# Chain Reductions for Multi-Signatures and the HBMS Scheme

Mihir Bellare and Wei Dai

Department of Computer Science and Engineering
University of California San Diego, USA
{mihir,wdai}@eng.ucsd.edu.

**Abstract.** Existing proofs for existing Discrete Log (DL) based multi-signature schemes give only weak guarantees if the schemes are implemented, as they are in practice, in 256-bit groups. This is because the underlying reductions, which are mostly in the standard model and from DL, are loose. We show that relaxing either the model or the assumption suffices to obtain tight reductions. Namely we give (1) tight proofs from DL in the Algebraic Group Model, and (2) tight, standard-model proofs from well-founded assumptions other than DL. We first do this for the classical 3-round schemes, namely BN and MuSig. Then we give a new 2-round multi-signature scheme, HBMS, as efficient as prior ones, for which we do the same. These multiple paths to security for a single scheme are made possible by a framework of chain reductions, in which a reduction is broken into a chain of sub-reductions involving intermediate problems. Overall our results improve the security guarantees for DL-based multi-signature schemes in the groups in which they are implemented in practice.

## 1 Introduction

Usage in cryptocurrencies has lead to interest in practical, Discrete-Log-based multi-signature schemes. Proposals exist, are efficient, and are supported by proofs, BUT, the bound on adversary advantage in the proofs is so loose that the proofs fail to support use of the schemes in the 256-bit groups in which they are implemented in practice. This leaves the security of in-practice schemes unclear.

We ask, is it possible to bridge this gap to give some valuable support, in the form of tight reductions, for in-practice schemes? As long as we stay in the current paradigm, namely standard-model proofs from DL, the answer is likely NO. To make progress, we need to be willing to change either the model or the assumption. We show that in fact changing either suffices. Our approach is to give, for any scheme, many different paths to security. In particular we give (1) tight reductions from DL in the Algebraic Group Model (AGM) [17], and (2) tight, standard-model reductions from well-founded assumptions other than DL. We obtain these results via a framework in which a reduction is "factored" into a chain of sub-reductions involving intermediate problems.

We implement this approach first with classical 3-round schemes, giving chain reductions yielding (1) and (2) above for the BN [7] and MuSig [25] schemes. Then, in the space of 2-round schemes, we give a new, efficient scheme, called HBMS, for which we do the same. We now look at all this in more detail.

BACKGROUND. A multi-signature $\sigma$ on a message $m$ can be thought of as affirming that "We, the members of this group, all, jointly, endorse $m$." The group is indicated by the vector $\mathbf{vk} = (\mathbf{vk}[1], \ldots, \mathbf{vk}[n])$ of individual public verification keys of its members, and can be dynamic, changing from one signature to another. Signing is done via an interactive protocol between group members; each member $i$ begins with its own public verification key $\mathbf{vk}[i]$, its matching private signing key $\boldsymbol{sk}[i]$, and the message $m$, and, at the end of the interaction, they output the multi-signature $\sigma$. The latter should be compact (of size independent of the size of the group), precluding the trivial solution in which $\sigma$ is a list of the individual signatures of the group members on $m$.

Following its suggestion in the 1980s [20], the primitive has seen much evolution [19, 22, 29, 26, 7]. Early schemes assumed all signers in the signing protocol picked their verification keys honestly. "Rogue-key attacks," in which a malicious signer picked its verification key as function of that of an honest signer, lead to an upgraded target, schemes that retain security even in the presence of adversarially-chosen verification keys. Towards this challenging end we first saw schemes either using interactive key-generation [26] or making the "knowledge of secret key" assumption [10, 23]. Finally, BN [7] gave an efficient, Schnorr-based scheme in the "plain public-key" model, where security was provided even in the face of maliciously-chosen verification keys, yet no more was assumed about these keys than their having certificates as per a standard PKI.

The BN model and definition have become the preferred target; it is the one used in the schemes we discuss next, and in our scheme as well. We denote the security goal as MS-UF. In Section 4 we define it via a game, and define the ms-uf advantage of an adversary as its probability of winning this game.

A NEW WAVE. Applications in blockchains and cryptocurrencies —see [11] for details— have fueled a resurgence of interest in multi-signatures. The desire here is MS-UF-secure, DL-based schemes that work over standard elliptic curves such as Secp256k1 or Curve25519. (Pairing-based schemes [11] are thus precluded.) The natural candidate is BN. But the new application arena has lead to a desire for the following further features, not possessed by BN: (1) Key aggregation. There should be a way to aggregate a set of verification keys into a single, short aggregate key, relative to which signatures are verified. (2) Two rounds. A signing protocol using only 2 rounds of interaction, as opposed to the 3 used by BN.

MuSig [25, 11] broke ground by adapting BN to add key aggregation. Now the effort moved to reducing the number of rounds. This proved challenging. Early proposals of two-round schemes —[2, 24, 35] as well as an early, two-round version of MuSig [25]— were broken by DEFKLNS [15]. To fill the gap, DEFKLNS gave a new two-round scheme, mBCJ. Other proposals followed: MuSig2 [27], MuSig-DN [28] and DWMS [1]. All these support key aggregation.

| | **Previous** | | **Ours** | |
|---|---|---|---|---|
| **Scheme** MS | $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p)$ | $p \approx 2^{256}$ | $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p)$ | $p \approx 2^{256}$ |
| BN [7] | $\sqrt{(q \cdot t^2)/p}$ | $2^{-8}$ | $t^2/p$ | $2^{-96}$ |
| MuSig [11, 25] | $\sqrt[4]{(q^3 \cdot t^2)/p}$ | 1 | $t^2/p$ | $2^{-96}$ |

**Fig. 1. Bounds on ms-uf advantage for the 3-round schemes BN and MuSig.** First we show prior bounds, then ours. In each case we first show the upper bound $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p)$ as a formula, where $t, q, q_s$ are, respectively the adversary running time, the number of its RO queries and the number of executions of the signing protocol, while prime $p$ is the size of the underlying group $\mathbb{G}$. We then show the evaluation with $t = q = 2^{80}$, $q_s = 2^{30}$ and $p \approx 2^{256}$, to capture security over 256-bit curves Secp256k1 or Curve25519.

All the schemes discussed here come with proofs of MS-UF security based on the hardness of the DL (Discrete Log) problem in the underlying group $\mathbb{G}$, up to variations in the model (standard or AGM [17]) or the type of DL problem (plain or OMDL [6]).

Current bounds. On being informed that a scheme has a proof of security based on the hardness of the DL problem in an underlying elliptic-curve group $\mathbb{G}$, the expectation of a practitioner is that the probability that a time $t$ attacker can violate MS-UF security is no more than the probability of successfully computing a discrete logarithm in $\mathbb{G}$, which, as per [34], is $t^2/p$, where $p$, a prime, is the size of $\mathbb{G}$. Concretely, with the 256-bit curves Secp256k1 or Curve25519 —$p \approx 2^{256}$— they would expect that a time $t \approx 2^{80}$ attacker has ms-uf advantage at most $2^{160-256} = 2^{-96}$.

But this expectation is only correct if the reduction in the proof is tight. Current proofs for DL-based multi-signature schemes are loose. With the 256-bit curves Secp256k1 or Curve25519, and for a $2^{80}$-time attacker, the proof of [7] for BN can preclude only a $2^{-8}$ ms-uf advantage, while the proof of [25, 11] for MuSig cannot even preclude a ms-uf advantage of 1, meaning there may be, per the proof, no security at all (cf. Figure 1). For 2-round schemes, the advantage precluded by current proofs is $2^{-16}$ in one case, and again just 1 for the others (cf. Figure 2). Overall, the proofs fail, by big margins, to support the parameter choices and expectations of practice.

Before continuing, let us expand on the above estimates. A proof of MS-UF security for a multi-signature scheme MS gives a formula $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p)$ that upper bounds the ms-uf advantage of an adversary as a function of its running time $t$, the number $q$ of its queries to the random oracle, and the number $q_s$ of executions of the signing protocol in the chosen-message attack in the ms-uf game. They are shown in Figures 1 and 2. We assume that $t \geq q \geq q_s$. To get these formulas, we first assume that the best attack against the DL problem is generic, so that a time $t$ attacker has success probability at most

| Scheme | Security | | Efficiency | |
|---|---|---|---|---|
| | $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p)$ | $p \approx 2^{256}$ | Sign | Vf |
| mBCJ [15] | $(q_s^3 \cdot q^2 \cdot t^2)/p$ | 1 | $T_2^{\text{me}} + T_3^{\text{me}}$ | $3T_2^{\text{me}}$ |
| MuSig-DN [28] | $\sqrt[4]{(q^3 \cdot t^2)/p}$ | 1 | NIZK | $T_2^{\text{me}}$ |
| MuSig2, $\nu \geq 4$ [27] | $\sqrt[4]{(q^3 \cdot t^2)/p}$ | 1 | $T_\nu^{\text{me}}$ | $T_2^{\text{me}}$ |
| MuSig2, $\nu = 2$ [27] | $(t^2 + q^3)/p$ | $2^{-16}$ | $T_2^{\text{me}}$ | $T_2^{\text{me}}$ |
| DWMS [1] | $t^2/p + q/\sqrt{p}$ | $2^{-48}$ | $T_2^{\text{me}} + T_{2N}^{\text{me}}$ | $T_2^{\text{me}}$ |
| HBMS | $t^2/p$ | $2^{-96}$ | $T_2^{\text{me}}$ | $T_3^{\text{me}}$ |

**Fig. 2. Bounds on ms-uf advantage for 2-round schemes.** First we show bounds for prior schemes, then the bounds for our new scheme HBMS. As before, we first show the upper bound formula $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p)$, where $t, q, q_s$ are, respectively the adversary running time, the number of its RO queries and the number of executions of the signing protocol, while prime $p$ is the size of the underlying group $\mathbb{G}$. We then show the evaluation with $t = q = 2^{80}$, $q_s = 2^{30}$ and $p \approx 2^{256}$, to capture security over 256-bit curves Secp256k1 or Curve25519. For MuSig2, results differ depending on a parameter $\nu$ of the scheme. We also show estimates of signing time (per signer) and verification time. Here $T_n^{\text{me}}$ is the time to compute one $n$-multi-exponentiation in $\mathbb{G}$. The "NIZK" for MuSig-DN indicates that signing requires computation and verification of a NIZKs, which is (much) more expensive then other operations shown.

$t^2/p$ [34]. Next, we use the concrete-security results, in theorems in the papers, that give reductions from the DL problem to the MS-UF security of their scheme. The square-roots in the formulas arise from uses of forking lemmas [31, 7, 2]; the fourth-roots from nested use. The bounds in our Figures are approximate, dropping negligible additive terms. The proofs on which the bounds of Figures 1 and 2 are based, are, for BN [7], MuSig [11, 25], mBCJ [15], MuSig-DN [28] and MuSig2 ($\nu \geq 4$) [27], in the standard model; and for MuSig2 ($\nu = 2$) [27], DWMS [1] and HBMS, in the AGM. See [4] for details.

Towards better bounds. Our thesis is that proofs should provide, not merely a qualitative guarantee, but one whose bounds quantitatively support parameter choices made in practice and the indications of cryptanalysis. Accordingly we want multi-signature schemes for which we can prove tight bounds on ms-uf advantage. How are we to reach this end? Impossibility results for Schnorr signatures [30, 21], on which the multi-signature schemes under consideration are based, indicate that a search for tight reductions that are both (1) in the standard model, and (2) from DL, is unlikely to succeed. We need to be flexible, and relax either (1) or (2). In fact we show that relaxing either suffices: We give (1) tight reductions from DL in the Algebraic Group Model (AGM) [17], and (2) tight, standard-model reductions from assumptions other than DL. Together, these provide valuable theoretical support for the use of practical multi-signature schemes in 256-bit groups.

AGM. The AGM considers a limited, but still large class of adversaries, called algebraic. When such an adversary queries a group element to an oracle, it provides also its representation in terms of prior group elements that the adversary has seen. Intuitively, the assumption is that the adversary "knows" how group elements it creates are represented. For elliptic curve groups, this appears to be a realistic assumption, and here the AGM captures natural and currently-known attack strategies.

When considering the merits of the AGM, an important one to keep in mind is that a proof in the AGM immediately implies a proof in the well-accepted Generic Group Model (GGM) of [34]. (So the AGM is only "better" than the GGM.) In more detail, a tight AGM reduction from DL to some problem X immediately yields a GGM bound on adversary advantage, for X, that matches the GGM bound for DL [17]. Thus, overall, tight AGM reductions provide a valuable guarantee. This is recognized by Fuchsbauer, Plouviez and Seurin [18] who use the AGM to give a tight reduction from DL to the UF security of the Schnorr signature scheme. Their result gives hope, realized here, that such reductions are possible for multi-signatures as well.

Chain reductions. We achieve the above ends, and more, as follows. For each multi-signature scheme $\mathsf{MS}$ we consider, we give a chain of reductions, starting from DL, that we depict as

$$\mathrm{DL} = \mathrm{P}_0 \to \mathrm{P}_1 \to \cdots \to \mathrm{P}_{m-1} \to \mathrm{P}_m = \mathsf{MS} \ ,$$

where $\mathrm{P}_1, \ldots, \mathrm{P}_{m-1}$ are intermediate computational problems. We refer to $m \geq 1$ as the length of the chain. For each step $\mathrm{P}_{i-1} \to \mathrm{P}_i$ we provide one of the following.

**1.** A tight, standard-model reduction. This is the ideal and done for as many steps as possible.

**2.** When **1.** is not possible, we give BOTH of the following:

**2.1** A tight AGM reduction, AND ALSO

**2.2** A non-tight standard-model reduction.

Since a tight standard-model reduction implies a tight AGM one, this yields a tight AGM reduction from DL to $\mathsf{MS}$, the first of our goals stated above. (A bit better, since some sub-reductions are standard-model.) For $i$ such that the chain $\mathrm{P}_i \to \cdots \to \mathsf{MS}$ consists only of tight standard-model reductions, we have a tight, standard model proof of $\mathsf{MS}$ from assumption $\mathrm{P}_i$, realizing our second goal, stated above, of tight standard-model reductions from assumptions other than DL. (Of course how interesting or valuable this is depends on the choice of $\mathrm{P}_i$, but as discussed below, we are able to make well-founded choices.)

Finally, something not yet mentioned, that follows from **1** and **2.2** of the chain reductions, is that we always have a standard model (even if non-tight) reduction DL $\to$ $\mathsf{MS}$. This means that, while adding tight AGM reductions that are valuable in practice, we are not lowering the theoretical or qualitative guarantees, these remaining as one would expect or desire.

$$\begin{tikzcd}
\mathrm{DL} \arrow[d, "1"'] & \\
\mathrm{IDL} \arrow[r, "2"] \arrow[d, "3"'] & \mathsf{BN} \\
\mathrm{XIDL} \arrow[r, "4"] \arrow[dr, "5"] & \mathsf{MuSig} \\
& \mathsf{HBMS}
\end{tikzcd}$$

| | Reduction | SM | AGM |
|---|---|---|---|
| 1 | DL → IDL | Th. 2 | **Th. 1** |
| 2 | IDL → BN | **Th. 5** | – |
| 3 | IDL → XIDL | Th. 4 | **Th. 3** |
| 4 | XIDL → MuSig | **Th. 7** | – |
| 5 | XIDL → HBMS | Th. 10 | **Th. 9** |

**Fig. 3. Chain reductions for multi-signatures.** SM stands for "Standard Model" and AGM for "Algebraic Group Model." An arrow P → Q means a reduction from P to Q; i.e. a proof that P implies Q. A boldface **Theorem Number** indicates the reduction is **tight**. A blank appears in the AGM column when a (tight) SM reduction to its left makes the AGM reduction unnecessary. Writing a MS scheme like BN, MuSig, HBMS as a point in a chain refers to MS-UF security of the scheme in question.

Chain reductions can be seen as a way to implement a modular proof framework in the style of [21], in which steps are reused across proofs for different schemes.

New bounds for classical schemes. We start by revisiting the classical 3-round schemes, namely BN and MuSig. Figure 3 illustrates our chains, that we now discuss.

IDL, formulated in [21] —they call it IDLOG, which we have abbreviated— is a purely group-based problem that is equivalent to the security against parallel impersonation under key-only attack (PIMP-KOA) of the Schnorr ID scheme. A tight GGM bound for IDL was shown by [21], but an AGM reduction DL → IDL does not seem to be in the literature; we fill this gap by providing it in Theorem 1. A (non-tight) standard model DL → IDL reduction is in [21], but we slightly improve it in Theorem 2.

Now our chain for BN is DL → IDL → BN. This chain has length 2. Our main result for BN is Theorem 5, which shows IDL → BN with a *tight, standard model* reduction. Putting this together with our above-mentioned tight DL → IDL AGM-reduction of Theorem 1, we get a tight DL → BN AGM-reduction. Also our tight, standard-model IDL → BN reduction says that BN is as secure as the Schnorr identification scheme, which is valuable in its own right since the latter has withstood cryptanalysis for many years.

We introduce an intermediate, purely group-based problem we call XIDL. We show IDL → XIDL with a tight AGM reduction (Theorem 3) and a (non-tight) standard-model reduction (Theorem 4).

Our chain for MuSig is DL → IDL → XIDL → MuSig. This chain has length 3. Our main result for MuSig is Theorem 7, which shows XIDL → MuSig with a *tight, standard model* reduction. Putting this together with the rest of the chain, we get a tight DL → MuSig AGM-reduction. If we are willing to view XIDL as an assumption extending IDL, we can also view MuSig as based tightly on that.

This means we show that $\mathbf{UB}_{\mathsf{MS}}^{\text{ms-uf}}(t, q, q_s, p) \leq t^2/p$ for both schemes, matching the DL bound. This is tight and optimal, since the multi-signature schemes can be broken by taking discrete-logs. Figure 1 compares our results with the prior ones.

New 2-round scheme. Turning to 2-round schemes, we give a new scheme, called HBMS. HBMS supports key aggregation, in line with other 2-round schemes. Our chain for our new 2-round HBMS scheme is DL → IDL → XIDL → HBMS. This chain has length 3. We show XIDL → HBMS with a tight AGM reduction (Theorem 9) and a (non-tight) standard-model reduction (Theorem 10). Putting this together with the rest of the chain, we get a tight DL → HBMS AGM-reduction, in particular showing $\mathbf{UB}_{\mathsf{MS}}^{\text{ms-uf}}(t, q, q_s, p) \leq t^2/p$, matching the DL bound. We also get a (non-tight) DL → HBMS standard-model-reduction.

Figure 2 compares HBMS with prior 2-round schemes. It shows that our improvement in security is not at the cost of efficiency. (Signing in HBMS is as efficient, or more so, than in prior schemes. For verification, MuSig-DN [28] is slightly faster, but signing in the latter is prohibitive due to the use of NIZKs.)

As the above shows, we reuse steps across different chains. Thus XIDL is an intermediate point for both MuSig and HBMS, and IDL for both BN and XIDL. This simplifies proofs and reduces effort. It also shows common elements and relations across schemes.

Equivalences. As discussed above, Theorem 5 shows IDL → BN with a tight, standard model reduction. We also give, in Theorem 6, a converse, namely a tight, standard-model reduction showing BN → IDL. This shows that IDL and BN are, security-wise, equivalent. Similarly, as discussed above, Theorem 7 shows MuSig → XIDL with a tight, standard model reduction, and we also give, in Theorem 8, a converse, namely a tight, standard-model reduction showing XIDL → MuSig. This shows that XIDL and MuSig are equivalent. Overall, this shows that IDL and XIDL are not arbitrary choices, but characterizations of the schemes whose consideration is necessary.

Definitional contributions. DEFKLNS [15] found subtle gaps in some prior proofs of security for some two-round multi-signature schemes [2, 24, 35]. This indicates a need for greater care in the domain of multi-signatures. We suggest that this needs to begin with *definitions*. The ones in prior work, stemming mostly from [7], suffer from some lack of detail and precision. In particular, the very *syntax* of a multi-signature scheme is not specified in detail. This results in scheme descriptions that lack in precision, and proofs that stay at a high level in part due to lack of technical language in which to give details. This in turn can lead to bugs.

To address these issues, we revisit the definitions. We start by giving a detailed syntax that formalizes the signing protocol as a stateful algorithm, run separately by each player. Details addressed include that a player knows its position in the signer list, that player identities are separate from public keys, and integration of the ROM through a parameter describing the type of ideal hash functions needed. Then we give a security definition written via a code-based game. See Section 4.

Related work. The interest for blockchains and cryptocurrencies, and thus our focus, is DL-based schemes over elliptic curves. There are many other multi-signature schemes, based on other hard problems. Aggregate signatures [12, 5] yield multi-signatures, but these use pairings (bilinear maps). A pairing-based multi-signature scheme is also given in [11]. Lattice-based multi-signature schemes include [16, 14].

As noted above, IDL [21] captures the security against parallel impersonation under key-only attack (PIMP-KOA) of the Schnorr ID scheme and thus, given the ZK property of the scheme, also its security against parallel impersonation under passive attack (PIMP-PA). “Parallel” means multiple impersonation attempts are allowed. IMP-PA, traditional security against impersonation under passive attack, is the case where just one impersonation attempt is allowed. The Reset Lemma [8] gives a standard model DL → IMP-PA reduction. This uses rewinding and is non-tight, with a square-root loss. BD [3] introduce the Multi-Base Discrete Logarithm (MBDL) problem, give a tight standard-model MBDL → IMP-PA reduction, and show that, in the GGM, the security of MBDL is the same as that of DL. An interesting open question is whether MBDL can be used as a starting point for tight reductions for multi-signature schemes. Rotem and Segev [32] give a standard model DL → IMP-PA reduction that improves the square-root-loss reduction but is still not tight.

## 2 Preliminaries

Notation. If $n$ is a positive integer, then $\mathbb{Z}_n$ denotes the set $\{0,\ldots,n-1\}$ and $[n]$ or $[1..n]$ denote the set $\{1,\ldots,n\}$. If $\boldsymbol{x}$ is a vector then $|\boldsymbol{x}|$ is its length (the number of its coordinates), $\boldsymbol{x}[i]$ is its $i$-th coordinate and $[\boldsymbol{x}] = \{\, \boldsymbol{x}[i] \;:\; 1 \le i \le |\boldsymbol{x}| \,\}$ is the set of all its coordinates. A string is identified with a vector over $\{0,1\}$, so that if $x$ is a string then $x[i]$ is its $i$-th bit and $|x|$ is its length. By $\varepsilon$ we denote the empty vector or string. The size of a set $S$ is denoted $|S|$.

Let $S$ be a finite set. We let $x \leftarrow\!\!{\scriptstyle\$}\; S$ denote sampling an element uniformly at random from $S$ and assigning it to $x$. We let $y \leftarrow A^{\mathrm{O}_1,\cdots}(x_1,\ldots;\rho)$ denote executing algorithm $A$ on inputs $x_1,\ldots$ and coins $\rho$ with access to oracles $\mathrm{O}_1,\ldots$, and letting $y$ be the result. We let $\rho \leftarrow\!\!{\scriptstyle\$}\; \mathrm{rand}(A)$ denote sampling random coins for algorithm $A$ and assigning it to variable $\rho$. We let $y \leftarrow\!\!{\scriptstyle\$}\; A^{\mathrm{O}_1,\cdots}(x_1,\ldots)$ be the result of $\rho \leftarrow\!\!{\scriptstyle\$}\; \mathrm{rand}(A)$ followed by $y \leftarrow A^{\mathrm{O}_1,\cdots}(x_1,\ldots;\rho)$. We let $[A^{\mathrm{O}_1,\cdots}(x_1,\ldots)]$ denote the set of all possible outputs of $A$ when invoked with inputs $x_1,\ldots$ and oracles $\mathrm{O}_1,\ldots$. Algorithms are randomized unless otherwise indicated. Running time is worst case.

Games. We use the code-based game playing framework of [9]. (See Fig. 4 for an example.) Games have procedures, also called oracles. Amongst these are Init and a Fin. In executing an adversary $\mathcal{A}$ with a game Gm, procedure Init is executed first, and what it returns is the input to $\mathcal{A}$. The latter may now call all game procedures except Init, Fin. When the adversary terminates, its output is viewed as the input to Fin, and what the latter returns is the game output. By $\mathrm{Gm}(\mathcal{A}) \Rightarrow y$ we denote the event that the execution of game Gm with adversary $\mathcal{A}$ results in output $y$. We write $\Pr[\mathrm{Gm}(\mathcal{A})]$ as shorthand for $\Pr[\mathrm{Gm}(\mathcal{A}) \Rightarrow \mathsf{true}]$, the probability that the game returns true. In writing game or adversary pseudocode, it is assumed that boolean variables are initialized to false, integer variables are initialized to 0 and set-valued variables are initialized to the empty set $\emptyset$.

A procedure (oracle) with a certain name O may appear in several games. (For example, Ch appears in two games in Figure 4.) To disambiguate, we may write Gm.O for the one in game Gm.

When adversary $\mathcal{A}$ is executed with game Gm, we consider the running time of $\mathcal{A}$ as the running time of the execution of $\mathrm{Gm}(\mathcal{A})$, which includes the time taken by game procedures. By $\mathrm{Q}^{\mathrm{O}}_{\mathcal{A}}$ we denote the number of queries made by $\mathcal{A}$ to oracle O in the execution. These counts are both worst case.

Groups. Throughout, $\mathbb{G}$ is a group whose order, assumed prime, we denote by $p$. We will use multiplicative notation for the group operation, and we let $1_{\mathbb{G}}$ denote the identity element of $\mathbb{G}$. We let $\mathbb{G}^* = \mathbb{G} \setminus \{1_{\mathbb{G}}\}$ denote the set of non-identity elements, which is the set of generators of $\mathbb{G}$ since the latter has prime order. If $g \in \mathbb{G}^*$ is a generator and $X \in \mathbb{G}$, then $\mathsf{DL}_{\mathbb{G},g}(X) \in \mathbb{Z}_p$ denotes the discrete logarithm of $X$ in base $g$.

Algebraic algorithms. We recall the definition of algebraic algorithms [17]. As above, fix a group $\mathbb{G}$ of prime order $p$, and let $g$ be a generator. In all of our security games involving $\mathbb{G}$ and $g$, we assume that any inputs and outputs of game oracles that are group elements (meaning, in $\mathbb{G}$) are distinguished. In particular, it will be clear from the game pseudocode definition which components of inputs and outputs are such group elements. We say that an adversary, against game Gm, is algebraic, if, whenever it submits a group element $Y \in \mathbb{G}$ as an oracle query, it also provides, alongside, a representation of $Y$ in terms of group elements previously returned by the game oracles (the latter including Init). Specifically, suppose during an execution of adversary $\mathcal{A}$ with game Gm, the adversary submits a group element $Y \in G$ to game oracle O. Then, alongside, it must provide a vector $(v_0, v_1, \ldots, v_m) \in \mathbb{Z}_p^m$, called a representation of $Y$, such that $Y = g^{v_0} \cdot h_1^{v_1} \cdots h_m^{v_m}$, where $h_1, \ldots, h_m$ are the group elements that have been returned to the adversary by game oracles of Gm, so far. When considering an execution of game Gm with an adversary $\mathcal{A}$ that is not algebraic, we omit the writing of representations in the oracle calls.

Hedging. Not all attacks are algebraic. The thesis of [17] is that natural ones are, and thus proving security relative to algebraic adversaries gives meaningful guarantees in practice. We adopt this here but add hedging. Recall this means that, for the same scheme, we seek both (1) A tight AGM reduction from DL, and (2) a standard-model (even if non-tight) reduction from DL. The former is used to guide and support parameter choices. The latter is viewed as at least qualitatively ruling out non-algebraic attacks.

Reductions. All our standard-model reductions are black-box and preserve algebraic-ness of adversaries, meaning, if the starting adversary is algebraic, so is the constructed one. This means that we can chain standard-model reductions with AGM-reductions to get overall AGM reductions.

## 3 Hardness of problems in groups

Our chain reductions exploit three computational problems related to groups: standard discrete log (DL); IDL [21]; and a new problem XIDL that we introduce. Here we give the definitions. We then show the length-2 chain DL → IDL → XIDL. We give reductions that are tight in the AGM and also give (non-tight) standard-model reductions, a total of four results. Referring to Figure 3, we are establishing the four theorems, shown in the table, that correspond to arrows 1 and 3. For the rest of the section, we fix a group $\mathbb{G}$ of prime order $p$, and a generator $g \in \mathbb{G}$.

DL. We recall the standard discrete logarithm (DL) problem via game $\mathrm{Gm}^{\mathrm{dl}}_{\mathbb{G},g}$ in Figure 4. Init provides the adversary, as input, a random challenge group element $X$, and to win it must output $x' = \mathsf{DL}_{\mathbb{G},g}(X)$ to Fin. We let $\mathbf{Adv}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A}) = \Pr[\mathrm{Gm}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A})]$ be the discrete-log advantage of adversary $\mathcal{A}$.

IDL. The identification discrete logarithm (IDL) problem, introduced by KMP [21], characterizes the hardness of parallel impersonation under key-only attack (PIMP-KOA) security [21] of the Schnorr identification scheme [33]. Formally, consider the game $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q}$ given in Fig. 4, where parameter $q$ is a positive integer. The IDL-adversary receives a random target point $X$ from Init. It is additionally given access to a challenge oracle Ch that can be called at most $q$ times. The oracle takes as query a group element $R$ (representing the commitment sent by the prover in Schnorr identification), stores it as $R_i$, and responds with a random challenge $c_i \leftarrow\!\!\$\; \mathbb{Z}_p$ (representing the one sent by the verifier). The adversary wins if it can produce the discrete log $z$ (representing the final prover response) of the group element $R_i \cdot X^{c_i}$, for a choice of $i$, denoted $I$, made by the adversary. We define the IDL-advantage of $\mathcal{A}$ to be $\mathbf{Adv}^{\mathrm{idl}}_{\mathbb{G},g,q}(\mathcal{A}) = \Pr[\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q}(\mathcal{A})]$.

KMP [21] study IDL in the Generic Group Model (GGM) [34] and prove a bound matching that for DL. Here, we strengthen this to give a tight AGM reduction DL → IDL. This could be seen as implicit in part of the AGM proof of security for the Schnorr signature scheme given in [18], although they make no connection to IDL.

**Theorem 1.** [DL → IDL, AGM] *Let $\mathbb{G}$ be a group of prime order $p$ with generator $g$. Let $q$ be a positive integer. Let $\mathcal{A}^{\mathrm{alg}}_{\mathrm{idl}}$ be an algebraic adversary against $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q}$. Then, adversary $\mathcal{A}_{\mathrm{dl}}$ can be constructed so that*

$$\mathbf{Adv}^{\mathrm{idl}}_{\mathbb{G},g,q}(\mathcal{A}^{\mathrm{alg}}_{\mathrm{idl}}) \le \mathbf{Adv}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A}_{\mathrm{dl}}) + \frac{q}{p}\ .$$

```
Game $\mathrm{Gm}^{\mathrm{dl}}_{\mathbb{G},g}$

INIT:
  $x \leftarrow\!\!\$\; Z_{|\mathbb{G}|}$ ; $X \leftarrow g^x$ ; Return $X$

FIN($x'$):
  Return $(x = x')$
```

```
Game $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q}$

INIT:
  $x \leftarrow \mathbb{Z}_{|\mathbb{G}|}$ ; $X \leftarrow g^x$
  Return $X$

CH($R$): // At most $q$ queries.
  $i \leftarrow i + 1$ ; $R_i \leftarrow R$
  $c_i \leftarrow\!\!\$\; \mathbb{Z}_{|\mathbb{G}|}$ ; Return $c_i$

FIN($I, z$):
  Return $(g^z = R_I \cdot X^{c_I})$
```

```
Game $\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}$

INIT:
  $x \leftarrow \mathbb{Z}_{|\mathbb{G}|}$ ; $X \leftarrow g^x$
  Return $X$

NWTAR($S$): // At most $q_1$ queries.
  $j \leftarrow j + 1$ ; $S_j \leftarrow S$
  $e_j \leftarrow\!\!\$\; \mathbb{Z}_{|\mathbb{G}|}$ ; $T_j \leftarrow S_j \cdot X^{e_j}$
  Return $e_j$

CH($j_{\mathrm{sel}}, R$): // At most $q_2$ queries.
  $i \leftarrow i + 1$ ; $R_i \leftarrow R$ ; $Y_i \leftarrow T_{j_{\mathrm{sel}}}$
  $c_i \leftarrow\!\!\$\; \mathbb{Z}_{|\mathbb{G}|}$ ; Return $c_i$

FIN($I, z$):
  Return $(g^z = R_I \cdot {Y_I}^{c_I})$
```

**Fig. 4.** Let $\mathbb{G}$ be a group of prime order $p = |\mathbb{G}|$, and let $g \in \mathbb{G}^*$ be a generator of $\mathbb{G}$. Let $q, q_1, q_2$ be positive integers. Top: Game defining discrete logarithm (DL) problem. Bottom left: Game defining identification logarithm (IDL) problem. Bottom right: Game defining random-target identification logarithm (XIDL) problem.

*Furthermore, the running time of $\mathcal{A}_{\mathrm{dl}}$ is about that of $\mathcal{A}^{\mathrm{alg}}_{\mathrm{idl}}$.*

The full proof is given in [4]. The idea of the proof is as follows. Since $\mathcal{A}^{\mathrm{alg}}_{\mathrm{idl}}$ is algebraic, its query $R$ to CH is accompanied by $(r_1, r_2)$ such that $R = g^{r_1} X^{r_2}$. Our adversary $\mathcal{A}_{\mathrm{dl}}$, who is running $\mathcal{A}^{\mathrm{alg}}_{\mathrm{idl}}$, records these as $R_i, r_{i,1}, r_{i,2}$, and responds with a random $c_i$. Eventually, $\mathcal{A}^{\mathrm{alg}}_{\mathrm{idl}}$ outputs $I, z$. Assuming it succeeds, we have $g^z = R_I \cdot X^{c_I} = g^{r_{I,1}} X^{r_{I,2}} X^{c_I}$, or $g^{z - r_{I,1}} = X^w$ where $w = (r_{I,2} + c_I) \bmod p$. Now $\mathsf{DL}_{\mathbb{G},g}(X)$ can be obtained as long as $w$ has an inverse modulo $p$, meaning is non-zero. But $c_I$ was chosen at random *after the adversary supplied* $r_{I,2}$, so the probability that $w$ is 0 is at most $1/p$. The factor of $q$ accounts for the adversary's having a choice of $I$ made after receiving challenges.

By $q$-IDL, we refer to IDL with parameter $q$. 1-IDL corresponds to IMP-KOA security of the Schnorr identification scheme, and a reduction DL → 1-IDL is obtained via the Reset Lemma of [8]. KMP show that 1-IDL → $q$-IDL. Overall this gives a standard model (very non-tight) DL → $q$-IDL reduction. However, a somewhat tighter (but still non-tight) result can be obtained when the forking lemma of [7] (which we recall as in [4].) is applied directly instead. Concretely, we give the following theorem, improving the prior reduction by a $\sqrt{q}$ factor. The proof is in [4].

**Theorem 2.** [DL → IDL, Standard Model] *Let* $\mathbb{G}$ *be a group of prime order* $p = |\mathbb{G}|$*, and let* $g \in \mathbb{G}^*$ *be a generator of* $\mathbb{G}$*. Let* $q$ *be a positive integer. Let* $\mathcal{A}_{\mathrm{idl}}$ *be an adversary against the game* $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q}$*. The proof constructs an adversary* $\mathcal{A}_{\mathrm{dl}}$ *such that*

$$\mathbf{Adv}^{\mathrm{idl}}_{\mathbb{G},g,q}(\mathcal{A}_{\mathrm{idl}}) \leq \sqrt{q \cdot \mathbf{Adv}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A}_{\mathrm{idl}})} + \frac{q}{p} \ . \tag{1}$$

*Additionally, the running time of* $\mathcal{A}_{\mathrm{dl}}$ *is approximately* $\mathrm{T}_{\mathcal{A}_{\mathrm{dl}}} \approx 2 \cdot \mathrm{T}_{\mathcal{A}_{\mathrm{idl}}}$*.*

Theorem 2 appears to yield a 1-IDL → $q$-IDL reduction with a bound that contradicts the lower bound claimed in [21, Corollary 4.4]. Our best guess as to an explanation is that our reduction does not meet the key and randomness preserving restrictions of [21, Corollary 4.4] or that their lower bound does not cover rewinding strategies.

XIDL. We define a new problem, random target identification discrete logarithm, abbreviated XIDL. It abstracts out the algebraic core of MuSig, and we will show that its security is equivalent to the MS-UF security of MuSig. It will also be an intermediate point in our reduction chain reaching our new HBMS scheme, thereby serving multiple purposes.

With $\mathbb{G}, p, g$ fixed as usual, XIDL is parameterized by positive integers $q_1, q_2$. Formally, consider the game $\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}$ given in Fig. 4. The adversary receives a randomly chosen group element $X$ from INIT. The game maintains a list $T_1, \ldots, T_{q_1}$ of “targets.” The adversary can create a target by querying the New Target oracle NWTAR with a group element $S$ of its choosing, whence $T_j = S \cdot X^{e_j}$ is added to the list of targets, for $e_j$ chosen randomly from $\mathbb{Z}_p$ by the game and returned to the adversary. The adversary can query the challenge oracle $\mathrm{CH}(j_{\mathrm{sel}}, R)$ by supplying an index $j_{\mathrm{sel}}$ and a group element $R$. The oracle records $T_{j_{\mathrm{sel}}}$ as $Y_i$, and $R$ as $R_i$, based on the counter $i$ it maintains. Intuitively, CH is similar to the challenge oracle CH in IDL game, besides that our adversary here needs to specify the target $T_{j_{\mathrm{sel}}}$ it is trying to impersonate against. The adversary wins the game if it can produce the discrete log $z$ of $R_I \cdot Y_I^{c_I}$, for an index $I$ of its choice. The oracles NWTAR and CH are allowed to be called at most $q_1$ and $q_2$ times, respectively. We define the XIDL advantage of $\mathcal{A}$ as $\mathbf{Adv}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}(\mathcal{A}) = \Pr[\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}(\mathcal{A})]$.

We show hardness of XIDL in both the AGM and the standard model, starting with the former. The theorem actually establishes the stronger DL → XIDL, tightly in the AGM.

**Theorem 3.** [DL → XIDL, AGM] *Let* $\mathbb{G}$ *be a group of order* $p$ *with generator* $g$*. Let* $q_1, q_2$ *be positive integers. Let* $\mathcal{A}^{\mathrm{alg}}_{\mathrm{xidl}}$ *be an algebraic adversary against* $\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}$*. Then, adversary* $\mathcal{A}_{\mathrm{dl}}$ *can be constructed so that*

$$\mathbf{Adv}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}(\mathcal{A}^{\mathrm{alg}}_{\mathrm{xidl}}) \leq \mathbf{Adv}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A}_{\mathrm{dl}}) + \frac{q_1 + q_2}{p} \ .$$

*Furthermore, the running time of* $\mathcal{A}_{\mathrm{dl}}$ *is about that of* $\mathcal{A}^{\mathrm{alg}}_{\mathrm{xidl}}$.

The full proof is given in [4]. Here we sketch the intuition. Since $\mathcal{A}^{\mathrm{alg}}_{\mathrm{xidl}}$ is algebraic, the $j$-th query to NwTar is of the form $S_j, s_{j,1}, s_{j,2}$ such that $S_j = g^{s_{j,1}} X^{s_{j,2}}$, and the $i$-th query to Ch is of the form $j_{\mathrm{sel}}, R_i, r_{i,1}, r_{i,2}$ such that $R_i = g^{r_{i,1}} X^{r_{i,2}}$. Let $e_j, c_i$ denote, respectively, the responses to the $j$-th query to NwTar and the $i$-th query to Ch. Eventually, $\mathcal{A}_{\mathrm{xidl}}$ outputs $I, z$. Assuming it succeeds, the equation $g^z = R_I \cdot T_J^{c_I} = R_I \cdot (S_J \cdot X^{e_J})^{c_I}$ must hold, where $J$ was the selected index $j_{\mathrm{sel}}$ in the $I$-th query to Ch. This means that $g^z = g^{r_{I,1}} X^{r_{I,2}} (g^{s_{J,1}} X^{s_{J,2}} X^{e_J})^{c_I}$, whence $g^{z - r_{I,1} - s_{J,1} \cdot c_I} = X^w$, where $w = r_{I,2} + (s_{J,2} + e_J) c_I$. As long as $w$ is non-zero modulo $p$, one can solve for the value of $\mathsf{DL}_{\mathbb{G},g}(X)$. But $e_J$ and $c_I$ were independently chosen after the adversary supplied $s_{J,2}$ and $r_{I,2}$, respectively. The probability that there exists $j$ such that $(s_{j,2} + e_j) = 0 \mod p$ is at most $q_1/p$ over $q_1$ queries to NwTar. Assuming there is no such $j$, the probability that $w = 0$ is at most $q_2/p$, due to the $q_2$ queries to Ch that $\mathcal{A}^{\mathrm{alg}}_{\mathrm{xidl}}$ can make.

In the standard model, techniques in the security proof of $\mathsf{MuSig}$ [11, 25] could be used to show DL → XIDL, which involves two applications of the Forking Lemma, leading to a fourth-root in the bound. We now show IDL → XIDL, using a single application of the forking lemma and thus with only a square-root in the bound. Combining this with Theorem 2 recovers the DL → XIDL reduction with its fourth-root.

**Theorem 4.** [IDL → XIDL, Standard Model] *Let* $\mathbb{G}$ *be a group of prime order* $p$ *with generator* $g$. *Let* $q_1, q_2$ *be positive integers. Let* $\mathcal{A}_{\mathrm{xidl}}$ *be an adversary against* $\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}$. *Then, an adversary* $\mathcal{A}_{\mathrm{idl}}$ *can be constructed so that*

$$\mathbf{Adv}^{\mathrm{xidl}}_{\mathbb{G},g,q_1,q_2}(\mathcal{A}_{\mathrm{xidl}}) \leq \sqrt{q_2 \cdot \mathbf{Adv}^{\mathrm{idl}}_{\mathbb{G},g,q_1}(\mathcal{A}_{\mathrm{idl}})} + \frac{q_2}{p} \ .$$

*Furthermore, the running time of* $\mathcal{A}_{\mathrm{idl}}$ *is about twice of that of* $\mathcal{A}_{\mathrm{xidl}}$.

The full proof is given in [4]. We now sketch the intuition. Adversary $\mathcal{A}_{\mathrm{idl}}$ receives $X$ from game $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}$ and runs adversary $\mathcal{A}_{\mathrm{xidl}}$, forwarding it $X$ as the target point. It answers queries to $\mathcal{A}_{\mathrm{xidl}}$'s NwTar oracle using its own $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}$.Ch oracle. Specifically, the $j$-th query $S$ to NwTar is responded to with $e_j \leftarrow\$ \mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}.\textsc{Ch}(S)$, and $\mathcal{A}_{\mathrm{idl}}$ additionally records the group element $T_j \leftarrow S \cdot X^{e_j}$. It simulates adversary $\mathcal{A}_{\mathrm{xidl}}$'s Ch oracle locally, meaning the $i$-th query $\textsc{Ch}(j_{\mathrm{sel}}, R)$ is responded to with a fresh challenge $c_i \leftarrow\$ \mathbb{Z}_p$. Eventually, adversary $\mathcal{A}_{\mathrm{xidl}}$ gives a response $I, z$. Our $\mathcal{A}_{\mathrm{idl}}$ adversary wins game $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}$ if it can produce the discrete log of $T_j$ for any $j$ of its choice. To do so, $\mathcal{A}_{\mathrm{idl}}$ uses rewinding, the analysis of which uses the Forking Lemma [7] that we recall in [4]. Rewinding is used to produce another response, $(I', z')$, from a forked execution of $\mathcal{A}_{\mathrm{xidl}}$. The Forking Lemma applies to an execution of an algorithm making queries to one oracle, but adversary $\mathcal{A}_{\mathrm{xidl}}$ has two oracles NwTar and Ch. We only "fork" $\mathcal{A}_{\mathrm{xidl}}$ on its queries to Ch. Specifically, we program oracle NwTar to behave identically compared to the first run (meaning we use previously recorded values of $e_1, \ldots$ as long as they are defined). In the second run, oracle Ch is replied with

$c_1, \ldots, c_{I-1}, c'_I, \ldots$, where $c'_I, \ldots$ are randomly chosen from $\mathbb{Z}_p$. Let us assume that $\mathcal{A}_{\text{idl}}$ has derived two valid responses from $\mathcal{A}_{\text{xidl}}$ using the Forking Lemma. Then it is guaranteed that $I = I'$ and $c_I \neq c'_I$. Moreover, we know the two executions of $\mathcal{A}_{\text{xidl}}$ only differ *after* the response of the $I$-th query to Ch, so the $I$-th query to Ch in both runs is some $J, R_I$. This allows our adversary to solve the equations $g^z = R_I \cdot T_J^{c_I}$ and $g^{z'} = R_I \cdot T_J^{c'_I}$ (which are guaranteed to be true if both runs succeed) to compute $\mathsf{DL}_{\mathbb{G},g}(T_J)$ and thus win the IDL game.

## 4 Definitions for multi-signatures

As discussed in Section 1, current definitions for multi-signatures, stemming mostly from [7], suffer from some lack of detail and precision, including lack of a precise syntax. This results in scheme descriptions that also lack somewhat in precision, and to proofs that stay at a high level in part due to lack of technical language in which to give details. This could be one of the contributors to bugs in these proofs [15].

To address this, we revisit the definitions. We give a detailed syntax that formalizes the signing protocol as a stateful algorithm, run separately by each player. (The state will be maintained by the overlying game.) Details addressed include that a player knows its position in the signer list, that player identities are separate from public keys, and integration of the ROM through a parameter describing the type of ideal hash functions needed. Then we give a security definition written via a code-based game.

SYNTAX. A multi-signature scheme $\mathsf{MS}$ specifies algorithms $\mathsf{MS.Kg}$, $\mathsf{MS.Vf}$, $\mathsf{MS.Sign}$, as well as a set $\mathsf{MS.HF}$ of functions, and an integer $\mathsf{MS.nr}$, whose intent and operation is as follows:

- *Key generation.* Via $(pk, sk) \leftarrow\$ \mathsf{MS.Kg}$, the key generation algorithm generates public signature-verification key $pk$ and secret signing key $sk$ for a user. (Each user is expected to run this independently to get its keys.)
- *Hash functions.* $\mathsf{MS.HF}$ is a set of functions, from which, via $\mathsf{h} \leftarrow\$ \mathsf{MS.HF}$, one is drawn and provided to scheme algorithms (except key generation) and the adversary as the random oracle. Specifying this as part of the scheme allows the domain and range of the random oracle to be scheme-dependent.
- *Verification.* Via $d \leftarrow \mathsf{MS.Vf}^{\mathrm{H}}(\boldsymbol{pk}, m, \sigma)$, the verification algorithm deterministically outputs a decision $d \in \{\mathsf{true}, \mathsf{false}\}$ indicating whether or not $\sigma$ is a valid signature on message $m$ under a vector $\boldsymbol{pk}$ of verification keys.
- *Signing.* The signing protocol is specified by signing algorithm $\mathsf{MS.Sign}$. In each round, each party, applies this algorithm to its current state $\mathsf{st}$ and the vector $\mathbf{in}$ of received messages from the other parties, to compute an outgoing message $\sigma$ (viewed as broadcast to the other parties) and an updated state $\mathsf{st}'$, written $(\sigma, \mathsf{st}') \leftarrow \mathsf{MS.Sign}^{\mathrm{H}}(\mathbf{in}, \mathsf{st})$. In the last round, $\sigma$ is the signature that this party outputs. (See Figure 5.)

- *Rounds.* The interaction consists of a fixed number $\mathsf{MS.nr}$ of rounds. (We number the rounds $0, \ldots, \mathsf{MS.nr}$. The final broadcast of the signature is not counted as in practice it is a local output.)

We say that a multi-signature scheme $\mathsf{MS}$ supports key aggregation if $\mathsf{MS}$ has two additional algorithms, $\mathsf{MS.Ag}$ and $\mathsf{MS.VfAg}$, such that the following hold: (1) Via $apk \leftarrow_\$ \mathsf{MS.Ag}^{\mathrm{H}}(pk_1, \ldots, pk_n)$, the key aggregation algorithm $\mathsf{MS.Ag}$ generates an aggregate public key, (2) Via $d \leftarrow \mathsf{MS.VfAg}^{\mathrm{H}}(apk, m, \sigma)$, the aggregate verification algorithm deterministically outputs a decision $d \in \{\mathsf{true}, \mathsf{false}\}$, and (3) the verification algorithm $\mathsf{MS.Vf}$ is defined exactly as $\mathsf{MS.Vf}^{\mathrm{H}}(\boldsymbol{pk}, m, \sigma) = \mathsf{MS.VfAg}^{\mathrm{H}}(\mathsf{MS.Ag}^{\mathrm{H}}(\boldsymbol{pk}), m, \sigma)$.

Some conventions will aid further definitions and scheme descriptions. A party's state $\mathsf{st}$ has several parts: $\mathsf{st.n}$ is the number of parties in the current execution of the protocol; $\mathsf{st.me} \in [1..\mathsf{st.n}]$ is the party's own identity; $\mathsf{st.rnd} \in [0..\mathsf{MS.nr}]$ is the current round number; $\mathsf{st.sk}$ is the party's own signing key; $\mathsf{st.pk}$ is the $\mathsf{st.n}$-vector of all verification keys; $\mathsf{st.msg}$ is the message being signed; $\mathsf{st.rej} \in \{\mathsf{true}, \mathsf{false}\}$ is the decision to reject (not produce a signature) or accept. It is assumed and required that each invocation of $\mathsf{MS.Sign}$ leaves all of these unchanged except for $\mathsf{st.rnd}$, which it increments by 1, and $\mathsf{st.rej}$, which is assumed initialized to $\mathsf{false}$ and may at some point be set to $\mathsf{true}$. The state can, beyond these, have other components that vary from protocol to protocol. (For example, Figure 6 describing the BN scheme has $\mathsf{st}.\boldsymbol{R}[j], \mathsf{st}.\boldsymbol{t}[j], \mathsf{st}.\boldsymbol{z}[j], \mathsf{st}.R, \ldots$.) We write $\mathsf{st} \leftarrow \mathsf{StInit}(j, sk, \boldsymbol{pk}, m)$ to initialize $\mathsf{st}$ by setting $\mathsf{st.n} \leftarrow |\boldsymbol{pk}|$ ; $\mathsf{st.me} \leftarrow j$ ; $\mathsf{st.rnd} \leftarrow 0$ ; $\mathsf{st.sk} \leftarrow sk$ ; $\mathsf{st.pk} \leftarrow \boldsymbol{pk}$ ; $\mathsf{st.msg} \leftarrow m$ ; $\mathsf{st.rej} \leftarrow \mathsf{false}$. If an execution $(\sigma, \mathsf{st}') \leftarrow \mathsf{MS.Sign}^{\mathrm{H}}(\mathbf{in}, \mathsf{st})$ returns $\sigma = \bot$ then it is assumed and required that further executions starting from $\mathsf{st}'$ all return $\bot$ as the output message.

<u>CORRECTNESS.</u> Algorithm $\mathsf{Exec}_{\mathsf{MS}}$, shown in the left column of Fig. 5, executes the signing protocol of $\mathsf{MS}$ on input a vector $\boldsymbol{sk}$ of signing keys, a vector $\boldsymbol{pk}$ of matching verification keys with $|\boldsymbol{sk}| = |\boldsymbol{pk}|$, and a message $m$ to be signed, and with access to random oracle $\mathsf{h} \in \mathsf{MS.HF}$. The number of parties $n$ at line 1 is the number of coordinates (length) of $\boldsymbol{pk}$. The state $\mathsf{st}_j$ of party $j$ at line 3 is initialized using the function $\mathsf{StInit}$ defined above. The loop at line 5 executes $\mathsf{MS.nr}$ rounds. Here $\boldsymbol{b}$ denotes the $n$-vector of currently-broadcast messages, meaning $\boldsymbol{b}[i]$ was broadcast by party $i$ in the prior round, and the entire vector is the input to party $j$ for the current round. At line 8, $\boldsymbol{b}$ now holds the next round of broadcasts.

The correctness game $\mathbf{G}^{\text{ms-cor}}_{\mathsf{MS},n}$ shown in the right column of Fig. 5 has only one procedure, namely FIN. We say that $\mathsf{MS}$ satisfies (perfect) correctness if for all positive integers $n$ we have $\Pr[\mathbf{G}^{\text{ms-cor}}_{\mathsf{MS},n}] = 1$.

<u>UNFORGEABILITY.</u> Game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ in Fig. 5 captures a notion of unforgeability for multi-signatures that slightly extends [7]. There is one honest player whose keys are picked at line 1, the adversary controlling all the other players. A new instance of the signing protocol is initialized by calling NS with an index $k$ and a vector $\boldsymbol{pk}$ of verification keys that the adversary can choose, possibly dishonestly, subject only to $\boldsymbol{pk}[k]$ being the verification key $pk$ of the honest player, as enforced by line 2. The first message of the honest player is sent out, and at

```
Algorithm $\mathsf{Exec}^{\mathsf{h}}_{\mathsf{MS}}(\boldsymbol{sk}, \boldsymbol{pk}, m)$:
$n \leftarrow |\boldsymbol{pk}|$
For $j = 1, \ldots, n$ do
    $\mathsf{st}_j \leftarrow \mathsf{StInit}(j, \boldsymbol{sk}[j], \boldsymbol{pk}, m)$
$\boldsymbol{b} \leftarrow (\varepsilon, \ldots, \varepsilon)$  // $n$-vector
For $i = 1, \ldots, \mathsf{MS.nr}$ do
    For $j = 1, \ldots, n$ do
        $(\sigma_j, \mathsf{st}_j) \leftarrow\!\!\$\ \mathsf{MS.Sign}^{\mathsf{h}}(\boldsymbol{b}, \mathsf{st}_j)$
    $\boldsymbol{b} \leftarrow (\sigma_1, \ldots, \sigma_n)$
Return $\sigma_1$
```

```
Game $\mathbf{G}^{\text{ms-cor}}_{\mathsf{MS},n}$

FIN:
$\mathsf{h} \leftarrow\!\!\$\ \mathsf{MS.HF}$
For $i = 1, \ldots, n$ do
    $(\boldsymbol{pk}[i], \boldsymbol{sk}[i]) \leftarrow\!\!\$\ \mathsf{MS.Kg}$
$\sigma \leftarrow\!\!\$\ \mathsf{Exec}^{\mathsf{h}}_{\mathsf{MS}}(\boldsymbol{sk}, \boldsymbol{pk}, m)$
$d \leftarrow \mathsf{MS.Vf}^{\mathsf{h}}(\boldsymbol{pk}, m, \sigma)$
Return $d$
```

```
Game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$

INIT:
$\mathsf{h} \leftarrow\!\!\$\ \mathsf{MS.HF}$ ; $(pk, sk) \leftarrow\!\!\$\ \mathsf{MS.Kg}$ ; Return $pk$

NS$(k, \boldsymbol{pk}, m)$:
$\boldsymbol{pk}[k] \leftarrow pk$ ; $u \leftarrow u + 1$ ; $\boldsymbol{pk}_u \leftarrow \boldsymbol{pk}$ ; $m_u \leftarrow m$ ; $\mathsf{st}_u \leftarrow \mathsf{StInit}(k, sk, \boldsymbol{pk}, m)$
$\boldsymbol{b} \leftarrow (\varepsilon, \ldots, \varepsilon)$ ; $(\sigma, \mathsf{st}_u) \leftarrow\!\!\$\ \mathsf{MS.Sign}^{\mathrm{H}}(\boldsymbol{b}, \mathsf{st}_u)$ ; Return $\sigma$

SIGN$_j(s, \boldsymbol{b})$:  // $1 \le j \le \mathsf{MS.nr}$
$(\sigma, \mathsf{st}_s) \leftarrow\!\!\$\ \mathsf{MS.Sign}^{\mathrm{H}}(\boldsymbol{b}, \mathsf{st}_s)$ ; Return $\sigma$

H$(x)$:
Return $\mathsf{h}(x)$

FIN$(k, \boldsymbol{pk}, m, \sigma)$:
If $(\boldsymbol{pk}[k] \neq pk)$ then Return false
If $(\boldsymbol{pk}, m) \in \{(\boldsymbol{pk}_i, m_i) : 1 \le i \le u\}$ then Return false
Return $\mathsf{MS.Vf}^{\mathrm{H}}(\boldsymbol{pk}, m, \sigma)$
```

**Fig. 5. Top left:** Procedure specifying an honest execution of the signing protocol associated with multi-signature scheme $\mathsf{MS}$. **Top right:** Correctness game. **Bottom:** Unforgeability game.

this point $\mathsf{st}_u.\mathsf{rnd} = 1$. Now the adversary can run multiple concurrent instances of the signing protocol with the honest signer. Oracle H is the random oracle, simply calling $\mathsf{h}$. Eventually the adversary calls FIN with a forgery index $k$, a vector of verification keys (subjected to $\boldsymbol{pk}[k]$ being the public key of the honest signer), a message and a claimed signature. It wins if verification succeeds and the forgery was non-trivial. The ms-uf-advantage of adversary $\mathcal{A}$ is $\mathbf{Adv}^{\text{ms-uf}}_{\mathsf{MS}}(\mathcal{A}) = \Pr[\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}(\mathcal{A})]$.

It is convenient for (later) proofs to have a separate signing oracle SIGN$_j$ for each round $j \in [1..\mathsf{MS.nr}]$. It is required that any SIGN$_j(s, \cdot)$ satisfy $s \in [1..u]$, and that the prior round queries $\textsc{Sign}_k(s, \cdot)$ for $k < j$ have already been made. It is required that for each $j, s$, at most one $\textsc{Sign}_j(s, \cdot)$ query is ever made.

Remarks. Our syntax and security notions for multi-signatures view a group of signers as captured by the vector (rather than the set) of their public keys. So for example, a forgery $((pk_1, pk_2), m, \sigma)$ is considered to be non-trivial even if there was a previous signing session under public keys $(pk_2, pk_1)$ and message $m$. This differs from previous formalizations that work instead with sets of public keys. However, previous definition can be recovered if a canonical encoding of sets of public keys into vectors of public keys is fixed in the usage of a scheme.

## 5 Analysis of the BN scheme

BN scheme. Let $\mathbb{G}$ be a group of prime order $p$. Let $g$ be a generator of $\mathbb{G}$ and let $\ell \geq 1$ be an integer. The associated BN [7] multi-signature scheme $\mathsf{MS} = \mathsf{BN}[\mathbb{G}, g, \ell]$ is shown in detail, in our syntax, in Fig. 6. The set $\mathsf{MS.HF}$ consists of all functions $\mathsf{h}$ such that $\mathsf{h}(0, \cdot) : \{0,1\}^* \rightarrow \{0,1\}^\ell$ and $\mathsf{h}(1, \cdot) : \{0,1\}^* \rightarrow \mathbb{Z}_p$. For $b \in \{0,1\}$ we write $\mathrm{H}_b(\cdot)$ for $\mathrm{H}(b, \cdot)$, so that scheme algorithms, and an ms-uf adversary, will have access to oracles $\mathrm{H}_0, \mathrm{H}_1$ rather than just H.

The signing protocol has 3 rounds. In round 0, player $j$ picks $r \leftarrow\!\!\$\ \mathbb{Z}_p$, stores $g^r$ in its state as $\mathsf{st}.\boldsymbol{R}[j]$, computes, and stores in its state, a value $\mathsf{st}.\boldsymbol{t}[j] \leftarrow \mathrm{H}_0((j, \mathsf{st}.\boldsymbol{R}[j]))$ that we call the BN-commitment, and broadcasts the BN-commitment. (Per our syntax, what is returned is the message to be broadcast and the updated state to be retained.) Since each player does this, in round 1, player $j$ receives the BN-commitments of the other players, storing them in vector $\mathsf{st}.\boldsymbol{t}$, and now broadcasting $\mathsf{st}.\boldsymbol{R}[j]$. In round 2, these broadcasts are received, so player $j$ can form the vector $\mathsf{st}.\boldsymbol{R}$. At line 20, it returns $\perp$ if one of the received values fails to match its commitment. As per our conventions, when this happens, this player will always broadcast $\perp$ in the future, so for round 3 we assume lines 21 and 22 are executed. These lines create the second component $\mathsf{st}.\boldsymbol{z}[j]$ of a Schnorr signature relative to the Schnorr-commitment $\mathsf{st}.\boldsymbol{R}[j]$ defined at line 13, and the player's own secret key, the computations being modulo $p$. This $\mathsf{st}.\boldsymbol{z}[j]$ is broadcast, so that, in round 3, our player receives the corresponding values from the other players. At line 27 it forms their modulo-$p$ sum $z$ and then forms the final signature $(\mathsf{st}.R, z)$.

Our description of the signing protocol differs, from that in [7], in some details that are brought out by our syntax, for example in using explicit party identities rather than seeing these as implicit in public keys.

Prior bounds. We recall the prior result of [7]. Let $\mathsf{MS} = \mathsf{BN}[\mathbb{G}, g, \ell]$ and let $\mathcal{A}_{\mathrm{ms}}$ be an adversary for game $\mathbf{G}^{\mathrm{ms\text{-}uf}}_{\mathsf{MS}}$. Assume the execution of game $\mathbf{G}^{\mathrm{ms\text{-}uf}}_{\mathsf{MS}}$ with $\mathcal{A}_{\mathrm{ms}}$ has at most $q$ distinct queries across $\mathrm{H}_0, \mathrm{H}_1$ and at most $q_{\mathrm{s}}$ queries to NS. Suppose the number of parties (length of verification-key vector) in queries to NS and $\textsc{Fin}$ is at most $n$. Let $a = 8q_{\mathrm{s}} + 1$ and $b = 2q + 16n^2q_{\mathrm{s}}$. Let $p = |\mathbb{G}|$.

```
$\mathsf{Kg}$:
    $sk \leftarrow\!\$\ \mathbb{Z}_p$ ; $pk \leftarrow g^{sk}$
    Return $(pk, sk)$
```

```
$\mathsf{Vf}^{\mathrm{H}}(\boldsymbol{pk}, m, \sigma)$:
    $(R, z) \leftarrow \sigma$ ; $(pk_1, \ldots, pk_n) \leftarrow \boldsymbol{pk}$
    $\mathsf{BN}$ :
        For $i = 1, \ldots, n$ do $c_i \leftarrow \mathrm{H}_1((i, R, \boldsymbol{pk}, m))$
        Return ( $g^z = R \cdot \prod_{i=1}^{n} pk_i^{c_i}$ )
    $\mathsf{MuSig}$ :
        $apk \leftarrow \prod_i^n pk_i^{\mathrm{H}_2((i,\boldsymbol{pk}))}$
        $c \leftarrow \mathrm{H}_1((R, apk, m))$
        Return ( $g^z = R \cdot apk^c$ )
```

```
$\mathsf{Sign}^{\mathrm{H}}(\boldsymbol{b}, \mathsf{st})$:
    $j \leftarrow \mathsf{st.me}$ ; $n \leftarrow \mathsf{st.n}$ ; $m \leftarrow \mathsf{st.msg}$ ; $sk \leftarrow \mathsf{st.sk}$ ; $\boldsymbol{pk} \leftarrow \mathsf{st.pk}$
    If $(\mathsf{st.rnd} = 0)$ then
        $\mathsf{st}.r \leftarrow\!\$\ \mathbb{Z}_p$ ; $\mathsf{st}.\boldsymbol{R}[j] \leftarrow g^r$ ; $\mathsf{st}.\boldsymbol{t}[j] \leftarrow \mathrm{H}_0((j, \mathsf{st}.\boldsymbol{R}[j]))$ ; $\mathsf{st.rnd} \leftarrow \mathsf{st.rnd} + 1$
        Return $(\mathsf{st}.\boldsymbol{t}[j], \mathsf{st})$
    If $(\mathsf{st.rnd} = 1)$ then
        For all $i \neq j$ do $\mathsf{st}.\boldsymbol{t}[i] \leftarrow \boldsymbol{b}[i]$
        $\mathsf{st.rnd} \leftarrow \mathsf{st.rnd} + 1$ ; Return $(\mathsf{st}.\boldsymbol{R}[j], \mathsf{st})$
    If $(\mathsf{st.rnd} = 2)$ then
        For all $i \neq j$ do $\mathsf{st}.\boldsymbol{R}[i] \leftarrow \boldsymbol{b}[i]$
        If ( $\exists\, i : \mathrm{H}_0((i, \mathsf{st}.R[i])) \neq \mathsf{st}.\boldsymbol{t}[i]$ ) then Return $(\bot, \mathsf{st})$
        $\mathsf{st}.R \leftarrow \prod_{i=1}^{n} \mathsf{st}.R[i]$
        $\mathsf{BN}$ : $c_j \leftarrow \mathrm{H}_1((j, R, \boldsymbol{pk}, m))$ ; $\mathsf{st}.\boldsymbol{z}[j] \leftarrow sk \cdot c_j + \mathsf{st}.r$
        $\mathsf{MuSig}$ :
            $apk \leftarrow \prod_{i=1}^{n} \boldsymbol{pk}[i]^{\mathrm{H}_2((i,\boldsymbol{pk}))}$ ; $c \leftarrow \mathrm{H}_1((R, apk, m))$
            $\mathsf{st}.\boldsymbol{z}[j] \leftarrow sk \cdot \mathrm{H}_2((\mathsf{st.me}, \boldsymbol{pk})) \cdot c + \mathsf{st}.r$
        $\mathsf{st.rnd} \leftarrow \mathsf{st.rnd} + 1$ ; Return $(\mathsf{st}.\boldsymbol{z}[j], \mathsf{st})$
    If $(\mathsf{st.rnd} = 3)$ then
        For all $i \neq j$ do $\mathsf{st}.\boldsymbol{z}[i] \leftarrow \boldsymbol{b}[i]$
        $z \leftarrow \sum_{i=1}^{n} \mathsf{st}.\boldsymbol{z}[i]$ ; Return $((\mathsf{st}.R, z), \mathsf{st})$
```

**Fig. 6.** Algorithms of the multi-signature scheme $\mathsf{BN}[\mathbb{G}, g, \ell]$ and $\mathsf{MuSig}[\mathbb{G}, g, \ell]$, where $\mathbb{G}$ is a group of prime order $p$ with generator $g$. Code that differs between the two schemes is marked explicitly. Oracle $\mathrm{H}_i(\cdot)$ is defined to be $\mathrm{H}(i, \cdot)$ for $i = 0, 1$ ($\mathsf{BN}$) and $i = 0, 1, 2$ ($\mathsf{MuSig}$).

Then BN [7] give a DL-adversary $\mathcal{A}_{\mathrm{dl}}$ such that

$$\mathbf{Adv}^{\mathrm{ms\text{-}uf}}_{\mathsf{MS}}(\mathcal{A}_{\mathrm{ms}}) \leq \sqrt{(q + q_{\mathrm{s}}) \cdot \left( \mathbf{Adv}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A}_{\mathrm{dl}}) + \frac{a}{p} + \frac{b}{2^{\ell}} \right)} \ . \qquad (2)$$

The running time of $\mathcal{A}_{\mathrm{dl}}$ is twice that of the execution of game $\mathbf{G}^{\mathrm{ms\text{-}uf}}_{\mathsf{MS}}$ with $\mathcal{A}_{\mathrm{ms}}$. BN obtain this result via their general forking lemma, which uses rewinding and accounts for the square-root in the bound.

Security of BN from IDL. We give a IDL → $\mathsf{BN}$ reduction that is *tight* and in the *standard model*. Combining this with our tight AGM reduction DL → IDL of Theorem 1 we conclude a tight AGM reduction DL → $\mathsf{BN}$. However, the standard model tight IDL → $\mathsf{BN}$ reduction is also interesting in its own right. It says that BN is just as secure as the Schnorr identification scheme. Since the latter has been around and resisted cryptanalysis for quite some time, this is good support for the security of BN.

**Theorem 5.** [IDL → $\mathsf{BN}$, Standard Model] *Let $\mathbb{G}$ be a group of prime order $p$. Let $g$ be a generator of $\mathbb{G}$ and let $\ell \geq 1$ be an integer. Let $\mathsf{MS} = \mathsf{BN}[\mathbb{G}, g, \ell]$ be the associated BN multi-signature scheme. Let $\mathcal{A}_{\mathrm{ms}}$ be an adversary for game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ of Figure 5. Assume the execution of game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ with $\mathcal{A}_{\mathrm{ms}}$ has at most $q_0, q_1, q_{\mathrm{s}}$ distinct queries to $\mathrm{H}_0, \mathrm{H}_1, \mathrm{NS}$, respectively, and the number of parties (length of verification-key vector) in queries to* NS *and* Fin *is at most $n$. Let $\alpha = q_{\mathrm{s}}(4q_0 + 2q_1 + q_{\mathrm{s}})$ and $\beta = q_0(q_0 + n)$. Then we construct an adversary $\mathcal{A}_{\mathrm{id}}$ for game $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}$ such that*

$$\mathbf{Adv}^{\text{ms-uf}}_{\mathsf{MS}}(\mathcal{A}_{\mathrm{ms}}) \leq \mathbf{Adv}^{\mathrm{idl}}_{\mathbb{G},g,q_1}(\mathcal{A}_{\mathrm{idl}}) + \frac{\alpha}{2p} + \frac{\beta}{2^{\ell}} \ . \tag{3}$$

*The running time of $\mathcal{A}_{\mathrm{idl}}$ is about that of the execution of game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ with $\mathcal{A}_{\mathrm{ms}}$. Furthermore, adversary $\mathcal{A}_{\mathrm{idl}}$ is algebraic if adversary $\mathcal{A}_{\mathrm{ms}}$ is.*

Above, $q_0$ is the number of distinct queries to $\mathrm{H}_0$ made, not directly by the adversary, but across the execution of the adversary in game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$, and similarly for $q_1$. A lower bound on $q_1$ is the length of $\boldsymbol{pk}$ in $\mathcal{A}_{\mathrm{ms}}$'s Fin query, so we can assume it is positive. With the above theorem, we can now derive an upperbound $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_{\mathrm{s}}, p)$ of the advantage of any MS adversary with running time $t$, making $q$ queries to H, and $q_{\mathrm{s}}$ signing interactions. We take $\ell \approx \log_2(p)$ and assume that $q_{\mathrm{s}} \leq q \leq t \leq p$. Additionally, we assume that the advantage of any IDL adversary with running time $t$ is at most $t^2/p$ (as justified by Theorem 2). We obtain $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p) \leq t^2/p$ as shown in Fig. 1.

The full proof of Theorem 5 is given in [4]. Here we give a sketch. The reduction adversary $\mathcal{A}_{\mathrm{idl}}$ receives a group element $X$ from $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}$ and forwards it to adversary $\mathcal{A}_{\mathrm{ms}}$ as the target public key. In order to run adversary $\mathcal{A}_{\mathrm{ms}}$, our adversary needs to be able to simulate the signing oracles $\mathrm{NS}, \textsc{Sign}_1, \textsc{Sign}_2$ as well as random oracles $\mathrm{H}_0$ and $\mathrm{H}_1$ without knowing $\mathsf{DL}_{\mathbb{G},g}(X)$. We first describe how the reduction proceeds if $\mathcal{A}_{\mathrm{ms}}$ makes no queries to $\mathrm{NS}, \textsc{Sign}_1$ or $\textsc{Sign}_2$, as this steps constitutes the main difference between our proof and the original proof of security for BN [7]. Adversary $\mathcal{A}_{\mathrm{idl}}$ uses the challenge oracle $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}.\textsc{Ch}$ to program the random oracle $\mathrm{H}_1$ (hence Ch needs to be able to be queried upto the number of times $\mathrm{H}_1$ is evaluated). In particular, for each query $\mathrm{H}_1((k, R, \boldsymbol{pk}, m))$ where $\boldsymbol{pk}[k] = X$, our adversary first computes $T \leftarrow R \cdot \prod_{j \neq k} \boldsymbol{pk}[j]^{\mathrm{H}_1((j,R,\boldsymbol{pk},m))}$, then obtains $c \leftarrow\!\!\$\ \textsc{Ch}(T)$ before returning $c$ as the return value for the query $\mathrm{H}_1((k, R, \boldsymbol{pk}, m))$. By construction, a valid forgery for $\boldsymbol{pk}, m$ is some signature

$\sigma = (R, z)$ such that

$$g^z = R \cdot \prod_{i=1}^{n} \boldsymbol{pk}[i]^{\mathrm{H}_1((i,R,\boldsymbol{pk},m))} = T \cdot X^c ,$$

where the first equality is by the verification equation of $\mathsf{BN}$ and the second equality is by the way $\mathrm{H}_1$ is programmed. This means that adversary $\mathcal{A}_{\mathrm{idl}}$ can simply forward the value of $z$ from a valid forgery, along with the index of the $\textsc{Ch}$ query corresponding to the $\mathrm{H}_1$ query of the forgery, to break game $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q_1}$. Moreover, adversary $\mathcal{A}_{\mathrm{idl}}$ succeeds as long as the forgery given by $\mathcal{A}_{\mathrm{ms}}$ is valid.

It remains to show that oracles $\mathrm{NS}, \textsc{Sign}_1, \textsc{Sign}_2$ can be simulated without knowledge of the secret key, $\mathsf{DL}_{\mathbb{G},g}(X)$. Roughly, this is done using the zero-knowledge property of the underlying Schnorr identification scheme as well as by programming the random oracles $\mathrm{H}_0$ and $\mathrm{H}_1$. The original proof by [7] constructs an adversary and argues that it simulates these oracles faithfully if certain bad events do not happen. We take a more careful approach and do this formally via a sequence of seven games and use the code-base game playing framework of [9]. This game sequence incurs the additive loss as indicated in Equation (3).

<u>Converse.</u> IDL is not merely some group problem that can be used to justify security of $\mathsf{BN}$ tightly; the hardness of IDL is, in fact, tightly equivalent to the MS-UF security of $\mathsf{BN}$. Formally, we give below a reduction turning any adversary against IDL into a forger $\mathcal{A}_{\mathrm{ms}}$ against $\mathsf{BN}$. This means that any security justification for $\mathsf{BN}$ must also justify the hardness of IDL.

**Theorem 6.** [$\mathsf{BN} \rightarrow$ IDL, Standard Model] *Let $\mathbb{G}$ be a group of prime order $p$. Let $g$ be a generator of $\mathbb{G}$ and let $\ell \geq 1$ be an integer. Let $\mathsf{MS} = \mathsf{BN}[\mathbb{G}, g, \ell]$ be the associated BN multi-signature scheme. Let $q$ be a positive integer and $\mathcal{A}_{\mathrm{idl}}$ be an adversary against $\mathrm{Gm}^{\mathrm{idl}}_{\mathbb{G},g,q}$. Then, we can construct an adversary $\mathcal{A}_{\mathrm{ms}}$ for game $\mathbf{G}^{\mathrm{ms\text{-}uf}}_{\mathsf{MS}}$, making no queries to* NS*, and at most $2q$ queries to* $\mathrm{H}_1$*, such that*

$$\mathbf{Adv}^{\mathrm{ms\text{-}uf}}_{\mathsf{MS}}(\mathcal{A}_{\mathrm{ms}}) \geq \mathbf{Adv}^{\mathrm{idl}}_{\mathbb{G},g,q}(\mathcal{A}_{\mathrm{idl}}) . \tag{4}$$

*The running time of $\mathcal{A}_{\mathrm{ms}}$ is about that of $\mathcal{A}_{\mathrm{idl}}$.*

*Proof (Theorem 6).* Consider the adversary given in Fig. 7. The adversary receives the target public key $pk$ from the MS-UF game and samples a key pair $(pk', sk') \leftarrow\!\!\$\ \mathsf{MS.Kg}$. The adversary will attempt to forge a signature against the vector of public keys $(pk, pk')$. Adversary $\mathcal{A}_{\mathrm{ms}}$ forwards $X = pk$ as the target point and runs IDL adversary $\mathcal{A}_{\mathrm{idl}}$. For each query $\textsc{Ch}(R)$ of $\mathcal{A}_{\mathrm{idl}}$, adversary $\mathcal{A}_{\mathrm{ms}}$ simulates the response as per line 4 to 6. If $\mathcal{A}_{\mathrm{idl}}$ succeeds, it must be that

$$g^z = R_I \cdot pk^{c_{I,1}} .$$

The value of $z$ can be used to construct a forgery signature (line 3). □

## 6 Analysis of the MuSig scheme

The current three-round version of MuSig has been proposed and analyzed by both [25] and [11]. Roughly, it is the $\mathsf{BN}$ scheme with added key aggregation.

```
$\underline{\mathcal{A}_{\mathrm{ms}}^{\mathrm{H}_1}(pk)}$:
$X \leftarrow pk$ ; $(pk', sk') \leftarrow\$ \ \mathsf{MS.Kg}()$
$(I, z) \leftarrow \mathcal{A}_{\mathrm{xidl}}^{\textsc{Ch}}(pk)$ // $g^z = R_I \cdot pk^{c_{I,1}}$
$\sigma \leftarrow (R_I, z + sk' \cdot c_{I,2} \mod p)$ ; Return $((pk, pk'), m_I, \sigma)$
$\underline{\textsc{Ch}(R)}$:
$i \leftarrow i + 1$ ; $R_i \leftarrow R$ ; $m_i \leftarrow \langle i \rangle$
$c_{i,1} \leftarrow\$ \ \mathrm{H}_1((1, R_i, (pk, pk'), m_i))$ ; $c_{i,2} \leftarrow\$ \ \mathrm{H}_1((2, R_i, (pk, pk'), m_i))$
Return $c_{i,1}$
```

**Fig. 7.** Adversary $\mathcal{A}_{\mathrm{ms}}$ for Theorem 7. For an integer $i$, $\langle i \rangle$ denote the binary representation of $i$.

Let $\mathbb{G}$ be a group of prime order $p$. And let $g$ be a generator of $g$ and $\ell \geq 1$ be an integer. The formal specification of $\mathsf{MS} = \mathsf{MuSig}[\mathbb{G}, g, \ell]$ in our syntax is shown in Fig. 6. There are minimal differences between $\mathsf{MuSig}$ and $\mathsf{BN}$ and we only highlight the differences. The set $\mathsf{MS.HF}$ consists of all functions $h$ such that $h(0, \cdot) : \{0,1\}^* \to \{0,1\}^\ell$ and $h(i, \cdot) : \{0,1\}^* \to \mathbb{Z}_p$ for $i = 1, 2$. Verification is done as follows. First, an aggregate key $apk$ for the list of keys $\boldsymbol{pk} = (pk_1, \ldots, pk_n)$ is computed as $apk \leftarrow pk_1^{\mathrm{H}_2((1, \boldsymbol{pk}))} \cdots pk_n^{\mathrm{H}_2((n, \boldsymbol{pk}))}$ (line 8). Next, a single challenge is derived from the commitment $R$ and aggregate key $apk$ (line 9). The signature $(R, z)$ is valid if $g^z = R \cdot apk^c$. The second round of signing also changes accordingly to generate a valid signature (line 24 and 25).

The following gives a tight, standard-model reduction XIDL → $\mathsf{MuSig}$. Combining this with our tight AGM chain DL → IDL → XIDL from Theorems 1 and 3, we get a tight AGM reduction DL → $\mathsf{MuSig}$.

**Theorem 7.** [XIDL → $\mathsf{MuSig}$, Standard Model] *Let $\mathbb{G}$ be a group of prime order $p$. Let $g$ be a generator of $\mathbb{G}$ and $\ell \geq 1$ be an integer. Let $\mathsf{MS} = \mathsf{MuSig}[\mathbb{G}, g, \ell]$ be the associated MuSig multi-signature scheme. Let $\mathcal{A}_{\mathrm{ms}}$ be an adversary for game $\mathbf{G}_{\mathsf{MS}}^{\mathrm{ms\text{-}uf}}$ of Figure 5. Assume the execution of game $\mathbf{G}_{\mathsf{MS}}^{\mathrm{ms\text{-}uf}}$ with $\mathcal{A}_{\mathrm{ms}}$ has at most $q_0, q_1, q_2, q_{\mathrm{s}}$ distinct queries to $\mathrm{H}_0, \mathrm{H}_1, \mathrm{H}_2, \mathrm{NS}$, respectively, and the number of parties (length of verification-key vector) in queries to* NS *and* Fin *is at most $n$. Let $\alpha = q_{\mathrm{s}}(4q_0 + 2q_1 + q_{\mathrm{s}}) + 2q_1q_2$ and $\beta = q_0(q_0 + n)$. Then we can construct an adversary $\mathcal{A}_{\mathrm{xidl}}$ for game $\mathrm{Gm}_{\mathbb{G},g,q_2,q_1}^{\mathrm{xidl}}$ such that*

$$\mathbf{Adv}_{\mathsf{MS}}^{\mathrm{ms\text{-}uf}}(\mathcal{A}_{\mathrm{ms}}) \leq \mathbf{Adv}_{\mathbb{G},g,q_2,q_1}^{\mathrm{xidl}}(\mathcal{A}_{\mathrm{xidl}}) + \frac{\alpha}{2p} + \frac{\beta}{2^\ell} \,. \tag{5}$$

*The running time of $\mathcal{A}_{\mathrm{xidl}}$ is about that of the execution of game $\mathbf{G}_{\mathsf{MS}}^{\mathrm{ms\text{-}uf}}$ with $\mathcal{A}_{\mathrm{ms}}$. Furthermore, adversary $\mathcal{A}_{\mathrm{xidl}}$ is algebraic if adversary $\mathcal{A}_{\mathrm{ms}}$ is.*

We remark that the values of $q_1$ and $q_2$ above arise from the number of queries to $\mathrm{H}_1$ and $\mathrm{H}_2$ made in the execution of $\mathbf{G}_{\mathsf{MS}}^{\mathrm{ms\text{-}uf}}(\mathcal{A}_{\mathrm{ms}})$. As a result, the appearance of $q_1$ and $q_2$ has their orders "switched" compared to in Section 3. With the above theorem, we can now derive an upperbound $\mathbf{UB}_{\mathsf{MS}}^{\mathrm{ms\text{-}uf}}(t, q, q_{\mathrm{s}}, p)$ of the advantage of any MS adversary with running time $t$, making $q$ queries to H, and $q_{\mathrm{s}}$ signing interactions. We take $\ell \approx \log_2(p)$ and assume that $q_{\mathrm{s}} \leq q \leq t \leq p$. Additionally, we assume that the advantage of any XIDL adversary with running time $t$ is at most $t^2/p$ (as justified by Theorem 4). We obtain $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_s, p) \leq t^2/p$ as shown in Fig. 1.

We again describe the reduction at a high level and defer the full proof to [4]. First, the reduction adversary $\mathcal{A}_{\text{xidl}}$ receives group element $X$ from game $\text{Gm}^{\text{xidl}}_{\mathbb{G}g,q_2,q_1}$ and runs $\mathcal{A}_{\text{ms}}$ with the target public key set to $X$. Similar to the proof of Theorem 5, our adversary needs to simulate the signing oracles $\text{NS}, \mathsf{Sign}_1, \mathsf{Sign}_2$ as well as $\text{H}_0, \text{H}_1, \text{H}_2$ without knowing $\mathsf{DL}_{\mathbb{G},g}(X)$ in order to run $\mathcal{A}_{\text{ms}}$. This again relies on the zero-knowledge property of the underlying Schnorr identification scheme and the programming of $\text{H}_0, \text{H}_1, \text{H}_2$. This step is done formally in a game sequence in the full proof and incurs the additive loss in Equation (5). To turn a forgery into a break against XIDL, our adversary programs $\text{H}_1$ and $\text{H}_2$ as follows. For the $j$-th query of $\text{H}_2((k, \boldsymbol{pk}))$ where $\boldsymbol{pk}[k] = X$, the adversary first computes $S \leftarrow \prod_{i \neq k} \boldsymbol{pk}[i]^{\text{H}_2((i,\boldsymbol{pk}))}$, then obtains $e_j \leftarrow\!\!\$\ \textsc{NwTar}(S)$ before returning $e_j$ as the response for the query. We remark that this particular query of $\text{H}_2$ have created an aggregate public key $apk = \prod_{i=1}^{|\boldsymbol{pk}|} \boldsymbol{pk}[i]^{\text{H}_2((i,\boldsymbol{pk}))} = S \cdot X^{e_j}$, which is also the value of $T_j$ that is recorded in the game $\text{Gm}^{\text{xidl}}_{\mathbb{G},g,q_2,q_1}$. For each $i$-th query of $\text{H}_1((R, apk, m))$, the adversary first finds the index $j_{\text{sel}}$ of the $\text{H}_2$-query that corresponds to the input $apk$, then obtains $c_i \leftarrow\!\!\$\ \textsc{Ch}(j_{\text{sel}}, R)$ before returning $c_i$ as the response for the query. If the eventual forgery is given for these two particular queries to $\text{H}_1$ and $\text{H}_2$, meaning forgery is $\boldsymbol{pk}, m, (R, z)$ for some $z$, then the verification equation of the signature scheme says that $g^z = R \cdot apk^{\text{H}_1((R,apk,m))}$. But this matches exactly the winning condition of $\text{Gm}^{\text{xidl}}_{\mathbb{G},g,q_2,q_1}$, since $apk = T_{j_{\text{sel}}}$ and $c_i = \text{H}_1((R, apk, m))$. Hence, our adversary $\mathcal{A}_{\text{xidl}}$ can simply return $(i, z)$ to break XIDL, as long as the forgery provided by $\mathcal{A}_{\text{ms}}$ is valid.

Similar to the relation between IDL and $\mathsf{BN}$, XIDL is also tightly equivalent to the MS-UF security of $\mathsf{MuSig}$. In particular, we turn any adversary breaking XIDL into a forger against $\mathsf{MuSig}$. This means that any security justification for $\mathsf{MuSig}$ must also justify the hardness of XIDL.

**Theorem 8.** [$\mathsf{MuSig} \rightarrow$ XIDL, Standard Model] *Let $\mathbb{G}$ be a group of prime order $p$. Let $g$ be a generator of $\mathbb{G}$ and let $\ell \geq 1$ be an integer. Let $\mathsf{MS} = \mathsf{MuSig}[\mathbb{G}, g, \ell]$ be the associated MuSig multi-signature scheme. Let $q_1, q_2$ be a positive integers and $\mathcal{A}_{\text{xidl}}$ be an adversary against $\text{Gm}^{\text{xidl}}_{\mathbb{G},g,q_2,q_1}$. Then, we can construct an adversary $\mathcal{A}_{\text{ms}}$ for game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$, making no queries to* NS*, and at most* $2q_1$ *and* $2q_2$ *queries to* $\text{H}_1$ *and* $\text{H}_2$ *respectively, such that*

$$\mathbf{Adv}^{\text{ms-uf}}_{\mathsf{MS}}(\mathcal{A}_{\text{ms}}) \geq \mathbf{Adv}^{\text{xidl}}_{\mathbb{G},g,q_2,q_1}(\mathcal{A}_{\text{xidl}}) \ . \tag{6}$$

*The running time of $\mathcal{A}_{\text{ms}}$ is about that of $\mathcal{A}_{\text{idl}}$.*

*Proof (Theorem 8).* Consider the adversary given in Fig. 8. The adversary receives the target publick key $pk$ from the MS-UF game. Adversary $\mathcal{A}_{\text{ms}}$ forwards $X = pk$ as the target point and runs XIDL adversary $\mathcal{A}_{\text{idl}}$. For each

```
$\mathcal{A}_{\mathrm{ms}}^{\mathrm{H}_1,\mathrm{H}_2}(pk)$:
1 $X \leftarrow pk$ ; $(I, z) \leftarrow \mathcal{A}_{\mathrm{xidl}}^{\textsc{NwTar},\textsc{Ch}}(pk)$ ; $J \leftarrow \mathrm{TI}[I]$
2 $\sigma \leftarrow (R_I, z)$ ; Return $((pk, S_J), m_I, \sigma)$
$\textsc{NwTar}(S)$:
3 $j \leftarrow j + 1$ ; $S_j \leftarrow S$
4 $e_{j,1} \leftarrow\$ \mathrm{H}_2((1, (pk, S)))$ ; $e_{j,2} \leftarrow\$ \mathrm{H}_2((2, (pk, S)))$ ; $e_j \leftarrow e_{j,2}/e_{j,1} \mod p$
5 $apk_j \leftarrow pk^{e_{j,1}} S^{e_{j,2}}$ ; $T_j \leftarrow pk \cdot S^{e_j}$ ; Return $e_j$
$\textsc{Ch}(j_{\mathrm{sel}}, R)$:
6 $i \leftarrow i + 1$ ; $R_i \leftarrow R$ ; $m_i \leftarrow \langle i \rangle$ ; $\mathrm{TI}[i] \leftarrow j_{\mathrm{sel}}$
7 $c_i \leftarrow \mathrm{H}_1((apk_{j_{\mathrm{sel}}}, R, m_i)) \cdot e_{j_{\mathrm{sel}},1}$ ; Return $c_i$
```

**Fig. 8.** Adversary $\mathcal{A}_{\mathrm{ms}}$ for Theorem 7. For an integer $i$, $\langle i \rangle$ denote the binary representation of $i$.

query $\textsc{NwTar}(S)$ of $\mathcal{A}_{\mathrm{xidl}}$, adversary $\mathcal{A}_{\mathrm{ms}}$ uses $S$ as a public key to generate the aggregate key $apk$ for the list $(pk, S)$. By construction, the $j$-th target $T_j$ for the XIDL game is related to $apk_j$ by $apk_j = T_j^{e_{j,1}}$. For each $\textsc{Ch}(j_{\mathrm{sel}}, R)$ query of $\mathcal{A}_{\mathrm{xidl}}$, adversary $\mathcal{A}_{\mathrm{ms}}$ programs in the $\mathrm{H}_1$ outputs corresponding to a forgery agaisnt the aggregate key $apk_{j_{\mathrm{sel}}}$ (line 6 and 7). By construction, if $\mathcal{A}_{\mathrm{xidl}}$ succeeds, it must be that

$$g^z = R_I \cdot T_J^{c_I} = R_I \cdot T_J^{\mathrm{H}_1((apk_J, R, m_i)) \cdot e_{J,1}} = R_I \cdot apk_J^{\mathrm{H}_1((apk_J, R, m_i))} .$$

Hence, adversary $\mathcal{A}_{\mathrm{ms}}$ produces a valid forgery at line 2. □

## 7 HBMS: Our new two-round multi-signature scheme

Recall that BN and MuSig are three-round schemes, and two-round schemes are desired due to blockchain applications. In this section, we introduce our new, efficient two-round multi-signature scheme supporting key-aggregation, HBMS. We first demonstrate its tight security against algebraic adversaries (Theorem 9), before justifying its security in the standard model (Theorem 10). Referring to Fig. 3, these results establish arrow 5. We refer to Fig. 2 for comparisons of HBMS against other two-round schemes.

Two-round MS scheme HBMS. The formal definition of our scheme is given in Fig. 9. HBMS has the same key generation algorithm Kg and key aggregation Ag algorithm as MuSig. We describe informally the process involved to sign a message $m$ under a vector of public keys $\boldsymbol{pk}$. In the first round, each signer $i$ samples $s_i$ and $r_i$ uniformly from $\mathbb{Z}_p$ and computes a commitment

$$T_i \leftarrow \mathrm{H}_0((\boldsymbol{pk}, m))^{s_i} \cdot g^{r_i} ,$$

which is sent to every other signer. In the second round, each signer receives the list of commitments $T_1, \ldots, T_n$ from each signer, and computes the ag-

```
$\underline{\mathsf{MS.Kg}}$:
  $sk \leftarrow_\$ \mathbb{Z}_p$ ; $pk \leftarrow g^{sk}$
  Return $(pk, sk)$
```

```
$\underline{\mathsf{MS.Vf}^{\mathrm{H}_0,\mathrm{H}_1,\mathrm{H}_2}(\boldsymbol{pk}, m, \sigma)}$:
  $(pk_1, \ldots, pk_n) \leftarrow \boldsymbol{pk}$ ; $apk \leftarrow \prod_i^n pk_i^{\mathrm{H}_2((i,\boldsymbol{pk}))}$
  $(T, s, z) \leftarrow \sigma$ ; $c \leftarrow \mathrm{H}_1((T, apk, m))$
  $h \leftarrow \mathrm{H}_0((\boldsymbol{pk}, m))$ ; Return $(g^z h^s = T \cdot apk^c)$
```

```
$\underline{\mathsf{MS.Sign}^{\mathrm{H}_0,\mathrm{H}_1,\mathrm{H}_2}(\boldsymbol{b}, \mathsf{st})}$:
  $j \leftarrow \mathsf{st.me}$ ; $n \leftarrow \mathsf{st.n}$ ; $m \leftarrow \mathsf{st.msg}$ ; $sk \leftarrow \mathsf{st.sk}$ ; $\boldsymbol{pk} \leftarrow \mathsf{st.pk}$
  $(pk_1, \ldots, pk_n) \leftarrow \boldsymbol{pk}$ ; $apk \leftarrow \prod_i^n pk_i^{\mathrm{H}_2((i,\boldsymbol{pk}))}$
  If $(\mathsf{st.rnd} = 0)$ then
      $\mathsf{st}.r[j] \leftarrow_\$ \mathbb{Z}_p$ ; $\mathsf{st}.s[j] \leftarrow_\$ \mathbb{Z}_p$
      $h \leftarrow \mathrm{H}_0((\boldsymbol{pk}, m))$ ; $\mathsf{st}.\boldsymbol{R}[j] \leftarrow g^{\mathsf{st}.r[j]}$ ; $\mathsf{st}.\boldsymbol{T}[j] \leftarrow \mathsf{st}.\boldsymbol{R}[j] \cdot h^{\mathsf{st}.s[j]}$
      $\mathsf{st.rnd} \leftarrow \mathsf{st.rnd} + 1$ ; Return $(\mathsf{st}.\boldsymbol{T}[j], \mathsf{st})$
  If $(\mathsf{st.rnd} = 1)$ then
      For all $i \neq j$ do $\mathsf{st}.\boldsymbol{T}[i] \leftarrow \boldsymbol{b}[i]$
      $\mathsf{st}.T \leftarrow \prod_{i=1}^n \mathsf{st}.\boldsymbol{T}[i]$ ; $\mathsf{st}.c \leftarrow \mathrm{H}_1((\mathsf{st}.T, apk, m))$ ; $e_j \leftarrow \mathrm{H}_2((j, \boldsymbol{pk}))$
      $\mathsf{st}.z[j] \leftarrow sk \cdot c \cdot e_j + \mathsf{st}.r[j]$ ; $\mathsf{st}.\boldsymbol{t}[j] \leftarrow (\mathsf{st}.s[j], \mathsf{st}.z[j])$
      $\mathsf{st.rnd} \leftarrow \mathsf{st.rnd} + 1$ ; Return $(\mathsf{st}.\boldsymbol{t}[j], \mathsf{st})$
  If $(\mathsf{st.rnd} = 2)$ then
      For all $i \neq j$ do $\mathsf{st}.\boldsymbol{t}[i] \leftarrow \boldsymbol{b}[i]$
      $(s, z) \leftarrow \sum_i^n \boldsymbol{t}[i]$ ; Return $((\mathsf{st}.T, s, z), \mathsf{st})$
```

**Fig. 9.** Two-round multi-signature scheme $\mathsf{MS} = \mathsf{HBMS}[\mathbb{G}, g]$ parameterized by a group $\mathbb{G}$ of prime order $p$ with generator $g$.

gregate value $T \leftarrow \prod_i T_i$. Each signer then computes the challenge value as $c \leftarrow \mathrm{H}_1((T, apk, m))$. To compute the reply, each signer $i$ computes $z_i \leftarrow r_i + sk \cdot c \cdot \mathrm{H}_2((i, \boldsymbol{pk}))$ and sends $(s_i, z_i)$ to every other signer. Finally, any signer can now compute the final signature as $(T, s, z)$ where $s = \sum_i s_i$ and $z = \sum_i z_i$. To verify a signature $(T, s, z)$ on $(\boldsymbol{pk}, m)$, the equation

$$g^z \cdot \mathrm{H}_0((\boldsymbol{pk}, m))^s = T \cdot apk^{\mathrm{H}_1((T, apk, m))} ,$$

must hold, where $apk = \prod_{i=1}^{|\boldsymbol{pk}|} \boldsymbol{pk}[i]^{\mathrm{H}_2((i,\boldsymbol{pk}))}$. Compared to $\mathsf{MuSig}$, the verification equation of $\mathsf{HBMS}$ involves an additional power of $\mathrm{H}((\boldsymbol{pk}, m))$ (hence the name $\mathsf{HBMS}$, or "Hash-Base Multi-Signature").

<u>Tight security against algebraic adversaries.</u> We first show that $\mathsf{HBMS}$ is tightly MS-UF-secure against algebraic adversaries.

**Theorem 9.** [DL → $\mathsf{HBMS}$, AGM] *Let $\mathbb{G}$ be a group of prime order $p$ with generator $g$. Let $\mathsf{MS}$ be the $\mathsf{HBMS}[\mathbb{G}, g]$ scheme. Let $\mathcal{A}_{\mathrm{ms}}^{\mathrm{alg}}$ be an algebraic adversary for game $\mathbf{G}_{\mathsf{MS}}^{\text{ms-uf}}$ of Figure 5. Assume the execution of game $\mathbf{G}_{\mathsf{MS}}^{\text{ms-uf}}$ with $\mathcal{A}_{\mathrm{ms}}$ has at most $q_1, q_2$ distinct queries to $\mathrm{H}_1, \mathrm{H}_2$, respectively. Then we can construct*

*an adversary $\mathcal{A}_{\mathrm{dl}}$ for game $\mathsf{DL}_{\mathbb{G},g}$ such that*

$$\mathbf{Adv}^{\text{ms-uf}}_{\mathsf{MS}}(\mathcal{A}^{\mathrm{alg}}_{\mathrm{ms}}) \leq \mathbf{Adv}^{\mathrm{dl}}_{\mathbb{G},g}(\mathcal{A}_{\mathrm{dl}}) + \frac{(q_1+1)q_2}{p} \ . \tag{7}$$

*The running time of $\mathcal{A}_{\mathrm{dl}}$ is about that of the execution of game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ with $\mathcal{A}^{\mathrm{alg}}_{\mathrm{ms}}$.*

Above, a reduction is given directly from DL, and there is no multiplicative loss. As before, assuming $q_{\mathrm{s}} \leq q \leq t \leq p$ and the generic hardness of DL (advantage of $t$-time adversary to be at most $t^2/p$), we derive that $\mathbf{UB}^{\text{ms-uf}}_{\mathsf{MS}}(t, q, q_{\mathrm{s}}, p) \leq t^2/p$, as shown in Fig. 2.

We give the highlevel proof sketch here and defer the full proof to [4]. Let $\mathcal{A}_{\mathrm{ms}}$ be the algebraic adversary against $\mathsf{HBMS}$. Our reduction adversary $\mathcal{A}_{\mathrm{dl}}$ sets its own target point $X$ (which it needs to obtain the discrete log of) as the target public key for $\mathcal{A}_{\mathrm{ms}}$. In order to run $\mathcal{A}_{\mathrm{ms}}$, our adversary $\mathcal{A}_{\mathrm{dl}}$ needs to be able to simulate oracles $\mathrm{NS}, \textsc{Sign}_1, \textsc{Sign}_2$ (oracles representing the honest signer) as well as random oracles $\mathrm{H}_0, \mathrm{H}_1, \mathrm{H}_2$. We first tackle the problem of simulating the honest signer without knowledge of the corresponding secret key. This is done by programming of random oracle $\mathrm{H}_0$. Suppose for $\boldsymbol{pk}, m$, we set $\mathrm{H}_0((\boldsymbol{pk}, m))$ to be $h = g^{\alpha} pk^{\beta}$ for some $\alpha, \beta \neq 0 \in \mathbb{Z}_p$ (whose exact distribution will be specified later). When the adversary interacts with the honest signer, the honest signer must first provide some commitment $T \in G$ (in the output of NS), then later produce $z, s \in \mathbb{Z}_p$ (in the output of $\textsc{Sign}_1$) such that

$$g^z h^s = T \cdot pk^c \ , \tag{8}$$

where $c \in \mathbb{Z}_p$ is some challenge value (that is derived using the random oracle and the responses of the adversary). To do this, our adversary set commitment $T = g^a h^b$ for $a, b \leftarrow\!\!\mathrm{s}\ \mathbb{Z}_p$. It shall be convenient to express $pk$ in terms of $g$ and $h$ as well. Note that as long as $\beta \neq 0$, $pk = h^{(\beta^{-1})} g^{-\alpha(\beta^{-1})}$. Since both $T$ and $pk$ are known to be of the form $g^{\star} h^{\star}$ (where $\star$ denotes some element of $\mathbb{Z}_p$), so is the group element $T \cdot pk^c$ (for any known value of $c$). Hence, the right-hand side of Equation (8) is of the form $g^z h^s$ for some values $z$ and $s$ that our adversary can compute, and our adversary can return them as response in the second round. Above, we noted that this works as long as $\beta \neq 0$. To guarantee this, we sample $\alpha \leftarrow\!\!\mathrm{s}\ \mathbb{Z}_p$ and $\beta \leftarrow\!\!\mathrm{s}\ \mathbb{Z}_p^*$ in $\mathrm{H}_0$. It remains to check that such way of simulating the honest signer is indistinguishable from the behavior of an honest signer holding the secrete key and executing the protocol. Roughly, this is because in both cases, the triple $(T, z, s)$ is uniformly distributed over $\mathbb{G} \times \mathbb{Z}_p^2$, subjected to the condition that Equation (8) holds.

Now, our adversary $\mathcal{A}_{\mathrm{dl}}$ can move onto turning a forgery from $\mathcal{A}_{\mathrm{ms}}$ into a discrete logarithm for target point $X$. Suppose adversary $\mathcal{A}_{\mathrm{ms}}$ returns forgery $(\boldsymbol{pk}, m, (T, s, z))$. Then,

$$g^z h^s = T \cdot apk^c \ , \tag{9}$$

where $apk = \prod_{i=1}^{|\boldsymbol{pk}|} \boldsymbol{pk}[i]^{\mathrm{H}_2((i, \boldsymbol{pk}))}$ and $c = \mathrm{H}_1((T, apk, m))$. Since $\mathcal{A}_{\mathrm{ms}}$ is algebraic, our adversary $\mathcal{A}_{\mathrm{dl}}$ can rewrite Equation (9) to the form $g^{\alpha_g} = X^{\alpha_X}$, which allows us to compute the discrete log of $X$ as $\alpha_g \alpha_x^{-1} \mod p$, as long as $\alpha_X$ is not zero. The full proof upperbounds the probability that $\alpha_X = 0$ to be at most

$q_1 q_2/p$. Outside of this bad event, our adversary $\mathcal{A}_{\mathrm{dl}}$ will successfully compute the value of $\mathsf{DL}_{\mathbb{G},g}(X)$ from a valid forgery.

STANDARD MODEL SECURITY OF $\mathsf{HBMS}$. We reduce the security of $\mathsf{HBMS}$ to the hardness of XIDL, with factor $q_{\mathrm{s}}$ loss. For applications, the number of signing queries $q_{\mathrm{s}}$ is much less than adversarial hash function evaluations. As a result, even though our reduction here is non-tight, the reduction loss is smaller compared to previous results for $\mathsf{BN}$, $\mathsf{MuSig}$ or other two round schemes (cf. Figure 1 and 2), at the expense of assuming the hardness of XIDL. Interestingly, due to Theorem 8, our results also state that $\mathsf{HBMS}$ is secure as long as $\mathsf{MuSig}$ is (via the reduction chain $\mathsf{MuSig} \rightarrow \mathrm{XIDL} \rightarrow \mathsf{HBMS}$), and this reduction again only losses a factor of $q_{\mathrm{s}}$ in the advantage.

**Theorem 10.** [XIDL → $\mathsf{HBMS}$, Standard Model] *Let $\mathbb{G}$ be a group of prime order $p$ with generator $g$. Let $\mathsf{MS}$ be the $\mathsf{HBMS}[\mathbb{G}, g]$ scheme given in Fig. 9. Let $\mathcal{A}_{\mathrm{ms}}$ be an adversary for game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ of Figure 5. Assume the execution of game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ with $\mathcal{A}_{\mathrm{ms}}$ has at most $q_0, q_1, q_2, q_{\mathrm{s}}$ distinct queries to $\mathrm{H}_0, \mathrm{H}_1, \mathrm{H}_2, \mathrm{NS}$, respectively. Then we can construct an adversary $\mathcal{A}_{\mathrm{xidl}}$ for game $\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_2,q_1}$ such that*

$$\mathbf{Adv}^{\text{ms-uf}}_{\mathsf{MS}}(\mathcal{A}_{\mathrm{ms}}) \leq e(q_{\mathrm{s}}+1) \cdot \mathbf{Adv}^{\mathrm{xidl}}_{\mathbb{G},g,q_2,q_1}(\mathcal{A}_{\mathrm{xidl}}) + \frac{q_1 q_2}{p} \ , \tag{10}$$

*where $e$ is the base of the natural logarithm. Adversary $\mathcal{A}_{\mathrm{xidl}}$ makes $q_2$ queries to* NWTAR *and $q_1$ queries to* CH. *The running time of $\mathcal{A}_{\mathrm{xidl}}$ is about that of the execution of game $\mathbf{G}^{\text{ms-uf}}_{\mathsf{MS}}$ with $\mathcal{A}_{\mathrm{ms}}$.*

Concretely, if we assume that XIDL is quantitatively as hard as DL, then against *any* adversary with running time $t$, making $q$ evaluations of the random oracle and making at most $q_{\mathrm{s}}$ signing queries, $\mathsf{HBMS}$ has security $(q_{\mathrm{s}}t^2+q^2)/p \approx q_{\mathrm{s}}t^2/p$.

We sketch the highlevel proof here and give the full proof in [4]. Our adversary receives the target point $X$ from the XIDL game and sets it as the target public key for adversary $\mathcal{A}_{\mathrm{ms}}$. As before, in order to run $\mathcal{A}_{\mathrm{ms}}$, we need to simulate oracles NWTAR, SIGN$_1$, SIGN$_2$ as well as $\mathrm{H}_0, \mathrm{H}_1, \mathrm{H}_2$. Recall that in the AGM proof, we can simulate the honest signer for $\boldsymbol{pk}, m$ if we set $\mathrm{H}_0((\boldsymbol{pk}, m)) = g^{\alpha}h^{\beta}$. However, this way of programming $\mathrm{H}_0$ does not facilitate in turning a forgery into a break for XIDL. Instead, we would like to program $\mathrm{H}_0((\boldsymbol{pk}, m)) = g^{\alpha}$ for the forgery $\boldsymbol{pk}, m$. To do this, we use a technique of Coron [13], which programs $\mathrm{H}_0((\boldsymbol{pk}, m))$ randomly in one of these two ways depending on a biased coin flip (with probability $\rho$ of giving 1). The reduction only succeeds if correct "guesses" are made. Specifically, we need that for every $\boldsymbol{pk}, m$ that is queried to the honest signer (in NS) then $\mathrm{H}_0((\boldsymbol{pk}, m))$ must have been programmed to be $g^{\alpha}pk^{\beta}$ (for some $\alpha$ and $\beta$), and for the forgery $\boldsymbol{pk}, m$, it must be that $\mathrm{H}_0((\boldsymbol{pk}, m)) = g^{\alpha}$ (for some $\alpha$). We can then optimize for the value of $\rho$, resulting in a multiplicative loss of $e(1+q_{\mathrm{s}})$.

Suppose adversary $\mathcal{A}_{\mathrm{ms}}$ returns a forgery $(\boldsymbol{pk}, m, (T, s, z))$ where we have previously programmed $\mathrm{H}_0((\boldsymbol{pk}, m)) = g^{\alpha}$. The verification equation say that $g^z h^s = T \cdot apk^c$. Since $h$ is just a power of $g$, the left-hand side of the verification equation is also a known power of $g$ (specifically $g^{z+\alpha\cdot s}$). This means that our adversary $\mathcal{A}_{\mathrm{xidl}}$ can proceed exactly as the reduction for $\mathsf{MuSig}$. In particular, for the $j$-th query of $\mathrm{H}_2((k, \boldsymbol{pk}))$ where $\boldsymbol{pk}[k] = X$, the adversary first computes $S \leftarrow \prod_{i \neq k} \boldsymbol{pk}[i]^{\mathrm{H}_2((i,\boldsymbol{pk}))}$, then obtains $e_j \leftarrow\!\!\$\ \textsc{NwTar}(S)$ before returning $e_j$ as the response for the query. We remark that this particular query of $\mathrm{H}_2$ have created an aggregate public key $apk = \prod_{i=1}^{|\boldsymbol{pk}|} \boldsymbol{pk}[i]^{\mathrm{H}_2((i,\boldsymbol{pk}))} = S \cdot X^{e_j}$, which is also the value of $T_j$ that is recorded in the game $\mathrm{Gm}^{\mathrm{xidl}}_{\mathbb{G},g,q_2,q_1}$. For each $i$-th query of $\mathrm{H}_1((T, apk, m))$, the adversary first finds the index $j_{\mathrm{sel}}$ of the $\mathrm{H}_2$-query that corresponds to the input $apk$, then obtains $c_i \leftarrow\!\!\$\ \textsc{Ch}(j_{\mathrm{sel}}, T)$ before returning $c_i$ as the response for the query. If the eventual forgery is given for these two particular queries to $\mathrm{H}_1$ and $\mathrm{H}_2$, meaning forgery is $\boldsymbol{pk}, m, (T, s, z)$, then the verification equation of the signature scheme says that $g^{z+\alpha\cdot s} = T \cdot apk^{\mathrm{H}_1((T,apk,m))}$ (if we programmed $\mathrm{H}_0((\boldsymbol{pk}, m))$ to be $g^{\alpha}$). Hence, our adversary $\mathcal{A}_{\mathrm{xidl}}$ can simply return $(i, z{+}\alpha{\cdot}s)$ to break XIDL, as long as the forgery provided by $\mathcal{A}_{\mathrm{ms}}$ is valid and we have made the right guesses in programming $\mathrm{H}_0$.

## Acknowledgments

We thank the ASIACRYPT 2021 reviewers for their careful reading and valuable comments.

Bellare was supported in part by NSF grant CNS-1717640 and a gift from Microsoft. Dai was supported in part by a Powell Fellowship and grants of the first author.

## References

1. H. K. Alper and J. Burdges. Two-round trip schnorr multi-signatures via delinearized witnesses. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 157–188, Virtual Event, Aug. 2021. Springer, Heidelberg. 2, 4
2. A. Bagherzandi, J. H. Cheon, and S. Jarecki. Multisignatures secure under the discrete logarithm assumption and a generalized forking lemma. In P. Ning, P. F. Syverson, and S. Jha, editors, *ACM CCS 2008*, pages 449–458. ACM Press, Oct. 2008. 2, 4, 7
3. M. Bellare and W. Dai. The multi-base discrete logarithm problem: Tight reductions and non-rewinding proofs for Schnorr identification and signatures. In K. Bhargavan, E. Oswald, and M. Prabhakaran, editors, *INDOCRYPT 2020*, volume 12578 of *LNCS*, pages 529–552. Springer, Heidelberg, Dec. 2020. 8
4. M. Bellare and W. Dai. Chain reductions for multi-signatures and the HBMS scheme. Cryptology ePrint Archive, Report 2021/404, 2021. 4, 11, 12, 13, 19, 22, 25, 26
5. M. Bellare, C. Namprempre, and G. Neven. Unrestricted aggregate signatures. In L. Arge, C. Cachin, T. Jurdzinski, and A. Tarlecki, editors, *ICALP 2007*, volume 4596 of *LNCS*, pages 411–422. Springer, Heidelberg, July 2007. 8
6. M. Bellare, C. Namprempre, D. Pointcheval, and M. Semanko. The one-more-RSA-inversion problems and the security of Chaum's blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003. 3

7. M. Bellare and G. Neven. Multi-signatures in the plain public-key model and a general forking lemma. In A. Juels, R. N. Wright, and S. De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, Oct. / Nov. 2006. 2, 3, 4, 7, 12, 13, 14, 15, 17, 18, 19, 20
8. M. Bellare and A. Palacio. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 162–177. Springer, Heidelberg, Aug. 2002. 8, 11
9. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. 9, 20
10. A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Y. Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 31–46. Springer, Heidelberg, Jan. 2003. 2
11. D. Boneh, M. Drijvers, and G. Neven. Compact multi-signatures for smaller blockchains. In T. Peyrin and S. Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 435–464. Springer, Heidelberg, Dec. 2018. 2, 3, 4, 8, 13, 20
12. D. Boneh, C. Gentry, B. Lynn, and H. Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 416–432. Springer, Heidelberg, May 2003. 8
13. J.-S. Coron. On the exact security of full domain hash. In M. Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 229–235. Springer, Heidelberg, Aug. 2000. 26
14. I. Damgård, C. Orlandi, A. Takahashi, and M. Tibouchi. Two-round $n$-out-of-$n$ and multi-signatures and trapdoor commitment from lattices. Cryptology ePrint Archive, Report 2020/1110, 2020. https://eprint.iacr.org/2020/1110. 8
15. M. Drijvers, K. Edalatnejad, B. Ford, E. Kiltz, J. Loss, G. Neven, and I. Stepanovs. On the security of two-round multi-signatures. In *2019 IEEE Symposium on Security and Privacy*, pages 1084–1101. IEEE Computer Society Press, May 2019. 2, 4, 7, 14
16. R. El Bansarkhani and J. Sturm. An efficient lattice-based multisignature scheme with applications to bitcoins. In S. Foresti and G. Persiano, editors, *CANS 16*, volume 10052 of *LNCS*, pages 140–155. Springer, Heidelberg, Nov. 2016. 8
17. G. Fuchsbauer, E. Kiltz, and J. Loss. The algebraic group model and its applications. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, Aug. 2018. 1, 3, 4, 5, 9
18. G. Fuchsbauer, A. Plouviez, and Y. Seurin. Blind schnorr signatures and signed ElGamal encryption in the algebraic group model. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 63–95. Springer, Heidelberg, May 2020. 5, 10
19. L. Harn. Group-oriented (t, n) threshold digital signature scheme and digital multisignature. *IEE Proceedings-Computers and Digital Techniques*, 141(5):307–313, 1994. 2
20. K. Itakura and K. Nakamura. A public-key cryptosystem suitable for digital multisignatures. *NEC Research & Development*, (71):1–8, 1983. 2
21. E. Kiltz, D. Masny, and J. Pan. Optimal security proofs for signatures from identification schemes. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 33–61. Springer, Heidelberg, Aug. 2016. 4, 6, 8, 10, 12

22. C.-M. Li, T. Hwang, and N.-Y. Lee. Threshold-multisignature schemes where suspected forgery implies traceability of adversarial shareholders. In A. D. Santis, editor, *EUROCRYPT'94*, volume 950 of *LNCS*, pages 194–204. Springer, Heidelberg, May 1995. 2
23. S. Lu, R. Ostrovsky, A. Sahai, H. Shacham, and B. Waters. Sequential aggregate signatures and multisignatures without random oracles. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 465–485. Springer, Heidelberg, May / June 2006. 2
24. C. Ma, J. Weng, Y. Li, and R. Deng. Efficient discrete logarithm based multi-signature scheme in the plain public key model. *Designs, Codes and Cryptography*, 54(2):121–133, 2010. 2, 7
25. G. Maxwell, A. Poelstra, Y. Seurin, and P. Wuille. Simple schnorr multi-signatures with applications to bitcoin. *Designs, Codes and Cryptography*, 87(9):2139–2164, 2019. 2, 3, 4, 13, 20
26. S. Micali, K. Ohta, and L. Reyzin. Accountable-subgroup multisignatures: Extended abstract. In M. K. Reiter and P. Samarati, editors, *ACM CCS 2001*, pages 245–254. ACM Press, Nov. 2001. 2
27. J. Nick, T. Ruffing, and Y. Seurin. MuSig2: Simple two-round Schnorr multi-signatures. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 189–221, Virtual Event, Aug. 2021. Springer, Heidelberg. 2, 4
28. J. Nick, T. Ruffing, Y. Seurin, and P. Wuille. MuSig-DN: Schnorr multi-signatures with verifiably deterministic nonces. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *ACM CCS 2020*, pages 1717–1731. ACM Press, Nov. 2020. 2, 4, 7
29. K. Ohta and T. Okamoto. A digital multisignature scheme based on the Fiat-Shamir scheme. In H. Imai, R. L. Rivest, and T. Matsumoto, editors, *ASIACRYPT'91*, volume 739 of *LNCS*, pages 139–148. Springer, Heidelberg, Nov. 1993. 2
30. P. Paillier and D. Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In B. K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 1–20. Springer, Heidelberg, Dec. 2005. 4
31. D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000. 4
32. L. Rotem and G. Segev. Tighter security for schnorr identification and signatures: A high-moment forking lemma for $\Sigma$-protocols. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 222–250, Virtual Event, Aug. 2021. Springer, Heidelberg. 8
33. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, Jan. 1991. 10
34. V. Shoup. Lower bounds for discrete logarithms and related problems. In W. Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997. 3, 4, 5, 10
35. E. Syta, I. Tamas, D. Visher, D. I. Wolinsky, P. Jovanovic, L. Gasser, N. Gailly, I. Khoffi, and B. Ford. Keeping authorities "honest or bust" with decentralized witness cosigning. In *2016 IEEE Symposium on Security and Privacy*, pages 526–545. IEEE Computer Society Press, May 2016. 2, 7