J Cryptol (2016) 29:729–774
DOI: 10.1007/s00145-015-9208-2

Journal of CRYPTOLOGY

Unconditionally Anonymous Ring and Mesh Signatures*

Xavier Boyen
Queensland University of Technology (QUT), Brisbane, QLD, Australia
xb@boyen.org

Communicated by Kenny Paterson.

Received 2 July 2008
Online publication 21 September 2015

Abstract. We generalize the ring signature primitive into the more general notion of mesh signature. Ring signatures are anonymous signatures made by someone who wishes to hide in the anonymity of a larger crowd. All that the signer needs to assemble such a virtual crowd is her own private key and the public keys of the other members. The crowd composition is all that the verifier will be able to see. In a sense, a ring signature expresses an anonymous endorsement of a message by a disjunction of signers. Mesh signatures generalize this notion by allowing the combination of "atomic" (i.e., regular) signatures, by one or multiple signers from an arbitrary larger crowd, into virtually any monotone "endorsement formula" with much more expressive power than a simple disjunction. The verifier sees only that the endorsement is valid for the stated formula, not how the formula is satisfied. As a special case, mesh signatures extend the ring signature functionality to certificate chains. This is useful when the anonymity-seeking signer wishes to hide in a crowd comprising uncooperative people who do not even have a published signature verification key on record. We give an efficient linear-size construction based on bilinear maps in the common random string model. Our mesh signatures achieve everlasting perfect anonymity—an imperative for the archetypical whistle-blowing use case of ring signatures—and, as a special case, yield the first unconditionally anonymous ring signatures without random oracles or trusted setup authorities. Non-repudiation is achieved from a mild extension of the SDH assumption, named Poly-SDH, which we introduce and justify meticulously.

Keywords. Cryptography, Bilinear maps, Digital signatures, Anonymous signatures, Everlasting privacy.

1. Introduction

Ring signatures, introduced in [37], are pseudonymous signatures that are issued in the name of a "ring" of users, and created by one of them without the participation of the

The author acknowledges funding under the ARC Discovery Project, grant number DP140103885. The author is presently supported under the ARC Future Fellowship, award number FT140101145.

*Expanded version of [10]

others, in a way that preserves the instigator's anonymity. The canonical application is for an individual "to leak a secret" non-repudiably on behalf of a crowd. Technically, ring signatures can be viewed as a witness-indistinguishable disjunction of regular signatures, but because of this, only people who have previously published a verification key are eligible to be conscripted into such a crowd. Ring signatures can thus only ever implicate individuals who, by the very act of publishing their key, are acquiescing to belonging in a ring scheme.

Mesh signatures generalize this notion from a mere disjunction to an arbitrarily complex monotone access structure, i.e., a logic formula with nested gates such as "And" ($\wedge$), "Or" ($\vee$), and "Threshold" ($\geq_t$), but without negation. The inputs to the formula are atomic statements of the form "User X says Y" to which we can assign the truth value "True" ($\top$) or "False" ($\bot$). We represent a monotone access structure formula as a tree, denoted Υ, where each leaf corresponds to one input atomic statement, each interior node including the root corresponds to one gate, and each node's output link indicates the truth value of the subformula corresponding to the subtree that it defines (hence, the root node's output link corresponds to the value of Υ). Such an access structure can be satisfied using different combinations of input values, which is to say, different truth-value assignments to the input statements. The mesh signature asserts that the entire tree evaluates to $\top$, without revealing anything else about the truth values of its inputs.

To create a mesh signature corresponding to a particular formula, one only needs a "satisfying set" of atomic signatures: Namely, a set of signatures on atomic statements that together suffice to satisfy the formula. Once created, a mesh signature does not reveal the particular set of atomic signatures that was used to create it. Furthermore, atomic signatures can be generated independently of each other and without regard to the formula(s) in which they are (or are not) intended to appear. In particular, atomic signatures need not be fresh: They can be reused indefinitely many times without the participation of the original signer—e.g., in a PKI, that would be the difference between merely having a signed certificate and having oracle access to a CA (certification authority).

The central result of this paper is thus a (constructive) proof of the following informal theorem: Suppose that a monotone formula Υ can be satisfied merely by setting to true all the input wires corresponding to atomic statements known to be true (because for each, we have an atomic signature saying so). Then, this set of atomic signatures can be efficiently transformed into a "mesh signatures," of size and complexity linear in the length of Υ, that signs Υ without leaking any information about its genesis. (Technically, we also require that all the clauses of the formula take distinct signers, meaning that no two atomic statements referencing the same verification key can appear the inputs. We can mitigate this technicality by defining for each user a virtual signature key pair, consisting of a number of distinct but equivalent real key pairs. Signers would sign using all their keys, and a signature is accepted if it verifies under any one of the keys.)

1.1. *Toy Examples*

We give two simple examples showing that the added expressiveness of mesh signatures is useful in the context of ring signatures proper—whose purpose, we recall, is to sign a message under cover of anonymity of a larger crowd without seeking its consent [37].

As a first illustration, we show a way to satisfy this lofty goal even if the members of the crowd are deliberately avoiding "conscription" by not publishing their keys. With traditional ring signatures, only users whose public keys are on the record can be made part of a signing ring. With mesh signatures, we sidestep that restriction, with the simple device of faking both the missing keys and their entire certificate chains, all the way up to known certification authorities, as needed. The technique will be more apparent on a concrete example.

Example 1. (Conscription of unwilling and unwitting ring members)

$$\sigma = \left[VK_{Alice}\colon Msg_1\right] \text{ or } \left(\left[VK_{CertAuth}\colon \left(\text{“}Bob,\text{”} VK_{Bob}\right)\right] \text{ and } \left[VK_{Bob}\colon Msg_2\right]\right).$$

Here, *Alice* is able to create σ using only the private key corresponding to VK_{Alice}, because the whole formula can be satisfied merely by satisfying the left-hand disjunct. In creating the right-hand disjunct, she will want to reference an actual CA public key to serve as a verifiable "anchor," but for *Bob*, she can make up a fake public key if she does not know his real one: The conjunction in σ's right-hand disjunct attributes a certified public key to *Bob* and then uses it to authenticate *Bob*'s endorsement of Msg_2. We see that *Alice* has convincingly conscripted *Bob* in a ring-like signature, not only without needing Bob's signature but even without Bob having ever had a private signing key to begin with.
Conversely, Bob could have created σ himself, by satisfying the right-hand disjunct by providing two atomic signatures, namely CertAuth's signature on VK_{Bob} and Bob's signature on Msg_2. Either way, our construction of σ ensures that it has the exact same distribution in either case, making it impossible for a third party to determine which one of Alice or Bob created it.

As a second illustration, we show how to facilitate the creation of anonymously signed messages with increased authority, e.g., for whistle-blowing purposes, by allowing endorsements that carry the weight of multiple signers.

Example 2. (Simple multiparty threshold ring signatures)

$$\sigma = \text{2-out-of-3 in } \{[CEO\colon \textit{secret-memo}]\,,\ [CFO\colon \textit{secret-memo}]\,,\ [COO\colon \textit{secret-memo}]\}\,.$$

The unconditional anonymity of mesh signatures guarantees that, as long as the whole signature σ is valid, there is no way to tell which two of the possible three atomic signatures were used to construct σ, thereby protecting the identity of the leakers.
Naturally, threshold gates like this can be fed entire certificate chains as in the previous example, allowing "keyless users" to be conscripted into this kind of multiparty ring signatures.

In general, the added expressiveness of mesh signatures over ring signatures will have useful benefits, even if we restrict ourselves to the typical applications of the latter. For instance, the crucial ability, demonstrated in Example 1, of mesh-based ring signature to conscript anyone, even users with keys of record, provides two desirable consequences:

1. For the whistle-blower, archetypal user of ring signatures, mesh signatures remove perhaps the biggest obstacle to their practical use, which is that in the real world, people will generally have neither a key of record nor the desire to acquire one—especially if doing so puts them at risk of being suspected of subversive activities.
2. For the average citizen, the mere theoretic practicality of mesh signatures paradoxically removes the reason stated above for refusing to embrace cryptography for mundane purposes (such as routine message signing, which requires publishing a key). Indeed, since with mesh signatures *anyone* is susceptible to be conscripted into a ring signature without either knowledge or consent, it does no longer help to shun cryptography as a personal choice to skirt such a possibility.

To make the use of certificate chains truly believable, it is important that mesh signatures be "modular," or constructible non-interactively from constituent atomic signatures reusable indefinitely. Indeed, if one plans to use an actual key certificate as part of a bigger signature (as Bob did in Example 1 above), one should be able (though not required) to reuse that same certificate more than once.

It is also important that the atomic signatures have nothing special about them specifically for mesh purposes, and furthermore be of compelling use *sui generis*, regardless of ring or mesh applications.

Anticipating on the subsequent sections, we note that our scheme satisfies both requirements. Specifically, atomic signatures are ordinary Boneh–Boyen "short signatures" [5] set in a common reference bilinear group (chosen at random, without secret or trapdoor, permanent or ephemeral). Naturally, each signer generates his or her own keys independently within that common group.

1.2. *Related Work*

The original ring signature primitive was defined in [37], to enable secret leaking that is at once authenticated (by a crowd) and anonymous (within the crowd). While that construction [37] was set in the ideal cipher model, a number of alternatives have subsequently been proposed, based on bilinear pairings [8], discrete logarithms [31], factoring (Strong RSA specifically) [24], or hybrids [1]; all these constructions are set in the random oracle model. Most have linear size in the ring membership count, except [24] which squeezes it all in constant size using accumulators in the random oracle model, and [15] which first managed to drop below the linear size in the standard model.

A number of existing protocols bear similarities with our new primitive. Perhaps the first such scheme is an anonymous authentication protocol of [23] that supports access structures and can be turned into a signature using the Fiat–Shamir heuristic. Another is an interactive anonymous authentication protocol, called deniable ring authentication [36], that combines the anonymity of ring signatures with the non-transferability of deniable authentication [26] and supports threshold and access structures. Among specific constructions in the random oracle model, we note the distributed ring signatures of [32] which let coalitions of users cooperate in an interactive signing protocol, and the hierarchical identity-based ring signatures of [42], which add signer ambiguity to the notion of hierarchical identity-based signature. Limited forms of identity-based ring signatures have also been studied in [3] and analyzed in [27].

In the category of ring signatures with expanded capabilities, we mention the threshold ring signatures of [12] the threshold identity-based signatures of [19] and the 1-out-of-n identity-based ring signature of [22]; the latter could be said to provide a mesh-like expressivity comparable to that of an n-ary disjunction of binary conjunctions. In general, identity-based ring signatures, by not requiring users to set up their key pairs in advance, do provide a flavor of forcible enrollment much closer to that of mesh signatures than that of (non-identity-based) ring signatures, as discussed in [20]. The main difference is that, in an identity-based ring scheme, all keys and certificates must emanate from the same central authority, which may stretch the limits of plausibility depending on the application. Additionally, we mention that mesh signatures could in principle be realized using signatures of knowledge [16], which allow the knowledge of a witness to an NP statement to serve as a signing key, in the common random string model.

Another related notion that has received much attention is that of group signatures, originally introduced in [17], which also provides for the anonymous creation of signatures on behalf of a crowd. The main difference is that group signatures require the anonymity to be revocable by a group manager, who also controls enrollment into the group. Group membership is often immutable although this restriction has been relaxed in [13]. There exist efficient constant-size group signature schemes, with random oracles [7], from interactive assumptions [2], and in the standard model [11]. See also [29] for a construction of theoretic interest with a strong proof of security.

Efficient ring signature constructions without random oracles have also been proposed recently, such as [4,21], and [38]. The construction of [21] uses bilinear groups and is efficient, but relies on a curious hardness assumption for which no justification is offered. The results of [4] include a scheme of theoretic interest from non-interactive Zaps [25], but also two efficient constructions (based on [14] or [41] signatures) for rings of size two, and a discussion of security models for ring signatures. Last but not least, [15] manages to combine square root size and full anonymity in a basic ring signature construction.

Probably, the most closely related to the present work is the ring signature scheme of [38] which can efficiently create linear-size ring signatures in the "trusted parameters" model; unforgeability is based on computational Diffie–Hellman and anonymity on the decisional subgroup [9] assumption. Because of the latter, the scheme requires a bilinear map in a group of composite order with a hidden factorization; such a group is set up explicitly by a central authority, which afterward must erase the factorization to ensure anonymity. It is possible to tweak their scheme, using ideas from [30], to base anonymity on the decisional linear [7] assumption, which would no longer require secret-coin *trusted parameters* (TP) but only a public-coin *common random string* (CRS), as in our scheme; however, anonymity would still remain computational. The main advantage of [38] over our ring scheme is that unforgeability rests on a weaker assumption.

2. Definitions and Security Models

Intuitively, a mesh signature is a non-interactive witness-indistinguishable proof that some monotone boolean expression Υ is true, where each input of Υ is notionally labeled with a key and message pair and is true only if the mesh signer is in possession of a valid atomic signature on the stated message under the stated key.

A mesh signature scheme should satisfy two security properties. First, it should be anonymous (ideally, unconditionally so), i.e., it should not reveal what assignment to the inputs of Υ caused it to be satisfied. Second, it should be unforgeable, i.e., the creation of a valid mesh signature must be predicated on the possession of a set of valid atomic signatures sufficient to satisfy Υ.

2.1. *Recursive Mesh Signature Specification*

We use ℓ to denote the number of atomic clauses allowed in any given formula (in a ring signature, this would be equal to the maximum number of users in any given ring). Let Υ be the expression generated by the following grammar, with propositional-logic semantics, under the restriction that, for each $i = 1, \ldots, \ell$, the production $\text{EXPR} ::= L_i$ corresponding to the symbol L_i be used at most once (in other words, no L_i may appear more than once in the written expression of Υ):

$$
\begin{array}{lll}
\text{EXPR} ::= L_1 \mid \cdots \mid L_\ell & & \text{input symbols (these productions are single-use each)} \\
\quad \mid \ \geq_t \{\text{EXPR}_1, \ldots, \text{EXPR}_m\} & & t\text{-out-of-}m \text{ threshold, with } 1 < t < m \\
\quad \mid \ \wedge\{\text{EXPR}_1, \ldots, \text{EXPR}_m\} & & m\text{-wise conjunction, with } 1 < m \\
\quad \mid \ \vee\{\text{EXPR}_1, \ldots, \text{EXPR}_m\} & & m\text{-wise disjunction, with } 1 < m
\end{array}
$$

Equivalently, we call Υ an "arborescent monotone threshold circuit" with ℓ Boolean inputs $L_1, \ldots, L_\ell$ and one Boolean output denoted $\Upsilon(L_1, \ldots, L_\ell)$. It is apparent by induction that Υ is always a non-trivial monotone function of its inputs and, in particular, $\Upsilon(\bot, \ldots, \bot) = \bot$ and $\Upsilon(\top, \ldots, \top) = \top$.

We use expressions of this form to state the meaning of mesh signatures. The signer specifies the circuit Υ and assigns to each symbol L_j an atomic proposition $\left[VK\colon Msg\right]$ to convey the meaning: "This is *Msg* signed under *VK*." The mesh signature then simply expresses that $\Upsilon(L_1, \ldots, L_\ell) = \top$ holds for the stated interpretation of the L_i (without revealing their individual truth values). For the example in the introduction, $\Upsilon = L_1 \vee (L_2 \wedge L_3)$ where L_1 denotes $\left[VK_{Alice}\colon Msg_1\right]$, *etc*.

Multiplicity of Keys. As mentioned, we require that no public key appears more than once in the clauses of Υ, i.e., for any two distinct $L_i = \left[VK_i\colon Msg_i\right]$ and $L_j = \left[VK_j\colon Msg_j\right]$ appearing in Υ, we have $VK_i \neq VK_j$.

To mitigate this technicality, we expressly allow users to own multiple keys, which means that expressions Υ with a multiple clauses involving the same signer can be constructed. This is perhaps primarily intended for certificate authorities, which could be asked to sign the same certificate under several published keys, any of which deemed sufficient for verification.

2.2. *Anonymity Model*

Since the motivating application of ring and mesh schemes is to leak secrets, it is crucial that anonymity be *unconditional* and *everlasting*, subsequently to the exposure of all secrets, for the long-term peace of mind of the signer. We thus insist on perfect (i.e., information-theoretic) anonymity, even upon prior disclosure of the signer's and every user's secret keys. Moreover, since a ring or mesh signature will normally refer to third-party keys (e.g., published keys from users conscripted into the ring), it is important that information-theoretic anonymity shall apply, even against an adversary who chooses third-party keys and knows the corresponding secret keys.

The strongest notion of anonymity defined in [4], "anonymity against full key exposure," in the context of ring signatures, requires that the signer remain anonymous following full exposure of all the private keys, after their use. It is, however, insufficient for our requirements because it does not allow the keys to be chosen by the adversary and provides anonymity when the private keys are only revealed *a posteriori*.

We remedy this situation by proposing the following, very simple but very strong definitions of anonymity. The first definition captures all that a user could normally wish for. The second definition is even stronger and captures what we can actually achieve.

Definition 3. (*Unconditional signer anonymity*) Formally, we say that a ring (resp., mesh) signature scheme is *unconditionally anonymous* if the identity of the signer (resp., the signing coalition) is, conditionally on the signature formula and all the public keys and messages referenced in its clauses, statistically independent of the corresponding private keys and the common reference string.

Definition 4. (*Ultimate signer anonymity*) As an extreme strengthening of the anonymity definition, we say that a ring (resp., mesh) signature scheme is *ultimately anonymous* if, conditionally on the signature formula and the information (i.e., public keys and messages) contained therein, the identity of the signer (resp., signing coalition) is statistically independent of all information instantiated in the scheme (i.e., public or secret, permanent or ephemeral).

The latter version is very strong. It also paradoxically entails that the signers' identities be (conditionally) independent of the very random coins used to make the signature—but could that be, as the coins do not merely leak but even "prove" who the signers are? The resolution is the information-theoretic nature of the definition. For each possible way to arrive at the observed signature, there exists a corresponding set of random coins "proving" it, either in reality or in counterfactually, that one could in principle divine—even though of course it may be computationally intractable to do so.

Another paradox, more germane to practical concerns, is that the strong information-theoretic definitions above only provide unassailable anonymity in the *asymptotically long term* and within the confines of the model. For example, leaked coins will always be deemed *prima facie* evidence of the true signers if it is indeed intractable to find coins providing a counterfactual explanation. More generally, side-channel evidence should be expected to remain convincing until such time as it is no longer unreasonable to consider that it might have been faked. These concerns, though important in practice, are out of scope of our discussion.

Anonymity, Unlinkability, and Randomization. A mesh signature σ is most generally constructed from a set of atomic signatures σ_i for the selected clauses $\left[VK_i : Msg_i\right]$ assigned the truth value $\top$. As we will see, mesh signatures are not unique, in part because of the "extrinsic" randomization due to the mesh signing process, but also because the atomic signatures themselves bring to the table their own "intrinsic" randomization (which will have to be faked for all clauses set to $\perp$). Therefore, it is a legitimate concern to wonder how the intrinsic randomization associated with the atomic signatures interacts with the requirements of unconditional anonymity.

To fix ideas, suppose that a mesh signature σ contains a (valid or invalid) atomic signature on a clause $\left[VK_1 : Msg_1\right]$ with intrinsic randomization t_1. Later, someone exposes an atomic signature σ_1 on the same clause and with the same randomization t_1. Is this evidence that σ_1 was used to construct σ, thereby putting the anonymity of the mesh signer in jeopardy? The answer is 'no,' as long as the owner of VK_1 could have created σ_1 after seeing t_1. Conversely, if σ had been revealed subsequently to σ_1, one could not infer that σ_1 was used in the creation of σ, as long as σ could plausibly have been created *ex post facto* to match the exposed randomization.

In general, if the intrinsic randomization of the atomic signature can be chosen freely and is conveyed in the clear, then matching randomization merely implies awareness and not linkability, i.e., it shows that the second signature was created with knowledge of the (randomization of the) first, but not that the two were actually created from each other, or by the same signer.

2.3. *Unforgeability Model*

The strongest notion of unforgeability defined in [4], "unforgeability with respect to insider corruption," for ring signatures, gives the adversary the ability to corrupt users dynamically and include its own public keys when making ring signature queries. Since the point of mesh signatures is to implicate uncooperative users, it is judicious to allow them to choose their keys maliciously.

However, as a compromise for unconditional anonymity, we relax the fully dynamic corruption model into an enhanced static one, in which the honest users are static and created ahead of time by a challenger, and the corrupted users are under the full control of an adversary who can bring them to life dynamically. We also need to specify what constitutes a valid forgery. For ring signatures, a forgery is any signature by a ring without adversarially controlled users. For mesh signatures, however, this would be overly restrictive, since it would exclude such forgeries as,

$$\Upsilon = ([U_1 : m_1] \wedge [U_3 : m_3]) \vee ([U_2 : m_2] \wedge [U_4 : m_4]),$$

where U_1 and U_2 are honest users and U_3 and U_4 are corrupted. Since Υ nominally entails $\Upsilon' = [U_1 : m_1] \vee [U_2 : m_2]$, a forger who signs Υ lacking the imprimatur of both U_1 and U_2 should be deemed successful. The same reasoning would continue to apply if the forger legitimately obtained an atomic signature on $[U_3 : m_3]$ even though U_3 were honest. We capture these circumstances by deeming admissible any forgery on a statement Υ if there exists a well-formed (and thus non-trivial) formula Υ' that contains no clause under the forger's control and such that $\Upsilon \Rightarrow \Upsilon'$.

To see where this comes from, for all corrupted users and all issued atomic signatures, let us set the corresponding literal $L_i \leftarrow \top$, which is the most that the adversary can do in legitimacy. If Υ then evaluates to $\top$, the forgery is inadmissible; otherwise, Υ will reduce to some well-formed formula Υ' that contains non-adversarial clauses exclusively. Hence, the existence of Υ' simply demands that Υ be unsatisfiable by the volition of the adversarial users alone. We distill all of this into the following existential unforgeability game and define the adversary's advantage as the probability of outputting an admissible valid forgery.

Definition 5. (*Existential unforgeability*) We define the existential mesh signature unforgeability game as the following interaction between a challenger and an adversary.

Challenger setup: The challenger designates a number ℓ of public keys, corresponding to the honest target users under the challenger's control.
Interaction: The following occurs interactively, in any order, driven by the adversary.

Adversary setup: The adversary reveals polynomially many public keys, one at a time, corresponding to the users under the adversary's control.
Mesh signature queries: The adversary makes up to q mesh signature queries on well-formed specifications Υ_j that involve zero or more adversarial users and at least one honest user (the latter condition being imposed to avoid queries that the adversary could trivially answer completely by itself).
Atomic signature queries: The adversary also makes up to q atomic signatures queries on clauses $\left[VK_i : Msg_j\right]$ for every honest user.

The challenger accepts or responds to each request before accepting the next one. The q mesh queries and the q ℓ atomic queries may be interleaved arbitrarily.

Signature forgery: the adversary produces a forged signature whose specification Υ satisfies $\forall j$, $\Upsilon \neq \Upsilon_j$ and implies a well-formed formula Υ' on the honest users, i.e., $\Upsilon(L_1, \ldots, L_\ell, \ldots) \Rightarrow \Upsilon'(L_1, \ldots, L_\ell)$, obtained by setting to "true" ($\top$) every literal L_i whose clause $\left[VK_i : Msg_i\right]$ involves an adversarial key or matches an atomic query.

The adversary's advantage at mesh unforgeability is the probability that it wins the foregoing game (for a random choice of common reference string (during/prior the challenger setup) if applicable).

In the adversary setup, one must recognize that the adversary might try to claim some of the challenger's keys as its own (perhaps re-randomized to make it less obvious). Since the same is possible in the real world, and is readily detectable by the challenger, we take no step to forbid it, other than to require that all specifications be well formed.

Mesh Unforgeability vs. Ring Unforgeability. The mesh security model allows the forger to make arbitrary atomic signature queries on behalf of the honest users: This is because mesh signatures must be constructible from any satisfying set of atomic signatures (such as PKI certificates) without requiring the private keys.

For ring signatures, atomic signature queries are superfluous, and we can obtain a tighter proof of security without them, mainly because we reduce the number of queries from $(\ell + 1)\, q$ to just q. Hence, we define existential unforgeability for ring signatures as for mesh signatures, but without atomic signature queries (also, regular signatures can always be emulated using rings of size one). We refer to [4] for ring signature unforgeability definitions with various security requirements.

3. Framework and Computational Assumption

We write $\mathbb{F}_p$ for the finite field of prime order p and $\mathbb{F}_p^\times = \mathbb{F}_p \setminus \{0\}$ for its multiplicative group of order $p-1$. We refer as a *bilinear context* to an algorithmically useful description

of an efficiently computable and non-degenerate bilinear map $\mathbf{e}$ between a set of groups $\mathbb{G}$ and $\hat{\mathbb{G}}$ of some prime order p and given by the respective generators g and $\hat{g}$, into a third group $\mathbb{G}_t$ of the same order. Let thus $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$ be a common bilinear context, where $\mathbf{e} : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_t$ is a pairing [35]. We use the "hat-notation" (as in $\hat{g}$) to indicate that an element belongs to $\hat{\mathbb{G}}$ rather than $\mathbb{G}$.

3.1. *Review of the SDH Assumption*

The complexity assumption we shall need is inspired by the Strong Diffie–Hellman assumption proposed in [5], which we now review. The q-SDH problem in a (bilinear) group $\mathbb{G}$ is stated:

> **(Original SDH)** Given elements $g,\ g^{\alpha},\ g^{\alpha^2},\ \ldots,\ g^{\alpha^q} \in \mathbb{G}$, choose $w \in \mathbb{F}_p$ and output $(w,\ g^{1/(\alpha+w)})$.

The SDH assumption then posits that the q-SDH problem above is intractable for $q = \mathrm{O}(\mathrm{poly}(\kappa))$. What makes this assumption special is that the problem admits not one but exponentially many "independent" solutions, which are all equally hard to find hence the modified q-SDH problem:

> **(Modified SDH)** Given $g,\ g^{\alpha} \in \mathbb{G}$ and $q-1$ pairs $(w_j,\ g^{1/(\alpha+w_j)})$, output another $(w,\ g^{1/(\alpha+w)})$.

It is known from [5] that if the original q-SDH problem is hard, then it is the modified problem.

Although the SDH problem statement does not require a bilinear group, it is because the bilinear map provides an efficient Decision Diffie–Hellman procedure [33] that the correctness of an SDH solution can be decided openly. Specifically, given g and g^{α}, deciding whether $(w,\ u) = (c,\ g^{1/(\alpha+w)})$ amounts to checking the equality $\mathbf{e}(u,\ \hat{g}^{\alpha}\,\hat{g}^{w}) = \mathbf{e}(g,\ \hat{g})$ is basically a DDH test that anyone can perform from public information. The short signature scheme of [5] relies on this.

3.2. *Poly-SDH: for Better Use of the Pairing*

The verifiability of SDH solutions with a simple DDH test suggests that more general assumptions could be made, based on the observation that the pairing is a powerful tool that can be used to decide more complex relations that are not efficiently reducible to DDH. For example, a natural generalization of the SDH problem is that of finding ℓ pairs $(w_i,\ u_i = g^{r_i/(\alpha+w_i)})$ for $i = 1, \ldots, \ell$, such that $\sum_{i=1}^{\ell} r_i = 1 \pmod{p}$. Purported solutions can then be verified using the equation,

$$\prod_{i=1}^{\ell} \mathbf{e}\left(u_i,\ \hat{g}^{\alpha}\,\hat{g}^{w_i}\right) = \mathbf{e}\left(g,\ \hat{g}\right). \tag{1}$$

Clearly, when $\ell = 1$, this is identical to the SDH problem. For larger values of ℓ, the adversary is given to spread the exponent inversion task across multiple pairs, by means of linear combination.

Unfortunately, for $\ell > 1$, the problem is in fact trivial, because Eq. (1) admits spurious solutions that do not require the solver to know the secret α and invert the exponent: For example, for $\ell = 2$ the solution $w_1 = 1,\ u_1 = g,\ w_2 = 0,\ u_2 = g^{-1}$ satisfies the equality regardless of α.

To remedy the preceding problem, we change the solver's task slightly and ask that the ℓ pairs to be output involve ℓ independent secrets $\alpha_1, \ldots, \alpha_\ell$ that appear once each, i.e., find,

$$\left(w_i,\ u_i = g^{\frac{r_i}{\alpha_i + w_i}} \right) : i = 1, \ldots, \ell, \quad \text{s.t.} \quad \sum_{i=1}^{\ell} r_i = 1 \pmod{p}.$$

To decide whether a solution $((w_1,\ u_1),\ \ldots,\ (w_\ell,\ u_\ell))$ to the new problem is correct, one needs, besides the generators g and $\hat{g}$, the ℓ group elements $(\hat{g}_1,\ \ldots,\ \hat{g}_\ell) = (\hat{g}^{\alpha_1},\ \ldots,\ \hat{g}^{\alpha_\ell})$. The verification equation is then,

$$\prod_{i=1}^{\ell} \mathbf{e}\left(u_i,\ \hat{g}_i\, \hat{g}^{w_i}\right) = \mathbf{e}\left(g,\ \hat{g}\right). \tag{2}$$

Notice that (1) is a special case of (2) where $\alpha_1 = \cdots = \alpha_\ell = \alpha$; however, for the security of the assumption it is important that the α_i be independently and uniformly distributed. Despite the added variables, Eq. (2) is no more expensive to verify (but necessitates large public parameters).

Based on the previous observations, the (q, ℓ)-Poly-SDH problem can be informally stated as:

> **(Poly-SDH)** Given $g,\ g^{\alpha_1},\ \ldots,\ g^{\alpha_\ell} \in \mathbb{G}$ and $q\,\ell$ pairs $(w_{i,j},\ g^{1/(\alpha_i + w_{i,j})})$ for $1 \le i \le \ell$ and $1 \le j \le q$, choose fresh $w_1,\ \ldots,\ w_\ell \in \mathbb{F}_p$ (i.e., such that $w_i \notin \{w_{i,1}, \ldots, w_{i,q}\}$) and output ℓ pairs $(w_i,\ g^{r_i/(\alpha_i + w_i)})$ such that $\sum_{i=1}^{\ell} r_i = 1$.

The α_i and $w_{i,j}$ in the instance are drawn from a uniform distribution. The w_i and r_i are chosen by the respondent. We require that $\forall i, \forall j,\ w_i \neq w_{i,j}$, lest the task be easy. The exponents r_i need not be revealed, since Eq. (2) can establish that a solution is correct and thus that $\sum_i r_i = 1$, without having to see the r_i.

We have chosen to state the (q, ℓ)-Poly-SDH problem in a form analog to Modified SDH, rather than Original SDH. There are several justifications for this:

- the modified form results in a weaker assumption (as Original SDH implies Modified SDH);
- it has a clear input/output symmetry which simplifies the security reductions;
- its instances are more concisely stated when more than one iterator is needed (i and j);
- the modified problem form is impervious to a generic analysis described in [18], which relies on the availability of g, g^{α}, and g^{α^d} for certain d, as in Original SDH instances.

The reason why there are no undesirably easy solutions to the (q, ℓ)-Poly-SDH problem will become apparent as we prove generic hardness in Sect. 3.3. See also "Formal Poly-SDH Definitions" of Appendix for formal definitions.

3.3. *Generic Hardness of Poly-SDH*

We now take some time to explain why the Poly-SDH assumption based on Eq. (2) is plausible, unlike our first attempt from Eq. (1) that was so easily broken. We give a heuristic argument based on the impossibility of efficient generic attacks. Specifically, we show that finding a solution to the (q, ℓ)-Poly-SDH problem will require, on expectation, $\Omega(\sqrt{p/q\,\ell})$ generic-group operations.

The generic-group model [39] assumes the lack of any structure beyond that of an (Abelian) cyclic group, restricting all manipulations on group elements to the group operation and its inverse (i.e., multiplication and division if the group is written multiplicatively). In the bilinear version of the model [5], one can also compute a pairing $\mathbf{e} : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_t$, as well as an isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$ (for "type-1" and "type-2" contexts) and its inverse $\psi^{-1} : \mathbb{G} \to \hat{\mathbb{G}}$ (for "type-1" only).

Let us assume that $\mathbb{G} = \hat{\mathbb{G}}$, which only makes the attack easier.[1] Recall that the Poly-SDH instance furnishes $g,\ g^{\alpha_1},\ \ldots,\ g^{\alpha_\ell}$, and a large number of pairs $(w_{i,j},\ u_{i,j} = g^{1/(\alpha_i + w_{i,j})})$. Based on this information, the attacker must output ℓ pairs $(w_i,\ u_i = g^{r_i/(\alpha_i + w_i)})$ such that $\sum_i r_i = 1$, where w_i is distinct from all $w_{i,j}$ with the same index i.

First, notice that the pairing $\mathbf{e}$ is useful to verify a solution, but not really to find one. This is because $\mathbf{e}$ maps to $\mathbb{G}_t$, and once we have landed in $\mathbb{G}_t$ we can never leave it. Also, ψ and ψ^{-1} just model the identity function since we have already assumed that $\mathbb{G} = \hat{\mathbb{G}}$. We can thus focus on multiplication and division in the multiplicative group $\mathbb{G}$ of prime order p.

Next, observe that all the group elements that can be created from g, $\{g^{\alpha_i}\}$, and $\{g^{1/(\alpha_i + w_{i,j})}\}$ are of the form $g^{\frac{\pi(\alpha_1,\ldots,\alpha_\ell)}{\Delta}}$, where $\pi \in \mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]_{q\ell+1}$ is any multivariate polynomial in $\alpha_1, \ldots, \alpha_\ell$ of total degree at most $q\,\ell + 1$, and where Δ is the common denominator $\Delta = \prod_{i=1}^{\ell} \prod_{j=1}^{q} (\alpha_i + w_{i,j})$. (Here, we use the notation $\mathbb{F}_p[x, y]$ to denotes the ring of polynomials in x and y over $\mathbb{F}_p$, and use the shorthand notation $\mathbb{F}_p[x]_d$ to denote the set of polynomials in x and y of total degree d or less.)

We need to produce ℓ elements $u_i = g^{r_i/(\alpha_i + w_i)}$ and the corresponding w_i. Our task is thus to find ℓ polynomials $\pi_1,\ \ldots,\ \pi_\ell \in \mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]_{q\ell+1}$ such that $\pi_i/\Delta = r_i/(\alpha_i + w_i)$ for some $\sum_i r_i = 1$, i.e., such that,

$$\sum_{i=1}^{\ell} (\alpha_i + w_i)\,\pi_i = \Delta = \prod_{i=1}^{\ell} \prod_{j=1}^{q} (\alpha_i + w_{i,j}).$$

[1] Recall that, in general, we require $\mathbb{G} \simeq \hat{\mathbb{G}}$ with an efficiently computable isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$. Assuming that the reverse isomorphism ϕ^{-1} is also efficiently computable only increases the power given to the adversary, thus making the attack easier. At the extreme, assuming $\mathbb{G} = \hat{\mathbb{G}}$ makes the attack even easier because it allows the adversary not only to move elements back and forth between the isomorphic groups, but also to mix them within the same algebraic expressions, which otherwise would not be allowed in the generic-group model. As noted earlier, the three cases exist in actual realizations of bilinear groups. For this proof, we place ourselves in the case where $\mathbb{G} = \hat{\mathbb{G}}$ in order to show the soundness of our hardness assumption in the most adversary-friendly setting, which will imply the weaker results. (This also helps to simplify the notation, by (temporarily) dropping all "hats" from the expressions.)

We show that there can be no such polynomials π_i using a linear change of variable. For all $i = 1, \ldots, \ell$ and $j = 1, \ldots, q$, we define $\alpha'_i = \alpha_i + w_i$ and $w'_{i,j} = w_{i,j} - w_i$. Notice that all $w'_{i,j} \neq 0$. Our new task becomes to find ℓ polynomials $\pi'_1, \ldots, \pi'_\ell$ of degree $\leq q\,\ell + 1$ in the variables $\alpha'_1, \ldots, \alpha'_\ell$, such that,

$$\sum_{i=1}^{\ell} \alpha'_i \pi'_i = \Delta = \prod_{i=1}^{\ell} \prod_{j=1}^{q} \left(\alpha'_i + w'_{i,j}\right).$$

Clearly, all the monomials in the left-hand side have degree in $\alpha'_1, \ldots, \alpha'_\ell$ at least 1. On the other hand, all $w'_{i,j}$ are nonzero, so the right-hand side yields a non-vanishing independent (degree-0) term equal to $\prod_i \prod_j w'_{i,j} = \prod_i \prod_j (w_{i,j} - w_i) \neq 0$, which is a contradiction.

The contradiction shows that the equations above cannot be satisfied identically in $\mathbb{F}_p[\alpha'_1, \ldots, \alpha'_\ell]$ or $\mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]$, which proves that the polynomials π'_i, and thus, π_i cannot exist. A standard argument then shows that the equations can only be satisfied in $\mathbb{F}_p$ for certain assignments of $\alpha_1, \ldots, \alpha_\ell \in \mathbb{F}_p$: the polynomial roots. Since the α_i are chosen at random, we can bound the probability of hitting those roots. We find that, if $q\,\ell < \mathrm{O}(\sqrt[3]{p})$, it will take $q_G = \Omega(\sqrt{\epsilon\, p/q\,\ell})$ operations to solve (q, ℓ)-Poly-SDH with probability ϵ in generic groups of order p.

We give a precise theorem and a complete proof based on this argument in "Generic-Group Complexity of Poly-SDH" of Appendix.

3.4. *Pluri-SDH: A Weaker Assumption*

Although we will need the Poly-SDH assumption to prove security of mesh signatures, ring signatures can be based on a slightly weaker assumption, due to the lack of atomic signature queries. Recall that in the (q, ℓ)-Poly-SDH problem, we are given ℓ generators g^{α_i} as well as ℓ series of q solution pairs $(w_{i,j},\ u_{i,j} = g^{1/\alpha_i + w_{i,j}})$. Our weaker assumption is similar, except that we only give out a single series of solution pairs, conventionally for an extra generator of index $i = 0$.

We define the $(q, \ell, 1)$-Pluri-SDH problem as a relaxed version of $(q, \ell + 1)$-Poly-SDH:

> **(Pluri-SDH)** Given generators $g, g^{\alpha_0}, \ldots, g^{\alpha_\ell} \in \mathbb{G}$ and q pairs $(w_{0,j},\ g^{1/(\alpha_0 + w_{0,j})})$ for $1 \leq j \leq q$, choose fresh $w_0, \ldots, w_\ell \in \mathbb{F}_p$ and output $\ell + 1$ pairs $(w_i,\ g^{r_i/(\alpha_i + w_i)})$ such that $\sum_{i=0}^{\ell} r_i = 1$.

See also "Formal Pluri-SDH Definitions" of Appendix for formal definitions, including that of the (q, ℓ, ℓ')-Pluri-SDH problem which is stated in an obvious way for $\ell' \geq 1$.

Generic Complexity. Regarding generic complexity, we can show that for $q\,\ell < \mathrm{O}(\sqrt[3]{p})$, a generic algorithm can solve the $(q, \ell, 1)$-Pluri-SDH problem with constant probability ϵ in a generic group of prime order p only by performing $q_G = \Omega(\sqrt{\epsilon\, p/q})$ generic-group operations on expectation.

A precise theorem for $(q, \ell, 1)$-Pluri-SDH and (q, ℓ, ℓ')-Pluri-SDH is given in "Generic-Group Complexity of Pluri-SDH" of Appendix.

3.5. *Comparing SDH with Pluri-SDH and Poly-SDH*

An interesting fact about the $(q, \ell, 1)$-Pluri-SDH problem in generic bilinear groups is that it is quantitatively as difficult as the (modified) q-SDH problem: In particular, the generic lower bounds are essentially the same as those found in [5] and do not strongly depend on ℓ. In other words, allowing the opponent to make ℓ-wise linear combinations has little adverse effect on generic security, provided that care has been taken to structure the problem to rule out all of the trivial solutions. A similar comparison can be made for the full (q, ℓ)-Poly-SDH problem, except that the relevant benchmark here is the $q\ell$-SDH problem. Although we appear to lose a factor ℓ in the number of allowed queries with respect to SDH, it will be a wash if the security reduction of interest allows ℓ times as many queries, which will be the case of our mesh unforgeability simulator.

The main difference between SDH and Pluri-SDH/Poly-SDH is thus not one of hardness. It is that the former is useful in any "Gap-DH" group where the Diffie–Hellman problem has a decision procedure, while Pluri-SDH and Poly-SDH require a group with an actually computable pairing (or at least an oracle for comparing products of pairings) in order to verify its solutions.

4. Special Case: Ring Signatures

We first describe a ring signature based on Pluri-SDH as a special case of our technique. It is more efficient than most other provably secure ring signature schemes without random oracles, and the first of those schemes to offer unconditional anonymity. It is set in the "public-coin" common random string model, i.e., requiring only minimal trust for setup. The scheme is in fact very close to a ring scheme from [21], but not the proof.

Initialization: Given a security parameter κ and a public random string $K \in \{0,1\}^{\mathrm{poly}(\kappa)}$, the parties generate from K a common bilinear instance $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e}) \leftarrow \mathcal{G}(1^\kappa; K)$ and a collision-resistant hash function $H : \{0,1\}^* \to \mathbb{F}_p$ shared by all. Since $\mathbf{G}$ has prime order and no hidden structure, it can safely be generated from public coins.

The string K is also used to generate three random elements $\hat{A}_0$, $\hat{B}_0$, and $\hat{C}_0$ in $\hat{\mathbb{G}}$. These elements define a public verification key "in the sky" whose matching signing key is undefined.

For notational convenience, we suppose for now that the isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$ is efficiently computable in the instance $\mathbf{G}$, and we let $A_0 = \psi(\hat{A}_0)$, $B_0 = \psi(\hat{B}_0)$, and $C_0 = \psi(\hat{C}_0)$ in $\mathbb{G}$. This temporary restriction will be lifted later in this section.

Key generation: To create a key pair, User #i draws a triple $(a_i, b_i, c_i) \in (\mathbb{F}_p^\times)^3$ as signing key and posts $(A_i, B_i, C_i, \hat{A}_i, \hat{B}_i, \hat{C}_i) = (g^{a_i}, g^{b_i}, g^{c_i}, \hat{g}^{a_i}, \hat{g}^{b_i}, \hat{g}^{c_i}) \in \mathbb{G}^3 \times \hat{\mathbb{G}}^3$ as verification key.

In case $\psi : \hat{\mathbb{G}} \to \mathbb{G}$ is easy to compute, users publish only $(\hat{A}_i, \hat{B}_i, \hat{C}_i)$ to avoid redundancy.

Ring signature: To create a ring signature on message $m_1, \ldots, m_\ell \in \mathbb{F}_p$ attributed to a ring of ℓ users, any member of the ring would proceed as follows. W.l.o.g., suppose that the signer is User #ℓ in the ring $R = (1, \ldots, \ell)$. The signer selects

$2\ell + 1$ random integers $s_0, s_1, \ldots, s_{\ell-1}, t_0, t_1, \ldots, t_\ell \in \mathbb{F}_p$ and outputs the signature,

$$\sigma = \left(g^{s_0},\ \ldots,\ g^{s_{\ell-1}},\ \left(g \cdot \prod_{i=0}^{\ell-1} \left(A_i\ B_i^{m_i}\ C_i^{t_i} \right)^{-s_i} \right)^{\frac{1}{a_\ell + b_\ell m_\ell + c_\ell t_\ell}},\ t_0,\ \ldots,\ t_\ell \right) \in \mathbb{G}^{\ell+1} \times \mathbb{F}_p^{\ell+1},$$

where $m_1, \ldots, m_\ell$ are the messages to be signed, and $m_0 = H((1, m_1), \ldots, (\ell, m_\ell))$, a collision-resistant hash of the statement expressed by the signature.

Ring verification: To verify a signature $\sigma = (S_1,\ \ldots,\ S_\ell,\ t_1,\ \ldots,\ t_\ell)$, test the equality,

$$\prod_{i=0}^{\ell} \mathbf{e}\left(S_i,\ \hat{A}_i\ \hat{B}_i^{m_i}\ \hat{C}_i^{t_i} \right) = \mathbf{e}\left(g,\ \hat{g} \right),$$

where $R = (1, \ldots, \ell)$ is the signature ring, $m_1, \ldots, m_\ell$ are the messages being signed, and $m_0 = H((1, m_1), \ldots, (\ell, m_\ell))$.

Consistency of the algorithms is readily verified. Note that the scheme is trivially modified to force all messages $m_1, \ldots, m_\ell$ to be the same, as in the traditional definition of ring signatures.

The purpose of including in the final signature a collision-resistant hash m_0 of the ring and all the messages, ostensibly binding m_0 to the public key "in the sky," is to prevent outsiders from appending new components to an existing signature, which would otherwise give an easy forgery (though perhaps a rather benign one). The second reason is that the key "in the sky" is useful in the security proof, and lets us rely on a weaker assumption.

4.1. *Anonymity*

Independently of setup assumptions, our ring signatures have irrevocable or everlasting, perfect, unconditional anonymity (i.e., with forward security against coerced disclosure of the long-term signing keys, and the randomness that created them, of all users in the system).

Theorem 6. *The ring signature has everlasting perfect anonymity.*

Proof. See "Anonymity of the Ring Scheme" of Appendix. □

4.2. *Unforgeability*

We then have existential unforgeability in the common random string model based on our computational assumption. More precisely, we can give two alternative reductions: One establishes security in the *ring* forgery game provided that the $(q, \ell, 1)$-Pluri-SDH problem is hard; the other proves security in the more demanding *mesh*

forgery game from the hardness of $(q, \ell + 1)$-Poly-SDH. Here, we recall from Sect. 2.3 that a mesh forger can also make atomic signature queries to the honest users in addition to mesh (or ring) queries, whereas a ring forger makes no atomic queries.

We now state the ring result, which is the most appropriate in the context of ring signatures. In "Unforgeability of the Ring Scheme" of Appendix, however, we shall state and prove the stronger result instead, because parts of that proof will be reused when proving security of the full mesh scheme of Sect. 5.

Theorem 7. *The ring signature is existentially unforgeable under an adaptive attack, against a static adversary that makes no more than q adaptive ring signature queries, provided that the* $(q, \ell, 1)$*-Pluri-SDH assumption holds in* $\mathbf{G}$*, in the common random string model.*

Proof. See "Unforgeability of the Ring Scheme" of Appendix □

4.3. *Bilinearity Without Isomorphism*

Since the most general types of bilinear instance $\mathbf{G}$ may fail to provide both an efficient isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$ and an efficient sampling procedure in $\hat{\mathbb{G}}$, it is useful to modify the ring scheme in order to relax these requirements. Although it is typically safe to rely on either one or the other [28], it is easy to eliminate both requirements at once in the following way.

- First, we redefine the random key "in the sky" to consist just of A_0, B_0, and C_0, to be sampled directly in $\mathbb{G}$ from the common random seed K (skipping $\hat{\mathbb{G}}$ altogether).
- Next, we modify the group element of index 0 in the signature, replacing $g^{s_0} \in \mathbb{G}$ with $\hat{g}^{s_0} \in \hat{\mathbb{G}}$. The signature becomes, e.g., with User #ℓ as the signer: $\sigma = (\hat{S}_0, \ldots, S_\ell, t_0, \ldots, t_\ell) =$

$$\left(\hat{g}^{s_0}, g^{s_1}, \ldots, g^{s_{\ell-1}}, \left(g \cdot \prod_{i=0}^{\ell-1} (A_i B_i^{m_i} C_i^{t_i})^{-s_i} \right)^{\frac{1}{a_\ell + b_\ell m_\ell + c_\ell t_\ell}}, t_0, \ldots, t_\ell \right)$$
$$\in \hat{\mathbb{G}} \times \mathbb{G}^\ell \times \mathbb{F}_p^{\ell+1},$$

- Last, we exchange the arguments under the pairing of index 0 and amend the verification equation into,

$$\mathbf{e}\left(A_0 B_0^{m_0} C_0^{t_0}, \hat{S}_0\right) \cdot \prod_{i=1}^{\ell} \mathbf{e}\left(S_i, \hat{A}_i \hat{B}_i^{m_i} \hat{C}_i^{t_i}\right) = \mathbf{e}\left(g, \hat{g}\right).$$

It is easy to see that the security theorems continue to hold in the modified ring signature scheme. On the one hand, anonymity is unconditional and thus insensitive to the existence of some efficient algorithm for ψ or for sampling in $\hat{\mathbb{G}}$. On the other hand, unforgeability relies no more on the presence of such algorithms than on their absence, as an inspection of the proof would show.

4.4. *The Key "In the Sky"*

A (tenuous) argument can be made that having a public key "in the sky" entails a stronger flavor of CRS than the mere sharing of a bilinear instance $\mathbf{G}$ and a collision-resistant hash function H.

The crux of the argument is that, for someone who controls the CRS, it is much easier to implant a trapdoor into the public key VK_0 than to prepare $\mathbf{G}$ for the subsequent efficient computation of discrete logarithms: The former can be done by constructing VK_0 from an explicit signing key (as the simulator does in the unforgeability proof), whereas the latter might involve the infeasible pre-computation of an exponential-size lookup table for the baby-step giant-step algorithm in $\mathbf{G}$. A counterargument is that if the CRS is truly random, then all of this is equally hard for everyone.

Either way, both flavors of the CRS model—with or without a plausible trapdoor—seem more palatable than the TP model—with its inescapable third-party secrets (ephemeral or permanent). We can even eliminate the "key in the sky" VK_0 altogether, but omit the details.

5. General Case: Mesh Signatures

We now describe our mesh signature scheme, based on the Poly-SDH assumption. We proceed in stages: We first define a few useful notions, which we then use to describe the actual system.

5.1. *Flattened Mesh Representation*

Recall that a mesh signature is characterized by an expression Υ generated by the grammar,

$$\Upsilon ::= N$$
$$N ::= L_1 \mid \ldots \mid L_\ell \mid \geq_t \{N_1, \ldots, N_m\} \mid \wedge \{N_1, \ldots, N_m\} \mid \vee \{N_1, \ldots, N_m\}.$$

To harmonize the notation with the scheme description, we need to consider an extra literal L_0 whose meaning is unimportant for now, and let $\tilde{\Upsilon}$ be as above with $\ell + 1$ input literals $L_0, \ldots, L_\ell$.

We show how to convert the recursive expression of $\tilde{\Upsilon}$ into a representation as a list of $\ell + 1$ polynomials in $\ell + 1$ variables (or fewer, depending on the structure of $\tilde{\Upsilon}$), akin to linear secret sharing structures [34,40].

The principle is as follows. To each input symbol L_i, we associate a degree-1 homogeneous polynomial $\pi_i = \sum_{j=0}^{\ell} y_{i,j} Z_j$, where the variables $Z_0, \ldots, Z_\ell$ are common to all polynomials and the coefficients $y_{i,j}$ are elements of $\mathbb{F}_p$. The polynomials are such that if the formula $\tilde{\Upsilon}$ is satisfied by setting some subset of symbols to $\top$, then the span of the corresponding polynomials will contain the pure monomial Z_0; conversely, any set of polynomials whose span contains the monomial Z_0 indicates a satisfying assignment.

The following algorithm computes such a representation from $\tilde{\Upsilon}$. Proceeding recursively, it assigns temporary polynomials to the interior nodes as it walks down the tree from the root to the leaves (i.e., from the output gate to the input symbols):

1. Initialize a counter $k_c \leftarrow 0$.
   The counter k_c is used for allocating new variables, so that each Z_{k+k_c} is always a "fresh" variable that is never used before or after in the algorithm.
2. Label the root node N_0 with the polynomial $\pi_{N_0} \leftarrow Z_0$.
3. Select a non-leaf node N with non-empty label $\pi_N \neq \emptyset$.
   (a) Denote by $N_1, \ldots, N_m$ the $m \geq 2$ children of N.
   (b) If N is $\vee\{N_1, \ldots, N_m\}$, then $\forall i = 1, \ldots, m$ let $\pi_{N_i} = \pi_N$.
   (c) If N is $\wedge\{N_1, \ldots, N_m\}$, then $\forall i = 1, \ldots, m$ let $\pi_{N_i} = \pi_N + \sum_{k=1}^{m-1} l_{i,k}\, Z_{k+k_c}$ where $l_{i,k} \in \mathbb{F}_p$. The selection of $l_{i,k}$ is explained below.
   (d) If N is $\geq_t \{N_1, \ldots, N_m\}$, then $\forall i = 1, \ldots, m$ let $\pi_{N_i} = \pi_N + \sum_{k=1}^{t-1} l_{i,k}\, Z_{k+k_c}$ where $l_{i,k} \in \mathbb{F}_p$. The selection of $l_{i,k}$ is explained below.
   (e) Label each child N_i with the polynomial π_{N_i}.
   (f) Unlabel node N, i.e., set $\pi_N \leftarrow \emptyset$.
   (g) Increment $k_c \leftarrow k_c + t - 1$ (using $t = 1$ for an $\vee$-gate, and $t = m$ for an $\wedge$-gate).
   (h) Continue at step 3 if an eligible node remains, otherwise skip to step 4.
4. Let $\vartheta \leftarrow k_c$ and output the polynomials $(\pi_0, \ldots, \pi_\ell)$ associated with the leaf nodes $L_0, \ldots, L_\ell$.
   Each polynomial π_i is represented as a vector of coefficients $(y_{i,0}, \ldots, y_{i,\vartheta}) \in \mathbb{F}_p^{\vartheta+1}$ such that $\pi_i = \sum_{k=0}^{\vartheta} y_{i,k}\, Z_k$ is the result of the sequence of operations in steps 3b, 3c, and 3d.

We note that the only variables with nonzero coefficients in the output polynomials are $Z_0, \ldots, Z_\vartheta$, where $\vartheta = k_c$ is the final counter value and may be equal to or lesser than ℓ.

In steps 3c and 3d, the coefficients $l_{i,k}$ need to ensure that no linear relation exists within any set of π_{N_i} of size $< m$ or $< t$. (By construction, m or t of them will always be linearly dependent.) To achieve this property, we let $\left(l_{i,k}\right)$ form a Vandermonde matrix in $\mathbb{F}_p^{m\times(m-1)}$ or $\mathbb{F}_p^{m\times(t-1)}$, i.e., set $l_{i,k} = a_i^k$ for distinct $a_i \in \mathbb{F}_p$; independence follows from the existence of polynomial interpolation. We also require that $\left(l_{i,k}\right)$ be constructed deterministically, so that anyone can verify that the π_i faithfully encode $\tilde{\Upsilon}$ simply by reproducing the process.

The following lemma shows the equivalence between the recursive specification of $\tilde{\Upsilon}$ and its flattened representation. It is adapted from a classic result [34] for linear secret sharing structures and proven by induction on the structure of $\tilde{\Upsilon}$. We refer to the literature [40] for further details.

Lemma 8. *[34] Let $\tilde{\Upsilon}$ be an arborescent monotone threshold circuit, and $\pi_0, \ldots, \pi_\ell$ a flattened representation of it per the above algorithm. A minimal truth assignment $\chi : \{L_0, \ldots, L_\ell\} \to \{\bot, \top\}$ satisfies $\tilde{\Upsilon}(\chi(L_0), \ldots, \chi(L_\ell)) = \top$ if and only if there exist in $\mathbb{F}_p$ coefficients $\nu_0, \ldots, \nu_\ell$ such that,*

$$\sum_{i=0}^{\ell} \nu_i\, \pi_i = Z_0, \quad \textit{and} \quad \forall i : \nu_i = 0 \iff \chi(L_i) = \bot.$$

In this context, a minimal assignment χ with respect to some monotone boolean function $\tilde{\Upsilon}$ is one that satisfies $\tilde{\Upsilon}$ but ceases to do so when any literal of χ is flipped from true to false.

Equivalently, if we expand the polynomials π_i into their coefficients $y_{i,k}$, and write $\delta_{0,k}$ for the Kronecker delta function, it holds that, $\forall k = 0, \ldots, \theta$,

$$\forall k = 0, \ldots, \theta : \sum_{i=0}^{\ell} \nu_i \, y_{i,k} = \delta_{0,k}.$$

5.2. *Information-Theoretic Blinding*

In the signature scheme (yet to be described), we use both the polynomials $(\pi_0, \ldots, \pi_\ell)$ and the linear combination $(\nu_0, \ldots, \nu_\ell)$ from Lemma 8: the latter to create a signature and the former to indicate how to verify it. However, since the linear coefficients ν_i reveal which of the L_i are true, they must be kept secret. In the actual signature, these coefficients appear not as integers but as exponents of elements of $\mathbb{G}$ and are thus already computationally hidden; however, this is not enough and we need to take an extra step to ensure information-theoretic hiding.

By Lemma 8, we know that $\sum_{i=0}^{\ell} \nu_i \, \pi_i = Z_0$, where each $\nu_i \in \mathbb{F}_p$ and each $\pi_i \in \mathbb{F}_p[Z_0, \ldots, Z_\vartheta]_1$. We hide the linear coefficients ν_i using random blinding terms $(h_0, \ldots, h_\ell)$ such that $\sum_{i=0}^{\ell} h_i \, \pi_i = 0$. Since $\sum_{i=0}^{\ell} (\nu_i + h_i) \, \pi_i = Z_0$, the blinded coefficients $\nu_i + h_i$ still bear witness that $\tilde{\Upsilon}(L_0, \ldots, L_\ell) = \top$. However, these witnesses have been rendered information-theoretically indistinguishable, because the distribution of $(\nu_0 + h_0, \ldots, \nu_\ell + h_\ell)$ is conditionally independent of the truth values of the L_i given that $\tilde{\Upsilon}(L_0, \ldots, L_\ell) = \top$.

The difficulty is that no scalar h_i will satisfy $\sum_{i=0}^{\ell} h_i \, \pi_i = 0$ when the π_i contain uninstantiated variables. However, given a specific set of π_i, it is easy to build h_i that have polynomial values.

1. Draw a random vector $\boldsymbol{s} = (s_1, \ldots, s_\ell) \in \mathbb{F}_p^\ell$ of scalar coefficients.
2. For $i = 1, \ldots, \ell$, define $h_i = -s_i \, \pi_0$, and set the remaining term $h_0 = \sum_{j=1}^{\ell} s_j \, \pi_j$.

In the actual scheme, these polynomials are evaluated "in the exponent" for unknown assignments to the Z_k, but regardless of their values, we have $\sum_{i=0}^{\ell} h_i \, \pi_i = (\sum_{j=1}^{\ell} s_j \, \pi_j) \, \pi_0 + \sum_{i=1}^{\ell} (-s_i \, \pi_0) \, \pi_i = 0$, and so the blinding terms $(h_0, \ldots, h_\ell)$ meet our requirements.

The random vector $\boldsymbol{s}$ can be chosen independently of the π_i. This is important for the actual signature scheme, where the relevant polynomials will have coefficients that involve discrete logarithms not known explicitly (in addition to the Z_k being instantiated as discrete logarithms of random group elements). In spite of this, we will be able to select a suitable vector $\boldsymbol{s}$ and compute the blinding terms h_i "in the exponent."

5.3. *Construction*

The full mesh signature scheme can now be described as follows. (In this description, we shall provide a somewhat "wasteful" construction and defer to Sect. 5.5 for a discussion of simple but effective ways to optimize it.)

Initialization: This step is parameterized by a security parameter κ and a bound λ on the number of clauses that can be incorporated into a mesh. It also assumes an agreed-upon public random string $K \in \{0,1\}^{\mathrm{poly}(\kappa)}$.
Given the security parameter κ and the reference string K, all the participants generate a common bilinear instance $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e}) \leftarrow \mathcal{G}(1^\kappa; K)$. Here, we require that the implied isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$ be efficiently computable.
The security parameter κ and the string K are also used to obtain a common hash function $H : \{0,1\}^* \to \mathbb{F}_p$ from a collision-resistant family.
Given the mesh size parameter λ and the string K, the participants then extract $\lambda + 1$ common elements $\hat{g}_0, \hat{g}_1, \ldots, \hat{g}_\lambda$ in $\hat{\mathbb{G}}$, and the corresponding images $g_0, g_1, \ldots, g_\lambda$ in $\mathbb{G}$ under ψ. The extraction process must ensure that the discrete logarithms of the g_i are unknown.
Finally, K defines $\lambda + 1$ random triples $(\hat{A}_{0,k},\ \hat{B}_{0,k},\ \hat{C}_{0,k}) \in \hat{\mathbb{G}}^3$ for $k \in \{0, \ldots, \lambda\}$; these elements together constitute a public verification key "in the sky" with no known signing key. Using the map ψ, everyone computes $A_{0,k} = \psi(\hat{A}_{0,k})$, $B_{0,k} = \psi(\hat{B}_{0,k})$, $C_{0,k} = \psi(\hat{C}_{0,k})$, in $\mathbb{G}$. We note that the public key "in the sky" is not well formed, in the sense that it satisfies none of the internal Diffie–Hellman relationships that regular user public keys, defined next, do.
Key generation: To create a key pair, User #i draws a triple $(a_i,\ b_i,\ c_i) \in (\mathbb{F}_p^\times)^3$ as signing key. User #i computes for each $k \in \{0, \ldots, \lambda\}$ the triple $(\hat{A}_{i,k},\ \hat{B}_{i,k},\ \hat{C}_{i,k}) = (\hat{g}_k^{a_i},\ \hat{g}_k^{b_i},\ \hat{g}_k^{c_i}) \in \hat{\mathbb{G}}^3$, and lets these $3(\lambda + 1)$ group elements constitute his or her verification key.
For simplicity, we write $(A_{i,k},\ B_{i,k},\ C_{i,k}) = (\psi(\hat{A}_{i,k}),\ \psi(\hat{B}_{i,k}),\ \psi(\hat{C}_{i,k})) = (g_k^{a_i},\ g_k^{b_i},\ g_k^{c_i}) \in \mathbb{G}^3$, which anyone can compute from the verification key of User #i thanks to ψ.
Mesh signature: Consider the following mesh signature prototype information:

- ℓ statements $\left[VK_i : Msg_i\right]$, assumed w.l.o.g. to involve the public keys of Users #$1, \ldots, \ell$, and whose propositional truth values are denoted by the literals L_i for $i = 1, \ldots, \ell$.
- an arborescent monotone threshold circuit Υ where each literal $L_1, \ldots, L_\ell$ is an input leaf; and an assignment $\chi : \{L_1, \ldots, L_\ell\} \to \{\bot, \top\}$ that satisfies $\Upsilon(L_1, \ldots, L_\ell) = \top$;
- $\forall i = 1, \ldots, \ell$ such that $\chi(L_i) = \top$, a valid Boneh–Boyen signature in $\mathbf{G}$, given as a pair,

$$\left(u_i = g^{\frac{1}{a_i + b_i\, w_i + c_i\, t_i}},\ t_i \right), \quad \text{for some} \quad t_i \in \mathbb{F}_p,$$

  where $w_i = Msg_i$ and (a_i, b_i, c_i) is the signing key for the statement $\left[VK_i : Msg_i\right]$.
- Optionally, a prescribed "random" value $t_i \in \mathbb{F}_p$ for any index i such that $\chi(L_i) = \bot$.

To create a mesh signature based on the preceding data, the signer firsts extends Υ into a new specification that involves the verification key "in the sky":

1. Hash the public mesh specification to get $Msg_0 = H(\lceil VK_1 : Msg_1 \rceil, \ldots, \lceil VK_\ell : Msg_\ell \rceil, \Upsilon)$, and implicitly associate the literal L_0 to the clause $\lceil VK_0 : Msg_0 \rceil$.
2. Construct $\tilde{\Upsilon} = L_0 \vee \Upsilon$, a well-formed arborescent monotone threshold circuit.
3. Extend χ so that $\chi(L_0) = \bot$, as we lack the corresponding atomic signature.

The signer then builds the mesh signature from the circuit $\tilde{\Upsilon}$, the assignment χ, and the atomic signatures (u_i, t_i) known for such i that $\chi(L_i) = \top$, as follows:

4. Create a flattened representation of $\tilde{\Upsilon}$ and χ as discussed in Sect. 5.1. Accordingly, let $\pi_0, \ldots, \pi_\ell \in \mathbb{F}_p[Z_0, \ldots, Z_\vartheta]$ be public degree-1 multivariate polynomials that encode $\tilde{\Upsilon}$, and $v_0, \ldots, v_\ell \in \mathbb{F}_p$ the secret scalar coefficients of a linear combination that expresses χ, as in Lemma 8. Compute the coefficients $y_{j,k} \in \mathbb{F}_p$ of the polynomials $\pi_j = \sum_{k=0}^{\vartheta} y_{j,k} Z_k$.
5. Create a random blinding vector $\boldsymbol{s} = (s_1, \ldots, s_\ell) \in \mathbb{F}_p^\ell$.
6. $\forall i \in \{0, \ldots, \ell\} : \chi(L_i) = \bot$, randomly draw $t_i \in \mathbb{F}_p$, and arbitrarily fix $u_i = g^0 = 1 \in \mathbb{G}$. Alternatively, instead of a random t_i, a prescribed value can be used.[2] (Recall that for $\chi(L_i) = \top$, the t_i and u_i are supplied with the atomic signatures.)
7. For all $j = 0, \ldots, \ell$ and $k = 0, \ldots, \vartheta$, calculate,

$$v_{j,k} = \left(A_{j,k}\, B_{j,k}^{m_j}\, C_{j,k}^{t_j} \right)^{y_{j,k}}, \quad \text{setting} \quad m_j = Msg_j.$$

(Note that if we instantiate $Z_k = \mathrm{dlog}_g(g_k)$, we get $v_j := \prod_{k=0}^{\vartheta} v_{j,k} = g^{(a_j + b_j m_j + c_j t_j)\pi_j}$ for all j except $j = 0$ since the key "in the sky" is ill formed.)

8. Compute, for $i = 1, \ldots, \ell$, and $k = 0, \ldots, \vartheta$, respectively,

$$S_i = {u_i}^{v_i}\, {v_0}^{-s_i}, \quad P_k = \prod_{j=1}^{\ell} {v_{j,k}}^{s_j}.$$

(The value of any intervening u_i such that $\chi(L_i) = \bot$ is unimportant since then $v_i = 0$; this is true in particular for the user "in the sky" of index 0.)

9. Output the mesh signature, consisting of the statement Υ and the tuple,

$$\sigma = (\, t_0,\ \ldots,\ t_\ell,\ S_1,\ \ldots,\ S_\ell,\ P_0,\ \ldots,\ P_\vartheta\,) \in \mathbb{F}_p^{\ell+1} \times \mathbb{G}^{\ell+\vartheta+1}.$$

Mesh verification: A fully qualified mesh signature package consists of:

- a list of $\ell + 1$ propositions $\lceil VK_0 : Msg_0 \rceil, \ldots, \lceil VK_\ell : Msg_\ell \rceil$ viewed as inputs to,
- an arborescent monotone threshold circuit $\tilde{\Upsilon} : \{\bot, \top\}^{\ell+1} \to \{\bot, \top\}$,
- a mesh signature $\sigma = (t_0, \ldots, t_\ell, S_1, \ldots, S_\ell, P_0, \ldots, P_\vartheta) \in \mathbb{F}_p^{\ell+1} \times \mathbb{G}^{\ell+\vartheta+1}$.

[2] The facility to use a prescribed value for t_i in clauses that are "false" is to give the appearance that such clauses are constructed from Boneh–Boyen atomic signatures with given t_i, as if they were "true." Without this provision, should a clause have the same t_i as that of a published signature (e.g., a certificate), it would be exposed as "true."

To verify such a signature, the verifier proceeds as follows:

1. Ascertain that $\tilde{\Upsilon}(\top, \star, \ldots, \star) = \top$, extract from $\tilde{\Upsilon}(L_0, \ldots, L_\ell)$ the subcircuit $\Upsilon(L_1, \ldots, L_\ell)$ such that $\tilde{\Upsilon} = \Upsilon \vee L_0$, and verify that $Msg_0 = H(\left[VK_1 : Msg_1\right], \ldots, \left[VK_\ell : Msg_\ell\right], \Upsilon)$.
2. Recompute the polynomials $(\pi_0, \ldots, \pi_\ell)$ representing the formula $\tilde{\Upsilon}$ by reproducing the deterministic conversion of Sect. 5.1.
3. For $i = 0, \ldots, \ell$, determine the coefficients $y_{i,k} \in \mathbb{F}_p$ of the polynomials $\pi_i = \sum_{k=0}^{\vartheta} y_{i,k} Z_k$.
4. For $i = 0, \ldots, \ell$ and $k = 0, \ldots, \vartheta$, retrieve $(\hat{A}_{i,k}, \hat{B}_{i,k}, \hat{C}_{i,k})$ from the key VK_i, and calculate,

$$\hat{v}_{i,k} = \left(\hat{A}_{i,k}\, \hat{B}_{i,k}^{m_i}\, \hat{C}_{i,k}^{t_i}\right)^{y_{i,k}}, \quad \hat{v}_i = \prod_{k=0}^{\vartheta} \hat{v}_{i,k}, \quad \text{setting} \quad m_i = Msg_i.$$

5 Using the pairing, verify the equalities, for all $k = 0, \ldots, \vartheta$,

$$\mathbf{e}\left(P_k,\ \hat{v}_0\right) \cdot \prod_{i=1}^{\ell} \mathbf{e}\left(S_i,\ \hat{v}_{i,k}\right) = \begin{cases} \mathbf{e}(g,\ \hat{g}_0) & \text{for } k = 0 \\ 1 & \text{otherwise} \end{cases}.$$

6 Accept the signature as valid if and only if all $\vartheta + 1$ preceding equalities hold in $\mathbb{G}_t$.

(Optional) **Probabilistic check:** Mesh signatures can be verified using fewer total pairings, at the cost of some additional random bits and exponentiations. In the same setting as above, it suffices to replace the end of the verification algorithm from step 5 onward by the following:

5′. Using the pairing, for $d_0 = 1$ and random $d_1, \ldots, d_\vartheta \in \mathbb{F}_p$, verify the single equality,

$$\mathbf{e}\left(\prod_{k=0}^{\vartheta} P_k^{d_k},\ \hat{v}_0\right) \cdot \prod_{i=1}^{\ell} \mathbf{e}\left(S_i,\ \prod_{k=0}^{\vartheta} \hat{v}_{i,k}^{d_k}\right) = \mathbf{e}(g,\ \hat{g}_0).$$

6′. Accept the signature as valid if and only if the preceding equality holds in $\mathbb{G}_t$.

The probabilistic verification incurs a negligible statistical error of accepting a signature that would not be accepted by the deterministic algorithm. It is however significantly faster.

5.4. *Security*

We state the correctness, anonymity, and unforgeability theorems for the mesh scheme. A corollary to the latter is also given, based on a weaker assumption, for the case where only a subset of the honest users are willing to answer atomic signature queries (e.g., certificate authorities).

Theorem 9. *The mesh signature is consistent.*

Proof. For any list of public polynomials $\pi_0, \ldots, \pi_\ell$ and secret coefficients $\nu_0, \ldots, \nu_\ell$ that, respectively, encode per Lemma 8 a well-formed mesh specification $\tilde{\Upsilon}$ and an assignment χ that satisfies it, we need to show that a signature created by the above algorithm will be accepted by the same. A straightforward sequence of substitutions in the scheme description shows this to be the case. □

Theorem 10. *The mesh signature has everlasting perfect anonymity.*

Proof. See "Anonymity of the Mesh Scheme" of Appendix □

Theorem 11. *The mesh signature is existentially unforgeable under an adaptive chosen message attack, against a static adversary that makes no more than q mesh signature queries, and no more than q atomic signature queries to each of the ℓ honest users, adaptively, provided that the $(q, \ell+1)$-Poly-SDH assumption holds in* **G***, in the common random string model.*

Proof. See "Unforgeability of the Mesh Scheme" of Appendix □

Corollary 12. *The mesh signature is existentially unforgeable under an adaptive chosen message attack, against a static adversary that makes no more than q mesh signature queries, and no more than q atomic signature queries to each of ℓ' among a total of $\ell + \ell'$ honest users, adaptively, provided that the $(q, \ell, \ell'+1)$-Pluri-SDH assumption holds in* **G***, in the common random string model.*

5.5. *Optimizations for Shorter Keys and CRS*

As previously mentioned, we can make both the ring and mesh signature schemes more compact and more efficient, by noting that there is no need for three secrets in the atomic signature triplets (a_i, b_i, c_i). As we shall see, two of them would suffice, though one is not enough. We exploit this by arbitrarily anchoring all instances of private keys' b_i to the constant 1 and accordingly fixing the corresponding public keys' B_i to known values that need no longer be published. The justification for this will become apparent in "Ring Scheme Security Proofs" and "Mesh Scheme Security Proofs" of Appendices, wherein the simulators that we construct are always allowed to know the value of the b_i, indicating that the latter do not actually contribute to security. Further anticipating from the security reductions, we note that we shall need to construct two different simulators, that will know either one of the remaining private-key secrets a_i and c_i, and that is the reason why we cannot shrink the private keys further. In summary, in both the mesh and the ring schemes, we can set $b_i = 1$ wherever it appears and omit the publication of any instance of $\hat{B}_{i,k} = \hat{g}_k^{b_i} = \hat{g}_k$. This results in public keys (including the key "in the sky") being shrunken to 2/3 of their original size. The scheme also becomes more computationally efficient as a result.[3]

[3] We note that with the B_i removed from the public keys, the ring scheme becomes syntactically very close to the ring signature scheme of [21]. If one temporarily ignores the generalization to the full mesh signature model, one way of looking at the ring signature scheme of Sect. 4 is a provably secure version of [21] from a provably sound hardness assumption (in the sense of being provably hard in the generic-group model).

A second way to achieve optimizations is further to compress the key "in the sky" to just two elements of $\hat{\mathbb{G}}$. This is based on the observation that, for $\tilde{\Upsilon} = \Upsilon \vee L_0$, the encoding algorithm of Sect. 5.1 always gives $\pi_0 = Z_0$, i.e., $y_{0,0} = 1$ and $y_{0,k} = 0$ for $k \neq 0$. This means that the tuples $(\hat{A}_{0,k}, \hat{B}_{0,k}, \hat{C}_{0,k})$ for $k \neq 0$ are in fact never used. Since it is safe to set $\hat{B}_{0,0} = \hat{g}$ as discussed above, the key "in the sky" can thus shrink to a mere pair $(\hat{A}_{0,0}, \hat{C}_{0,0})$ of random elements.

6. Conclusion

We have introduced mesh signatures as a generalization of ring signatures with a richer language for expressing signer ambiguity. Mesh signatures scale to large crowds with many cosigners and independent certificate authorities; they can even implicate unwilling individuals who, by withholding their ring public key, would have otherwise remained out of reach. Because in principle mesh signatures require no central authority and only a minimal-trust CRS, they provide a credible answer to the question of how to leak a secret authoritatively.

We have constructed a simple and practical mesh signature scheme in prime-order bilinear groups, which achieves everlasting unconditional anonymity, and existential unforgeability in the common random string model. To obtain this result, we introduced a new complexity assumption, which we prove sound in the generic model; it is in the spirit of the SDH assumption, but better exploits the group structure of the values computed by pairing. Incidentally, we obtain a very efficient and the first unconditionally anonymous ring signature without random oracles as a special case of our construction.

Appendix 1: The Poly-SDH and the Pluri-SDH Assumptions

We gave informal definitions of the Poly-SDH and Pluri-SDH assumptions in Sects. 3.2 and 3.4. Next, we give concrete and asymptotic definitions of the Poly-SDH and Pluri-SDH assumptions, whose generic hardness we prove in "Generic-Group Complexity of Poly-SDH" and "Generic-Group Complexity of Pluri-SDH" of Appendices.

Formal Poly-SDH Definitions

Our formal statement of the (q, ℓ)-Poly-SDH problem applies to canonical bilinear groups of all types, whether $\mathbb{G} = \hat{\mathbb{G}}$ or $\mathbb{G} \neq \hat{\mathbb{G}}$. In general, most of the group elements that appear in the problem instance will belong in $\mathbb{G}$, and the solution is required to be in group. Several elements must also be given in $\hat{\mathbb{G}}$, mainly to enable the pairing-based verification of the solutions; but rather than require the isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$ to be efficiently computable and let the user do the translation, we shall give these elements in both groups explicitly.

Definition 13. In a bilinear context $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$, the (q, ℓ)-Poly-SDH problem is:

> Given, $\forall i = 1, \ldots, \ell$, $\forall j = 1, \ldots, q$, $g^{\alpha_i} \in \mathbb{G}$, $\hat{g}^{\alpha_i} \in \hat{\mathbb{G}}$, and $(w_{i,j},\ g^{\frac{1}{\alpha_i + w_{i,j}}}) \in \mathbb{F}_p \times \mathbb{G}$,
>
> output $\forall i$, $(w_i,\ g^{\frac{r_i}{\alpha_i + w_i}})$, subject to: $\sum_{i=1}^{\ell} r_i = 1 \pmod{p}$, $\forall i, \forall j, w_i \neq w_{i,j}$.

The advantage of an algorithm $\mathcal{A}$ in solving the (q, ℓ)-Poly-SDH problem is,

$$\mathbf{Adv}_{\mathcal{A}}^{\text{PolySDH}} = \Pr\left[\mathcal{A}\left(\begin{array}{l} g,\ (g^{\alpha_i})_{1 \le i \le \ell},\ \left(w_{i,j},\ g^{\frac{1}{\alpha_i + w_{i,j}}} \right)_{\substack{1 \le i \le \ell \\ 1 \le j \le q}} \\ \hat{g},\ (\hat{g}^{\alpha_i})_{1 \le i \le \ell} \end{array} \right) = \left(w_i,\ g^{\frac{r_i}{\alpha_i + w_i}} \right)_{1 \le i \le \ell} : w_i \neq w_{i,j},\ \sum_{i=1}^{\ell} r_i = 1 \right]$$

The probability is over the random choice of generators $g \in \mathbb{G} \setminus \{1\}$ and $\hat{g} \in \hat{\mathbb{G}} \setminus \{1\}$, of exponents $\alpha_1, \ldots, \alpha_\ell \in \mathbb{F}_p^\times$, of integers $w_{i,j} \in \mathbb{F}_p \setminus \{-\alpha_i\}$, and the random bits consumed by $\mathcal{A}$. (Nevertheless, we allow g and/or $\hat{g}$ to be given externally, as it is convenient in type-1 and type-2 bilinear contexts to take $g = \psi(\hat{g})$ under a fixed isomorphism ψ.)

Definition 14. We say that the (q, ℓ, t, ϵ)-Poly-SDH assumption holds in $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$ if no t-time and ϵ-advantage randomized algorithm solves the (q, ℓ)-Poly-SDH problem in $\mathbf{G}$.

Formal Pluri-SDH Definitions

As in the previous case, we formally define the $(q, \ell, 1)$-Pluri-SDH problem for bilinear groups of all types, whether symmetric ($\mathbb{G} = \hat{\mathbb{G}}$) or asymmetric ($\mathbb{G} \neq \hat{\mathbb{G}}$), with or without efficient isomorphism ($\psi : \hat{\mathbb{G}} \to \mathbb{G}$), at the cost of some possible redundancy in the instance statement in some cases.

Definition 15. In a bilinear context $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$, the $(q, \ell, 1)$-Pluri-SDH problem is:

> Given: $\forall i = 0, \ldots, \ell$, $(g^{\alpha_i}, \hat{g}^{\alpha_i}) \in \mathbb{G} \times \hat{\mathbb{G}}$, and: $\forall j = 1, \ldots, q$, $(w_{0,j},\ g^{\frac{1}{\alpha_0 + w_{0,j}}}) \in \mathbb{F}_p \times \mathbb{G}$,
>
> output: $\forall i$, $(w_i,\ g^{\frac{r_i}{\alpha_i + w_i}})$, subject to: $\sum_{i=0}^{\ell} r_i = 1 \pmod{p}$, and: $\forall j$, $w_0 \neq w_{0,j}$.

The advantage of an algorithm $\mathcal{A}$ in solving the $(q, \ell, 1)$-Pluri-SDH problem is,

$$\mathbf{Adv}^{\text{PluriSDH}}_{\mathcal{A}} = \Pr\left[\mathcal{A}\left(\begin{array}{l} g,\ \left(g^{\alpha_i}\right)_{0\le i\le \ell}\ ,\ \left(w_{0,j},\ g^{\frac{1}{\alpha_0+w_{0,j}}}\right)_{1\le j\le q} \\ \hat{g},\ \left(\hat{g}^{\alpha_i}\right)_{0\le i\le \ell} \end{array}\right) = \left(w_i,\ g^{\frac{r_i}{\alpha_i+w_i}}\right)_{0\le i\le \ell} : w_0 \neq w_{0,j}\ ,\ \sum_{i=1}^{\ell} r_i = 1\right]$$

The probability is over the random choice of $g \in \mathbb{G}\backslash\{1\}$ and $\hat{g} \in \hat{\mathbb{G}}\backslash\{1\}$, of $\alpha_0, \ldots, \alpha_\ell \in \mathbb{F}_p^\times$, of $w_{0,j} \in \mathbb{F}_p \setminus \{-\alpha_0\}$, and over the random bits consumed by $\mathcal{A}$. (We allow g and/or $\hat{g}$ to be externally given, as it is convenient in some bilinear contexts to take $g = \psi(\hat{g})$ under a fixed isomorphism ψ.)

Definition 16. We say the $(q, \ell, 1, t, \epsilon)$-Pluri-SDH assumption to hold in $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$ if no t-time and ϵ-advantage randomized algorithm solves the $(q, \ell, 1)$-Pluri-SDH problem in $\mathbf{G}$.

A Spectrum of Assumptions from Pluri-SDH to Poly-SDH. Given the clear similarity between the $(q, \ell + 1)$-Poly-SDH and $(q, \ell, 1)$-Pluri-SDH problems, it is natural to consider a generalized notion of Pluri-SDH with $\ell + \ell'$ random generators in total, only ℓ' of which would be given with accompanying series of solution pairs. (In the basic Pluri-SDH problem, we had $\ell' = 1$.)

Definition 17. In a bilinear context $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$, the (q, ℓ, ℓ')-Pluri-SDH problem is (using positive and negative indices $i \neq 0$ to differentiate the two types of generators):

Given: $\forall i' = 1, \ldots, \ell,\ (g^{\alpha_{i'}}, \hat{g}^{\alpha_{i'}}) \in \mathbb{G} \times \hat{\mathbb{G}}$,
and: $\forall i'' = -\ell', \ldots, -1,\ \forall j = 1, \ldots, q,\quad (g^{\alpha_{i''}}, \hat{g}^{\alpha_{i''}}) \in \mathbb{G} \times \hat{\mathbb{G}},\ (w_{i'',j},\ g^{\frac{1}{\alpha_{i''}+w_{i'',j}}}) \in \mathbb{F}_p \times \mathbb{G}$,
output: $\forall i = -\ell', \ldots, -1, 1, \ldots, \ell,\ (w_i,\ g^{\frac{r_i}{\alpha_i+w_i}})$,
subject to: $\sum_{i\in\{-\ell',\ldots,\ell\}\backslash\{0\}} r_i = 1 \pmod p$, and also: $\forall i < 0, \forall j,\ w_i \neq w_{i,j}$.

The advantage of an algorithm $\mathcal{A}$ in solving the (q, ℓ, ℓ')-Pluri-SDH problem is,

$$\mathbf{Adv}^{\text{PluriSDH}}_{\mathcal{A}} = \Pr\left[\mathcal{A}\left(\begin{array}{l} g,\ \left(g^{\alpha_i}\right)_{\substack{i=-\ell',\ldots,\\ -1,1,\ldots,\ell}}\ ,\ \left(w_{i,j},\ g^{\frac{1}{\alpha_i+w_{i,j}}}\right)_{\substack{1\le j\le q\\ i=-\ell',\ldots,\\ -1,1,\ldots,\ell}} \\ \hat{g},\ \left(\hat{g}^{\alpha_i}\right)_{\substack{i=-\ell',\ldots,\\ -1,1,\ldots,\ell}} \end{array}\right)$$

$$= \left(w_i,\ g^{\frac{r_i}{\alpha_i + w_i}} \right)_i \ :\ w_i \neq w_{i,j}\ ,\ \sum_{i=1}^{\ell} r_i = 1 \Bigg]$$

The probability is over the random choice of $g \in \mathbb{G} \setminus \{1\}$ and $\hat{g} \in \hat{\mathbb{G}} \setminus \{1\}$, of all $\alpha_i \in \mathbb{F}_p^\times$, of all $w_{i,j} \in \mathbb{F}_p \setminus \{-\alpha_i\}$, and over the random bits consumed by $\mathcal{A}$. (We allow g and/or $\hat{g}$ to be externally given, as it is convenient in some bilinear contexts to take $g = \psi(\hat{g})$ under a fixed isomorphism ψ.)

Definition 18. We say the $(q, \ell, \ell', t, \epsilon)$-Pluri-SDH assumption to hold in $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$ if no t-time and ϵ-advantage randomized algorithm solves the (q, ℓ, ℓ')-Pluri-SDH problem in $\mathbf{G}$.

The generalized (q, ℓ, ℓ')-Pluri-SDH problem is useful to prove security of the mesh scheme in the case where there are $\ell + \ell'$ honest users, but only ℓ' of them are willing to answer queries for atomic signatures (e.g., certificate authorities willing to issue certificates on demand, *vs.* ordinary users only occasionally making opportunistic signatures).

Asymptotic Definitions

For completeness, we also give a complexity-theoretic statement of either assumption, based on the asymptotic limit in an infinite family of bilinear groups. All our security theorems are easy to restate in terms of the asymptotic definition.

Definition 19. We say that the Poly-SDH assumption holds for a bilinear instance generaror $\mathcal{G}$ if, for any polynomially bounded functions $q(\cdot)$, $\ell(\cdot)$, $t(\cdot)$, $\epsilon(\cdot)$, and any algorithm $\mathcal{A}$, there exists a threshold κ^*, beyond which for all values of the security parameter $\kappa > \kappa^*$, the $(q(\kappa), \ell(\kappa), t(\kappa), \epsilon(\kappa))$-Poly-SDH assumption holds against $\mathcal{A}$ in all bilinear instances $\mathbf{G}$ generated by $\mathcal{G}(1^\kappa)$.

Definition 20. We say that the Pluri-SDH assumption holds for a bilinear instance generaror $\mathcal{G}$ if, for any polynomial functions $q(\cdot)$, $\ell(\cdot)$, $\ell'(\cdot)$, $t(\cdot)$, $\epsilon(\cdot)$, and any algorithm $\mathcal{A}$, there exists a threshold κ^*, beyond which for all values of the security parameter $\kappa > \kappa^*$, the $(q(\kappa), \ell(\kappa), \ell'(\kappa), t(\kappa), \epsilon(\kappa))$-Pluri-SDH assumption holds against $\mathcal{A}$ in all bilinear instances $\mathbf{G}$ generated by $\mathcal{G}(1^\kappa)$.

Generic-Group Complexity of Poly-SDH

We now give a complete proof that the Poly-SDH assumption (and hence, also the weaker Pluri-SDH assumption) is sound in the generic group model in the sense of Shoup [39]. The proof is based on the argument given in Sect. 3.3.

In this model, slightly extended to incorporate bilinearity, the groups $\mathbb{G}$, $\hat{\mathbb{G}}$, and $\mathbb{G}_t$ are assumed to have a generic presentation, i.e., only canonical operations may be performed on their elements (the group operations in each of $\mathbb{G}$, $\hat{\mathbb{G}}$, and $\mathbb{G}_t$, the iso-

morphism from $\hat{\mathbb{G}}$ to $\mathbb{G}$ and its inverse, and of course the pairing). Specifically, the solver $\mathcal{A}$ performs the canonical operations by interacting with an oracle $\mathcal{O}$, in such a way that $\mathcal{A}$ only sees arbitrary representations of the group elements. This is modeled using arbitrary injective encoding functions f, $\hat{f}$, and f_t, one for each group, and by representing any group element $h \in \mathbb{G}$ as the string $f(h)$ when interacting with $\mathcal{A}$ (and similarly for elements of the other groups).[4] $\mathcal{A}$ is otherwise computationally unbounded.

The following theorem gives an upper bound on the success probability of any generic Poly-SDH attacker $\mathcal{A}$, or equivalently, a lower bound on the complexity of solving Poly-SDH generically.

Theorem 21. *Let $\mathcal{A}$ be a computationally unbounded algorithm for the (q, ℓ)-Poly-SDH problem in generic bilinear groups of prime order p, that makes at most q_G generic-group oracle queries in total (i.e., for $\mathbf{e}$, ψ, ψ^{-1}, and $*$ in $\mathbb{G}$, $\hat{\mathbb{G}}$, and $\mathbb{G}_t$). Then,*

$$\Pr\left[\mathcal{A}^{\mathcal{O}}\left(\begin{array}{l} f(g),\ \ f(g^{\alpha_i})_{1\le i\le \ell}\ ,\ (\,w_{i,j},\ f(g^{\frac{1}{\alpha_i+w_{i,j}}})\,)_{\substack{1\le i\le \ell\\ 1\le j\le q}} \\ \hat{f}(\hat{g}),\ \ \hat{f}(\hat{g}^{\alpha_i})_{1\le i\le \ell} \end{array}\right) = \left(w_i,\ f(g^{\frac{r_i}{\alpha_i+w_i}})\right)_{1\le i\le \ell} : \ w_i \neq w_{i,j}\ ,\ \sum_{i=1}^{\ell} r_i = 1\right]$$
$$= \mathbf{Adv}^{\text{PolySDH}}_{\mathcal{A}} \le \frac{(q_G + q\,\ell + 2\,\ell + 2)^2\,(q\,\ell)}{p-1} = \mathrm{O}\left(\frac{q_G^2\,q\,\ell + (q\,\ell)^3}{p}\right).$$

In order to be as general as possible, we state and prove the theorem assuming the solver is given access to both ψ and ψ^{-1}, as in type-1 bilinear groups. Since the withholding of one or both of these capabilities can only hurt $\mathcal{A}$, the same theorem gives us valid bounds for type-2 and type-3 groups without needing another proof (albeit, at the cost of a small slack factor).

Proof. We merely sketch the oracle simulation, which follows a similar routine as in [5]; the subsequent analysis will be the core of the argument.

We construct an algorithm $\mathcal{B}$ that simulates the generic-group oracle $\mathcal{O}$ without committing to values for the α_i, and analyze what $\mathcal{A}$ can extract from it. Internally, $\mathcal{B}$ keeps track of the group elements by their discrete logarithms to the generators $g \in \mathbb{G}$, $\hat{g} \in \hat{\mathbb{G}}$, and $g_t = \mathbf{e}(g, \hat{g}) \in \mathbb{G}_t$. Since the variables $\alpha_1, \ldots, \alpha_\ell$ are left undetermined, in all generality these discrete logs will be multivariate expressions in $\mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]$, which we denote by ρ_k, $\hat{\rho}_k$, or ρ'_k, depending on the group. Externally, $\mathcal{B}$ maps the correspond-

[4] The encodings $f(h)$, *etc.*, of group elements can be thought of providing unique but intrinsically meaningless "handles" for manipulating "opaque" group elements. The handles can be assigned at random, or sequentially in the order of queries, as long as the handles remain independent of any arithmetic representation of the group elements.

ing group elements to random strings it gives to $\mathcal{A}$: in the group $\mathbb{G}$ it associates ρ_i to $f_i = f(g^{\rho_i})$, in $\hat{\mathbb{G}}$ it maps $\hat{\rho}_i$ to $\hat{f}_i = \hat{f}(\hat{g}^{\hat{\rho}_i})$, and in $\mathbb{G}_t$ it maps ρ'_i to $f'_i = f_t(g_t^{\rho'_i})$.

To initiate the interaction, $\mathcal{B}$ must provide $\mathcal{A}$ with an instance of the (q, ℓ)-Poly-SDH problem. To do so, $\mathcal{B}$ picks random $w_{i,j} \in \mathbb{F}_p$ for $i = 1, \ldots, \ell$ and $j = 1, \ldots, q$, and creates:

- two strings, f_0 and $\hat{f}_0$, which it binds to the constants $\rho_0 = 1$ and $\hat{\rho}_0 = 1$ respectively;
- $2\,\ell$ strings, f_i and $\hat{f}_i$ for $i = 1, \ldots, \ell$, bound to the expressions $\rho_i = \alpha_i$ and $\hat{\rho}_i = \alpha_i$ respectively;
- $q\,\ell$ strings, f_k for $k = i + j\,\ell$ where $i = 1, \ldots, \ell$ and $j = 1, \ldots, q$, bound to the terms $\rho_k = \frac{1}{\alpha_i + w_{i,j}}$.

For simplicity, and to avoid dealing with ratios, we reduce all the expressions to the common denominator $\Delta = \prod_{i=1}^{\ell} \prod_{j=1}^{q} (\alpha_i + w_{i,j})$: for all k, we define $\pi_k = \rho_k\,\Delta$, $\hat{\pi}_k = \hat{\rho}_k\,\Delta$, and $\pi'_k = \rho'_k\,\Delta$. Observe that all of these are polynomials in $\mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]$ of degree $\leq q\,\ell + 1$.

$\mathcal{B}$ gives to $\mathcal{A}$ all the strings f_k and $\hat{f}_k$ created above (but not the corresponding polynomials). $\mathcal{B}$ also initializes three counters: $n \leftarrow (q\,\ell + \ell + 1)$, $\hat{n} \leftarrow (\ell + 1)$, and $n_t \leftarrow 0$.

$\mathcal{A}$ then makes a total of q_G adaptive queries to the generic-group oracle, of the following kinds:

> **Group operations:** Suppose $\mathcal{A}$ wants to compute the product of two operands in the group $\mathbb{G}$ represented as f_i and f_j. To answer, $\mathcal{B}$ retrieves the polynomials π_i and π_j that these strings correspond to (the same string may appear under multiple indices i, but in all cases the associated polynomial will be the same, so there is no inconsistency). It calculates the polynomial sum $\pi_n = \pi_i + \pi_j \in \mathbb{F}_p[\alpha]$. If the result π_n is already present in the $\mathbb{G}$-mapping, the corresponding string is copied into f_n, otherwise a new string is created; then the entry (π_n, f_n) is added to the mapping. Finally, $\mathcal{B}$ gives f_n to $\mathcal{A}$, and then increments $n \leftarrow n + 1$.
> If $\mathcal{A}$ had wanted the ratio in lieu of the product, $\mathcal{B}$ would have let $\pi_n = \pi_i - \pi_j \in \mathbb{F}_p[\alpha]$.
> Group operation queries in $\hat{\mathbb{G}}$ or $\mathbb{G}_t$ are serviced analogously, using the relevant mappings.
> **Isomorphisms:** Suppose $\mathcal{A}$ wants to map an operand $\hat{f}_i$ in $\hat{\mathbb{G}}$ to its image in $\mathbb{G}$ by the isomorphism. To answer, $\mathcal{B}$ proceeds as above to retrieve the polynomial $\hat{\pi}_i$, assigns $\pi_n = \hat{\pi}_i$, and lets f_n be the copy of an existing string or a new string depending on whether π_n was already present in the $\mathbb{G}$-mapping. It adds (π_n, f_n) to the mapping, gives f_n to $\mathcal{A}$, then increments $n \leftarrow n + 1$.
> Inverse isomorphism queries are answered similarly, after exchanging the roles of $\mathbb{G}$ and $\hat{\mathbb{G}}$.
> **Pairing:** Suppose $\mathcal{A}$ wants to compute the bilinear map between f_i in $\mathbb{G}$ and $\hat{f}_j$ in $\hat{\mathbb{G}}$. To answer, $\mathcal{B}$ retrieves the polynomials π_i and $\hat{\pi}_j$, computes the polynomial product $\pi_{n_t} = \pi_i \cdot \pi_j \in \mathbb{F}_p[\alpha]$, determines whether π_{n_t} already exists in the $\mathbb{G}_t$-mapping, and accordingly lets f_{n_t} be a clone of the corresponding string or a new string. It adds (π_{n_t}, f_{n_t}) to the mapping, gives f_{n_t} to $\mathcal{A}$, then increments $n_t \leftarrow n_t + 1$.

W.l.o.g., we have assumed that $\mathcal{A}$ only queried $\mathcal{B}$ on legitimate strings that were previously revealed, and similarly assume that it outputs its solution in terms of such strings exclusively.

After q_G queries, $\mathcal{A}$ returns a solution to the Poly-SDH instance, in the form of ℓ pairs $(w_1,\ f_{k_1}),\ \ldots,\ (w_\ell,\ f_{k_\ell})$ where $0 \le k_i < n$. We denote by π_{k_i} the formal polynomials that correspond to the generic representations f_{k_i}.

To verify the solution, $\mathcal{B}$ randomly selects $\alpha_i^* \in \mathbb{F}_p^\times$ for each i, and evaluates the formal polynomials under the assignment $(\alpha_1, \ldots, \alpha_\ell) = (\alpha_1^*, \ldots, \alpha_\ell^*)$. The validation equation is,

$$\sum_{i=1}^{\ell} (\alpha_i + w_i)\, \rho_{k_i}(\alpha_1, \ldots, \alpha_\ell) = 1,$$

$$i.e., \quad \sum_{i=1}^{\ell} (\alpha_i + w_i)\, \pi_{k_i}(\alpha_1, \ldots, \alpha_\ell) = \Delta. \qquad (\star)$$

$\mathcal{A}$ wins the game outright if Eq. ($\star$) holds in $\mathbb{F}_p$ under the random assignment.

$\mathcal{A}$ wins by default if any two distinct polynomials representing elements of the same group (e.g., $\pi_i \neq \pi_j$, or $\hat{\pi}_i \neq \hat{\pi}_j$, or $\pi'_i \neq \pi'_j$) evaluate to the same value in $\mathbb{F}_p$ under the assignment (*viz.*, $\pi_i(\alpha_1^*, \ldots, \alpha_\ell^*) = \pi_j(\alpha_1^*, \ldots, \alpha_\ell^*)$, *etc.*): if this is the case, $\mathcal{B}$'s simulation is flawed since it portrayed as distinct two instances of the same group element.

$\mathcal{A}$ loses the game barring the two scenarios above. We shall bound the probability of either event, below.

Notice that all polynomials π used by $\mathcal{B}$ to represent an element in $\mathbb{G}$ or $\hat{\mathbb{G}}$ have degree $\le q\,\ell + 1$, since they are constructed using sums and differences from an initial set of polynomials with this property (the polynomials of highest degrees are $\pi_i = \rho_i\,\Delta = \alpha_i\,\Delta$ for $i = 1, \ldots, \ell$, of degree $q\,\ell + 1$). In the target group $\mathbb{G}_t$, the polynomials can be of degree $\le 2\,q\,\ell + 2$, which is the highest degree that the product of two polynomials of degree $\le q\,\ell + 1$ can attain.

To bound the probability that $\mathcal{A}$ makes a correct answer, we bound the probability that Eq. ($\star$) will be satisfied. First, we show that it cannot hold identically in $\mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]$. Consider the following change of variables: for all $i = 1, \ldots, \ell$ and $j = 1, \ldots, q$, we define $\alpha'_i = \alpha_i + w_i$ and $w'_{i,j} = w_{i,j} - w_i$, where we note that $w'_{i,j} \neq 0$ because of the constraints on the w_i. The change of variable is well-defined and unambiguous, and lets us rewrite Eq. ($\star$) as,

$$\sum_{i=1}^{\ell} \alpha'_i\, \pi_{k_i}\left(\alpha'_1 - w_1, \ldots, \alpha'_\ell - w'_1\right) = \Delta$$

$$= \prod_{i=1}^{\ell} \prod_{j=1}^{q} (\alpha_i + w_{i,j}) = \prod_{i=1}^{\ell} \prod_{j=1}^{q} \left(\alpha'_i + w'_{i,j}\right). \qquad (\star\star)$$

Eq. ($\star\star$) cannot hold identically in $\mathbb{F}_p[\alpha'_1, \ldots, \alpha'_\ell]$ since the left-hand side has no independent term while the right-hand side contains the independent term $\prod_i \prod_j w'_{i,j} = \prod_i \prod_j (w_{i,j} - w_i) \neq 0$. Equation ($\star$) then does not hold identically either, in $\mathbb{F}_p[\alpha_1, \ldots, \alpha_\ell]$, which means that it is a non-trivial polynomial equation of degree

$\leq q\,\ell + 2$. For a random assignment of $(\alpha_1, \ldots, \alpha_\ell) \in (\mathbb{F}_p^\times)^\ell$, Eq. ($\star$) will thus be satisfied in $\mathbb{F}_p$ with probability $\leq \frac{q\,\ell+2}{p-1}$.

To bound the probability that $\mathcal{A}$ wins by default, recall that all the polynomials representing elements in $\mathbb{G}$ and $\hat{\mathbb{G}}$ have degree $\leq q\,\ell + 1$, and those in $\mathbb{G}_t$ have degree $\leq 2\,q\,\ell + 2$. Again using the standard argument on the probability of satisfying polynomial equations, and combining all of these with the union bound, we find that the probability that there is at least one flaw in $\mathcal{B}$'s simulation is $\leq \binom{n}{2}\frac{q\,\ell+1}{p-1} + \binom{\hat{n}}{2}\frac{q\,\ell+1}{p-1} + \binom{n_t}{2}\frac{2\,q\,\ell+2}{p-1}$.

(The denominators are $p-1$ rather than p to reflect that the α_i are random in $\mathbb{F}_p^\times$ and not in $\mathbb{F}_p$: conceivably, the adversary could exploit this by forcing all polynomial solutions to be nonzero.) Now, if we take the union bound over the two events, and bring into account the loop invariant, $(n + \hat{n} + n_t) = (q\,\ell + 2\,\ell + 2 + q_G)$, we thus establish, $\mathbf{Adv}_{\mathcal{A}}^{\text{PolySDH}} \leq \binom{n}{2}\frac{q\,\ell+1}{p-1} + \binom{\hat{n}}{2}\frac{q\,\ell+1}{p-1} + \binom{n_t}{2}\frac{2\,q\,\ell+2}{p-1} + \frac{q\,\ell+1+2}{p-1} \leq (q_G + q\,\ell + 2\,\ell + 2)^2\,(q\,\ell)/(p-1)$, i.e., $\mathbf{Adv}_{\mathcal{A}}^{\text{PolySDH}} = \mathrm{O}(q_G^2\,q\,\ell/p + (q\,\ell)^3/p)$. □

We can state the generic hardness of the asymptotic Poly-SDH assumption much more simply, as follows.

Corollary 22. *All algorithms that solve (q, ℓ)-Poly-SDH problems with constant probability $\epsilon > 0$ in generic bilinear groups of order p such that $q\,\ell < \mathrm{O}(\sqrt[3]{p})$ require $\Omega(\sqrt{\epsilon\,p/q\,\ell})$ generic operations.*

Generic-Group Complexity of Pluri-SDH

Essentially the same proof as in "Generic-Group Complexity of Poly-SDH" od Appendix gives the following generic-group complexity lower bounds for the Pluri-SDH assumption. We give the bounds for the general (q, ℓ, ℓ')-Pluri-SDH problem, and note that they also apply if we set $\ell' = 1$, which is the relevant assumption for the ring signature proof.

Theorem 23. *Let $\mathcal{A}$ be a computationally unbounded algorithm for the (q, ℓ, ℓ')-Pluri-SDH problem in generic bilinear groups of prime order p, that makes q_G queries to a generic-group oracle. Then,*

$$\Pr\left[\mathcal{A}^{\mathcal{O}}\left(\begin{array}{l} f(g),\ \ f(g^{\alpha_i})_{\substack{i=-\ell',\ldots,\\ -1,1,\ldots,\ell}},\ \left(w_{i,j},\ f(g^{\frac{1}{\alpha_i+w_{i,j}}})\right)_{\substack{1\leq j\leq q\\ i=-\ell',\ldots,\\ -1,1,\ldots,\ell}} \\ \hat{f}(\hat{g}),\ \ \hat{f}(\hat{g}^{\alpha_i})_{\substack{i=-\ell',\ldots,\\ -1,1,\ldots,\ell}} \end{array}\right) = \left(w_i,\ f\left(g^{\frac{r_i}{\alpha_i+w_i}}\right)\right)_i \ :\ w_i \neq w_{i,j}\ ,\ \sum_{i=1}^{\ell} r_i = 1\right]$$

$$= \mathbf{Adv}^{\text{PluriSDH}}_{\mathcal{A}} \leq \frac{(q_G + q\,\ell' + 2\,\ell' + 2\,\ell)^2\,(q\,\ell' + \ell)}{p-1}$$
$$= \mathrm{O}\left(\frac{q_G^2\,(q\,\ell' + \ell) + (q\,\ell' + \ell)^3}{p}\right).$$

Proof. The proof is analogous to that of Theorem 21. □

We can restate the generic hardness of the asymptotic Pluri-SDH assumption more concisely:

Corollary 24. *Any algorithm that solves the* (q, ℓ, ℓ')*-Pluri-SDH problem with constant probability* $\epsilon > 0$ *in generic bilinear groups of prime order* p *such that* $\ell \leq q\,\ell' < \mathrm{O}(\sqrt[3]{p})$ *requires* $\Omega(\sqrt{\epsilon\,p/q\,\ell'})$ *generic-group operations.*

On Selecting the Group Order

A recent analysis of SDH shows that it may be helpful (but not required) to ensure that $p \pm 1$ be "rough" or "non-smooth" (i.e., have large prime factors). Specifically, in [18] it is shown how to recover the secret α generically from an Original SDH instance, in time as low as $\Theta(\log p\,\sqrt{p/d})$, given a smooth divisor d of $p-1$ or $p+1$ not exceeding $\sqrt[3]{p}$ or $\sqrt[4]{p}$ respectively. Conversely, for p "safe" in the reciprocal sense, it is conjectured that the generic time complexity of breaking SDH is the same "square root" $\Theta(\sqrt{p})$ generic lower bound proven for Discrete Log in [39]. We note that the analysis from [18] does not violate the Original SDH "cube root" $\Theta(\sqrt[3]{p})$ generic lower bound proven in [5], but can come as close as within a logarithmic factor of that "prediction."

For mesh signatures, like for regular signatures, it is enticing to parametrize security based on optimistic conjectures rather than formal generic bounds, in the hope of improving efficiency for a prescribed security level. To benefit from the "square root" conjecture from [18] while avoiding known attacks, requires the simultaneous satisfaction of extra constraints such as the primality of $(p \pm 1)/2$, in conjunction to p itself being prime and amenable to the construction of pairing-friendly groups of order p. Undoubtedly this can complicate the setup.

A safer and more principled approach is to stick to the methodology advocated in [5,6], which, in our context, means to abide by the "cube root" generic bounds given in Sect. 3.3, and select the size of p accordingly. Since our generic bounds are valid for all primes p, this opens up the possibility of working with bilinear groups that might impose their own constraints on p.

All in all, it is a judgment call—but one whose long-term security only affects the non-repudiation security aspect of mesh signatures; recall that anonymity is unconditional.

Appendix 2: Ring Scheme Security Proofs

This section focuses on the security properties of the special-case ring signature scheme of Sect. 4.

Anonymity of the Ring Scheme

We prove Theorem 6 using an information-theoretic argument.

Proof of Theorem 6. We need to show that it is impossible to determine which ring member produced a signature, even if all the secret keys are revealed. Observe that the distribution of σ is uniform over the $(2\,\ell+1)$-dimensional variety of $\mathbb{G}^{\ell+1} \times \mathbb{F}_p^{\ell+1}$ defined by the ring verification equation, i.e.,

$$\mathbf{P}(\sigma) = \mathbf{U}\Bigg((S_0,\ \ldots,\ S_\ell,\ t_0,\ \ldots,\ t_\ell) \in \mathbb{G}^{\ell+1} \times \mathbb{F}_p^{\ell+1} : \prod_{i=0}^{\ell} \mathbf{e}\left(S_i,\ \hat{A}_i\ \hat{B}_i^{m_i}\ \hat{C}_i^{t_i}\right) = \mathbf{e}(g,\ \hat{g}) \Bigg).$$

Conditionally on the public information, the random variable σ is jointly independent of the signer identity and all the secret data (including the random coins used to generate the signing keys). The theorem follows immediately. □

The reader will notice that the independence does not extend conditionally on the ephemerals s_i that were used to create the signature. Even though it is in the interest of the signer to erase those immediately, a powerful adversary could possibly recompute them (by solving a discrete logarithm, as infeasible as it sounds), so one may ask how this affects anonymity. The answer is that is does not, because there are many possible sets of ephemerals, one for each member of the ring, and by Theorem 6 none of them can be preferred over any other based on public or secret information.[5]

Unforgeability of the Ring Scheme

We now prove a stronger version of Theorem 7 where the forger is allowed to make atomic signature queries, based on the stronger $(q, \ell+1)$-Poly-SDH assumption in $\mathbf{G}$. (The simpler Theorem 7 can be proven along the same lines.) The precise statement of the theorem we prove is as follows.

Theorem 25. *The ring signature is existentially unforgeable under an adaptive attack, against a static adversary that makes no more than q ring signature queries, and q atomic signature queries to each one of the ℓ honest users, adaptively, provided that the $(q, \ell+1)$-Poly-SDH assumption holds in $\mathbf{G}$, in the common random string model.*

[5] There remains the possibility that the signer could reveal the ephemerals *voluntarily* to revoke *her own* anonymity, but this falls outside of the purview of Theorem 6: The same outcome can be achieved generically in all ring signatures. To do so, the signer would append to the message being ring-signed, a one-time signature verification key, an encrypted signature of the same under her public key, and a signature of a hash of the decryption key under the one-time key; she could then provably de-anonymize herself by revealing the decryption key.

Proof. Let $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$ be some bilinear instance with a computable isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$. Suppose that $H : \{0,1\}^* \to \mathbb{F}_p$ is a collision-resistant hash function.

We are given a random instance of the $(q, \ell+1)$-Poly-SDH problem in $\mathbf{G}$, stated as $\ell+1$ pairs $(g_i = g^{\alpha_i},\ \hat{g}_i = \hat{g}^{\alpha_i})$ for $i = 0, \ldots, \ell$, and $(\ell+1)\,q$ pairs $(w_{i,j},\ u_{i,j} = g^{1/\alpha_i + w_{i,j}})$ for $i = 0, \ldots, \ell$ and $j = 1, \ldots, q$. Our task is to output $\ell+1$ pairs $(w_i^*,\ u_i^* = g^{r_i^*/\alpha_i + w_i^*})$ for $i = 0, \ldots, \ell$, for some public choice of w_i^* such that $w_i^* \notin \{w_{i,1}, \ldots, w_{i,q}\}$ for all $i = 0, \ldots, \ell$, and some secret choice of r_i^* such that $\sum_{i=0}^{\ell} r_i^* = 1 \pmod{p}$.

We construct an algorithm $\mathcal{B}$ that solves such instances of the Poly-SDH problem by interacting with a black-box forger $\mathcal{A}$ for the ring scheme. For simplicity, we give sequential numbers in $\{1, 2, \ldots\}$ to all users (i.e., potential ring members). Since the adversary is static, we suppose w.l.o.g. that the target users will consist of the set $\{1, \ldots, \ell\}$. For simplicity of notation, we further suppose that the target ring for the forgery is the whole set $R^* = \{1, \ldots, \ell\}$, rather than any subset. Indeed, in the model of Sect. 2.3, it is clear that a signature by a subset ring "implies" a signature by the whole ring; we find it more convenient for this first proof to assume a forgery with a fixed form. (The fully general case of mesh signatures will be addressed in a later section).

In this setting, we use the index $i = 0$ for the key "in the sky", and assume that $\mathcal{B}$ has ownership of the first ℓ users with $1 \le i \le \ell$, and that $\mathcal{A}$ controls all the others with $\ell+1 \le i \le I_{\max}$ for some polynomial bound $I_{\max}$. Each player reveals the public keys of the users in its custody, first $\mathcal{B}$, and then $\mathcal{A}$, without revealing the private keys.

For generality and homogeneity of notation, we will allow each ring signature component to bear on an arbitrary message in $\mathbb{F}_p$, i.e., we do not require that the messages be the same.

Our proof makes use of two distinct simulations, one of which will be chosen at random, to account for two possible behaviors of the adversary. □

Regular Simulation. We first describe a "regular" simulation that $\mathcal{B}$ can use when the forgery returned by the adversary $\mathcal{A}$ never causes the associated parameter w_i^* to fall in the set $\{w_{i^*,1}, \ldots, w_{i^*,q}\}$. The reduction for this case follows.

> To start, the simulator $\mathcal{B}$ must fix the common random string from the distribution expected by $\mathcal{A}$. First, $\mathcal{B}$ publishes the bilinear instance $\mathbf{G}$ and the isomorphism ψ. Next, $\mathcal{B}$ chooses random $b_0,\ c_0 \in \mathbb{F}_p^\times$ and publishes $(\hat{A}_0,\ \hat{B}_0,\ \hat{C}_0) = (\hat{g}_0,\ \hat{g}^{b_0},\ \hat{g}^{c_0})$ as the public key "in the sky"; this implicitly reveals $(A_0,\ B_0,\ C_0) = (g_0,\ g^{b_0},\ g^{c_0})$. Last, $\mathcal{B}$ publishes the description of H.
>
> $\mathcal{B}$ gives $\mathcal{A}$ the public keys of the first ℓ users. To do so, for each $i = 1, \ldots, \ell$, it draws random $b_i,\ c_i \in \mathbb{F}_p^\times$ and publishes the tuple $(A_i,\ B_i,\ C_i,\ \hat{A}_i,\ \hat{B}_i,\ \hat{C}_i) = (g_i,\ g^{b_i},\ g^{c_i},\ \hat{g}_i,\ \hat{g}^{b_i},\ \hat{g}^{c_i})$.
>
> $\mathcal{A}$ gives $\mathcal{B}$ the public keys $(A_i,\ B_i,\ C_i,\ \hat{A}_i,\ \hat{B}_i,\ \hat{C}_i)$ of the users it controls, $i = (\ell+1), \ldots, I_{\max}$. Here, $I_{\max}$ is the total number of users, which must be polynomially bounded.
>
> Per the unforgeability model, $\mathcal{A}$ is allowed to give those values to $\mathcal{B}$ interactively, interleaved with the signature queries described below. We only require that $\mathcal{A}$ reveal a public key before it makes any signature query that involves that key.

$\mathcal{A}$ makes q_S distinct ring signature queries to $\mathcal{B}$, one at a time, proceeding adaptively. To exhaust the quota of queries that are available to $\mathcal{A}$, we assume w.l.o.g. that $q_S = q$.

For $j = 1, \ldots, q_S$, the j-th query is a pair $(M_j,\ R_j)$ where R_j is a ring and M_j is a vector of messages to be signed by that ring. Let n be the number of users in the ring. Let thus $R_j = \{i_1, \ldots, i_n\} \subseteq \{1, \ldots, I_{\max}\}$, and $M_j = (m_1, \ldots, m_n) \in \mathbb{F}_p^n$. We require that $R_j \cap \{1, \ldots, \ell\} \neq \emptyset$; otherwise the simulator is not supposed to respond and the query will be denied.

$\mathcal{B}$ responds to a well-formed j-th query as follows.

Select n random exponents $s_1, \ldots, s_n \in \mathbb{F}_p$, and n random integers $t_1, \ldots, t_n \in \mathbb{F}_p$. Compute $m_0 = H((i_1, m_1), \ldots, (i_n, m_n))$ and set $t_0 = (w_{0,j} - b_0\, m_0)/c_0$. Here, b_0 and c_0 are the secret exponents chosen during setup, and $w_{0,j}$ is taken from $(w_{0,j},\ u_{0,j})$ in the problem instance.

$\mathcal{B}$ replies to $\mathcal{A}$ by returning the signature $\sigma_j = (S_0,\ \ldots,\ S_n,\ t_0,\ \ldots,\ t_n)$, which is given by,

$$\sigma_j = \left(\left(u_{0,j} \cdot \prod_{k=1}^{n} \left(A_{i_k}\, B_{i_k}^{m_k}\, B_{i_k}^{t_k} \right)^{-s_k} \right),\ \left(A_0\, B_0^{m_0}\, C_0^{t_0} \right)^{s_1},\ \ldots,\ \left(A_0\, B_0^{m_0}\, C_0^{t_0} \right)^{s_n},\ t_0,\ t_1,\ \ldots,\ t_n \right).$$

Note that σ_j is properly randomized and passes the validity test: $\prod_{k=0}^{n} \mathbf{e}(S_k,\ \hat{A}_{i_k}\, \hat{B}_{i_k}^{m_k}\, \hat{C}_{i_k}^{t_k}) = \mathbf{e}\left(g^{\frac{1}{\alpha_0 + w_{0,j}}},\ \prod_{k=1}^{n} (A_{i_k}\, B_{i_k}^{m_k}\, C_{i_k}^{t_k})^{-s_k},\ \hat{A}_0\, \hat{B}_0^{m_0}\, \hat{C}_0^{t_0} \right) \cdot \prod_{k=1}^{n} \mathbf{e}\left((A_0\, B_0^{m_0}\, C_0^{t_0})^{s_k},\ \hat{A}_{i_k}\, \hat{B}_{i_k}^{m_k}\, \hat{C}_{i_k}^{t_k} \right) = \mathbf{e}\left(g^{\frac{1}{\alpha_0 + w_{0,j}}},\ \hat{A}_0\, \hat{B}_0^{m_0}\, \hat{C}_0^{t_0} \right) = \mathbf{e}\left(g^{\frac{1}{\alpha_0 + w_{0,j}}},\ \hat{g}^{\alpha_0 + w_{0,j}} \right) = \mathbf{e}(g,\ \hat{g})$, as required.

$\mathcal{A}$ can also make $q_{S,i}$ distinct atomic signature queries for each user $i = 1, \ldots, \ell$ controlled by $\mathcal{B}$. These queries can be adaptive and interleaved with the ring signature queries. Again to exhaust the quota of queries that are available to $\mathcal{A}$, we assume w.l.o.g. that $q_{S,i} = q$ for $i = 1, \ldots, \ell$.

$\mathcal{B}$ responds to the j-th query (i, m) to the i-th user, by retrieving the pair $(w_{i,j},\ u_{i,j})$ from the Poly-SDH instance, computing $t = (w_{i,j} - b_i\, m)/c_i$, and returning the Boneh–Boyen signature $(u_{i,j},\ t)$ to $\mathcal{A}$.

$\mathcal{A}$ eventually outputs a forgery σ^* on some list of messages $M^* = (m_1^*, \ldots, m_\ell^*)$ attributed to the target ring $R^* = \{1, \ldots, \ell\}$. To preclude trivial forgeries, we demand that $(M^* \neq M_j) \vee (R^* \neq R_j)$ for all queries (M_j, R_j) previously made by $\mathcal{A}$.

The forgery is a signature of the form $\sigma^* = (S_0^*,\ \ldots,\ S_\ell^*,\ t_0^*,\ \ldots,\ t_\ell^*)$ and must necessarily satisfy, for $m_0^* = H((1, m_1^*), \ldots, (\ell, m_\ell^*))$,

$$\prod_{i=0}^{\ell} \mathbf{e}\left(S_i^*,\ \hat{A}_i\, \hat{B}_i^{m_i^*}\, \hat{C}_i^{t_i^*} \right) = \mathbf{e}(g,\ \hat{g}).$$

It follows that there exist ℓ unknown exponents $r_1^*, \ldots, r_\ell^* \in \mathbb{F}_p$ such that,

$$S_0 = \left(g^{\frac{1}{\alpha_0 + b_0\, m_0^* + c_0\, t_0^*}} \right)^{1 - \sum_{k=1}^{\ell} r_k^*}, \quad \text{and} \quad S_i = \left(g^{\frac{1}{\alpha_i + b_i\, m_i^* + c_i\, t_i^*}} \right)^{r_i^*} \quad \text{for } i = 1, \ldots, \ell.$$

Thus, if we compute $w_i^* = b_i\, m_i^* + c_i\, t_i^*$ for $i = 0, \ldots, \ell$, and formally define $r_0^* = 1 - \sum_{k=1}^{\ell} r_k^*$, we obtain the following solution to the Poly-SDH problem instance,

$$\begin{aligned} &((w_0^*,\ S_0^*),\ \ldots,\ (w_\ell^*,\ S_\ell^*)) \\ &\quad = \left(\left(w_0^*,\ u_0^* = g^{\frac{r_0^*}{\alpha_0 + w_0^*}} \right),\ \ldots,\ \left(w_\ell^*,\ u_\ell^* = g^{\frac{r_\ell^*}{\alpha_\ell + w_\ell^*}} \right) \right). \end{aligned}$$

$\mathcal{B}$ can compute all the w_i^* and u_i^*, even though it does not know any of the r_i^*.

The preceding reduction will succeed to produce a solution to the given Poly-SDH instance, when the adversary's forgery is such that $w_i^* \notin \{w_{i,1}, \ldots, w_{i,q}\}$ for all $i = 0, \ldots, \ell$.

However, $\mathcal{A}$ could produce an anomalous forgery where $w_{i^*}^* \in \{w_{i^*,1}, \ldots, w_{i^*,q}\}$ for some $i^* \leq \ell$, either by luck or by design. The preceding simulation is useless in this case, since it would produce an easy solution to the Poly-SDH problem that is explicitly forbidden. We use an "alternative" simulation to deal with this anomaly.

Alternative Simulation. We construct an "alternative" reduction for the case where $\mathcal{A}$'s forgery causes the parameter $w_{i^*}^*$ to land in the set $\{w_{i^*,1}, \ldots, w_{i^*,q}\}$ of parameters explicitly listed in the problem instance given to $\mathcal{B}$. The simulation is as follows.

To start, the simulator $\mathcal{B}$ posts a common random string from the distribution expected by $\mathcal{A}$. To do so, $\mathcal{B}$ publishes the bilinear instance $\mathbf{G}$ and the isomorphism ψ; it chooses random $a_0,\ b_0 \in \mathbb{F}_p^\times$ and publishes $(\hat{A}_0,\ \hat{B}_0,\ \hat{C}_0) = (\hat{g}^{a_0},\ \hat{g}^{b_0},\ \hat{g}_0)$; it also publishes a description of H.

$\mathcal{B}$ gives $\mathcal{A}$ the public keys of the first ℓ users. To do so, for each $i = 1, \ldots, \ell$, it draws random $a_i,\ b_i \in \mathbb{F}_p^\times$ and publishes the tuple $(A_i,\ B_i,\ C_i,\ \hat{A}_i,\ \hat{B}_i,\ \hat{C}_i) = (g^{a_i},\ g^{b_i},\ g_i,\ \hat{g}^{a_i},\ \hat{g}^{b_i},\ \hat{g}_i)$.

$\mathcal{A}$ gives $\mathcal{B}$ the public keys $(A_i,\ B_i,\ C_i,\ \hat{A}_i,\ \hat{B}_i,\ \hat{C}_i)$ of the users it controls, $i = (\ell + 1), \ldots, I_{\max}$. Again, $I_{\max}$ is the total number of users, which must be polynomially bounded.
As before, $\mathcal{A}$ may keep introducing new keys after it has started making signature queries, as long as public keys are revealed before queries that depend on them.

$\mathcal{A}$ makes q distinct ring signature queries to $\mathcal{B}$, proceeding adaptively. As before, the j-th query is given as a pair $(M_j,\ R_j)$ with $M_j = (m_1, \ldots, m_n) \in (\mathbb{F}_p)^n$ for some n and $R_j = \{i_1, \ldots, i_n\} \subseteq \{1, \ldots, I_{\max}\}$ such that $R_j \cap \{1, \ldots, \ell\} \neq \emptyset$. Upon receiving this query, $\mathcal{B}$ responds as follows.
$\mathcal{B}$ computes $m_0 = H((i_1, m_1), \ldots, (i_n, m_n))$ and sets $t_0 = (a_0 + b_0\, m_0)/w_{0,j}$ using the j-th pair $(w_{0,j}, u_{0,j})$ from the problem instance. $\mathcal{B}$ also computes $V_j =$

$B_0^{m_0}\, C_0^{t_0}$ and stores the tuple $(0,\ V_j,\ m_0,\ t_0,\ M_j,\ R_j)$ in some searchable database for future use.

$\mathcal{B}$ then chooses n random exponents $s_1,\ \ldots,\ s_n \in \mathbb{F}_p$, and n random integers $t_1,\ \ldots,\ t_n \in \mathbb{F}_p$, and answers the query by giving to $\mathcal{A}$ the signature,

$$\sigma_j = \left(\left(u_{0,j}^{1/t_0} \cdot \prod_{k=1}^{n} (A_{i_k}\, B_{i_k}^{m_k}\, C_{i_k}^{t_k})^{-s_k} \right),\ \left(A_0\, B_0^{m_0}\, C_0^{t_0}\right)^{s_1},\ \ldots,\ \left(A_0\, B_0^{m_0}\, C_0^{t_0}\right)^{s_n},\ t_0,\ t_1,\ \ldots,\ t_n \right).$$

We show that the signature $\sigma_j = (S_0,\ \ldots,\ S_n,\ t_0,\ \ldots,\ t_n)$ has the correct distribution. First, it has the requisite $2\,n + 1$ degrees of freedom, and so is adequately randomized. Second, in the verification equation, all the factors under the pairings cancel out, except for $u_{0,j}^{1/t_0}$. Therefore, writing $i_0 = 0$, we find, $\prod_{k=0}^{n} \mathbf{e}(S_k,\ \hat{A}_{i_k}\, \hat{B}_{i_k}^{m_k}\, \hat{C}_{i_k}^{t_k}) = \mathbf{e}((u_{0,j}^{1/t_0}),\ \hat{A}_0\, \hat{B}_0^{m_0}\, \hat{C}_0^{t_0}) = \mathbf{e}(g^{\frac{w_{0,j}/(a_0+b_0\, m_0)}{\alpha_0+w_{0,j}}},\ \hat{g}^{a_0+b_0\, m_0}\, \hat{g}_0^{t_0}) = \mathbf{e}(g^{\frac{w_{0,j}}{\alpha_0+w_{0,j}}},\ \hat{g})\, \mathbf{e}(g^{\frac{1}{\alpha_0+w_{0,j}}},\ \hat{g}^{\alpha_0}) = \mathbf{e}(g,\ \hat{g})$, as required.

$\mathcal{A}$ also makes q distinct atomic signature queries for each of the users $i = 1, \ldots, \ell$ controlled by $\mathcal{B}$. These queries are adaptive and interleaved with the ring signature queries.

$\mathcal{B}$ responds to the j-th query (i, m) on behalf of the i-th user, as follows.

It retrieves the fresh pair $(w_{i,j},\ u_{i,j})$ from the Poly-SDH instance, and sets $t = (a_i + b_i\, m)/w_{i,j}$. It also computes $\mathcal{B}$ also computes $V = B_i^m\, C_i^t$ and stores the tuple $(i,\ V,\ m,\ t,\ \emptyset,\ \emptyset)$ the database for future use. It then returns the Boneh–Boyen signature $(u_{i,j}^{1/t},\ t)$ to $\mathcal{A}$.

$\mathcal{A}$ finally outputs a forgery σ^* bearing on a message vector $M^* = (m_1^*, \ldots, m_\ell^*)$ and the target ring $R^* = \{1, \ldots, \ell\}$, provided that $(M^* \neq M_j) \vee (R^* \neq R_j)$ for all queries (M_j, R_j) made earlier. We let $m_0^* = H((1, m_1^*), \ldots, (\ell, m_\ell^*))$.

The forgery $\sigma^* = (S_0^*,\ \ldots,\ S_\ell^*,\ t_0^*,\ \ldots,\ t_\ell^*)$ is not valid unless $\prod_{i=0}^{\ell} \mathbf{e}(S_i,\ \hat{A}_i\, \hat{B}_i^{m_i^*}\, \hat{C}_i^{t_i^*}) = \mathbf{e}(g,\ \hat{g})$, i.e., there must exist $r_1,\ \ldots,\ r_\ell \in \mathbb{F}_p$ and $r_0 = 1 - \sum_{k=1}^{\ell} r_k$ such that,

$$S_i = \left(g^{\frac{1}{a_i + b_i\, m_i^* + \alpha_i\, t_i^*}} \right)^{r_i} \quad \text{for } i = 0, \ldots, \ell.$$

For $i = 0, \ldots, \ell$, let us define $w_i^* = b_i\, m_i^* + c_i\, t_i^*$ for the value $c_i = \mathrm{dlog}_g(C_i) = \alpha_i$ that would have given $\mathcal{A}$ the same view in the regular simulation. Let us also define $w_i^{**} = (a_i + b_i\, m_i^*)/t_i^*$ with a_i and b_i as chosen in the present simulation. Observe that $\mathcal{B}$ is unable to compute any of the w_i^*, but it can and does compute all the w_i^{**}.

$\mathcal{B}$ exploits all of this as follows. It computes $V_i^* = (B_i)^{m_i^*}\, (C_i)^{t_i^*} = g^{w_i^*}$ for $i = 0, \ldots, \ell$, and then searches the database for an entry $(i,\ V_j,\ m_i^{(j)},\ t_i^{(j)},\ M_j,\ R_j)$ such that $V_j = V_i^*$, or equivalently, such that $b_i\, m_i^{(j)} + \alpha_i\, t_i^{(j)} = w_i^*$.

There are three (possibly overlapping) possibilities for success, and one for failure:

1. An entry was found with $V_j = V_i^*$ and $m_i^{(j)} \neq m_i^*$ for some i: In this case, $\mathcal{B}$ can resolve the Poly-SDH instance explicitly by recovering the secret exponent $\alpha_i \in \mathbb{F}_p$,

$$\alpha_i = \frac{\left(t_i^* - t_i^{(j)}\right)}{b_i \, (m_i^{(j)} - m_i^*)}.$$

2. An entry was found with $V_j = V_0^*$ and $m_0^{(j)} = m_0^*$: In this case, $\mathcal{B}$ has found a collision in the supposedly collision-resistant hash function H, since we have,

$$\overbrace{H\left(\underbrace{f(M_j, R_j)}\right)}^{m_0^{(j)}} = \overbrace{H\left(\underbrace{f(M^*, R^*)}\right)}^{m_0^*}.$$
$$\left(\left(i_1, m_1^{(j)}\right), \ldots, \left(i_n, m_n^{(j)}\right)\right) \neq \left(\left(1, m_1^*\right), \ldots, \left(\ell, m_\ell^*\right)\right)$$

3. An entry was found with $V_j = V_i^*$ and $m_i^{(j)} = m_i^*$ for $i \neq 0$: In this case, the ring signature forgery includes a clause $\left[VK_i : m_i^*\right]$ on which $\mathcal{A}$ had previously made an atomic signature query (it being the j-th query to the i-th user). The forgery is therefore inadmissible and this case may be discounted (regardless of whether $V_j = V_i^*$ or $V_j \neq V_i^*$).
4. No entry that matches the conditions was found in the database: This corresponds to the event that $\forall i, \forall j, (w_i^* \neq b_i \, m_i^{(j)} + c_i \, t_i^{(j)})$ where $c_i = \mathrm{dlog}_g(C_i)$. In this case $\mathcal{B}$ is stuck, because $w_i^* \notin \{w_{i^*,1}, \ldots, w_{i^*,q}\}$, but from the point of view of $\mathcal{A}$ this is precisely when the "regular" simulation will be able to proceed to completion.

Fair Prior Apportionment. Even though $\mathcal{B}$ will not know in advance which of the "regular" and "alternative" simulation should be used, we can see that they all appear identical to the adversary, so by a standard argument, if $\mathcal{B}$ makes a fair random choice at the onset, the final reduction will be successful with probability $\frac{1}{2}$.

For completeness, we mention that the alternative simulation may succeed not by solving the Poly-SDH instance but by finding a hash collision. Since it is possible to build (keyed) collision-resistant hash function families from CDH, and thus from SDH or Poly-SDH, the theorem still holds without the need for a hashing assumption.

Appendix 3: Mesh Scheme Security Proofs

This section focuses on the security properties of the general mesh signature scheme of Sect. 5.

Anonymity of the Mesh Scheme

We prove Theorem 10 using an information-theoretic argument.

Proof of Theorem 10. We need to show that it is impossible to determine which assignment caused $\tilde{\Upsilon}$ to be satisfied in the signature σ. By design of the blinding factors, the distribution of σ is the uniform distribution $\mathbf{U}$ over the $(2\,\ell+1)$-dimensional variety of $\mathbb{F}_p^{\ell+1} \times \mathbb{G}^{\ell+\vartheta+1}$ defined by the mesh verification equation, i.e.,

$$\mathbf{P}(\sigma) = \mathbf{U}\left(\begin{array}{l}(t_0,\ \ldots,\ t_\ell,\ S_1,\ \ldots,\ S_\ell,\ P_0,\ \ldots,\ P_\vartheta) \in \mathbb{F}_p^{\ell+1} \times \mathbb{G}^{\ell+\vartheta+1} : \\ \quad \prod_{k=0}^{\vartheta} \mathbf{e}\left(P_0,\ \hat{A}_{0,k}\,\hat{B}_{0,k}^{m_0}\,\hat{C}_{0,k}^{t_0}\right)^{y_{0,k}} \cdot \prod_{i=1}^{\ell} \mathbf{e}\left(S_i,\ \hat{A}_{i,0}\,\hat{B}_{i,0}^{m_i}\,\hat{C}_{i,0}^{t_i}\right)^{y_{i,0}} = \mathbf{e}(g,\ \hat{g}_0) \\ \quad \wedge \prod_{k=0}^{\vartheta} \mathbf{e}\left(P_1,\ \hat{A}_{0,k}\,\hat{B}_{0,k}^{m_0}\,\hat{C}_{0,k}^{t_0}\right)^{y_{0,k}} \cdot \prod_{i=1}^{\ell} \mathbf{e}\left(S_i,\ \hat{A}_{i,1}\,\hat{B}_{i,1}^{m_i}\,\hat{C}_{i,1}^{t_i}\right)^{y_{i,1}} = 1 \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \vdots \\ \quad \wedge \prod_{k=0}^{\vartheta} \mathbf{e}\left(P_\vartheta,\ \hat{A}_{0,k}\,\hat{B}_{0,k}^{m_0}\,\hat{C}_{0,k}^{t_0}\right)^{y_{0,k}} \cdot \prod_{i=1}^{\ell} \mathbf{e}\left(S_i,\ \hat{A}_{i,\vartheta}\,\hat{B}_{i,\vartheta}^{m_i}\,\hat{C}_{i,\vartheta}^{t_i}\right)^{y_{i,\vartheta}} = 1\end{array}\right),$$

where $\forall i = 0, \ldots, \ell$ the exponents $y_{i,k}$ come from the polynomial coefficients of $\pi_i = \sum_{k=0}^{\vartheta} y_{i,k}\, Z_k$, whereas $\forall k = 0, \ldots, \vartheta$ the elements $\hat{A}_{i,k}$, $\hat{B}_{i,k}$, $\hat{C}_{i,k}$ come from the public verification keys VK_i.

First, we observe that $\mathbf{P}(\sigma)$ is indeed uniform over the stated $(2\,\ell+1)$-dimensional variety. Indeed, σ lives in a $(2\,\ell+\vartheta+2)$-dimensional space and is subject to $\vartheta+1$ independent linear constraints in the verification algorithm, so it has at most $2\,\ell+1$ degrees of freedom. Conversely, for each joint random assignment to $t_0, \ldots, t_\ell$ and $s_1, \ldots, s_\ell$ there exists a distinct valid signature based on that assignment, and so σ has at least $2\,\ell+1$ degrees of freedom, which are then as stated.

Next, we verify that $\mathbf{P}(\sigma)$ is indeed independent of the true signer, and more generally of the secret linear combination $\nu_0, \ldots, \nu_\ell$ that was used in the creation of σ (conditionally on Υ being satisfied). This is immediate since the polynomials π_i and their coefficients $y_{i,k}$ form a public encoding of Υ determined independently of the truth-value assignment χ from which the ν_i derive.

Last, we note that the hash value $Msg_0 = H(\left[VK_1 : Msg_1\right], \ldots, \left[VK_\ell : Msg_\ell\right], \Upsilon)$ and all ancillary information adjoined to σ is itself a function of public information only. The theorem follows. □

Unforgeability of the Mesh Scheme

We prove Theorem 11 from the $(q, \ell+1)$-Poly-SDH assumption in $\mathbf{G}$.

Proof of Theorem 11. Let $\mathbf{G} = (p, \mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_t, g, \hat{g}, \mathbf{e})$ be a bilinear instance with a computable isomorphism $\psi : \hat{\mathbb{G}} \to \mathbb{G}$, and suppose that $H : \{0,1\}^* \to \mathbb{F}_p$ is a collision-resistant hash function.

As before, we are given a random instance of the $(q, \ell+1)$-Poly-SDH problem in $\mathbf{G}$, stated as $\ell+1$ pairs $(g^{\alpha_i}, \hat{g}^{\alpha_i})$ for $i = 0, \ldots, \ell$, and $(\ell+1)\,q$ pairs $(w_{i,j},\ u_{i,j} = g^{1/\alpha_i + w_{i,j}})$ for $i = 0, \ldots, \ell$ and $j = 1, \ldots, q$. Our task is to output

$\ell + 1$ pairs $(w_i^*,\ u_i^* = g^{r_i^*/\alpha_i + w_i^*})$ for $i = 0, \ldots, \ell$, for some public choice of w_i^* such that $w_i^* \notin \{w_{i,1}, \ldots, w_{i,q}\}$ for all i, and some secret choice of r_i^* such that $\sum_{i=0}^{\ell} r_i^* = 1 \pmod p$. □

We construct an algorithm $\mathcal{B}$ that solves such instances of the Poly-SDH problem by interacting with a black-box forger $\mathcal{A}$ for the mesh scheme. For simplicity, we give sequential numbers in $\{1, 2, \ldots\}$ to all users (that is, to all individuals that can be potentially named as mesh members).

Since the adversary is static, we suppose w.l.o.g. that the target users will consist of the set $\{1, \ldots, \ell\}$. Per the mesh security model, a successful forgery must bear on a mesh expression Υ over an arbitrary set of users, as long as it logically implies a weaker expression Υ' that involves only the users in $\{1, \ldots, \ell\}$, and no component that matches an earlier atomic query. Recall that this is to ensure that Υ remains falsifiable even if the adversary sets all the literals under its control to $\top$, lest we accept a trivial forgery. We remark that a maximally weak such Υ' representable in the language is a disjunction over all target users in $\{1, \ldots, \ell\}$, i.e., the ring signature disjunction $\Upsilon'' = \vee_{i=1}^{\ell} \big[VK_i : Msg_i\big]$. Here, the clauses $\big[VK_i : Msg_i\big]$ will be the same as in the original forgery, except for those that were the object of earlier atomic queries in which case they are replaced by $\big[VK_i : Msg_i'\big]$ for some arbitrary Msg_i'. Observe that Υ'' specifies a ring signature forgery analogous to that in the proof of Theorem 25.

Our strategy will thus be, first, to accept from $\mathcal{A}$ a forgery σ with formula Υ on some arbitrary set of users. Next, we transform σ into a pseudo-signature σ' with a weaker formula Υ' defined as a disjunction of the clauses $\big[VK_i : Msg_i\big]$ from Υ that involve no adversarial user and match no prior atomic signature query, i.e., only clauses under the simulator's control are permitted in Υ'; the result σ' is a ring signature that is technically invalid because the clause $[VK_0 : m_0]$ for the "key in the sky" pertains to the original hash value $m_0 = H(\ldots, \Upsilon)$. Then, we transform σ' into an even weaker ring pseudo-signature σ'' whose formula Υ'' is a disjunction over the entire ring $R^* = \{1, \ldots, \ell\}$ of all honest users: for the latter step we simply add a clause $\big[VK_i : Msg_i'\big]$ for each user of index $i \le \ell$ that is missing from Υ'; the messages Msg_i' may be arbitrary and match Msg_i from the forgery, as long as we steer clear from all messages used in atomic queries made to VK_i. From σ'', which still contains the original hash m_0 but would otherwise be a valid ring signature, $\mathcal{B}$ can compute a solution to the Poly-SDH instance exactly as in Theorem 25.

It suffices to show how to respond to well-formed mesh signatures and atomic signature queries during the adaptive query phase, and then how to effect the final transformation $\sigma \mapsto \sigma' \mapsto \sigma''$ corresponding to $\Upsilon \mapsto \Upsilon' \mapsto \Upsilon''$ at the end of the game.

Regular Simulation. We merely show how the regular mesh simulation generalizes that of the basic ring signature.

> To start, $\mathcal{B}$ publishes the bilinear instance $\mathbf{G}$, the isomorphism ψ, and the hash function H. Furthermore, $\mathcal{B}$ chooses $\lambda + 1$ random exponents $z_0, \ldots, z_\lambda \in \mathbb{F}_p^\times$, and for each k publishes the reference element $\hat{g}_k = (\hat{g}^{\alpha_0})^{z_k}$, based on the value of $\hat{g}^{\alpha_0}$ given in the Poly-SDH instance. This also induces $g_k = (g^{\alpha_0})^{z_k} = \psi(\hat{g}_k)$, for $k = 0, \ldots, \lambda$.

Additionally, $\mathcal{B}$ publishes the verification key "in the sky," consisting of $3\,(\lambda+1)$ elements $\hat{A}_{0,k} = (\hat{g}^{\alpha_0})^{a_{0,k}}$, $\hat{B}_{0,k} = \hat{g}^{b_{0,k}}$, $\hat{C}_{0,k} = \hat{g}^{c_{0,k}}$ with random $a_{0,k}, b_{0,k}, c_{0,k} \in \mathbb{F}_p^\times$ for $k = 0, \ldots, \lambda$. (Notice that the $\hat{A}_{0,k}$ are powers of $\hat{g}^{\alpha_0}$ given in the Poly-SDH instance, while the $\hat{B}_{0,k}$ and $\hat{C}_{0,k}$ are mere powers of $\hat{g}$.)

$\mathcal{B}$ gives $\mathcal{A}$ the public keys of the first ℓ users. To do so, for each $i = 1, \ldots, \ell$, it draws random exponents $b_i, c_i \in \mathbb{F}_p^\times$ and publishes $\hat{A}_{i,k} = (\hat{g}^{\alpha_i})^{z_k}$, $\hat{B}_{i,k} = \hat{g}^{b_i\, z_k}$, $\hat{C}_{i,k} = \hat{g}^{c_i\, z_k}$ for $k = 0, \ldots, \lambda$. (Here, the $\hat{A}_{i,k}$ derive from the $\hat{g}^{\alpha_i}$ from the Poly-SDH instance, whereas the $\hat{B}_{i,k}$ and $\hat{C}_{i,k}$ are known powers of $\hat{g}$.)

$\mathcal{A}$ gives $\mathcal{B}$ the public keys of the users it controls, of indices $i = (\ell+1), \ldots, I_{\max}$, during the course of queries. For each such i, a key consists of $3\,(\lambda+1)$ elements $\hat{A}_{i,k}$, $\hat{B}_{i,k}$, $\hat{C}_{i,k}$, for $k = 0, \ldots, \lambda$.
It must be the case that $\mathrm{dlog}_{\hat{g}_0}(\hat{A}_{i,0}) = \ldots = \mathrm{dlog}_{\hat{g}_\lambda}(\hat{A}_{i,\lambda})$, and similarly for the $\hat{B}_{i,k}$ and the $\hat{C}_{i,k}$, which the simulator can easily verify using the pairing.

$\mathcal{A}$ makes $q_S = q$ distinct mesh signature queries to $\mathcal{B}$, one at a time, proceeding adaptively.
Each query is a well-formed mesh statement Υ to be signed by any coalition of users that can satisfy Υ. Let n be the number of mesh literals $L_1, \ldots, L_n$, and let the corresponding clauses be $\left[VK_{i_1} : m_1\right], \ldots, \left[VK_{i_n} : m_n\right]$. We require that Υ be unsatisfiable using only clauses with user indices $i_j > \ell$, otherwise the adversary should be able to create a signature without having to query for it. (We remark that the simulator is still able to respond even in this case, using the signing key "in the sky," but chooses not to do so.)
The adversary may also specify, as part of the query, an atomic signature (u_j, t_j) on m_j for any number of users $i_j > \ell$ in its control. This is to simulate the scenario where the signer wishes to make use of atomic signatures on messages it knows, e.g., PKI certificates on designated users' keys, but lacks the ability to create different signatures on those messages. (In the simulation, $\mathcal{B}$ simply ignores the given u_j and uses the t_j instead of random values.)
$\mathcal{B}$ responds to a well-formed j-th query very much as in the ring signature simulation of Theorem 25. The differences are as follows.
First, $\mathcal{B}$ creates $\tilde{\Upsilon} = \Upsilon \vee L_0$ where L_0 corresponds to the proposition $[VK_0 : m_0]$ with $m_0 = H(\left[VK_1 : Msg_1\right], \ldots, \left[VK_n : Msg_n\right], \Upsilon)$, per the mesh scheme. $\mathcal{B}$ computes a representation of $\tilde{\Upsilon}$ as a list of degree-1 polynomials $\pi_0, \ldots, \pi_n \in \mathbb{F}_p[Z_0, \ldots, Z_\vartheta]$ with coefficients $y_{i,k} \in \mathbb{F}_p$.
Next, $\mathcal{B}$ satisfies $\tilde{\Upsilon}$ with a truth assignment such that $\chi(L_0) = \top$ and $\chi(L_i) = \bot$ everywhere else. Accordingly, $\mathcal{B}$ uses the $(0, j)$-th pair $(w_{0,j},\ u_{0,j})$ from the Poly-SDH instance to obtain an atomic signature on m_0 under a combination of the keys "in the sky" expressed by the polynomial π_0. Precisely, it defines $a_0 =$

$\sum_{k=0}^{\vartheta} a_{0,k}\, y_{0,k}, b_0 = \sum_{k=0}^{\vartheta} b_{0,k}\, y_{0,k}, c_0 = \sum_{k=0}^{\vartheta} c_{0,k}\, y_{0,k}$, and builds the signature as: $(u_0 = {u_{0,j}}^{1/a_0},\ t_0 = \frac{a_0\, w_{0,j} - b_0\, m_0}{c_0})$.[6,7]

Then, $\mathcal{B}$ runs the mesh signing algorithm as in the real scheme, based on the π_i and ν_i it has. Specifically, $\mathcal{B}$ takes t_0, chooses $t_1, \ldots, t_n$ at random (or uses the $\mathcal{A}$-specified values for them), and computes $S_{i_1}, \ldots, S_{i_n}$ and $P_0, \ldots, P_\vartheta$ as in the real scheme. It can do so without knowing any signing key because $\nu_i = 0$ for all users $i \neq 0$; however, the mesh prototype it obtains is invalid precisely for that reason. (The only nonzero coefficient in the linear combination is the coefficient ν_0 that applies to the polynomial π_0, but the signing algorithm of Sect. 5.3 purposely ignores it since in real life it is always zero. For $\tilde{\Upsilon} = \Upsilon \vee L_0$, the algorithm of Sect. 5.1 gives $\pi_0 = Z_0$, and thus $\nu_0 = 1$ per Lemma 8 since all other coefficients are zero. The simulator needs to incorporate ν_0 manually into the prototype signature.)

Thus, $\mathcal{B}$ has to amend the prototype to turn it into a valid mesh signature. This is done by multiplying u_0 into the component P_0. Indeed, because the prototype so far contains only blinding factors and no actual signature, the left-hand sides of all the verification equations resolve to $1 \in \mathbb{G}_t$, including the main equation (the one involving P_0), which should equate to $\mathbf{e}(g, \hat{g}_0) \in \mathbb{G}_t$ instead. A substitution of $u_0^{\nu_0} P_0 = u_0 P_0$ for the prototype's P_0 effects the desired correction without disturbing the other equations.

The resulting mesh signature is given to $\mathcal{A}$. It is valid, and uniformly distributed over the correct space, as we count $2n+1$ degrees of freedom among its $2n+\vartheta+1$ components.

$\mathcal{A}$ also makes $q_{S,i} = q$ distinct atomic signature queries for each user $i = 1, \ldots, \ell$ controlled by $\mathcal{B}$; these can be arbitrarily interleaved with the ring signature queries. $\mathcal{B}$ responds to the j-th query (i, m) to the i-th user, by retrieving the (i, j)-th pair $(w_{i,j},\ u_{i,j})$ from the Poly-SDH instance, computing $t = (w_{i,j} - b_i\, m)/c_i$, and returning the Boneh–Boyen signature $(u_{i,j},\ t)$ to $\mathcal{A}$.

$\mathcal{A}$ finally outputs a forgery $\sigma = (t_0, \ldots, t_n, S_1, \ldots, S_n, P_0, \ldots, P_\vartheta)$ bearing on a well-formed mesh statement Υ such that $\Upsilon \Rightarrow \Upsilon'$ for some other well-formed statement Υ' that involves only users of indices $i \leq \ell$ and no clause that matches an earlier atomic query. W.l.o.g., we can assume that Υ' is a disjunction, since every well-formed formula in the language can be weakened into a disjunction of its inputs.

$\mathcal{B}$ performs the first transformation $\Upsilon \mapsto \Upsilon'$ by eliminating from σ the components of user indices $i > \ell$, or those that match an atomic query, which will produce a ring pseudo-signature σ' on the disjunction Υ'. The process amounts to performing

[6] Notice that, unlike the user keys which have a lot of internal structure (exhibited by the many obvious discrete log relations), the various components of the key "in the sky" are independently distributed and so a signature that will verify under one triple $(\hat{A}_{0,k_1}, \hat{B}_{0,k_1}, \hat{B}_{0,k_1})$ will not verify under another $(\hat{A}_{0,k_2}, \hat{B}_{0,k_2}, \hat{B}_{0,k_2})$. What we need is a signature that will verify for the particular combination of such triples given by the coefficients of π_0.

[7] It can also be shown that for $\tilde{\Upsilon} = \Upsilon \vee L_0$, the algorithm of Sect. 5.1 always gives $\pi_0 = Z_0$, and so we have $y_{0,k} = 0$ for $k \neq 0$. It follows that the public key "in the sky" only needs one triple $(\hat{A}_{0,0}, \hat{B}_{0,0}, \hat{B}_{0,0})$ instead of $\lambda + 1 \geq \vartheta + 1$ of them. We omit this optimization from the present proof to avoid further complicating the argument, but since it greatly shortens the CRS we explicitly recommend its use in Sect. 5.5.

Gaussian elimination of every variable S_i that corresponds to an undesirable clause, in the linear system that lives "in the exponents" of the verification equations. Each step of the Gaussian elimination will "consume" exactly one of the ϑ supplemental verification equation (i.e., those that involve some P_k for $k \neq 0$), with the aim of eliminating one of the remaining offending S_i, thus chosen as our "pivot." As in classical Gaussian elimination, only an S_i whose coefficient $y_{i,k} \neq 0$ at the end of a previous step can serve as pivot and hence be eliminated by consuming the next equation in P_k. Observe that we do not eliminate P_0. A straightforward argument shows that if $\Upsilon \Rightarrow \Upsilon'$ where Υ' represents a disjunction over $L_0, \ldots, L_\ell$, then a pseudo-signature corresponding to Υ' can be obtained by the Gaussian elimination process ("pseudo" because, once again, the hash value m_0 is not updated in the process). We now describe Gaussian elimination "in the exponent".

Suppose w.l.o.g. that $\mathcal{B}$ seeks to eliminate S_n from σ, where S_n appears with exponent $y_{n,\vartheta} \neq 0$ in the equation that involves P_ϑ, i.e.,

$$\mathbf{e}\left(P_\vartheta,\ \hat{v}_0\right) \cdot \prod_{i=1}^{n} \mathbf{e}\left(S_i,\ \left(\hat{A}_{i,\vartheta}\ \hat{B}_{i,\vartheta}^{m_i}\ \hat{C}_{i,\vartheta}^{t_i}\right)^{y_{i,\vartheta}}\right) = 1.$$

The idea is to find a linear combination that will cancel out S_n in the main equation, i.e.,

$$\mathbf{e}\left(P_0,\ \hat{v}_0\right) \cdot \prod_{i=1}^{n} \mathbf{e}\left(S_i,\ \left(\hat{A}_{i,0}\ \hat{B}_{i,0}^{m_i}\ \hat{C}_{i,0}^{t_i}\right)^{y_{i,0}}\right) = \mathbf{e}(g,\ \hat{g}_0).$$

This is easy to do once we observe that $(\hat{A}_{i,\vartheta}\ \hat{B}_{i,\vartheta}^{m_i}\ \hat{C}_{i,\vartheta}^{t_i})^{z_0} = (\hat{A}_{i,0}\ \hat{B}_{i,0}^{m_i}\ \hat{C}_{i,0}^{t_i})^{z_\vartheta}$ for all $i = 1, \ldots, n$. We raise both sides of the P_ϑ equation (on top) to the power of $\rho = \frac{y_{n,0}\ z_0}{y_{n,\vartheta}\ z_\vartheta}$, and divide the result into the main equation, causing the pairing with S_n to vanish. We then consolidate the new terms into the existing ones to preserve the form of the equation. That is, we let $P_0' = P_0 / P_\vartheta^\rho$ replace the ratio of P_0 and P_ϑ^ρ, and for all $i = 1, \ldots, n-1$, we merge the ratio of $\mathbf{e}(S_i,\ (\hat{A}_{i,0}\ \hat{B}_{i,0}^{m_i}\ \hat{C}_{i,0}^{t_i})^{y_{i,0}})$ and $\mathbf{e}(S_i,\ (\hat{A}_{i,\vartheta}\ \hat{B}_{i,\vartheta}^{m_i}\ \hat{C}_{i,\vartheta}^{t_i})^{y_{i,\vartheta}})^\rho = \mathbf{e}(S_i,\ (\hat{A}_{i,0}\ \hat{B}_{i,0}^{m_i}\ \hat{C}_{i,0}^{t_i})^{\rho\ \frac{z_\vartheta}{z_0}\ y_{i,\vartheta}})$, into the one pairing $\mathbf{e}(S_i,\ (\hat{A}_{i,0}\ \hat{B}_{i,0}^{m_i}\ \hat{C}_{i,0}^{t_i})^{y'_{i,0}})$ by letting $y'_{i,0} = y_{i,0} - \rho\ \frac{z_\vartheta}{z_0}\ y_{i,\vartheta} = y_{i,0} - \frac{y_{n,0}}{y_{n,\vartheta}}\ y_{i,\vartheta}$. Notice that $y'_{n,0} = 0$ is evidence that S_n has vanished from the equation.

We similarly eliminate S_n from each of the remaining verification equations, for $k = 1, \ldots, \vartheta - 1$, using the same technique. This produces new points P'_k to replace the old P_k, as well as new exponents $y'_{i,k}$ for all $i = 0, \ldots, n-1$ to replace the $y_{i,k}$. The signature σ has been transformed into $(t_0, \ldots, t_{n-1}, S_1, \ldots, S_{n-1}, P'_0, \ldots, P'_{\vartheta-1})$, which has three fewer components.

After ϑ iterations of this process, we obtain a signature $(t_0, \ldots, t_{n-\vartheta}, S_1, \ldots, S_{n-\vartheta}, P''_0)$, which corresponds to a mesh statement Υ' which is a flat disjunction of $n - \vartheta + 1$ inputs. We have thus obtained a ring signature over a ring of size $n - \vartheta$ plus the user "in the sky." We can rename the component P''_0 as S_0 and rearrange the signature into the familiar form $\sigma' = (S_0, \ldots, S_{n-\vartheta}, t_0, \ldots, t_{n-\vartheta})$,

which represents a disjunction, and should only contain target users and honest clauses if the original forgery was admissible in the first place.
$\mathcal{B}$ then performs the second transformation $\Upsilon' \mapsto \Upsilon''$. Starting from σ', it expands the ring to cover all ℓ users, by adding dummy signatures by the missing users on arbitrary messages: this is done by adding S_j components for randomly chosen t_j. Since these dummy signature components need not (and must not) contribute to the final output, they consist only of a blinding factor that is easy to cancel in an existing S_i, using the usual reciprocity trick under the pairing. We finally obtain a ring pseudo-signature $\sigma'' = (t_0, \ldots, t_\ell, S_0, \ldots, S_\ell)$ for the full ring $R^* = \{1, \ldots, \ell\}$ of honest users plus the verification key "in the sky."
The final step of the reduction is to turn σ'' into a solution to the Poly-SDH problem. We omit the description of this step as it is exactly the same as in the regular simulation in Theorem 25, once $\mathcal{B}$ has removed whichever (known) exponent $y''_{i,0}$ still remains in each S_i.

We already mentioned that the ring pseudo-signature σ'' is technically invalid since the message m_0 borne by the key "in the sky" is the hash value $m_0 = H(\lceil VK_1 : Msg_1 \rceil, \ldots, \lceil VK_n : Msg_n \rceil, \Upsilon)$ from the original mesh specification Υ, and not Υ''. This does not affect the final reduction.

However, just as in the proof of Theorem 25, for some $i \in \{0, \ldots, \ell\}$ the forgery message m_i could induce a value $w_i^* \in \{w_{i,1}, \ldots, w_{i,q}\}$ that was already given in the Poly-SDH instance (using the w_i^* notation from Theorem 25). If the component S_i can be eliminated by the above process, then all is well. Otherwise, we need an alternative reduction from the same mold as in Theorem 25.

Alternative Simulation. The alternative reduction works only on a final forgery that "matches" one of the Poly-SDH pairs that was given to the simulator, in the sense that it corresponds to a value $w_0^* \in \{w_{0,1}, \ldots, w_{0,q}\}$. In this case, the simulator can either recover the Poly-SDH secret exponent α_0, or find a collision under the hash function H.

The simulation proceeds as in the ring signatures of Theorem 25, based on the same judicious choice of known and unknown discrete logarithms to let the final reduction go through. As in the regular simulation above, there are two main difficulties compared to the ring signature case:

1. We need to answer queries not for ring signatures but for more complicated mesh signatures. This is easy to do by using the signing key "in the sky" and the method described in the regular mesh simulation above, except for how the atomic signature of index 0 is created. Here, the simulator knows the discrete logarithm of A_0 and B_0 instead of B_0 and C_0, and so the atomic signature is constructed as in the alternative simulation for ring signatures.
   Queries on atomic signatures for the ℓ honest users are answered using the lists of pairs for values of $i = 1, \ldots, \ell$, as the alternative simulator in the ring proof.
2. We need to transform the final mesh forgery (provided it is admissible) into a ring forgery that exactly covers the target users $\{1, \ldots, \ell\}$. The transformation operates in two steps analogous to those of the regular mesh simulation, except for a minor modification: in order for the final reduction to work, the simulator must arrange to know the discrete logarithms of the private key components $\hat{A}_i$ and $\hat{B}_i$, instead

of $\hat{B}_i$ and $\hat{C}_i$. However, this does not affect the transformation, which exploits a different set of discrete logarithms, namely the z_k.

Success Probability. The important point that justifies that either the regular or the alternative simulation will give the desired result, is that in both cases the final reduced ring signature σ'' still contains a component of index 0 that bears on a hash m_0 of the full unretouched mesh specification Υ given by the forger (and with the pristine original randomizer t_0, too). Since it is those values and the given Poly-SDH instance that determine which of the regular and alternative simulations will work, we have the desired perfect dichotomy.

We conclude that if $\mathcal{B}$ chooses one at random at the onset, the overall success probability of the reduction (conditionally on $\mathcal{A}$'s success) will be $\frac{1}{2}$. □

References

[1] M. Abe, M. Ohkubo, K. Suzuki, 1 signatures from a variety of keys, in *Proceedings of AsiaCrypt 2002*. LNCS, vol. 2501 (Springer, 2002), pp. 415–432

[2] G. Ateniese, J. Camenisch, S. Hohenberger, B. de Medeiros, Practical group signatures without random oracles. Cryptology ePrint Archive, Report 2005/385 (2005). http://eprint.iacr.org/

[3] M.H. Au, J.K. Liu, T.H. Yuen, D.S. Wong, ID-based ring signature scheme secure in the standard model, in *Proceedings of IWSEC 2006.* LNCS, vol. 4266 (2006), pp. 1–16

[4] A. Bender, J. Katz, R. Morselli, Ring signatures: stronger definitions, and constructions without random oracles. *J. Cryptol.***22**(1):114–138 (2009)

[5] D. Boneh, X. Boyen, Short signatures without random oracles, in *Advances in Cryptology—EUROCRYPT 2004*. LNCS, vol. 3027 (Springer, 2004), pp. 56–73

[6] D. Boneh, X. Boyen, Short signatures without random oracles and the SDH assumption in bilinear groups. *J. Cryptol.***21**(2):149–177 (2008)

[7] D. Boneh, X. Boyen, H. Shacham, Short group signatures, in *Advances in Cryptology—CRYPTO 2004*. LNCS, vol. 3152 (Springer, 2004), pp. 41–55

[8] D. Boneh, C. Gentry, B. Lynn, H. Shacham, Aggregate and verifiably encrypted signatures from bilinear maps, in *Advances in Cryptology—EUROCRYPT 2003*. LNCS, vol. 2656 (Springer, 2003), pp. 416–432

[9] D. Boneh, E.-J. Goh, K. Nissim, Evaluating 2-DNF formulas on ciphertexts, in *Proceedings of TCC 2005*, Lecture Notes in Computer Science (Springer, 2005)

[10] X. Boyen, Mesh signatures: how to leak a secret with unwitting and unwilling participants, *in Advances in Cryptology—EUROCRYPT 2007*. LNCS, vol. 4515 (Springer-Verlag, 2007), pp. 210–227

[11] X. Boyen, B. Waters, Full-domain subgroup hiding and constant-size group signatures, in *Public Key Cryptography—PKC 2007.* LNCS, vol. 4450 (Springer, 2007), pp. 1–15

[12] E. Bresson, J. Stern, M. Szydlo, Threshold ring signatures and applications to ad-hoc groups, in *Advances in Cryptology—CRYPTO 2002.* LNCS, vol. 2442 (2002), pp. 465–480

[13] J. Camenisch, A. Lysyanskaya, Signature schemes with efficient protocols, in *Proceedings of SCN 2002*. LNCS (Springer, 2002)

[14] J. Camenisch, A. Lysyanskaya, Signature schemes and anonymous credentials from bilinear maps, in *Advances in Cryptology—CRYPTO 2004*. LNCS, vol. 3152 (Springer, 2004)

[15] N. Chandran, J. Groth, A. Sahai, Ring signatures of sub-linear size without random oracles, in *Proceedings of ICALP 2007*, vol. 4596 (2007), pp. 423–443.

[16] M. Chase, A. Lysyanskaya, Signatures of knowledge, in *Advances in Cryptology—CRYPTO 2006*. LNCS, vol. 4117 (Springer, 2006)

[17] D. Chaum, E. van Heyst, Group signatures, in *Advances in Cryptology—EUROCRYPT 1991*. LNCS, vol. 547 (Springer, 1991), pp. 257–265

[18] J.H. Cheon, Security analysis of the strong Diffie–Hellman problem, in *Advances in Cryptology—EUROCRYPT 2006*. LNCS, vo. 4004 (Springer, 2006), pp. 1–13

[19] S.S.M. Chow, L.C.K. Hui, S.-M. Yiu. Identity based threshold ring signature, in *Proceedings of ICISC 2004*. LNCS, vol. 3506 (2004), pp. 218–232
[20] S.S.M. Chow, R.W.C. Lui, L.C.K. Hui, S.-M. Yiu, Identity based ring signature: why, how and what next, in *Proceedings of EuroPKI 2005*. LNCS, vol. 3545 (2005), pp. 144–161
[21] S.S.M. Chow, V.K.-W. Wei, J.K. Liu, T.H. Yuen, Ring signatures without random oracles, in *Proceedings of AsiaCCS 2006* (ACM Press, 2006), pp. 297–302
[22] S.S.M. Chow, S.-M. Yiu, L.C.K. Hui, Efficient identity based ring signature, in *Proceedings of ACNS 2005*. LNCS, vol. 3531 (2005) pp. 499–512
[23] R. Cramer, I. Damgård, B. Schoenmakers, Proofs of partial knowledge and simplified design of witness hiding protocols, in *Advances in Cryptology—CRYPTO 1994*. LNCS, vol. 839 (Springer, 1994), pp. 174–187
[24] Y. Dodis, A. Kiayias, A. Nicolosi, V. Shoup, Anonymous identification in ad hoc groups, in *Advances in Cryptology—EUROCRYPT 2004*. LNCS, vol. 3027 (Springer, 2004), pp. 609–626.
[25] C. Dwork, M. Naor, Zaps and their applications. *SIAM J. Comput.***36**(6):1513–1543 (2007)
[26] C. Dwork, M. Naor, A. Sahai, Concurrent zero-knowledge or the timing model for designing concurrent protocols. *J. ACM***51**(6):851–898 (2004)
[27] A.L. Ferrara, M. Green, S. Hohenberger, M.O. Pedersen, Practical short signature batch verification. Cryptology ePrint Archive, Report 2008/015 (2008). http://eprint.iacr.org/
[28] S.D. Galbraith, K.G. Paterson, N.P. Smart, Pairings for cryptographers. Cryptology ePrint Archive, Report 2006/165 (2006). http://eprint.iacr.org/
[29] J. Groth, Fully anonymous group signatures without random oracles, in *Proceedings of ASIACRYPT 2007*. LNCS, vol. 4833 (2007), pp. 164–180
[30] J. Groth, R. Ostrovsky, A. Sahai, Non-interactive Zaps and new techniques for NIZK, in *Advances in Cryptology—CRYPTO 2006*. LNCS (Springer, 2006)
[31] J. Herranz, G. Sáez, Forking lemmas for ring signature schemes, in *Proceedings of IndoCrypt 2003*. LNCS, vol. 2904 (Springer, 2003), pp. 266–279
[32] J. Herranz, G. Sáez, New distributed ring signatures for general families of signing subsets. Cryptology ePrint Archive, Report 2004/377 (2004). http://eprint.iacr.org/
[33] A. Joux, K. Nguyen, Separating decision Diffie–Hellman from computational Diffie–Hellman in cryptographic groups. *J. Cryptol.*, **16**(4) (2003)
[34] M. Karchmer, A. Wigderson, On span programs, in *Annual Conference on Structure in Complexity Theory* (1993)
[35] V. Miller, The Weil pairing, and its efficient calculation. *J. Cryptol.***17**(4) (2004)
[36] M. Naor, Deniable ring authentication, in *Advances in Cryptology—CRYPTO 2002*. LNCS, vol. 2442 (Springer, 2002), pp. 481–498
[37] R. Rivest, A. Shamir, Y. Tauman, How to leak a secret, in *Proceedings of AsiaCrypt 2001*. LNCS, vol. 2248 (Springer, 2001), pp. 552–565
[38] H. Shacham, B. Waters, Efficient ring signatures without random oracles, in *Public Key Cryptography—PKC 2007*. LNCS, vol. 4450 (Springer, 2007)
[39] V. Shoup, Lower bounds for discrete logarithms and related problems, in *Advances in Cryptology—EUROCRYPT 1997*. LNCS, vol. 1233 (Springer, 1997)
[40] M. van Dijk, A linear construction of secret sharing schemes. *Des. Codes Cryptogr.***12**(2):161–201 (1997)
[41] B. Waters, Efficient identity-based encryption without random oracles, in *Advances in Cryptology—EUROCRYPT 2005*. LNCS, vol. 3494 (Springer, 2005)
[42] V.K. Wei, T.H. Yuen, (Hierarchical identity-based) threshold ring signatures. Cryptology ePrint Archive, Report 2006/193 (2006). http://eprint.iacr.org/