{\$} \{0,1\}$
122 $b \leftarrow \mathrm{B}[X^*\|m]$
123 Return $b\|\mathrm{HT}[b\|X^*\|m]^x$

Finalize$((X_1, m_1), \ldots, (X_n, m_n), (b_1, \ldots, b_n, \sigma))$

130 $\mathrm{B}[X^*\|m_1] \xleftarrow{\$} \{0,1\}$
131 If $\mathrm{B}[X^*\|m_1] = b_1$ then $\mathtt{bad} \leftarrow \mathtt{true}$
132 Return $\mathbf{e}(\sigma, g) = \prod_{i=1}^{n} \mathbf{e}(\mathrm{HT}[b_i\|X_i\|m_i], X_i)$

Figure 9: Games G_0, G_1 for the proof of Lemma 5.3.

$$= \mathbf{e}(h^s, X^*) \cdot \prod_{i=1}^{n} \mathbf{e}(\psi(X_i)^{y_i}, g) \tag{21}$$

$$= \mathbf{e}(h^s, X^*) \cdot \mathbf{e}(\tau, g) \;. \tag{22}$$

Equation (17) is from the verification algorithm of the scheme and the assumption that A's forgery is valid. Equation (18) is true because $X_i = X^*$ for $i \in I$ and the sets I, J partition $\{1, \ldots, n\}$. Equation (19) is due to the way G_0 responds to hash queries. Specifically, for $i \in I$ we have $\mathrm{HT}[b_i\|X^*\|m_i] = \psi(g)^{y_i}h$ because $b_i \neq \mathrm{B}[X^*\|m_i]$, while for $i \in J$ we have $\mathrm{HT}[b_i\|X_i\|m_i] = \psi(g)^{y_i}$ both when $X_i = X^*$ and when $X_i \neq X^*$. Equation (20) uses the bilinearity of $\mathbf{e}$. Equation (21) uses two things: that $s = |I|$, and that $\mathbf{e}(\psi(A), B) = \mathbf{e}(\psi(B), A)$ for all $A, B \in \mathbb{G}_2$. Equation (22) is true because $\tau = \prod_{i=1}^{n} \psi(X_i)^{y_i}$ due to line 034 of G_0. Now, from (22), we have

$$\mathbf{e}(\sigma\tau^{-1}, g) \;=\; \mathbf{e}(h^s, X^*) \;=\; \mathbf{e}(h^s, g^x) \;=\; \mathbf{e}(h^{xs}, g) \;.$$

So it must be that $h^{xs} = \sigma\tau^{-1}$. We now claim that if $b_1 \neq \mathrm{B}[X^*\|m_1]$ —meaning $\mathtt{bad}$ is $\mathtt{false}$— then $s \in \mathbb{Z}_p^*$. This means s has an inverse modulo p, and hence the value γ computed at line 036 of G_0 equals h^x. We now justify the above claim. By property **1** of the simplified adversary A, we know that m_1 was not a sign query. So if $b_1 \neq \mathrm{B}[X^*\|m_1]$, then $1 \in I$, meaning $I \neq \emptyset$ and $s > 0$. However, we have written the code of B with the assumption that $s < p$. (We may assume without loss of generality that $n < p$, because otherwise B can exhaustively search for x in time $O(nt_{\mathrm{exp}})$ and compute h^x.) Thus, $s \in \mathbb{Z}_p^*$ as claimed. We have just argued that if A's forgery is valid (meaning the output of G_0 is $\mathtt{true}$) and

$b_1 \neq \mathrm{B}[X^* \| m_1]$ (meaning GOOD holds) then $\gamma = h^x$. That is, we have justified Equation (16).

We now claim that

$$\Pr\left[\, G_0^{\mathsf{A}} \Rightarrow \texttt{true} \;\wedge\; \mathrm{GOOD}_0 \,\right] \;=\; \Pr\left[\, G_1^{\mathsf{A}} \Rightarrow \texttt{true} \;\wedge\; \mathrm{GOOD}_1 \,\right] \tag{23}$$

$$\;=\; \Pr\left[\, G_1^{\mathsf{A}} \Rightarrow \texttt{true} \,\right] \cdot \Pr\left[\, \mathrm{GOOD}_1 \,\right] . \tag{24}$$

Game G_0 answers hash oracle queries correctly, so these answers are distributed just as those of G_1. The same is true for sign queries. However, G_1 only defines $\mathrm{B}[X^* \| m]$ when m is a sign query. Property **1** says that m_1 was not a sign query, so the sign procedure does not pick $\mathrm{B}[X^* \| m_1]$. Thus, its choice is delayed until the **Finalize** procedure. Thus, Equation (23) is true. Now, it is clear that in G_1 the game output is determined before $\mathrm{B}[X^* \| m_1]$ is chosen, so that the events GOOD_1 and "$G_1^{\mathsf{A}} \Rightarrow \texttt{true}$" are independent. This justifies Equation (24). However,

$$\Pr\left[\, G_1^{\mathsf{A}} \Rightarrow \texttt{true} \,\right] = \mathbf{Adv}^{\mathrm{agg\text{-}uf}}_{\mathcal{AS}\text{-}4}(\mathsf{A}) \quad \text{and} \quad \Pr\left[\, \mathrm{GOOD}_1 \,\right] = \frac{1}{2} . \tag{25}$$

Combining Equations (16), (24) and (25) yields Equation (15).

For the running time analysis, recall that B runs A answering A's hash and sign queries using the code of the corresponding procedures of G_0 and then, once A outputs a forgery, B executes lines 030–036. Thus, in the worst case, each hash-oracle query incurs two exponentiations from lines 013–018 (thus, the $2q_{\mathrm{H}}^*$ term), and each sign-oracle query incurs one exponentiation from line 021 (thus, the q_{S} term). Additionally, each loop iteration on lines 033–034 incurs one exponentiation (thus, the $n^*_{\max}$ term), and line 036 incurs one exponentiation. Lemma 5.3 follows. ▌

5.2 A variant of $\mathcal{SAS}$-*1* with a tight reduction

The security result of $\mathcal{SAS}$-*1* can be improved. Similar to the case of GAS schemes, we can apply the Katz-Wang technique of [16] to make the reduction tight.

Ideas. We use ideas similar to the design of the $\mathcal{AS}$-*4* GAS scheme. Specifically, the key generation of the new scheme is the same as that of $\mathcal{SAS}$-*1*. To generate a sequential aggregate on $(\pi, \pi^{-1}), m, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n)$, the signer picks a random bit b and returns $\mathrm{H}(b \| (\pi_1, m_1) \| \ldots \| (\pi_n, m_n) \| \pi \| m)$ where $\mathrm{H}\colon \{0,1\}^* \to \mathbb{G}$ is a random oracle. Similar to $\mathcal{AS}$-*4*, to avoid having the verification algorithm try each value of the bit for each constituent message, which would take 2^n steps for a sequential aggregate signature comprising n permutation-message pairs, we output the bit b as part of the signature. This increase in the signature length is small. See Section 5.1 for discussions on this issue.

The $\mathcal{SAS}$-*2* scheme. Let $\mathcal{SAS}\text{-}2 = (\mathsf{Kg}, \mathsf{SASign}, \mathsf{SAVf})$ be the SAS scheme defined as follows. The key generation algorithm has each signer generate a key pair (pk, sk) by computing $(\pi, \overline{\pi}, \pi^{-1}) \xleftarrow{\$} \mathsf{Gen}\,;\, pk \leftarrow \pi\,;\, sk \leftarrow (\pi, \pi^{-1})$. Let $\mathrm{H} : \{0,1\}^* \to \mathbb{G}$ be a random oracle. The sequential aggregation and verification algorithms are described in Figure 10. We have presented the signing algorithm as stateful, but it can be rendered stateless as discussed in Section 5.1. Namely, rather than having it pick the bit b at random, it can compute it by applying a PRF to the message, where the key for the PRF is part of the secret signing key.

Security of $\mathcal{SAS}$-*2*. The following theorem states that $\mathcal{SAS}$-*2* scheme is secure in the random oracle model assuming that the underlying family of permutations is claw-free.

Theorem 5.4 *If the family of certified claw-free trapdoor permutations Π is (t', ϵ')-claw-free, then the $\mathcal{SAS}$-2 sequential aggregate signature scheme is $(t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon)$-secure for any $t, q_{\mathrm{S}}, n_{\max}, q_{\mathrm{H}}, \epsilon$ satisfying*

$$\epsilon \;\geq\; 2 \cdot \epsilon' \;\; \text{and} \;\; t \;\leq\; t' - n_{\max} \cdot (q_{\mathrm{H}} + 2q_{\mathrm{S}} + 1) \cdot O(T_{\Pi}) . \;▌$$

Proof of Theorem 5.4: We do not give a full proof here, but rather sketch the differences with the proof of Theorem 4.1. The new algorithm B works largely in the same way. However, in its simulation of A in Game G_0, it maintains an extra table $\mathrm{B}[\cdot]$ in which it associates a random bit to each Q_n that is

Algorithm $\mathsf{SASign}((\pi, \pi^{-1}), m, \sigma_n, (\pi_1, m_1), \ldots, (\pi_n, m_n))$: If $n = 0$ then $s_0 \leftarrow 1$ Else $Q_n \leftarrow (\pi_1, m_1) \| \ldots \| (\pi_n, m_n)$ If $\mathsf{SAVf}(Q_n, \sigma_n) \neq 1$ then return $\perp$ If $\mathrm{ST}[m \| \sigma_n \| Q_n] \neq \perp$, then return $\mathrm{ST}[m \| \sigma_n \| Q_n]$ Parse σ_n as $(b_1, \ldots, b_n, s_n)$ $b_{n+1} \xleftarrow{\$} \{0,1\}$; $h \leftarrow \mathrm{H}(b_{n+1} \| Q_n \| (\pi, m))$ $s_{n+1} \leftarrow \pi^{-1}(h \cdot s_n)$; $\sigma_{n+1} \leftarrow (b_1, \ldots, b_{n+1}, s_{n+1})$ $\mathrm{ST}[m \| \sigma_n \| Q_n] \leftarrow \sigma_{n+1}$ Return σ_{n+1}	Algorithm $\mathsf{SAVf}((\pi_1, m_1), \ldots, (\pi_n, m_n), \sigma)$: Parse σ as $(b_1, \ldots, b_n, s_n)$ For $i = n, \ldots, 1$ do If $\mathsf{Test}(\pi_i) = 0$ then return 0 $h_i \leftarrow \mathrm{H}(b_i \| (\pi_1, m_1) \| \ldots \| (\pi_i, m_i))$ $s_{i-1} \leftarrow \pi_i(s_i) \cdot h_i^{-1}$ If $s_0 = 1$ then return 1 else return 0 ▌

Figure 10: The aggregation and verification algorithms for the $\mathcal{SAS}$-2 SAS scheme.

part of a $H(b \| Q_n \| (\pi, M))$ query. Also, the table HT now uses $b \| Q_n$ as the key rather than simply Q_n. On a hash query $b \| Q_n \| (\pi, M)$ where $Q_n = \pi_1 \| m_1 \| \cdots \| \pi_n \| m_n$, if $\pi_n \neq \pi$ (recall that π is the public key given as input to A), then B computes $\mathrm{HT}[b \| Q_n] \leftarrow \pi_n(\sigma[Q_n]) \cdot \sigma[Q_{n-1}]^{-1}$ as before. If $\pi_n = \pi$ and $b = \mathrm{B}[Q_n]$, then B computes $\mathrm{HT}[b \| Q_n] \leftarrow \pi_n(\sigma[Q_n]) \cdot \sigma[Q_{n-1}]^{-1}$ as in the case for $\pi_n \neq \pi$. If $\pi_n = \pi$ and $b \neq \mathrm{B}[Q_n]$, however, B computes $\mathrm{HT}[b \| Q_n]$ as $\overline{\pi}(\sigma[Q_n]) \cdot \sigma[Q_{n-1}]^{-1}$. ▌

Acknowledgments

We thank the ICALP 2007 anonymous referees for their valuable comments.

References

[1] M. Bellare, A. Boldyreva, and J. Staddon. Randomness re-use in multi-recipient encryption schemes. In Y. Desmedt, editor, *PKC 2003: 6th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 85–99, Miami, USA, Jan. 6–8, 2003. Springer-Verlag, Berlin, Germany.

[2] M. Bellare and O. Goldreich. On defining proofs of knowledge. In E. F. Brickell, editor, *Advances in Cryptology – CRYPTO'92*, volume 740 of *Lecture Notes in Computer Science*, pages 390–420, Santa Barbara, CA, USA, Aug. 16–20, 1992. Springer-Verlag, Berlin, Germany.

[3] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, Nov. 3–5, 1993. ACM Press.

[4] M. Bellare and P. Rogaway. Code-based game-playing proofs and the security of triple encryption. In S. Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, St.-Petersburg, Russia, May 29 – June 1, 2006. Springer-Verlag, Berlin, Germany. Available as Cryptology ePrint Report 2005/334.

[5] M. Bellare and M. Yung. Certifying permutations: Noninteractive zero-knowledge based on any trapdoor permutation. *Journal of Cryptology*, 9(3):149–166, 1996.

[6] A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Y. Desmedt, editor, *PKC 2003: 6th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 31–46, Miami, USA, Jan. 6–8, 2003. Springer-Verlag, Berlin, Germany.

[7] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In E. Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 416–432, Warsaw, Poland, May 4–8, 2003. Springer-Verlag, Berlin, Germany.

[8] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. A survey of two signature aggregation techniques. *RSA's CryptoBytes*, 6(2), Summer 2003.

[9] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In C. Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 514–532, Gold Coast, Australia, Dec. 9–13, 2001. Springer-Verlag, Berlin, Germany.

[10] D. Catalano, D. Pointcheval, and T. Pornin. Trapdoor hard-to-invert group isomorphisms and their application to password-based authentication. *Journal of Cryptology*, 20(1):115–149, 2006.

[11] J.-S. Coron. On the exact security of full domain hash. In M. Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235, Santa Barbara, CA, USA, Aug. 20–24, 2000. Springer-Verlag, Berlin, Germany.

[12] S. Galbraith, K. Paterson, and N. Smart. Pairings for cryptographers. Cryptology ePrint Archive, Report 2006/165, 2006. `http://eprint.iacr.org/`.

[13] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, Apr. 1988.

[14] R. Hayashi, T. Okamoto, and K. Tanaka. An RSA family of trap-door permutations with a common domain and its applications. In F. Bao, R. Deng, and J. Zhou, editors, *PKC 2004: 7th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 291–304, Singapore, Mar. 1–4, 2004. Springer-Verlag, Berlin, Germany.

[15] Intergovernmental Oceanographic Commission of UNESCO. Towards the establishment of a tsunami warning and mitigation system for the Indian Ocean. Available at `http://ioc3.unesco.org/indotsunami/`. Last accessed April 13, 2007.

[16] J. Katz and N. Wáng. Efficiency improvements for signature schemes with tight security reductions. In *ACM CCS 03: 10th Conference on Computer and Communications Security*, pages 155–164, Washington D.C., USA, Oct. 27–30, 2003. ACM Press.

[17] S. Lu, R. Ostrovsky, A. Sahai, H. Shacham, and B. Waters. Sequential aggregate signatures and multisignatures without random oracles. In S. Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, St.-Petersburg, Russia, May 29 – June 1, 2006. Springer-Verlag, Berlin, Germany.

[18] A. Lysyanskaya, S. Micali, L. Reyzin, and H. Shacham. Sequential aggregate signatures from trapdoor permutations. In C. Cachin and J. Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 74–90, Interlaken, Switzerland, May 2–6, 2004. Springer-Verlag, Berlin, Germany.

[19] H. Shacham. *New Paradigms in Signature Schemes*. PhD thesis, Stanford University, 2005.