

It wasn't me!

Repudiability and Unclaimability of Ring Signatures

Sunoo Park
MIT and Harvard

Adam Sealfon
MIT

Abstract

Ring signatures, introduced by [RST01], are a variant of digital signatures which certify that *one among a particular set* of parties has endorsed a message while hiding *which* party in the set was the signer. Ring signatures are designed to allow *anyone* to attach anyone else's name to a signature, as long as the signer's own name is also attached.

But what guarantee do ring signatures provide if a purported signatory wishes to denounce a signed message — or alternatively, if a signatory wishes to later come forward and claim ownership of a signature? Prior security definitions for ring signatures do not give a conclusive answer to this question: under most existing definitions, the guarantees could go either way. That is, it is consistent with some standard definitions that a non-signer might be able to *repudiate* a signature that he did not produce, or that this might be impossible. Similarly, a signer might be able to later convincingly claim that a signature he produced is indeed his own, or not. Any of these guarantees might be desirable. For instance, a whistleblower might have reason to want to later claim an anonymously released signature, or a person falsely implicated in a crime associated with a ring signature might wish to denounce the signature that is framing them and damaging their reputation. In other circumstances, it might be desirable that even under duress, a member of a ring cannot produce proof that he did or did not sign a particular signature. In any case, a *guarantee* one way or the other seems highly desirable.

In this work, we formalize definitions and give constructions of the new notions of *repudiable*, *unrepudiable*, *claimable*, and *unclaimable* ring signatures. Our repudiable construction is based on VRFs, which are implied by several number-theoretic assumptions (including strong RSA or bilinear maps); our claimable construction is a black-box transformation from any standard ring signature scheme to a claimable one; and our unclaimable construction is derived from the lattice-based ring signatures of [BK10], which rely on hardness of SIS. Our repudiable construction also provides a new construction of standard ring signatures.

1 Introduction

Ring signatures, introduced by [RST01], are a variant of digital signatures which certify that *one among a particular set* of parties has signed a particular message, without revealing which specific party is the signer. This set is called a “ring.” Ring signatures can be useful, for example, to certify that certain leaked information comes from a privileged set of government or company officials without revealing the identity of the whistleblower, to issue important orders or directives without setting up the signer to be a scapegoat for repercussions,¹ or to enable untraceable transactions in

¹When it comes to national security issues, for instance, there may be some reluctance among law-makers to “roll back” existing laws or reduce checking or surveillance measures, which could be due in part to the risk of ending up a scapegoat upon any future national security incident like a terrorist attack.

cryptocurrencies (as in Monero [Mon]).

In a ring signature scheme, just as in a traditional digital signature scheme, any party can create a key pair for signing and verification, and publish the verification key. Signers can produce signatures that verify with respect to any set of verification keys that includes their own, and unforgeability guarantees that no party can produce a valid signature with respect to a set of verification keys without possessing a corresponding secret key.

But what guarantee does a ring signature scheme provide if a purported signatory wishes to denounce a signed message — or alternatively, if a signatory wishes to later come forward and *claim* ownership of a signature? Given the motivation of anonymity behind the notion of a ring signature, a natural first intuition might be that parties should be able neither to denounce nor to claim a signature in a convincing way. However, depending on the threat model, we believe that the opposite guarantees — that is, to guarantee the ability to denounce or claim signatures — may be useful too, as elaborated below. Furthermore, whatever one’s preference, a guarantee one way or the other seems more desirable than no guarantee either way.

Prior security definitions for ring signatures do not conclusively provide these guarantees one way or the other. That is, a non-signer might be able to *repudiate* a signature that he did not produce (“repudiability”), or this might be impossible (“unrepudiability”). Similarly, a signer might be able to later convincingly claim that a signature he produced is indeed his own (“claimability”), or be unable to do so (“unclaimability”).

The most detailed taxonomy of security definitions for ring signatures was given by [BKM09], which presents a series of anonymity guarantees of increasing strength. A natural anonymity guarantee defined by [BKM09], called “anonymity against adversarially chosen keys,” is informally described as follows: an adversary who controls all but $t \geq 2$ parties in a ring, and who may produce his own malformed key pairs as well as corrupt honest parties’ keys, must have negligible advantage at guessing which of the t honest parties produced a given signature. This anonymity definition might allow a party to ascertain whether a given signature was produced by her own signing key, and perhaps also to convince others of this fact — but it does not *guarantee* or *prohibit* either of these capabilities.

On the other hand, the strongest of the anonymity definitions of [BKM09] (called “anonymity against full key exposure”) requires that even if an adversary compromises every single party in a ring, the adversary cannot identify the signers of past signatures. It is relatively straightforward to see that under such a strong anonymity guarantee, Alice would have no way to convince anyone that she did not produce the objectionable message; indeed, she herself cannot tell the difference between a signature produced using her own signing key and one produced using someone else’s.

The ability to identify whether one’s own signing key was used to produce a particular signature can be a feature or a bug. To protect anonymity of past signatures against a very strong adversary who might compromise all the secret keys in a ring, it seems desirable to prevent distinguishing one’s own signatures from those generated by someone else. On the other hand, without the ability to distinguish, it would be virtually impossible to tell if someone had stolen your signing key. Moreover, as discussed below, it could be beneficial in certain circumstances for members of a ring to have the ability to disown signatures of messages that they have strong reasons to denounce; and conversely, in some circumstances the signer of a message might later wish to prove to the world that he was the one who produced a particular signature in the past.

We have now identified four potentially useful notions for ring signatures: *repudiability*, *unrepudiability*, *claimability*, and *unclaimability*. The main contributions in this paper consist both of

new definitions and constructions of each of these notions. Before diving into an overview of definition and constructions, we provide some discussion of why each of these notions — some of which directly oppose each other — may be meaningful and desirable: the following scenarios explore a few of the circumstances in which various of the above guarantees might be appropriate. Though some of the scenarios are phrased somewhat whimsically, we believe that each scenario illustrates a meaningful threat model motivating the definition concerned.

Scenario 1 (Repudiability) Let us consider a hypothetical tale, wherein two candidates Alice and Bob are running for president in the land of Oz. Oz is notorious for its petty partisan politics and its tendency to prefer whomever appears friendlier in a series of nationally televised grinning contests between the main-party candidates. At the peak of election season, a disgruntled citizen Eve decides to help out her preferred candidate Bob by publishing the following message, which goes viral on the social networks of Bob supporters:

*I created a notorious terrorist group and laundered lots of money!
Signed: Alice or Eve or Alice's campaign chairman.*

Of course, the virally publicized message does not actually incriminate Alice at all, since any one of the signatories could have produced it. However, perhaps there is nothing that Alice can do to allay the doubt in the minds of her suspicious detractors. As mentioned above, ring signatures are deliberately designed to allow *anyone* to attach anyone else's name to a signature, without the latter's knowledge or consent. Despite this, there could be realistic situations in which non-signing members of a ring associated with a particular message could suffer serious consequences through no fault of their own, perhaps due to the real signer adversarially trying to damage their reputation. In light of this, perhaps it would be desirable in some contexts for the owner of a verification key to be able to denounce messages, e.g., to clear her name of a crime or hate speech accusation that might otherwise impact her life in terms of reputation, job prospects, or incarceration.

Scenario 2 (Claimability) Our next story concerns a talented brewery employee who developed new statistical techniques to test the quality of beers. Naturally, his employer was protective of its competitive advantage since other breweries at the time may not have been using similar statistical methods. Yet, in the interest of science, they allowed him to publish his results — on condition of anonymity.² A credible way to prove authorship at a later date, after the need for anonymity has ceased to exist, might be very useful — especially in case of competing claims by impostors. As we see here, claiming authorship of an anonymous work may become appropriate after a passage of time. The next example illustrates quite a different type of situation in which claimability at the signer's discretion may be valuable.

Consider an employee Emily who is concerned about unethical practices at her company, and takes it upon herself to expose what is going on and publish a critical commentary. Concerned about her job security and possible retribution, as well as the credibility of her allegations, she maintains her anonymity using ring signatures. It emerges, in fact, that similar practices are prevalent across the industry: related revelations drive a wider movement of reform. Some time later, after her company has substantially reformed its practices and her fears of retribution have been allayed

²This is the true story of William Sealy Gosset's invention of the Student's t -test at Guinness Brewery in 1908 [Man00].

— perhaps by her promotion, or by a change in leadership — Emily seeks to reveal her identity and add her voice to the growing movement, providing her solidarity, legitimation, and follow-up story. In addition, if following the reforms, those involved in the earlier unethical practices were subject to stigma or even prosecution, claimability of her earlier ring signatures would allow Emily to exculpate herself.

Scenario 3 (Unrepudiability and unclaimability) Let us return to the government of the fictional country of Oz. The parliament of Oz is mired in partisan gridlock, with legislators from each party ruthlessly voting down any bills, however reasonable, proposed by members of the opposing party — preventing any laws at all from being enacted and effectively shutting down the government, which is in no party’s interest. Suppose that instead of directly proposing a new law, a legislator of Oz anonymously publishes the text of the proposed bill using a ring signature scheme:

*Proposed: that free ice cream shall be provided every Tuesday.*³

Signed: a member of the Parliament of Oz.

If the signer used an unclaimable ring signature scheme, then she could not decide to reveal her identity upon a later change of heart, allowing legislators of both parties to support or oppose the bill on its merits without worrying about purely political considerations.

Unclaimability and unrepudiability may be particularly useful guarantees in scenarios where the placement of whole groups of people under duress is a substantial concern. For instance, in circumstances where an employer or authoritarian government may coercively compel individuals to provide a repudiation or proof of authorship (e.g. signing randomness) for a signature, the provable inability to do so convincingly may be essential. Unrepudiability may also be desirable in situations in which members of a ring are likely to have conflicting individual incentives but there is a possibility of collective benefit in case of cooperation, as in a prisoner’s dilemma scenario.

Summary of technical contributions. We formalize *repudiability*, *unrepudiability*, *claimability*, and *unclaimability* of ring signatures, as well as strengthened anonymity and unforgeability definitions which are compatible with each of these notions. We show that *unclaimability* implies *unrepudiability* (intuitively, because a failed repudiation can be used as a claim). Anonymity against adversarially chosen keys is the strongest anonymity notion compatible with *repudiability* and *claimability*, and anonymity against full key exposure is implied by *unclaimability* and equivalent to *unrepudiability*.

We provide three constructions based on different assumptions, one for each of the three notions of *repudiability*, *claimability*, and *unclaimability*. Perhaps the most surprising of these is *unclaimability*, which guarantees that the signer cannot later credibly convince others that she produced a particular signature. A natural first intuition is that meaningful notions of unclaimability might be impossible to achieve, since a signer can always remember the signing randomness (and later present it as “proof” of having produced a signature). The key insight for our definition and construction of unclaimable ring signatures is that the signing randomness does not constitute a convincing claim if *anyone in the ring can also produce credible signing randomness* for any signature in which they are implicated. Our construction of *unclaimable* ring signatures is an augmentation of the

³Even if each party might support this legislation, they may be unwilling to do so if it were proposed by the other party, decrying their respective opponents as either fiscally irresponsible or in the pocket of Big Ice Cream.

lattice-based ring signature scheme of [BK10] that adds additional algorithms allowing anyone in the ring to generate credible signing randomness; this capability is achieved via lattice trapdoors.

Our construction of *repudiable* ring signatures is based on verifiable random functions (VRFs), which are implied by either the (strong) RSA assumption, assumptions on bilinear maps, or NIWIs and commitments; see [Bit17; GHKW17] and references therein for more detailed discussion of the assumptions that imply VRFs.⁴ Our construction does not use standard ring signatures as a building block, and as such can also be viewed as a new construction of standard ring signatures. Our construction of *claimable* ring signatures, on the other hand, is a simple and generic black-box transformation from any standard ring signature scheme to a claimable one. We overview our contributions in more detail below.

1.1 Definitional contributions

Repudiability. We define a *repudiable ring signature scheme* as a ring signature scheme that is equipped with additional algorithms `Repudiate` and `VerRepud` as follows. `Repudiate` takes as input a signing key sk , a ring signature σ , and a “ring” R (i.e., a set of verification keys), and outputs a *repudiation* ξ . `VerRepud` takes as input a ring R , a signature σ , a repudiation ξ , and a verification key vk , and outputs a single bit indicating whether or not ξ is a valid repudiation attesting that σ was not produced by vk . The two requirements for a ring signature scheme to be *repudiable* are, informally, as follows.

1. *Correctness*: Any member of a ring must be able to produce valid repudiations of any signature that he did not produce.
2. *Soundness*: A cheating signer must not be able to produce a valid signature with respect to a ring, and also be able to produce valid repudiations of that signature under every verification key in that ring that he owns.

Once a ring signature scheme is equipped with these additional repudiation algorithms, the standard definitions of *unforgeability* and *anonymity* against adversarially chosen keys are insufficient to capture the natural guarantees that would be desired for a repudiable ring signature scheme: we need the release of repudiations not to compromise the unforgeability or anonymity of any future signatures. Accordingly, we modify the definitions of unforgeability and anonymity for repudiable ring signatures (Definitions 3.4 and 3.5), by additionally giving the adversary access to a *repudiation oracle*. This ensures that repudiations of past signatures do not affect the security guarantees of future signatures. See Section 3.1 for formal definitions of repudiability.

Correction of an earlier version. As pointed out by Hao Lin and Mingqiang Wang [Lin19; LW19], previous versions of this paper (including the conference version) had an error in the anonymity definition for repudiable ring signatures (Definition 3.4).⁵ This has been corrected in the current version. See Remark 7 in Section 3.1.1 for details.

Claimability. We define a *claimable ring signature scheme* as a ring signature scheme equipped with additional algorithms `Claim` and `VerClaim` as follows. `Claim` takes as input a signing key sk , a signature σ , and a ring R , and outputs a claim ζ . `VerClaim` takes as input a ring R , a verification

⁴Note that the ring signature construction of [BKM09] assumes the closely related assumption of NIZKs, which are implied by VRFs and also imply VRFs if the NIZK is subexponentially hard under a standard derandomization assumption; see [Bit17] for details.

⁵Specifically, the difference between the old and new definitions is in the details of puncturing an oracle in the security game.

key vk , a signature σ , and a claim ζ , and outputs a single bit indicating whether or not ζ is a valid claim attesting that σ was produced by vk . The three requirements for a claimable ring signature scheme are, informally, as follows.

1. *Correctness*: Any honest signer must be able to produce a valid claim with respect to any signature that he produced.
2. *Soundness*: No adversary can produce a valid claim with respect to a signature produced by an honest signer, even if the adversary can choose the message and ring with respect to which the signature is produced, and can insert malformed verification keys into the ring.
3. *No framing*: No adversary can produce a signature together with a valid claim of that signature on behalf of an honest (non-signing) party.

As above, once a ring signature scheme is equipped with these additional claiming algorithms, the standard definitions of *unforgeability* and *anonymity* against adversarially chosen keys are insufficient. We modify the definitions of anonymity and unforgeability for claimable ring signatures (Definitions 3.10 and 3.11), by additionally giving the adversary access to a *claim oracle*. See Section 3.3 for formal definitions of repudiability.

Repudiability and *claimability* are compatible, i.e., a ring signature scheme can be both repudiable and claimable. Indeed, our repudiable and claimable constructions together give rise to such a scheme. Notably, the unforgeability and anonymity definitions corresponding to the natural notion of a repudiable-and-claimable ring signature scheme are *not* the conjunction of unforgeability and anonymity for repudiable ring signatures and for claimable ring signatures. Rather, the unforgeability and anonymity definitions for a repudiable-and-claimable ring signature scheme involve a stronger adversary which is simultaneously given access to both a repudiation oracle and a claim oracle. See Section 3.5 for further discussion on repudiable-and-claimable schemes.

Unclaimability We also introduce *unclaimable* ring signature schemes, in which the signer *provably cannot* convincingly claim that she was the one who produced the signature. As briefly mentioned above, while the signer can always save the signing randomness and reveal it along with her secret key in an attempt to claim authorship of a signature, it is not always true that this constitutes a convincing claim. In particular, such a claim is not credible if *any* member of the ring can take a valid signature and produce fake randomness that produces the desired signature using her own signing key.

The idea that a non-signer can adaptively produce fake randomness is reminiscent of deniable encryption [CDNO97], in which an encryptor and/or recipient is required to produce fake randomness “explaining” that a particular ciphertext is an encryption of an adversarially chosen message.

We define an *unclaimable* ring signature scheme to capture just this requirement: that is, any member of the ring must be able to produce fake signing randomness for a signature that is distributed indistinguishably from real signing randomness. Intuitively, the only information potentially possessed by a signer but not by the other members of the ring is the signing randomness, so non-signers that can generate convincing simulated signing randomness can also convincingly simulate any additional information that might be released by the signer in an attempt to claim the signature. We consider a strong flavor of this definition in which the indistinguishability property, described informally below, is statistical.

1. *Indistinguishability*: Any member of a ring must be able to produce fake signing randomness given a signature. The signature and fake signing randomness must be distributed statistically close to an honestly generated signature and corresponding signing randomness used by that

Repudiable	VRF (Section 4)
Unrepudiable	RS anonymous against FKE (Section 3.2)
Claimable	Transformation from any RS (Section 5)
Unclaimable	SIS (Section 6)

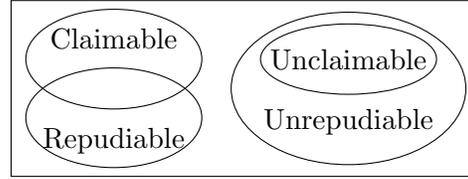


Figure 1: Summary of our results and assumptions relied on. VRF = verifiable random function, RS = ring signature, FKE = full key exposure, SIS = short integer solution problem.

individual to sign the same message, even given all verification keys and signing keys.

We formally define unclaimability in Section 3.4.

Remark 1. Even under this definition, if the signer chooses a message to sign that corresponds to a secret known only to herself, then she may still be able to convince others that she was the signer. For instance, if the signed message is the output of a one-way function, she may be able to convince others that she was the signer by subsequently revealing the preimage. Even more flagrantly, the signed message could contain a signature using a standard (non-ring) signature scheme, directly identifying the signer. This property is rather inherent: if knowledge of the contents of the message itself at the time of signing are enough to identify the signer, then no security property on the signature scheme can enforce that the signer remains hidden, since the identification of the signer is unrelated to the signature and based only on the signed message.

Indeed, ring signatures were not designed to provide anonymity for signers who *want* to identify themselves, but rather for those who desire anonymity. Similarly, our unclaimability definition does not guarantee unclaimability for those who *want* to identify themselves, but rather provides credibility for a signer who *wants* to later be able to claim (e.g., under duress) that she could not convincingly claim the signature even if she wanted to. In particular, even an adversary with unlimited computational power who obtains the secret keys belonging to every member of the ring and a purported signing randomness from an alleged signer, he still will not be convinced of the identity of the signer, since fake signing randomness from the right distribution can be produced for every member of the ring.

Unrepudiability Unclaimability intuitively guarantees that no member of the ring can convincingly prove that she was the signer. A related, weaker notion that might be desirable in some circumstances is that of *unrepudiability*, which guarantees that no member of the ring can convincingly prove that she was *not* the signer. Unrepudiability is equivalent to anonymity against full key exposure and is implied by unclaimability.

1.2 Overview of our constructions

Our repudiable construction Our construction relies on ZAPs (two-round public-coin witness-indistinguishable proofs) and verifiable random functions (VRFs) as building blocks.⁶ Our building blocks have some overlap with those of the ring signature construction of [BKM09], which uses ZAPs, public-key encryption (PKE), and a digital signature scheme. Both our scheme and theirs

⁶VRFs imply ZAPs, so it suffices to assume VRFs. [GO92; DN07]

use ZAPs to achieve anonymity of the ring signatures, but with different approaches: the statements proven by the ZAPs are quite unrelated in the two constructions. Moreover, in our scheme, we do not need PKE or signature schemes, and instead use VRFs directly to achieve unforgeability and repudiability. The structure of our construction is thus very different from that of [BKM09].

At a very high level, each signing key in our construction contains a tuple of four VRF keys. A signature consists of the output of each of the signer’s VRFs on the message, along with a ZAP proof that (several of) the VRF values in the signature are correct w.r.t. the VRF verification key of some member of the ring. A repudiation for individual i consists of a ZAP proof that some of the VRF values in the signature are different from the correct values for party i ’s VRFs evaluated at the message. One complication arises because we must guarantee that the release of a repudiation for individual i on a message does not subsequently allow a different member of the ring to produce a signature on the message that cannot be repudiated by individual i . We overcome this difficulty by relying on the witness indistinguishability property of the ZAP and ensuring that the repudiation does not reveal the actual VRF outputs of the repudiator; that is, the ZAP proof is produced with the VRF proof as a *witness*. The specific statement proven by the ZAPs is that some specific combination of at least two of the purported VRF outputs is correct. Although in the honest usage of the scheme, all four are produced correctly, we design the specific structure of the statements proved in order to allow a hybrid argument to argue indistinguishability between signatures of different signers in a ring. This scheme of proving the correctness of VRF outputs turns out also to imply unforgeability, not only repudiability, so we do not need to rely on any underlying signature scheme as building block. (In other words, our scheme can also be seen as a new construction of standard ring signatures based on VRFs.)

Our claimable construction We give a generic transformation from any standard ring signature scheme RS to a claimable one. The transformation uses commitment schemes, standard signatures, and PRFs (which are all achievable from one-way functions). The basic idea is to take a signature σ_{RS} under RS and append to it a *commitment* c to (vk, σ_{RS}) where vk is the verification key of the signer. The verification algorithm simply checks whether σ_{RS} verifies. The claim consists of a decommitment revealing that c is a commitment to (vk, σ_{RS}) . Intuitively, by the hiding property of the commitment scheme, the identity of the signer is hidden until he chooses to publish a claim.

The simple transformation just described runs into a couple of problems when examined in detail. First, what if a signer commits to $(\sigma_{\text{RS}}, vk')$ where vk' is not his own key but that of someone else in the ring? This ability would violate equation (6) of Definition 3.9 (claimability). To prevent such behavior, our construction actually commits to a *standard (non-ring) signature* on (vk, σ_{RS}) . The unforgeability property of standard signatures then guarantees, intuitively, that a signer cannot convincingly make a claim with respect to any verification key unless he knows a corresponding signing key.

A second issue encountered by the scheme thus far described is that the signer must remember the commitment randomness in order to produce a claim. It is preferable that the signer not be stateful between signing and claiming; indeed, Definition 3.9 requires this. To resolve this, our construction derives commitment randomness from a PRF. For similar reasons, the signing randomness for the standard (non-ring) signature in our construction is also derived from a PRF.

Remark 2. Among the constructions presented in this paper, claimability is by far the simplest. Moreover, as a generic transformation, it has the advantage of adding minimal efficiency overhead to the existing state of the art in ring signatures. The simplicity of achieving claimability is perhaps unsurprising in light of the natural intuition that claiming should be possible simply by

remembering the signing randomness. As evidenced by *unclaimability*, this intuition is not strictly true in general, as in certain schemes, producing signing randomness may not prove authorship. In a nutshell, our generic transformation ensures that signing randomness is indeed a convincing proof of authorship in the resulting scheme, and moreover builds into the scheme a simple method of efficiently recovering the signing randomness without storing it explicitly.

Our unclaimable construction Our construction of unclaimable ring signatures is an extension of the SIS-based ring signature scheme of Brakerski and Kalai [BK10]. The construction is based on trapdoor sampling. In this overview, we describe a simplified version of the scheme. The full scheme is described in Section 6. The basic idea for obtaining unclaimability is that each identity corresponds to a public matrix $A_i \in \mathbb{Z}_q^{n \times m}$ sampled together with a secret trapdoor T_i . A signature will consist of short vectors $x_i \in \mathbb{Z}_q^m$ such that

$$\sum_i A_i x_i = y,$$

where y is a target value. For this overview, we can think of y as the output of a random oracle on the message; in the actual construction, y will be obtained as the sum of additional matrix-vector products. In order to sign the message, signer i first samples short vectors x_j for each $j \neq i$. Then, using the lattice trapdoor T_i , he samples a short vector x_i such that the equation

$$x_i = y - \sum_{j \neq i} A_j x_j$$

is satisfied. The signature is the list of vectors $\sigma = (x_i)_i$. Using properties of lattice trapdoors, it follows that the distribution over $(x_i)_i$ can be made statistically close no matter which trapdoor was used to produce the signature. Moreover, given a vector x^* to be produced, we can sample random coins that will yield that vector under either the ordinary sampling algorithm or the trapdoor sampling algorithm. Consequently, we obtain an algorithm that can produce explanatory randomness for a signature under any identity in the ring.

Removing the random oracle to obtain ring signatures in the plain model (and unclaimable ones) requires several complications. [BK10] first describes a basic ring signature scheme with weaker unforgeability properties, in which the target vector y is determined using additional matrix-vector products for matrices that depend on the bits of the message. They then amplify the security of the scheme through a sequence of transformations that ultimately yield a scheme with full unforgeability. In Section 6, we first define an algorithm for producing explanatory randomness for their basic scheme, and then describe how to modify this algorithm for each modification of the basic scheme, ultimately yielding an unclaimable ring signature scheme based on the SIS assumption.

Remark 3. The idea that a non-signer of a given signature can adaptively produce fake signing randomness is reminiscent of deniable encryption [CDNO97], in which an encryptor of a given ciphertext can adaptively produce fake randomness consistent with it being an encryption of a different message. In this context, it may seem somewhat surprising that our construction relies on a relatively standard assumption (SIS) while many natural definitions of deniable encryption are not known to be achievable without heavier assumptions such as indistinguishability obfuscation [SW14; CPP18]. A subtle difference that is significant here is that a deniably encrypted message must still be recoverable by the honest decryptor, while in the unclaimable ring signature setting, the signer’s identity need not be recoverable by anyone.

1.3 Other related work

Several constructions of ring signatures based on lattice assumptions have been proposed (e.g., [BK10; Mel+13; BLO18]). The only other construction of ring signatures based on ZAPs is [BKM09], to our knowledge. Numerous other ring signature constructions have been proposed, mostly based on various assumptions on bilinear maps, many but not all of which are in the random oracle model (e.g., [Ngu05; SS10; Boo+15]).

Two additional works in the lattice trapdoor literature bear mentioning: the seminal [Ajt99], and the more recent [MP12]. The latter is more recent than [GPV08], whose trapdoors our unclaimable construction relies on (this reliance is carried over from the [BK10] construction).

Ring signatures with additional guarantees Since the original proposal of ring signatures by [RST01], various variant definitions have been proposed. For example, *linkable* ring signatures [LWW04] allow identification of signatures that were produced by the same signer, without compromising the anonymity of the signer within the ring. An enhancement to this notion called *designated linkability* [LSW06] does not allow linkability by default, but instead allows links to be revealed at will by a designated party. A related notion is that of *designated identity verifier* ring signatures [SP14], which allow identification of the signer by a special party who is publicly designated by the signer at the time of signing. Another notion called *traceable* ring signatures [FS07] considers a setting where signatures are generated with respect to “tags” and each member may sign at most a single message (say, a vote) with respect to a particular tag, or else his identity will be revealed. *Accountable* ring signatures [XY04; Boo+15] allow a signer to assign the power to de-anonymize her signature to a specific publicly identified party.

It may seem that some of these variants of ring signature schemes have properties that would be useful for constructing *claimable* ring signatures as introduced in this paper. This implication is unsurprising in the context of our results: *all* of the above types of ring signature schemes in fact imply claimable ring signatures, since our construction of claimable ring signatures is a generic transformation from *any* ring signature scheme. It is unclear if leveraging the additional features of variant schemes would be more desirable than applying our generic transformation, which has very low overhead and moreover can be applied to a simpler, more efficient ring signature scheme that may lack these additional properties.

Group signatures Group signatures [CH91] are a different type of signature that allow signing w.r.t. a set of verification keys and provide anonymity of the signer within that set. This concept differs most strikingly from ring signatures in that there is a central authority that (1) sets up the group (i.e., set of signers) and issues keys to members of the group and (2) has the power to revoke the anonymity of the signer of a signature. Notions such as (un)linkability, described above, have been applied to the group signature setting as well. Notably, there has also been proposed a notion of *deniable group signatures* [Ish+16], in which the group manager may issue proofs that a particular group member did *not* sign a particular signature. This bears a little resemblance to our notion of *repudiability* in ring signatures; however, the presence of a central authority in the group signature setting means these problems are technically rather disparate. [LNWX17] construct lattice-based deniable group signatures; however, their technique for deniability is very different from ours, and relies on zero-knowledge proofs of plaintext inequality for LWE ciphertexts, which do not suffice in our setting.

2 Anonymity and unforgeability of ring signatures

This section overviews standard ring signature definitions: syntax, correctness, anonymity, and unforgeability. We express the anonymity and unforgeability definitions differently from prior work, as explained in their respective subsections. However, our definitions are equivalent to the correspondingly named definitions from prior work. Throughout the paper, k denotes the security parameter.

Definition 2.1 (Ring signature). A *ring signature scheme* is a triple of PPT algorithms $RS = (\text{Gen}, \text{Sign}, \text{Verify})$, satisfying the three properties of *correctness* (Definition 2.2), *anonymity* (Definitions 2.5–2.6), and *unforgeability* (Definition 2.8). The syntax of Gen , Sign , and Verify follows.

- $\text{Gen}(1^k)$ takes k as input and outputs a *verification key* vk and a *signing key* sk .
- $\text{Sign}(R, sk, m)$ takes as input a signing key sk , a message m , and a set of verification keys $R = \{vk_1, \dots, vk_N\}$, and outputs a signature σ . The set R is also known as a “ring.”
- $\text{Verify}(R, \sigma, m)$ takes as input a set R of verification keys, a signature σ , and a message m , and outputs a single bit indicating whether or not σ is a valid signature on m w.r.t. R .

Where it may not be clear from context, we sometimes write $RS.\text{Gen}$, $RS.\text{Sign}$, $RS.\text{Verify}$ to denote the Gen , Sign , Verify algorithms belonging to RS .

Definition 2.2 (Correctness). A ring signature scheme $RS = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies *correctness* if there is a negligible function ε such that for any $N = \text{poly}(k)$, any N key pairs $(vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k)$, any $i \in [N]$, and any message m ,

$$\Pr[\text{Verify}(R, \text{Sign}(R, sk_i, m), m) = 1] = 1 - \varepsilon(k), \quad (1)$$

where $R = \{vk_1, \dots, vk_N\}$. RS satisfies *perfect correctness* if (1) holds for $\varepsilon = 0$.

Remark 4. Definition 2.2 considers only for honestly generated keys. One could also consider a stronger requirement that verification be successful for honestly generated signatures with respect to rings containing adversarial keys. Any scheme satisfying Definition 2.2 can be transformed into one satisfying the stronger definition, by modifying original signature algorithm Sign to a new algorithm Sign' that operates as follows. On input R, sk, m , for a sufficiently large polynomial p :

1. $\sigma \leftarrow \text{Sign}(R, sk, m)$
2. $b_1, \dots, b_p \leftarrow \text{Verify}(R, \sigma, m)$
3. if $\forall i \in [p], b_i = 1$, output σ ; else, go back to step 1

2.1 Anonymity

Prior work, notably [RST01; BKM09], has presented a variety of anonymity definitions for ring signatures. Two of the definitions from prior work are relevant to this paper: anonymity against adversarially chosen keys, and anonymity against full key exposure.

This section presents a new, generalized anonymity definition which is parametrized by *oracle sets*, and expresses the two relevant anonymity definitions as instantiations of the generalized definition. This generalized definition is useful to consolidate the existing definitions and make clear their relationship to one another; it captures not only the two definitions we rely on here, but also others from prior work. Moreover, the generalized definition will be essential to concisely express the new anonymity definitions that we introduce in later sections for anonymity of *repudiable*

and *claimable* ring signature schemes (in Sections 3.1.1 and 3.3.1 respectively). In a nutshell, this is because the new definitions need to allow the adversary access to additional oracles related to repudiation and/or claiming.

The generalized definition follows. It is parametrized by sets of oracles $\mathcal{O}_1, \mathcal{O}_2$ and an additional parameter $\alpha \in \{0, 1, 2\}$ that limits the adversary's corruptions.

Definition 2.3 ($(\mathcal{O}_1, \mathcal{O}_2, \alpha)$ -anonymity). Let $\mathcal{O}_1, \mathcal{O}_2$ be sets of oracles, where each oracle in the set is parametrized by a list of key-pairs. Define $\text{Corr}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$ to take as input $i \in [N]$ and output $\omega_i \leftarrow \text{Gen}^{-1}(vk_i, sk_i)$.⁷

A ring signature scheme $\text{RS} = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies $(\mathcal{O}_1, \mathcal{O}_2, \alpha)$ -anonymity if for any PPT adversary \mathcal{A} and any polynomial $N = \text{poly}(k)$, $\Pr[b' = b]$ in the above game is negligibly close to $1/2$. That is, formally, \forall PPT $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, $N = \text{poly}(k)$, there is a negligible function ε such that

$$\Pr \left[\begin{array}{l} (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\ ((m^*, i_0^*, i_1^*, R^*), \mathfrak{s}) \leftarrow \mathcal{A}_1^{\mathcal{O}_1, \text{Corr}}(vk_1, \dots, vk_N) \\ b \leftarrow \{0, 1\} \\ \sigma \leftarrow \text{Sign}(R^* \cup \{vk_{i_0^*}, vk_{i_1^*}\}, sk_{i_b^*}, m^*) \\ b' \leftarrow \mathcal{A}_2^{\mathcal{O}_2, \text{Corr}}(\mathfrak{s}, \sigma) \end{array} : b' = b \wedge |\{i_0^*, i_1^*\} \cap I| \leq \alpha \right] < \frac{1}{2} + \varepsilon(k), \quad (2)$$

where I is the set of queries to the corruption oracle; and the notation $\mathcal{A}^{\mathcal{O}, \text{Corr}}$ means that for each oracle O in \mathcal{O} , \mathcal{A} has oracle access to $O_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$, and \mathcal{A} also has oracle access to $\text{Corr}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$.

Definitions 2.5 and 2.6 are instantiations of Definition 2.3. They are equivalent to the correspondingly named definitions in [BKM09].

Definition 2.4 (Signing oracle OSign). For a ring signature scheme RS , the oracle $\text{OSign}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$ is defined to take as set R , input $i \in [N]$, and a message m , and output $\text{RS.Sign}(R \cup \{vk_i\}, sk_i, m)$. When the oracle is invoked with respect to a single key pair (i.e., $\text{OSign}_{(vk, sk)}$), we treat the oracle as taking only two inputs, m and R , since i is superfluous in this case.

Definition 2.5 (Anonymity against adversarially chosen keys). A ring signature scheme $\text{RS} = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies *anonymity against adversarially chosen keys* if it is $(\{\text{OSign}\}, \emptyset, 0)$ -anonymous. Moreover, RS satisfies *adaptive anonymity against adversarially chosen keys* if it is $(\{\text{OSign}\}, \{\text{OSign}\}, 0)$ -anonymous.

Definition 2.5 captures the guarantee that as long as there are at least two honest parties in a ring (represented by i_0^*, i_1^*), even if all other parties in the ring are corrupted by an adversary, the adversary cannot tell which of the honest parties produced a signature. One can also consider an even stronger definition where the adversary may corrupt *all but one* or even *all* of the parties in the ring, as in Definition 2.6.

⁷The function Gen^{-1} takes as input a verification key vk and signing key sk produced by Gen , and produces the randomness used by Gen to produce this key pair. That is, it samples from the set $\{\omega : \text{Gen}(1^k; \omega) = (vk, sk)\}$. In practice we will only ever invoke Gen^{-1} on a key pair produced by Gen , so we could invert efficiently by simply remembering the randomness used by Gen , but for the purposes of this definition we will describe it as a sampling procedure. Upon the first invocation on an input i , Corr samples $\omega_i \leftarrow \text{Gen}^{-1}(vk_i, sk_i)$, stores it, and outputs it. If Corr is queried twice on the same input i then it outputs the same ω_i that was previously stored.

Definition 2.6 (Anonymity against full key exposure). A ring signature scheme $RS = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies *anonymity against full key exposure* if it is $(\{\text{OSign}\}, \emptyset, 2)$ -anonymous.

Remark 5. Adaptive variants of anonymity were not given or discussed in prior work.⁸ In this paper, we refer primarily to *adaptive anonymity against adversarially chosen keys*: this is the strongest notion compatible with repudiability and claimability. To see this, observe that knowledge of a single one (say, $sk_{i_0}^*$) of the two challenge secret keys is sufficient to violate anonymity in a repudiable or claimable scheme, since the challenge signature σ was produced by $sk_{i_0}^*$ if and only if repudiating (resp. claiming) σ using $sk_{i_0}^*$ yields an invalid repudiation (resp. valid claim).

Definition 2.6 does not include an adaptive version because adaptivity does not give the adversary any additional power when he can corrupt all the keys.⁹

2.2 Unforgeability

The first unforgeability definition that follows is parametrized by an oracle set, taking a similar approach to our anonymity definitions above. In this section, we only give one instantiation of the parametrized definition of unforgeability. We will give other instantiations of Definition 2.7 later in the paper, in defining unforgeability for *repudiable* and *claimable* ring signature schemes (in Sections 3.1.1 and 3.3.1 respectively).

Definition 2.7 (\mathcal{O} -unforgeability). Let \mathcal{O} be a set of oracles, where each oracle in the set is parametrized by a list of key-pairs. A ring signature scheme $RS = (\text{Gen}, \text{Sign}, \text{Verify})$ is \mathcal{O} -unforgeable if for any PPT \mathcal{A} and any $N = \text{poly}(k)$, there is a negligible function ε such that

$$\Pr \left[\begin{array}{l} (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\ (R^*, m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}, \text{OSign}, \text{Corr}}(vk_1, \dots, vk_N) \\ b \leftarrow \text{Verify}(R^*, \sigma^*, m^*) \end{array} : \begin{array}{l} b = 1 \wedge R^* \subseteq \{vk_1, \dots, vk_N\} \setminus I \\ \wedge Q \cap \{(\cdot, m^*, R^*)\} = \emptyset \end{array} \right] < \varepsilon(k),$$

where the notation $\mathcal{A}^{\mathcal{O}, \text{OSign}, \text{Corr}}$ is defined as in Definition 2.3, and I and Q are the sets of queries made to the corruption and signing oracles respectively.

We refer to the event that the conditions on the right-hand side of the colon in the above probability expression are met as a “successful forgery.”

Definition 2.8 (Unforgeability of ring signatures). A ring signature scheme $RS = (\text{Gen}, \text{Sign}, \text{Verify})$ is *unforgeable* if it is \emptyset -unforgeable.¹⁰

⁸This notwithstanding, the prior constructions of [RST01; BKM09; BK10] all achieve adaptive anonymity. The constructions of [RST01; BK10] achieve perfect/statistical anonymity, respectively — the former, in the random oracle model — so adaptivity follows. The proof of anonymity of [BKM09] construction essentially suffices to prove adaptive anonymity, though the argument was not made in that paper since there was no adaptive definition of anonymity.

⁹An intermediate notion between Definitions 2.5 (which sets $\alpha = 0$) and 2.6 (which sets $\alpha = 2$) is *anonymity against attribution attacks*, defined in [BKM09], which effectively sets $\alpha = 1$. Adaptive anonymity against attribution attacks is not equivalent to the non-adaptive variant of the same.

¹⁰This is the definition described as *unforgeability with respect to insider corruption* in [BKM09], and is the strongest of the three unforgeability definitions considered therein.

3 New definitions: (un)repudiability and (un)claimability

3.1 Repudiable ring signatures

Repudiability addresses the question of whether or not members of a ring can prove that they did *not* sign a particular message (when they in fact did not sign it).

Definition 3.1 (Repudiable ring signature). A *repudiable ring signature scheme* is a ring signature scheme with an additional pair of algorithms (Repudiate , VerRepud), satisfying the four properties of *correctness* (Definition 2.2), *repudiability* (Definition 3.3), *anonymity* (Definition 3.4), and *unforgeability* (Definition 3.5). The syntax of Repudiate and VerRepud follows.

- $\text{Repudiate}(R, sk, \sigma, m)$ takes as input a set of verification keys $R = \{vk_1, \dots, vk_N\}$, a signing key sk , a ring signature σ , and a message m , and outputs a repudiation ξ .
- $\text{VerRepud}(R, vk, \sigma, m, \xi)$ takes as input a set R of verification keys, an identity vk , a signature σ , a message m , and a repudiation ξ , and outputs a single bit indicating whether or not ξ is a valid repudiation of signature σ for identity vk .

Definition 3.2 (Repudiation oracle ORpd). For a repudiable ring signature scheme RS , the oracle $\text{ORpd}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$ is defined to take as input a set R , $i \in [N]$, a signature σ , and a message m , and output $\text{RS.Repudiate}(R \cup \{vk_i\}, sk_i, \sigma, m)$. When the oracle is invoked with respect to a single key pair (i.e., $\text{ORpd}_{(vk, sk)}$), we treat the oracle as taking only three inputs, R , σ , and m , since i is superfluous in this case.

Additionally, we define the oracle $\text{ORpd}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}^{(R^*, m^*)}$ to output \perp when it receives both the ring R^* and message m^* as inputs, and otherwise to give the same response as $\text{ORpd}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$. We refer to this oracle as “punctured at R^* and m^* .”

Repudiability requires two conditions, expressed by equations (3) and (4) below. Intuitively, (3) captures the requirement “good people can repudiate,” i.e., that for any (possibly maliciously generated) signature, an honest party who did not produce it should be able to successfully repudiate. (4) captures the requirements that “bad people cannot repudiate a signature they produced,” i.e., addressing the case where the malicious signature and repudiation are both produced using the key being verified, and thus we want the signer to be unable to produce a valid repudiation.

Definition 3.3 (Repudiability). A ring signature scheme $\Sigma = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies *repudiability* if equipped with algorithms (Repudiate , VerRepud) such that the following conditions hold.

1. (*Non-signers can repudiate*) Let $\mathcal{O} = \{\text{OSign}\}$. For any (possibly adversarial) PPT signing algorithm $\mathcal{A}_{\text{Sign}}$, there exists a negligible function ε such that

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{Gen}(1^k) \\ (\sigma, m, R') \leftarrow \mathcal{A}_{\text{Sign}}^{\mathcal{O}, \text{ORpd}_{(vk, sk)}}(vk) \\ \xi \leftarrow \text{Repudiate}(R', sk, \sigma, m) \\ b \leftarrow \text{VerRepud}(R', vk, \sigma, m, \xi) \\ b' \leftarrow \text{Verify}(R', \sigma, m) \end{array} : \begin{array}{l} b = 1 \vee b' = 0 \\ \forall Q \cap \{(\cdot, m, R')\} \neq \emptyset \end{array} \right] > 1 - \varepsilon(k). \quad (3)$$

2. (*Signer cannot repudiate*) For any (possibly adversarial) sign-and-repudiate algorithm $\mathcal{A}_{\text{S\&R}}$,

there is a negligible function ε such that for any $N = \text{poly}(k)$,

$$\Pr \left[\begin{array}{l} (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\ (\sigma, R', m, \{\xi_{vk}\}_{vk \in R' \setminus R}) \leftarrow \mathcal{A}_{\text{S\&R}}^{\mathcal{O}}(R) \\ \forall vk \in R' \setminus R, b_{vk} \leftarrow \text{VerRepub}(R', vk, \sigma, m, \xi_{vk}) \\ b' \leftarrow \text{Verify}(R', \sigma, m) \end{array} : \begin{array}{l} R' \cap R = \emptyset \vee \bigvee_{vk \in R' \setminus R} b_{vk} = 0 \\ \forall b' = 0 \vee Q \cap \{(\cdot, m, R')\} \neq \emptyset \end{array} \right] > 1 - \varepsilon(k), \quad (4)$$

where $R = \{vk_1, \dots, vk_N\}$, $\mathcal{O} = \{\text{OSign}, \text{ORpd}\}$, and Q is the set of queries to OSign .

Remark 6. Equation 4 guarantees that a party possessing a set of signing keys cannot repudiate under all of these keys, *as long as some key in the ring is honestly generated*. If the adversary generates all keys in the ring, then he may be able to produce a repudiation under every key in the ring. However, this does not undermine the purpose of repudiability: indeed, if presented with repudiations under every key in a ring, one can confidently conclude that all keys in the ring were generated dishonestly, and thus that all parties in the ring effectively colluded to produce each signature under that ring. Similarly, given repudiations for a subset of the identities in a ring, one can conclude that *either* one of the remaining identities in the ring produced the signature *or* all of the remaining identities in the ring colluded maliciously to produce the signature. That is, either way, at least one of the remaining identities is responsible for the signature.

3.1.1 Anonymity and unforgeability of repudiable ring signatures

The definitions of anonymity and unforgeability need to be adapted for repudiable ring signature schemes, to incorporate a repudiation oracle as described next.

Definition 3.4 (Anonymity of repudiable ring signatures). A repudiable ring signature scheme

$$(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Repudiate}, \text{VerRepub}))$$

satisfies *anonymity against adversarially chosen keys* if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $(\{\text{OSign}, \text{ORpd}\}, \emptyset, 0)$ -anonymous (Definition 2.3). Moreover, the repudiable ring signature satisfies *adaptive anonymity against adversarially chosen keys* if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $(\{\text{OSign}, \text{ORpd}\}, \{\text{OSign}, \text{ORpd}^{(R^*, m^*)}\}, 0)$ -anonymous, where R^* and m^* are the challenge ring and message in the anonymity experiment (Equation 2).¹¹

Recall from Remark 5 that adaptive anonymity against adversarially chosen keys is the strongest anonymity notion compatible with repudiability.

Definition 3.5 (Unforgeability of repudiable ring signatures). A repudiable ring signature scheme

$$(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Repudiate}, \text{VerRepub}))$$

is unforgeable if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $\{\text{ORpd}\}$ -unforgeable (Definition 2.7).

¹¹ One may ask whether the oracle ORpd in the second phase should be punctured more specifically only at the challenge identities i_0^*, i_1^* , since signatures of the message m^* by other parties in the ring R^* seem like they should be admissible repudiation queries. But puncturing at just the challenge identities is unnecessary: the adversary in Definition 2.3 additionally receives a corruption oracle, so without loss of generality, the adversary can corrupt all parties in the ring except for i_0^*, i_1^* , and can run the repudiation algorithm for these parties without invoking the oracle. The repudiation oracle in the second phase is thus only necessary for repudiations with respect to the challenge identities i_0^*, i_1^* .

Remark 7. Some previous versions of this paper (including the conference version) presented a slightly different definition of anonymity of repudiable ring signatures: puncturing the second ORpd oracle at the challenge signature σ instead of at R^*, m^* . These previous versions erred in presenting the same construction as the current one, R-RS, which in fact satisfies the current definition but not the old one. Thanks to Hao Lin and Mingqiang Wang for pointing out this issue and correction [Lin19; LW19].

It turns out that the old version of the definition implies that any repudiable ring signature scheme must satisfy a strong and unintuitive non-malleability property: namely, it must be infeasible, given an honest party’s signature, to produce a different signature on the same message (valid for the same honest party). In other words, the old version of the definition is undesirably brittle in that any repudiable signature scheme can be generically transformed to a very similar scheme which no longer technically satisfies repudiability under the old definition, by appending a uniformly random bit to each signature: this yields an attack where the adversary simply flips the last bit of the challenge signature and violates anonymity by querying the ORpd oracle on the “flipped” signature and one of the challenge identities. In light of this, we believe the current Definition 3.4 is the natural one.

3.2 Unrepudiable ring signatures

We next consider a notion where it is *not* possible for a party to prove to others that he did not produce a particular signature. In fact, though it may not be immediately apparent, a natural formalization of this notion is expressed by the definition of *anonymity against full key exposure* (Definition 2.6): that is, the strongest of the anonymity definitions given in Section 2. The following paragraphs justify this claim with detailed intuition.

Recall that anonymity against full key exposure (FKE) preserves signer anonymity even against an adversary that obtains all of the secret keys of all members of a ring. A ring signature scheme that satisfies repudiability could not also satisfy anonymity against FKE, because of the following attack: the adversary obtains all secret keys in the ring, attempts to repudiate using each secret key, and identifies as the signer the one secret key with respect to which the repudiation algorithm does not produce a valid repudiation. With overwhelming probability, by definition of repudiability, there is exactly one such secret key.

This informal argument establishes that anonymity against FKE must imply any reasonable notion of unrepudiability. The question then arises: are the two notions equivalent? While there arguably exist meaningful definitions of unrepudiability that are weaker than anonymity against FKE, we believe anonymity against FKE is the most reasonable definition of unrepudiability, as explained next.

Any reasonable definition of unrepudiability should capture the intuitive requirement that non-signers cannot behave distinguishably from signers. A little more precisely, for any protocol that could be executed by a non-signer Nancy with respect to a signature σ and her verification key vk' , the signer Sigmund of that signature must be able to engage in the same protocol with respect to his own verification key vk and behave indistinguishably from Nancy. In other words, we require that if Nancy’s secret key were stolen, the thief would be unable to tell whether σ was produced by Nancy or by someone else. Indeed, if Nancy were stateless and did not remember what signatures she had produced in the past, or simply lent her secret key to someone else who used it to produce signatures, then she herself would not be able to tell. The definition of anonymity against FKE embodies almost exactly this requirement — but instead of requiring anonymity against the thief

who steals just Nancy’s key, the definition makes the stronger requirement that anonymity must hold even against a thief who has every secret key in the ring corresponding to σ .

Is a weaker definition, which only rules out *unilateral* repudiations by a single party, a meaningful definition of unrepudiability? Perhaps. However, it is more in keeping with the intuitive goals and standard properties of ring signatures to protect against adversaries that may have many or all secret keys in a ring: that is, to rule out even the possibility of multiple ring members *colluding* to produce a repudiation for some ring member. Thus we arrive at the following definition.

Definition 3.6 (Unrepudiable ring signature scheme). A ring signature scheme satisfies *unrepudiability* if it satisfies anonymity against full key exposure (Definition 2.6).

3.3 Claimable ring signatures

Claimability addresses whether the actual signer can prove later that they were the signer, without remembering the signing randomness.

Definition 3.7 (Claimable ring signature). A *claimable ring signature scheme* is a ring signature scheme with an additional pair of algorithms ($\text{Claim}, \text{VerClaim}$), satisfying the four properties of *correctness* (Definition 2.2), *claimability* (Definition 3.9), *anonymity* (Definition 3.10), and *unforgeability* (Definition 3.11). The syntax of Claim and VerClaim follows.

- $\text{Claim}(R, sk, \sigma)$ takes as input a signing key sk , a ring signature σ , and a set of verification keys $R = \{vk_1, \dots, vk_N\}$, and outputs a claim ζ .¹²
- $\text{VerClaim}(R, vk, \sigma, \zeta)$ takes as input a set R of verification keys, a signature σ , a claim ζ , and an identity vk , and outputs a single bit indicating whether or not ζ is a valid claim of signature σ for identity vk .

Definition 3.8 (Claim oracle OClaim). For a claimable ring signature scheme RS , the oracle $\text{OClaim}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$ is defined to take as input $i \in [n]$, a set R , and a signature σ , and output $\text{RS.Claim}(R, sk, \sigma)$. When the oracle is invoked with respect to a single key pair (i.e., $\text{OClaim}_{(vk, sk)}$), we treat the oracle as taking only two inputs, R and σ , since i is superfluous in this case.

Additionally, we define the oracle $\text{OClaim}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}^{\langle \sigma^* \rangle}$ to output \perp when it receives the signature σ^* as input, and otherwise to give the same response as $\text{OClaim}_{(vk_1, sk_1), \dots, (vk_N, sk_N)}$.

Claimability requires three conditions, expressed by equations (5), (6), and (7) below. Informally, (5) requires that honest signers can successfully claim their signatures, (6) requires that adversarial parties cannot successfully claim a signature that they did not produce, and (7) requires that adversarial parties cannot produce a signature along with a claim that appears to be produced by an honest party (that is, falsely framing the honest party as the signer).

Definition 3.9 (Claimability). A ring signature scheme ($\text{Gen}, \text{Sign}, \text{Verify}$) is *claimable* if equipped with algorithms ($\text{Claim}, \text{VerClaim}$) such that the following conditions hold.

1. (*Honest signer can claim*) There exists a negligible function ε such that for any $N = \text{poly}(k)$ and $(vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k)$ and any $i \in [N]$, it holds for any message m that

$$\Pr[\sigma \leftarrow \text{Sign}(R, sk_i, m) : \text{VerClaim}(R, vk_i, \sigma, \text{Claim}(R, sk_i, \sigma)) = 1] > 1 - \varepsilon(k), \quad (5)$$

¹² Claim could also take the message m as an additional input (analogously to Repudiate). We omit m to simplify notation, because our construction of Claim does not use m . However, definitionally speaking, m may be thought of as an additional implicit input to Claim (which could be used in other constructions).

where $R = \{vk_1, \dots, vk_N\}$.¹³

2. (*Non-signers cannot claim*) Let $\mathcal{O} = \{\text{OSign}\}$. For any (possibly adversarial) PPT sampling-and-claiming algorithm $\mathcal{A}_{\text{Claim}} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function ε such that

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{Gen}(1^k) \\ (R', m, \mathfrak{s}) \leftarrow \mathcal{A}_1^{\mathcal{O}, \text{OClaim}_{(vk, sk)}}(vk) \\ \sigma \leftarrow \text{Sign}(R' \cup \{vk\}, sk, m) \\ (\zeta, vk') \leftarrow \mathcal{A}_2^{\mathcal{O}, \text{OClaim}_{(vk, sk)}}(R' \cup \{vk\}, \sigma, \mathfrak{s}) \\ b \leftarrow \text{VerClaim}(R' \cup \{vk\}, vk', \sigma, \zeta) \\ b' \leftarrow \text{Verify}(R' \cup \{vk\}, \sigma, m) \end{array} : \begin{array}{l} b = 1 \wedge b' = 1 \\ \wedge vk' \neq vk \end{array} \right] < \varepsilon(k). \quad (6)$$

3. (*Malicious signer cannot frame an honest party*) For any PPT adversary $\mathcal{A}_{\text{S\&C}}$, there exists a negligible function ε such that

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{Gen}(1^k) \\ (R', m, \sigma, \zeta) \leftarrow \mathcal{A}_{\text{S\&C}}^{\mathcal{O}, \text{OClaim}_{(vk, sk)}}(vk) \\ b \leftarrow \text{VerClaim}(R' \cup \{vk\}, vk, \sigma, \zeta) \\ b' \leftarrow \text{Verify}(R' \cup \{vk\}, \sigma, m) \end{array} : \begin{array}{l} b = 1 \wedge b' = 1 \\ \wedge Q \cap \{(\cdot, \sigma)\} = \emptyset \end{array} \right] < \varepsilon(k). \quad (7)$$

where $\mathcal{O} = \{\text{OSign}\}$ and Q is the set of queries made to oracle $\text{OClaim}_{vk, sk}$.

Remark 8. Our definition and construction guarantee that all honestly generated signatures can be claimed by the signer, and even malicious signers cannot frame non-signers. Our definition does not guarantee that all signatures that verify can be claimed by someone; requiring this would be a reasonable alternative definition, but we believe our claimability definition is reasonable and sufficient in many settings. Since claimability is a notion designed for the signer’s benefit, it is not an “attack” that the signer can choose to waive the ability to claim. If two ring members collude to produce a signature, then they are effectively both signers; therefore, it is not an attack if two signers collude to produce a signature claimable by one of them.

Our definition has a property that may be counterintuitive at first glance: namely, a single party might be able to claim and repudiate the same signature. However, this possibility is in fact fully consistent with our security guarantees, as follows. If the party does both actions, we (the observers) know she is adversarial. If she claims, we know that she either signed or was part of an adversarial group that signed. If she repudiates, we know that either she didn’t sign, or she was part of an adversarial group of which another member must be unable to repudiate.

3.3.1 Anonymity and unforgeability of claimable ring signatures

The definitions of anonymity and unforgeability must be adapted for claimable ring signature schemes, to allow the adversary a claim oracle as described next.

Definition 3.10 (Anonymity of claimable ring signatures). A claimable ring signature scheme $(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Claim}, \text{VerClaim}))$ satisfies anonymity against adversarially chosen keys if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $(\{\text{OSign}, \text{OClaim}\}, \emptyset, 0)$ -anonymous (Definition 2.3). Moreover, the repudiable ring signature satisfies *adaptive* anonymity against adversarially chosen keys if $(\text{Gen}, \text{Sign}, \text{Verify})$ is

$$(\{\text{OSign}, \text{OClaim}\}, \{\text{OSign}, \text{OClaim}^{(\sigma)}\}, 0)\text{-anonymous},$$

¹³Like Definition 2.2, Equation 5 considers only honestly generated keys. See Remark 4 for further discussion.

where σ is the challenge signature in the anonymity experiment (Equation (2)).

Recall from Remark 5 that adaptive anonymity against adversarially chosen keys is the strongest anonymity notion compatible with claimability.

Definition 3.11 (Unforgeability of claimable ring signatures). A claimable ring signature scheme $(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Claim}, \text{VerClaim}))$ is *unforgeable* if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $\{\text{OClaim}\}$ -unforgeable (Definition 2.7).

3.4 Unclaimable ring signatures

An unclaimable ring signature scheme has the property that the signer cannot later convince anyone of her identity. That is, for any function that the true signer can compute given the signing randomness and the secret key, any other member of the ring can compute an indistinguishable function. The result is that even an adversary holding all ring members under duress cannot figure out who produced a given signature. This is true even if the ring members under duress attempt to cooperate with the adversary.

To achieve this, it suffices for any member of the ring to be able to extract signing randomness distributed indistinguishably from true signing randomness, that would produce the given signature under their secret key. More formally, the following guarantee should hold.

Definition 3.12 (Unclaimable ring signatures). A *unclaimable* ring signature scheme is a ring signature scheme augmented with an additional algorithm `ExtractRandomness` as follows.

- `ExtractRandomness` (R, sk, σ, m) takes as input a ring R , a secret key sk , a signature σ and a message m . If sk is one of the secret keys for ring R , and σ is a signature on m with respect to R , then it outputs randomness ρ .

`ExtractRandomness` must satisfy the following condition.

- (*Statistical unclaimability*) Let \mathcal{R} be the distribution of signing randomness. For any $N = \text{poly}(k)$ there is a negligible function ϵ such that the following holds. Let $(vk_1, sk_1), (vk_2, sk_2) \leftarrow \text{Gen}(1^k)$. For any message m and any vk_3, \dots, vk_N and sk_3, \dots, sk_N , let $R = \{vk_1, \dots, vk_N\}$ and $S = \{(i, vk_i, sk_i)\}_{i \in [N]}$. Let $\rho \leftarrow \mathcal{R}$, $\sigma_1 \leftarrow \text{Sign}(R, sk_1, m; \rho)$, and $\rho_1 \leftarrow \text{ExtractRandomness}(R, sk_2, \sigma_1, m)$. Let $\rho_2 \leftarrow \mathcal{R}$ and $\sigma_2 \leftarrow \text{Sign}(R, sk_2, m; \rho_2)$. Then

$$(S, \rho_1, \sigma_1) \approx_\epsilon (S, \rho_2, \sigma_2).$$

Definition 3.12 is unusual among the definitions in this paper, in that it gives a statistical rather than a computational guarantee. We opted to give the statistical definition because it is simpler, it is a stronger guarantee, and our construction in this case achieves the statistical guarantee. One could also consider a computational definition.

Remark 9 (Claimability is not the opposite of unclaimability). According to these definitions, unclaimability is *not* technically the opposite of claimability (even when ignoring the fact that the formal definitions give a statistical guarantee for unclaimability but a computational guarantee for claimability). Claimability requires the ability to “voluntarily claim” a signature *without remembering the signing randomness*, whereas unclaimability rules out the ability to “claim under duress” *even given the signing randomness*. For voluntary claims, the natural and stronger definition is to guarantee the ability to claim adaptive, without “planning ahead” and without the storage requirement of remembering the signing randomness. In contrast, when considering attempts to claim

under duress, the natural and stronger definition is to rule out the possibility of successful claims even in the presence of the signing randomness.

Remark 10 (Unclaimability protects *honest* signers). An adversarial signer who *wants* to claim can devise ways of credibly later claiming a ring signature, even when using an unclaimable ring signature scheme.¹⁴ This does not decrease the utility of an unclaimable ring signature scheme for honest signers who *want* their signatures to be unclaimable.

Remark 11 (Secure erasure and rubber-hose adversaries). Our unclaimability definition does not protect signers against adversaries that can unexpectedly compromise the internal state of ring members, in the following sense: either the signer will have securely erased the signing randomness, in which case the adversary will not be able to determine the signer’s identity; or she won’t have erased the signing randomness, in which case the adversary will be able to distinguish the signer from other ring members. In the first case, the signer’s anonymity comes from her diligent erasure, rather than from unclaimability; and the first case is only possible assuming secure erasure is feasible.

We believe our unclaimability definition still has twofold value. First, even assuming secure erasure, a system is stronger if it eliminates reliance on users to erase diligently. Users often fail to use systems as they are supposed to; this is a major source of real-world security failures. Our definition’s security does not rely on any user actions beyond running the honest signing algorithm.

Secondly, unclaimability is useful beyond the context of such powerful, real-time adversaries. Consider an honest signer who does not intend to claim at the time of signing, but later is corrupted or changes her mind. Our definition guarantees that the signer’s intent at the time of signing is what matters: if she intended unclaimability when signing, she cannot later claim that signature.

3.4.1 Unclaimability implies unrepudiability

Any unclaimable ring signature scheme is also unrepudiable. Recall that the definition of unclaimability captures the idea that for any function that the true signer can compute given the signing randomness and the secret key, any other member of the ring can compute an indistinguishable function. Intuitively, the implication follows from the fact that repudiation would require a non-signer to behave in a way that *distinguishable* from any possible behavior of the actual signer.

Theorem 3.13. *Any unclaimable ring signature scheme is also unrepudiable.*

Proof. Recall that unrepudiability (Definition 3.6) is defined as anonymity against full key exposure (Definition 2.6), which is $(\{\text{OSign}\}, \emptyset, 2)$ -anonymity under the framework of Definition 2.3. Thus, a ring signature scheme is unrepudiable if for any adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and polynomial N , it satisfies (2) of Definition 2.3 for $(\mathcal{O}_1, \mathcal{O}_2, \alpha) = (\{\text{OSign}\}, \emptyset, 2)$. Note that since $\alpha = 2$, we can consider without loss of generality only adversaries that use the corruption oracle to learn all N secret keys in (2), and do not make any queries to the signing oracle (since the adversary can produce signatures himself, using the secret keys).

To establish the theorem, it suffices to show that for any unclaimable ring signature scheme, the view of the adversary \mathcal{A} is indistinguishable between the cases $b = 0$ and $b = 1$. Unclaimability (Definition 3.12) directly implies that these two views are indistinguishable. To see this, recall

¹⁴For example, an adversarial signer might use a PRG output as his signing randomness, or append it to his message, and remember the preimage. If he later revealed the preimage, it would likely serve as a credible claim to authorship of the signature.

that the adversary’s inputs in (2) of Definition 2.3 are an honestly generated set of verification keys vk_1, \dots, vk_N and a signature σ produced by the honest signing algorithm using a secret key corresponding to some $vk_i, i \in [N]$. The theorem follows. \square

Remark 12. It is unclear that the reverse implication holds even for a computational definition of unclaimability. The main complication is that unclaimability must hold even if the signing randomness is saved, while this is not an issue for unrepudiability. For instance, an algorithm that appends a commitment to the signing randomness (or to a random nonce) could be unrepudiable, but could be claimed by a signer who remembered the signing randomness.

3.5 Repudiable-and-claimable ring signatures

Suppose that $(\text{Gen}, \text{Sign}, \text{Verify})$ is a ring signature scheme, and there are algorithms Repudiate , VerRepud , Claim , and VerClaim such that

$$(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Repudiate}, \text{VerRepud})) \quad \text{and} \quad (\text{Gen}, \text{Sign}, \text{Verify}, (\text{Claim}, \text{VerClaim}))$$

are a repudiable ring signature scheme and a claimable ring signature scheme respectively (Definitions 3.1 and 3.7). The seven algorithms together do *not* necessarily satisfy the natural notion of a “repudiable-and-claimable” ring signature scheme.

The reason, in a nutshell, is that Definition 3.1 (repudiability) does *not* allow the adversary a claim oracle, and likewise Definition 3.7 (claimability) does *not* allow the adversary a repudiation oracle. Indeed, it would not make sense even syntactically for the “other oracles” to be provided: since each of Definitions 3.1 and 3.7 is defined with respect to a *quintuple* of algorithms either containing Repudiate but not Claim , or vice versa, the concept of the “other oracle” is undefined within the scope of each definition.

Thus, it could be that when an adversary has access to both a claim and a repudiation oracle, the resulting scheme is no longer secure. Indeed, there are simple (though arguably unnatural) examples of schemes where this happens, such as the following.

Example 1. Given *any* ring signature scheme, augment the signing key sk to a new signing key $sk' = (sk, \eta_0, \eta_1)$ that additionally contains a pair η_0, η_1 such that η_0 is sampled uniformly randomly and $\eta_0 \oplus \eta_1 = sk$. Sign works just as in the original scheme, using only sk and ignoring η_0, η_1 . Repudiate produces repudiations just as in the original scheme, but additionally appends η_0 to every repudiation. Claim produces claims just as in the original scheme, but additionally appends η_1 to every repudiation. This modified scheme would be repudiable if the original scheme was, and also claimable if the original scheme was. However, an adversary that could see both a repudiation and a claim would straightforwardly be able to recover sk and thereby forge signatures.

The natural security definition for a repudiable-and-claimable ring signature scheme is to include both repudiation and claim oracles throughout the repudiability, claimability, anonymity, and unforgeability definitions. As the resulting formal definitions are somewhat repetitive, we defer them to Appendix A.

4 Repudiable construction

4.1 Building blocks

ZAPs ZAPs are two-message public coin witness indistinguishable proofs [DN07].

Definition 4.1 (ZAP). A ZAP for an NP language L with witness relation \mathcal{R}_L is a triple of algorithms $\text{ZAP}_L = (\text{ZAP.Setup}_L, \text{ZAP.Prove}_L, \text{ZAP.Verify}_L)$, where ZAP.Setup and ZAP.Prove are PPT and ZAP.Verify is polynomial-time and deterministic, satisfying the following properties.

Public coin. For some polynomial $\ell = \ell(k)$, ZAP.Setup is the algorithm that on input 1^k , outputs a uniformly random element of $\{0, 1\}^\ell$.

Completeness. For $(x, w) \in \mathcal{R}_L$ and any $\rho \in \{0, 1\}^{\ell(k)}$ we have

$$\Pr_{\pi \leftarrow \text{ZAP.Prove}(\rho, x, w)} [\text{ZAP.Verify}(\rho, \pi, x) = 1] = 1.$$

Adaptive soundness. There exists a negligible function ϵ such that

$$\Pr_{\rho \leftarrow \text{ZAP.Setup}(1^k)} [\exists(x, \pi) : x \notin L \wedge \text{ZAP.Verify}(\rho, \pi, x)] \leq \epsilon(k).$$

Witness indistinguishability. For any sequences $\{\rho_k\}_{k \in \mathbb{N}}$, $\{x_k\}_{k \in \mathbb{N}}$, $\{w_{0,k}\}_{k \in \mathbb{N}}$, $\{w_{1,k}\}_{k \in \mathbb{N}}$, where for all k , $\rho_k \in \{0, 1\}^{\ell(k)}$, $x_k \in L$ and $(x_k, w_{0,k}), (x_k, w_{1,k}) \in \mathcal{R}_L$, the following pair of ensembles is computationally indistinguishable:

$$\{\text{ZAP.Prove}(\rho_k, x_k, w_{0,k})\}_{k \in \mathbb{N}} \stackrel{c}{\approx} \{\text{ZAP.Prove}(\rho_k, x_k, w_{1,k})\}_{k \in \mathbb{N}}.$$

In this work, for simplicity, we will assume use of a ZAP for some NP-complete language L_{NP} (with witness relation $\mathcal{R}_{L_{\text{NP}}}$) and for any $L \in \text{NP}$ with witness relation \mathcal{R}_L , we define ZAP.Prove_L and ZAP.Verify_L as follows.

- ZAP.Prove_L takes as input a triple (ρ, x, w) . If $(x, w) \notin \mathcal{R}_L$, then output \perp . Otherwise, use an NP reduction on (x, w) to get a pair $(x_{\text{NP}}, w_{\text{NP}}) \in \mathcal{R}_{L_{\text{NP}}}$, and output $\text{ZAP.Prove}(\rho, x, w)$.
- ZAP.Verify_L takes as input a triple (ρ, π, x) , uses the same NP reduction to obtain x_{NP} (which is in L_{NP} iff $x \in L$), and outputs $\text{ZAP.Verify}(\rho, \pi, x)$.

Verifiable random functions (VRFs) [MRV99] are another main building block of our construction. The important property of VRFs that we rely on is *residual pseudorandomness*, i.e., that VRF outputs on inputs for which the adversary has not received proofs remain indistinguishable from random.

Definition 4.2 (VRF). A *verifiable random function* (VRF) is a tuple of algorithms $\text{VRF} = (\text{VRF.Gen}, \text{VRF.Eval}, \text{VRF.Prove}, \text{VRF.Verify})$, where Gen and Verify are PPT and Eval and Prove are polynomial time and deterministic, satisfying:

Complete provability With probability at least $1 - 2^{-\Omega(k)}$ over $(pk, sk) \leftarrow \text{VRF.Gen}(1^k)$, we have for all inputs x that

$$\Pr[\text{VRF.Verify}(pk, x, \text{VRF.Eval}(sk, x), \text{VRF.Prove}(sk, x)) = 1] > 1 - 2^{-\Omega(k)}.$$

Unique provability For all $pk, x, y_1, y_2, \tau_1, \tau_2$ with $y_1 \neq y_2$, for either $i = 1$ or $i = 2$ it holds that

$$\Pr[\text{VRF.Verify}(pk, x, y_i, \tau_i) = 1] < 2^{-\Omega(k)}.$$

Residual pseudorandomness Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a probabilistic polynomial-time adversary, where both \mathcal{A}_1 and \mathcal{A}_2 get oracle access to the VRF evaluation and prove algorithms. Let $(pk, sk) \leftarrow \text{VRF.Gen}(1^k)$, and let

$$(x, \mathfrak{s}) \leftarrow \mathcal{A}_1^{\text{VRF.Eval}(sk, \cdot), \text{VRF.Prove}(sk, \cdot)}(1^k, pk).$$

Let $b \leftarrow \{0, 1\}$, and let v be either $\text{VRF.Eval}(sk, x)$ or uniformly random, depending on the choice bit b . Let

$$b' = \mathcal{A}_2^{\text{VRF.Eval}(sk, \cdot), \text{VRF.Prove}(sk, \cdot)}(1^k, v, \mathfrak{s}).$$

Then there is a negligible function ϵ such that $\Pr[b = b' \text{ and } x \notin Q] < 1/2 + \epsilon(k)$, where Q is the set of oracle queries made by \mathcal{A} to either oracle.

For simplicity, we assume that Eval takes inputs x of any length, i.e., $x \in \{0, 1\}^*$.

Definition 4.3. The *verification failure probability* of a VRF VRF is

$$\Pr \left[\begin{array}{l} (pk, sk) \leftarrow \text{VRF.Gen}(1^k) \\ b \leftarrow \text{VRF.Verify}(pk, x, \text{VRF.Eval}(sk, x), \text{VRF.Prove}(sk, x)) \end{array} : b = 0 \right].$$

The residual pseudorandomness property of the VRF still holds even if the adversary gets to query many key pairs at once, and gets to adaptively choose to learn some of the secret keys (in this case, residual pseudorandomness holds for the uncorrupted keys only).

Lemma 4.4 (Parallel VRF Game). *Let VRF be a VRF. Then \forall PPT $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and all $N = \text{poly}(k)$, there is a negligible function ϵ such that*

$$\Pr \left[\begin{array}{l} (pk_1, sk_1), \dots, (pk_N, sk_N) \leftarrow \text{VRF.Gen}(1^k) \\ (m^*, \mathfrak{s}) \leftarrow \mathcal{A}_1^{\mathcal{V}, \text{Corr}}(vk_1, \dots, vk_N) \\ \forall i \in [N], y_{i,0} \leftarrow \text{VRF.Eval}(sk_i, m^*) \\ \forall i \in [N], y_{i,1} \leftarrow \$ \\ b \leftarrow \{0, 1\} \\ b' \leftarrow \mathcal{A}_2^{\mathcal{V}^{(m^*)}}(\mathfrak{s}, (y_{i,b})_{i \in [N] \setminus C}) \end{array} : b = b' \wedge \forall i \in [N] \setminus C, (i, m^*) \notin Q \right] < 1/2 + \epsilon(k), \quad (8)$$

where \mathcal{V} is an oracle that takes as input a pair (i, m) and outputs

$$(y, \tau) = (\text{VRF.Eval}(sk_i, m), \text{VRF.Prove}(sk_i, m)),$$

$\mathcal{V}^{(m^*)}$ is the same oracle punctured at $m = m^*$ (i.e. on input (i, m) , the oracle outputs \perp if $m = m^*$ and behaves the same as \mathcal{V} if $m \neq m^*$) and Corr is an oracle that takes as input an index $i \in [N]$ and outputs sk_i , and C denotes the set of queries made to the corruption oracle, and Q denotes the set of the queries to \mathcal{V} .

We will refer to the game described by (8) as the Parallel VRF Game.

4.2 Construction

Construction 4.5. Our construction R-RS is parametrized by ZAP, VRF, and M , where

- ZAP is a ZAP;
- VRF is a VRF with input domain $\{0,1\}^*$, whose Verify algorithm takes ν bits of randomness and whose verification failure probability (Definition 4.3) is ε ,¹⁵
- M is a polynomial satisfying $M \geq (\nu + k)/\log_2(1/\varepsilon)$.¹⁶

We first present the Gen algorithm of our ring signature scheme R-RS.

R-RS.Gen(1^k)

1. $(vk_{\text{VRF}}^1, sk_{\text{VRF}}^1), \dots, (vk_{\text{VRF}}^4, sk_{\text{VRF}}^4) \leftarrow \text{VRF.Gen}(1^k)$.
Let $\vec{vk}_{\text{VRF}} = (vk_{\text{VRF}}^1, \dots, vk_{\text{VRF}}^4)$ and $\vec{sk}_{\text{VRF}} = (sk_{\text{VRF}}^1, \dots, sk_{\text{VRF}}^4)$.
2. $\rho \leftarrow \text{ZAP.Setup}(1^k)$.
3. $\vec{\alpha} = (\alpha_1, \dots, \alpha_M) \leftarrow (\{0,1\}^\nu)^M$.
4. Output $vk = (\vec{vk}_{\text{VRF}}, \rho, \vec{\alpha})$ and $sk = (\vec{sk}_{\text{VRF}}, vk)$.¹⁷

In the rest of the section, we (implicitly) use the following convention to parse a ring R .

$$\begin{aligned} & \text{Write } R = \{vk_1, \dots, vk_N\}. \\ & \text{For each } i \in [N], \text{ write } vk_i = (vk_{\text{VRF}}^i = (vk_{\text{VRF}}^{i,1}, \dots, vk_{\text{VRF}}^{i,4}), \rho_i, \vec{\alpha}_i = (\alpha_1^i, \dots, \alpha_M^i)). \end{aligned} \tag{9}$$

Definition 4.6. Let L be the following NP language.

$$\begin{aligned} & \{(R, m, \varphi, (y_1, y_2, y_3, y_4)) : \exists i^*, \tau_1, \tau_2, \tau_3, \tau_4, \gamma \text{ s.t. } (b_1 \vee b_2) \wedge (b_3 \vee b_4) \\ & \text{where } \forall \eta \in \{1, 2, 3, 4\}, b_\eta = \bigwedge_{i \in [N], j \in [M]} \text{VRF.Verify}(vk_{\text{VRF}}^{i^*, \eta}, (R, m, \varphi), y_\eta, \tau_\eta; \alpha_j^i \oplus \gamma)\} . \end{aligned}$$

We now present the Sign and Verify algorithms of our construction.

R-RS.Sign(R, sk, m)

1. Parse R as described above and $sk = ((sk_{\text{VRF}}^1, \dots, sk_{\text{VRF}}^4), vk)$.
2. If $vk \notin R$ output \perp and halt.
3. Define $i^* \in [N]$ such that $vk_{i^*} = vk$.
4. $\gamma \leftarrow \{0,1\}^\nu$. (This is used as part of the ZAP witness in Step 6.)
5. $\varphi \leftarrow \{0,1\}^k$. (This is used as a salt for the VRF input in Step 7, and output in Step 8.)
6. For $\eta \in \{1, 2, 3, 4\}$, let $y_\eta = \text{VRF.Eval}(sk_{\text{VRF}}^\eta, (R, m, \varphi))$ and $\tau_\eta = \text{VRF.Prove}(sk_{\text{VRF}}^\eta, (R, m, \varphi))$.
Let $\vec{y} = (y_1, \dots, y_4)$.
7. For each $i \in [N]$, let $\pi_i \leftarrow \text{ZAP.Prove}_L(\rho_i, (R, m, \varphi, \vec{y}), (i^*, \tau_1, \perp, \tau_3, \perp, \gamma))$.
Let $\vec{\pi} = (\pi_1, \dots, \pi_N)$.
8. Output $\sigma = (\vec{\pi}, \vec{y}, \varphi)$.

¹⁵ ε need not be explicitly known, as discussed in footnote 16.

¹⁶This inequality is required in order to invoke Corollary 4.11. As explained in Remark 14, a satisfactory value of M can be set even without knowledge of ε . If ε happens to be known, a smaller value of M can be chosen.

¹⁷We include the verification key in sk so that the Sign procedure can identify the verification key in the ring corresponding to the signing key.

R-RS.Verify(R, σ, m)

1. Parse R as above and $\sigma = ((\pi_1, \dots, \pi_N), \vec{y}, \varphi)$.
2. Output $\bigwedge_{i \in [N]} \text{ZAP.Verify}_L(\rho_i, \pi_i, (R, m, \varphi, \vec{y}))$.

Now that we have described the main algorithms of R-RS, we proceed to describe the repudiation algorithms for R-RS.

Definition 4.7. Let L' be the following NP language:

$$\left\{ (R, m, \varphi, (y_1, \dots, y_4), vk = (v\vec{k}_{\text{VRF}}, \rho, \vec{\alpha})) : \exists i^*, y'_1, \dots, y'_4, \tau'_1, \dots, \tau'_4, \gamma \text{ s.t.} \right. \\ \left. ((b'_1 \wedge b'_2) \vee (b'_3 \wedge b'_4)) \wedge vk = vk_{i^*}, \text{ where } \forall \eta \in \{1, 2, 3, 4\}, \right. \\ \left. b'_\eta = \left(y'_\eta \neq y_\eta \wedge \bigwedge_{i \in [N], j \in [M]} \text{VRF.Verify}(vk_{\text{VRF}}^{i^*, \eta}, (R, m, \varphi), y'_\eta, \tau'_\eta; \alpha_j^i \oplus \gamma) \right) \right\}.$$

R-RS.Repudiate(R, sk, σ, m)

1. If R-RS.Verify(R, σ, m) = 0 output \perp and halt.
2. Parse R as above, $sk = ((sk_{\text{VRF}}^1, \dots, sk_{\text{VRF}}^4), vk)$, and $\sigma = (\vec{\pi}, \vec{y}, \varphi)$.
3. If $vk \notin R$ output \perp and halt.
4. Define $i^* \in [N]$ such that $vk_{i^*} = vk$.
5. For $\eta \in \{1, 2\}$: let $y'_\eta = \text{VRF.Eval}(sk_{\text{VRF}}^\eta, (R, m, \varphi))$ and let $\tau'_\eta = \text{VRF.Prove}(sk_{\text{VRF}}^\eta, (R, m, \varphi))$.
6. $\gamma \leftarrow \{0, 1\}^\nu$. (This is used as part of the ZAP witness in Step 7.)
7. For each $i \in [N]$, let $\xi_i \leftarrow \text{ZAP.Prove}_{L'}(\rho_i, (R, m, \varphi, \vec{y}, vk), (i^*, y'_1, y'_2, \perp, \perp, \tau'_1, \tau'_2, \perp, \perp, \gamma))$.
8. Output $\xi = (\xi_1, \dots, \xi_N)$.

R-RS.VerRepud(R, vk, σ, m, ξ)

1. Parse R as above. If $vk \notin R$, output 1 and halt.
2. Parse $\sigma = (\vec{\pi}, \vec{y}, \varphi)$, and $\xi = (\xi_1, \dots, \xi_N)$.
3. Output $\bigwedge_{i \in [N]} \text{ZAP.Verify}_{L'}(\rho_i, \xi_i, (R, m, \varphi, \vec{y}, vk))$.

Remark 13. As written, the size of the VRF input (R, m, φ) scales with the size of the ring R , and we have assumed that the VRF has input domain $\{0, 1\}^*$, i.e., can take variable-length inputs. When this is not the case, or when it is desirable for efficiency reasons to evaluate the VRF on a smaller input, the scheme can be straightforwardly modified by employing a collision-resistant hash function, and evaluating the VRF on the *hash of* (R, m, φ) rather than on (R, m, φ) directly. We have presented the version of the scheme without the hash function, for simplicity of exposition.

4.3 Security proof

Theorem 4.8. *Let VRF be a VRF and ZAP be a ZAP. Then R-RS is a repudiable ring signature scheme.*

Proof. Follows directly from Lemmata 4.9 (correctness), 4.12 (repudiability), 4.13 (unforgeability), and 4.14 (anonymity). \square

Lemma 4.9 (Correctness of R-RS). R-RS *satisfies correctness* (Definition 2.2).

Correctness is immediate so we omit the proof. Before presenting the proof of repudiability, we establish the following supporting lemma and corollary, which proceed according to an argument of [DN07].

Lemma 4.10. *Let \mathcal{V} be a randomized algorithm that takes ν bits of randomness and outputs one bit. Let $\beta \in \{0, 1\}$ be a bit, and let x be an input such that for some negligible ε ,*

$$\Pr \left[\mathcal{V}(1^k, x) = \beta \right] \geq 1 - \varepsilon . \quad (10)$$

Let M be a polynomial such that $M \geq (\nu + k)/\log_2(1/\varepsilon)$. (Note that the right-hand side is at most polynomial since the numerator is polynomial and the denominator is super-constant.) Then the following probability is overwhelming:

$$\Pr_{(\alpha_1, \dots, \alpha_M) \leftarrow (\{0, 1\}^\nu)^M} \left[\forall \gamma \in \{0, 1\}^\nu, \exists i \in [M] \text{ s.t. } \mathcal{V}(1^k, x; \alpha_i \oplus \gamma) = \beta \right] . \quad (11)$$

Proof. Fix any $\gamma \in \{0, 1\}^\nu$. Let $\psi_{i,\gamma} = \alpha_i \oplus \gamma$ for each $i \in [M]$. Since the α_i are distributed randomly and independently, the distribution of $(\psi_{i,\gamma})_{i \in [M]}$ is uniform over $(\{0, 1\}^\nu)^M$ even when conditioned on γ . Therefore, conditioned on any given γ ,

$$\Pr \left[\forall i \in [M], \mathcal{V}(1^k, x; \psi_{i,\gamma}) \neq \beta \right] < \varepsilon^M .$$

There are 2^ν possible values of γ , so by a union bound,

$$\Pr \left[\exists \gamma \in \{0, 1\}^\nu \text{ s.t. } \forall i \in [M], \mathcal{V}(1^k, x; \psi_{i,\gamma}) \neq \beta \right] < 2^\nu \cdot \varepsilon^M .$$

Since $M \geq (\nu + k)/\log_2(1/\varepsilon) = \log_\varepsilon(2^{-(\nu+k)})$ by assumption, the right-hand side is at most 2^{-k} , which is negligible. The lemma follows. \square

The next corollary states that the implication established in Lemma 4.10 in fact goes both ways.

Corollary 4.11. *Let \mathcal{V} be a randomized algorithm takes ν bits of randomness and outputs one bit. Let $\beta \in \{0, 1\}$ be a bit, and let ε be a negligible function. Let M be a polynomial such that $M \geq (\nu + k)/\log_2(1/\varepsilon)$. Then (10) holds if and only if (11) is overwhelming.*

Proof. Follows from applying Lemma 4.10 for both $\beta = 0$ and $\beta = 1$. \square

Remark 14. For all large enough k , Corollary 4.11 holds for any $M \geq \nu + k$. This is because if (10) holds for some negligible ε , then for all large enough k , (10) also holds for $\varepsilon = 1/2$ (or indeed, any constant ε). Substituting $\varepsilon = 1/2$ into $M \geq (\nu + k)/\log_2(1/\varepsilon)$ yields $M \geq \nu + k$. In particular, this means that a satisfactory M can be chosen without knowledge of ε .

Next, we give the proof of repudiability of R-RS.

Lemma 4.12 (Repudiability of R-RS). R-RS *is repudiable* (Definition 3.3).

Proof. Suppose, for contradiction, that R-RS is not repudiable. Then by Definition 3.3, it must be that either (3) or (4) does not hold. We consider these two possibilities in turn.

Suppose first that (3) does not hold for R-RS. Then there is a PPT $\mathcal{A}_{\text{Sign}}$ that generates a valid signature σ with respect to some ring R , so that σ is not repudiable by some honest party in the ring. That is, the following probability is non-negligible:

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{Gen}(1^k) \\ (\sigma, m, R) \leftarrow \mathcal{A}_{\text{Sign}}^{\mathcal{O}}(vk) \\ \xi \leftarrow \text{Repudiate}(R, sk, \sigma, m) \\ b \leftarrow \text{VerRepud}(R, vk, \sigma, m, \xi) \\ b' \leftarrow \text{Verify}(R, \sigma, m) \end{array} : \begin{array}{l} b = 0 \wedge b' = 1 \\ \wedge Q \cap \{(\cdot, m, R')\} = \emptyset \end{array} \right], \quad (12)$$

where Q is the set of queries to OSign .

Based on $\mathcal{A}_{\text{Sign}}$, we build an adversary \mathcal{A} for Parallel VRF Game (defined in Lemma 4.4), as follows. \mathcal{A} first invokes $\text{R-RS.Gen}(1^k)$ to obtain (vk, sk) . The vk is a tuple $(\vec{vk}_{\text{VRF}}, \rho, \vec{\alpha})$, where \vec{vk}_{VRF} can be parsed further as $(vk_{\text{VRF}}^1, \dots, vk_{\text{VRF}}^4)$.

\mathcal{A} obtains two verification keys $vk_{\text{VRF}}^{3,*}, vk_{\text{VRF}}^{4,*}$ from the VRF challenger, and replaces vk_{VRF}^3 and vk_{VRF}^4 with these keys, setting

$$vk' = ((vk_{\text{VRF}}^1, vk_{\text{VRF}}^2, vk_{\text{VRF}}^{3,*}, vk_{\text{VRF}}^{4,*}), \rho, \vec{\alpha}).$$

Let $sk_{\text{VRF}}^{3,*}, sk_{\text{VRF}}^{4,*}$ be the VRF secret keys corresponding to $vk_{\text{VRF}}^{3,*}, vk_{\text{VRF}}^{4,*}$, respectively.¹⁸

Next, \mathcal{A} runs $\mathcal{A}_{\text{Sign}}(vk')$, answering $\mathcal{A}_{\text{Sign}}$'s oracle queries as follows.

- On query (R'', m'') to OSign : \mathcal{A} runs the honest signing algorithm R-RS.Sign on input $(R'' \cup \{vk'\}, sk, m'')$, with the following modification: in step 6, instead of using sk_{VRF}^3 and sk_{VRF}^4 to generate y_3, τ_3 , and y_4, τ_4 , \mathcal{A} invokes its VRF oracle.
- On query (R'', σ'', m'') to ORpd : \mathcal{A} runs the honest repudiation algorithm R-RS.Repudiate on input $(R'' \cup \{vk'\}, sk, \sigma'', m'')$. (Note that sk_{VRF}^3 and sk_{VRF}^4 are not used in algorithm R-RS.Repudiate , so we don't need to invoke the VRF oracle here.)

Let (σ, m, R) be the output of $\mathcal{A}_{\text{Sign}}$. \mathcal{A} parses $\sigma = (\vec{\pi}, \vec{y}, \varphi)$ and $\vec{y} = (y_1, \dots, y_4)$. Then \mathcal{A} sends (R, m, φ) to the VRF challenger, receiving responses y'_3 and y'_4 . If $y_3 = y'_3$ or $y_4 = y'_4$, \mathcal{A} outputs 0. Otherwise, \mathcal{A} outputs a random bit. Let us now consider \mathcal{A} 's behavior in the two cases where the VRF challenger's bit b is equal to 0 and equal to 1.

Case $b = 0$ in Parallel VRF Game. In this case, the view of $\mathcal{A}_{\text{Sign}}$ is identical to the view in (12), so by assumption $\mathcal{A}_{\text{Sign}}$ will win the game — i.e., produce a signature that verifies but is not repudiable by an honest party — with non-negligible probability. Note that whenever $\mathcal{A}_{\text{Sign}}$ wins the game, the condition $Q \cap \{(\cdot, m, R')\} \neq \emptyset$ in (12) implies that \mathcal{A} has not previously made an oracle query on the VRF challenge message (R, m, φ) during the query phase. Let us suppose that $\mathcal{A}_{\text{Sign}}$ wins the game described in (12), and consider the implications.

By definition, if R-RS.VerRepud rejects (with non-negligible probability) on an honestly generated repudiation $\xi = (\xi_i)_{i \in [R]}$ generated with respect to vk' , then

$$\exists i \in [R] \text{ s.t. } \text{ZAP.Verify}_L(\rho_i, \xi_i, (R, m, \varphi, \vec{y}, vk')) = 0. \quad (13)$$

By the completeness of the ZAP and the complete provability of the VRF, since ξ is honestly generated with respect to vk' , the statement $\neg b'_3 \vee \neg b'_4$ then holds with overwhelming probability,

¹⁸Note that $sk_{\text{VRF}}^{3,*}, sk_{\text{VRF}}^{4,*}$ are generated by the VRF challenger and not accessible by \mathcal{A} .

where b'_3, b'_4 are as defined in Definition 4.7. Expanding the definition of b'_3, b'_4 , and again using that ξ is honestly generated, we have that

$$\begin{aligned} & \exists \eta \in \{3, 4\}, i' \in [|R|], j' \in [M], \gamma \text{ s.t.} \\ & \text{VRF.Verify}(vk_{\text{VRF}}^{\eta,*}, (R, m, \varphi), y_\eta, \text{VRF.Prove}(sk_{\text{VRF}}^{\eta,*}, (R, m, \varphi)); \alpha_{j'}^{i'} \oplus \gamma) = 0 . \end{aligned} \quad (14)$$

Since $vk' \in R$ is honestly generated by assumption, we have that the $\vec{\alpha} = (\alpha_1, \dots, \alpha_M)$ within vk' is distributed uniformly over $(\{0, 1\}^\nu)^M$. Then applying Corollary 4.11 (setting the algorithm \mathcal{V} to be VRF.Verify): (14) implies *either*

$$\begin{aligned} & \exists \eta \in \{3, 4\} \text{ and } \tau \text{ s.t.} \\ & \Pr [\text{VRF.Verify}(vk_{\text{VRF}}^{\eta,*}, (R, m, \varphi), y_\eta, \tau) = 1] \text{ is overwhelming,} \end{aligned} \quad (15)$$

or a negligible probability event occurred. By the complete and unique provability of the VRF, (15) implies that

$$y_3 = \text{VRF.Eval}(sk_{\text{VRF}}^{3,*}, (R, m, \varphi)) \text{ or } y_4 = \text{VRF.Eval}(sk_{\text{VRF}}^{4,*}, (R, m, \varphi)) . \quad (16)$$

Chaining together the implications, we conclude that (16) holds with all but negligible probability conditioned on the non-negligible-probability event of $\mathcal{A}_{\text{Sign}}$ winning the game described in (12).

Finally, by definition of Parallel VRF Game, when $b = 0$,

$$y'_3 = \text{VRF.Eval}(sk_{\text{VRF}}^{3,*}, (R, m, \varphi)) \text{ and } y'_4 = \text{VRF.Eval}(sk_{\text{VRF}}^{4,*}, (R, m, \varphi)) . \quad (17)$$

From (16) and (17): when $b = 0$, there is a non-negligible probability that

$$y_3 = y'_3 \text{ or } y_4 = y'_4 . \quad (18)$$

Recall that (18) is the trigger condition for \mathcal{A} to output 0. Therefore, when $b = 0$, \mathcal{A} outputs 0 with non-negligible probability (and outputs a random bit the rest of the time).

Case $b = 1$ in Parallel VRF Game. In this case, y'_3 and y'_4 are uniformly random and independent of the rest of the experiment, so they will be distinct from y_3 and y_4 with overwhelming probability. Consequently, in this case, with all but negligible probability \mathcal{A} outputs a random bit.

Thus, \mathcal{A} wins Parallel VRF Game with non-negligible probability. This contradicts the security of the VRF. Therefore, R-RS satisfies (3).

It remains to show that R-RS satisfies (4). The argument for this part of the proof follows a somewhat similar outline to the argument already presented. The rest of the proof is deferred to Appendix B. \square

Lemma 4.13 (Unforgeability of R-RS). *R-RS is unforgeable (in the sense of Definition 3.5).*

Proof. This proof is very similar to the second half of the proof of Lemma 4.12.

Suppose that this is not the case. Then there exists some $N = \text{poly}(k)$ and some PPT \mathcal{B} such that the following probability is non-negligible, where I and Q are the sets of queries made to the corruption and signing oracles respectively:

$$\Pr \left[\begin{array}{l} (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\ (R^*, m^*, \sigma^*) \leftarrow \mathcal{B}^{\text{OSign, ORpd, Corr}}(vk_1, \dots, vk_N) \\ b \leftarrow \text{Verify}(R^*, \sigma^*, m^*) \end{array} : \begin{array}{l} b = 1 \wedge R^* \subseteq \{vk_1, \dots, vk_N\} \setminus I \\ \wedge Q \cap \{(\cdot, m^*, R^*)\} = \emptyset \end{array} \right] . \quad (19)$$

We build an adversary \mathcal{A} to the Parallel VRF Game of Lemma 4.4. \mathcal{A} first samples

$$(vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k)$$

and parses each vk_i and sk_i as follows:

$$\begin{aligned} vk_i &= (\vec{vk}_{\text{VRF}}^i = (vk_{\text{VRF}}^{i,1}, vk_{\text{VRF}}^{i,2}, vk_{\text{VRF}}^{i,3}, vk_{\text{VRF}}^{i,4}), \rho_i, \vec{\alpha}_i) \\ sk_i &= (\vec{sk}_{\text{VRF}}^i = (sk_{\text{VRF}}^{i,1}, sk_{\text{VRF}}^{i,2}, sk_{\text{VRF}}^{i,3}, sk_{\text{VRF}}^{i,4}), vk_i) \end{aligned}$$

Then \mathcal{A} chooses a random $i^* \leftarrow [N]$, obtains a pair of verification keys $vk_{\text{VRF}}^{*,3}, vk_{\text{VRF}}^{*,4}$ from the VRF challenger, and lets

$$vk^* = (\vec{vk}_{\text{VRF}}^i = (vk_{\text{VRF}}^{i,1}, vk_{\text{VRF}}^{i,2}, vk_{\text{VRF}}^{*,3}, vk_{\text{VRF}}^{*,4}), \rho_i, \vec{\alpha}_i).$$

Let $sk_{\text{VRF}}^{*,3}, sk_{\text{VRF}}^{*,4}$ denote the VRF secret keys corresponding to $vk_{\text{VRF}}^{*,3}, vk_{\text{VRF}}^{*,4}$, respectively.

Let $vk_{i^*}^* = vk^*$ and let $vk_i^* = vk_i$ for every $i \neq i^*$. Let $R^* = \{vk_1^*, \dots, vk_N^*\}$. \mathcal{A} then runs \mathcal{B} on input R^* , answering \mathcal{B} 's oracle queries as follows.

- On query (R'', i'', m'') to **OSign**: \mathcal{A} runs the honest signing algorithm **R-RS.Sign** on input $(R'' \cup \{vk_{i''}^*\}, sk_{i''}, m'')$, with the following modification if $i = i^*$: in step 6, instead of using $sk_{\text{VRF}}^{i^*,3}$ and $sk_{\text{VRF}}^{i^*,4}$ to generate y_3, τ_3 , and y_4, τ_4 , \mathcal{A} invokes its VRF oracle.
- On query $(R'', i'', \sigma'', m'')$ to **ORpd**: \mathcal{A} runs the honest repudiation algorithm **R-RS.Repudiate** on input $(R'' \cup \{vk_{i''}^*\}, sk_{i''}, \sigma'', m'')$. (sk_{VRF}^3 and sk_{VRF}^4 are not used by **R-RS.Repudiate**, so \mathcal{A} does not need to invoke the VRF oracle here.)
- On each invocation of **Corr** on some index $i \in [N]$: if $i = i^*$ then \mathcal{A} outputs a random bit and aborts; otherwise, \mathcal{A} responds to \mathcal{B} with sk_i .

Let (R', m, σ) be the output of \mathcal{B} . \mathcal{A} parses $\sigma = (\vec{\pi}, \vec{y}, \varphi)$ and $\vec{y} = (y_1, \dots, y_4)$, and submits (R', m, φ) to the VRF challenger and then receive responses y'_3 and y'_4 . If $y'_3 = y_3$ or $y'_4 = y_4$, \mathcal{A} outputs 0. Else, \mathcal{A} outputs a random bit.

It remains to show that \mathcal{A} distinguishes with non-negligible advantage, in the parallel VRF security game, between VRF outputs and random values.

Let us consider the behavior of \mathcal{A} in the two cases where the VRF challenger's bit is equal to 0 and equal to 1.

Case $b = 0$ in Parallel VRF Game. In this case, the response that \mathcal{A} receives to the challenge (R', m) consists of VRF outputs on input (R', m) with respect to the keys $vk_{\text{VRF}}^{*,3}, vk_{\text{VRF}}^{*,4}$. In particular, whenever \mathcal{B} does not query **Corr** on i^* , the view of \mathcal{B} is identical to the view in (19). So, conditioned on \mathcal{B} not corrupting i^* , \mathcal{B} will win the game — i.e., produce a valid signature — with non-negligible probability, by assumption.

Let us consider the probability that \mathcal{B} queries **Corr** for input i^* . Recall that this event causes \mathcal{A} to abort and output a random bit. The distribution of the view (i.e., verification keys and oracle responses) of \mathcal{B} is unaffected by \mathcal{A} 's choice of i^* , until the point at which \mathcal{B} submits an oracle query to **Corr** for input i^* (if at all). The condition $R^* \subseteq \{vk_1, \dots, vk_N\} \setminus I$ in (19) ensures that if \mathcal{B} wins the game with non-negligible probability, then \mathcal{B} leaves one or more keys uncorrupted with at least that non-negligible probability. Since i^* is chosen at random by \mathcal{A} , it follows that $\Pr[i^* \notin I]$ is non-negligible.

Let E' denote the event that \mathcal{A} does not abort (i.e., $i^* \notin I$) and \mathcal{B} 's output signature verifies (i.e., $\text{R-RS.Verify}(R', \sigma, m) = 1$). We have established that E' occurs with non-negligible probability.

Then, by the same argument given from (28) to (32) in the proof of Lemma 4.12, we have that *either*

$$\begin{aligned} & \exists j^* \in [|R'|], \eta \in \{3, 4\}, \text{ and } \tau \text{ s.t.} \\ & \Pr \left[\text{VRF.Verify}(vk_{\text{VRF}}^{j^*, \eta}, (R', m, \varphi), y_\eta, \tau) = 1 \right] \text{ is overwhelming,} \end{aligned}$$

or a negligible probability event occurred. When $j^* = i^*$, this moreover implies

$$y_3 = \text{VRF.Eval}(sk_{\text{VRF}}^{*,3}, (R', m, \varphi)) \text{ or } y_4 = \text{VRF.Eval}(sk_{\text{VRF}}^{*,3}, (R', m, \varphi)) . \quad (20)$$

Let us consider the probability that $j^* = i^*$. As also observed above, the distribution of the view (i.e., verification keys and oracle responses) of \mathcal{B} is unaffected by \mathcal{A} 's choice of i^* , until the point at which \mathcal{B} submits an oracle query to Corr for input i^* (if at all). Since i^* is chosen at random by \mathcal{A} (and is thus independent of j^*), and $i^*, j^* \in [|R'|]$, $\Pr[i^* = j^*]$ must be non-negligible. Therefore, (20) holds with non-negligible probability.

Finally, by definition of the Parallel VRF Game, whenever the challenger's bit b is 0,

$$y'_3 = \text{VRF.Eval}(sk_{\text{VRF}}^{*,3}, (R', m, \varphi)) \text{ or } y'_4 = \text{VRF.Eval}(sk_{\text{VRF}}^{*,3}, (R', m, \varphi)) . \quad (21)$$

From (20) and (21) we have that with non-negligible probability,

$$y_3 = y'_3 \text{ or } y_4 = y'_4 . \quad (22)$$

Recall that (22) is the trigger condition for \mathcal{A} to output 0. We conclude that when the VRF challenger's bit $b = 0$, the trigger condition for \mathcal{A} to output 0 is met with non-negligible probability; and by construction, \mathcal{A} outputs a random bit the rest of the time (i.e., when the trigger condition is not met).

Case $b = 1$ in Parallel VRF Game. In this case, the response that \mathcal{A} receives to the challenge message (R, m, φ) consists of truly random strings instead of VRF outputs, and so $\Pr[y'_3 = y_3 \vee y'_4 = y_4]$ is negligible. Thus, \mathcal{A} outputs a random bit with overwhelming probability.

We have shown that \mathcal{A} 's output is non-negligibly different depending on the VRF challenger's bit, and so \mathcal{A} must win the Parallel VRF Game with non-negligible probability. This contradicts the security of the VRF. Therefore, R-RS is unforgeable. \square

Lemma 4.14 (Anonymity of R-RS). *R-RS satisfies adaptive anonymity against adversarially chosen keys (Definition 3.4).*

Proof sketch. The proof is a hybrid argument using the security of VRFs and ZAPs to change the values y_2 and y_4 within the signature first to truly random values, then to VRF outputs w.r.t. party i_1^* rather than i_0^* . Then the same procedure is applied to y_1 and y_3 , so that finally y_1, \dots, y_4 are all VRF outputs w.r.t. party i_1^* . Details of the proof are in Appendix C. \square

5 Claimable transformation

In this section, we give a simple black-box transformation from any ring signature to a claimable ring signature scheme. The transformation relies on one-way functions. If the original scheme was repudiable, the resulting scheme is moreover *claimable-and-repudiable*.

5.1 Building blocks

We assume familiarity with the standard notions of commitment schemes, standard signatures, and PRFs, and simply establish syntax in this subsection.¹⁹ For simplicity, we denote a commitment scheme by a single algorithm Com , and assume that the decommitment simply consists of the commitment randomness.

Definition 5.1 (Commitment scheme). A *commitment scheme* Com has the following syntax: Com takes as input a message μ and randomness r and outputs a commitment c . Sometimes we leave r implicit and write $\text{Com}(\mu)$ instead of $\text{Com}(\mu; r)$. The *decommitment* of c is the randomness r .

A commitment scheme must satisfy the following properties.

- *Hiding*: For all PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, \exists negligible ε s.t. $\forall k \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} (\mu_0, \mu_1, \mathfrak{s}) \leftarrow \mathcal{A}_1(1^k) \\ b \leftarrow \{0, 1\} \\ c \leftarrow \text{Com}(\mu_b) \\ b' \leftarrow \mathcal{A}_2(c, \mathfrak{s}) \end{array} : b' = b \right] \leq 1/2 + \varepsilon(k). \quad (23)$$

- *Binding*: For all PPT adversaries \mathcal{A} , \exists negligible ε s.t. $\forall k \in \mathbb{N}$,

$$\Pr \left[(c, \mu, r, \mu', r') \leftarrow \mathcal{A}(1^k) : \mu \neq \mu' \wedge \text{Com}(\mu; r) = c = \text{Com}(\mu'; r') \right] \leq \varepsilon(k).$$

Definition 5.2 (Standard signature (syntax)). A *signature scheme* is a triple of PPT algorithms $\Sigma = (\Sigma.\text{Gen}, \Sigma.\text{Sign}, \Sigma.\text{Verify})$ with the following syntax:

- $\Sigma.\text{Gen}(1^k)$ takes as input the security parameter k and outputs a *verification key* vk and a *signing key* sk .
- $\Sigma.\text{Sign}(sk, m)$ takes as input a signing key sk and a message m , and outputs a signature σ .
- $\Sigma.\text{Verify}(vk, \sigma, m)$ takes as input a verification key vk , a signature σ , and a message m , and outputs a single bit indicating whether or not σ is a valid signature on m w.r.t. vk .

Definition 5.3 (PRF (syntax)). A *pseudorandom function (PRF)* is a pair of algorithms $\text{PRF} = (\text{PRF}.\text{Gen}, \text{PRF}.\text{Eval})$, where:

- $\text{PRF}.\text{Gen}$ is a PPT algorithm that takes as input 1^k and outputs a *PRF key* sk_{PRF} , and
- $\text{PRF}.\text{Eval}$ is a polynomial-time deterministic algorithm that takes as input a PRF key sk_{PRF} and $x \in \{0, 1\}^*$ and outputs a string r .

For simplicity, we assume PRFs that take arbitrary-length inputs (i.e., $\{0, 1\}^*$).

5.2 The transformation

Our transformation builds on any ring signature scheme, RS , to construct a claimable ring signature scheme C-RS . The basic idea is to take a signature σ_{RS} under RS and append to it a *commitment* c to (vk, σ_{RS}) where vk is the verification key of the signer. The verification algorithm simply checks whether σ_{RS} verifies. The claim consists of a decommitment revealing that c is a commitment to (vk, σ_{RS}) . Intuitively, by the hiding property of the commitment scheme, the identity of the signer is hidden until he chooses to publish a claim.

The simple transformation just described runs into a couple of problems when examined in detail. First, what if a signer commits to $(\sigma_{\text{RS}}, vk')$ where vk' is not his own key but that of

¹⁹We refer to any standard textbook (e.g., [KL14]) for the relevant security definitions.

someone else in the ring? This ability would violate equation (6) of Definition 3.9 (claimability). To prevent such behavior, our construction actually commits to a *standard (non-ring) signature* on (vk, σ_{RS}) . The unforgeability property of standard signatures then guarantees, intuitively, that a signer cannot convincingly make a claim with respect to any verification key unless he knows a corresponding signing key.

A second hurdle encountered by the scheme thus far described is that the signer must remember the commitment randomness in order to produce a claim. It is preferable that the signer need not be stateful in between signing and claiming; and indeed, recall that Definition 3.9 formalizes this property. To resolve this, our construction derives commitment randomness from a PRF. For similar reasons, the signing randomness for the standard (non-ring) signature in our construction is also derived from a PRF.

The formal description of the transformation follows.

Construction 5.4. Our transformation C-RS is parametrized by the following:

- RS, a ring signature scheme,
- Σ , a standard signature scheme,
- Com, a commitment scheme, and
- PRF, a PRF.

For convenience, and without loss of generality, we assume that the commitment randomness of Com, the signing randomness of Σ , and the output of PRF.Eval all have the same length of ν bits.

C-RS.Gen(1^k)

1. Let $(vk_{RS}, sk_{RS}) \leftarrow \text{RS.Gen}(1^k)$.
2. Let $(vk_{\Sigma}, sk_{\Sigma}) \leftarrow \Sigma.\text{Gen}(1^k)$.
3. Let $sk_{\text{PRF}} \leftarrow \text{PRF.Gen}(1^k)$.
4. Output $vk = (vk_{RS}, vk_{\Sigma})$ and $sk = (vk, sk_{RS}, sk_{\Sigma}, sk_{\text{PRF}})$.

In the rest of the construction, we implicitly parse verification keys and signing keys of C-RS as $vk = (vk_{RS}, vk_{\Sigma})$ and $sk = (vk, sk_{RS}, sk_{\Sigma}, sk_{\text{PRF}})$ respectively. Also, for a ring

$$R = (vk_1 = (vk_{RS}^1, vk_{\Sigma}^1), \dots, vk_N = (vk_{RS}^N, vk_{\Sigma}^N)) ,$$

we write $\text{RS}(R)$ to denote $(vk_{RS}^1, \dots, vk_{RS}^N)$.

C-RS.Sign(R, sk, m)

1. Let $\sigma_{RS} \leftarrow \text{RS.Sign}(\text{RS}(R), sk_{RS}, m)$.
2. Let $r_{\Sigma} = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{RS}, 0))$.
3. Let $\sigma_{\Sigma} = \Sigma.\text{Sign}(sk_{\Sigma}, (vk, \sigma_{RS}); r_{\Sigma})$.
4. Let $r_{\text{Com}} = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{RS}, 1))$.
5. Let $c = \text{Com}((vk, \sigma_{\Sigma}); r_{\text{Com}})$.
6. Let $\sigma = (\sigma_{RS}, c)$.
7. If $\text{C-RS.VerClaim}(R, vk, \sigma, \text{C-RS.Claim}(R, sk, \sigma)) = 1$, output σ .
8. Otherwise, output (\perp, \perp) .

C-RS.Verify($R, \sigma = (\sigma_{RS}, c), m$)

1. If $\sigma_{RS} = \perp$, output 0.

2. Otherwise, output $\text{RS.Verify}(\text{RS}(R), \sigma_{\text{RS}}, m)$.

C-RS.Claim($R, sk, \sigma = (\sigma_{\text{RS}}, c)$)

1. Let $r'_{\Sigma} = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{\text{RS}}, 0))$.
2. Let $r'_{\text{Com}} = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{\text{RS}}, 1))$.
3. Let $\sigma'_{\Sigma} = \Sigma.\text{Sign}(sk_{\Sigma}, (vk, \sigma_{\text{RS}}); r'_{\Sigma})$.
4. If $c \neq \text{Com}(\sigma'_{\Sigma}, r'_{\text{Com}})$, output $\zeta = \perp$.
5. Otherwise, output $\zeta = (r'_{\text{Com}}, \sigma'_{\Sigma})$.

C-RS.VerClaim($R, vk, \sigma = (\sigma_{\text{RS}}, c), \zeta = (r'_{\text{Com}}, \sigma'_{\Sigma})$)

1. Let $c' = \text{Com}((vk, \sigma'_{\Sigma}); r'_{\text{Com}})$.
2. Output $(c = c') \wedge \Sigma.\text{Verify}(vk_{\Sigma}, \sigma'_{\Sigma}, (vk, \sigma_{\text{RS}}))$.

If RS is a *repudiable* ring signature scheme equipped with algorithms (*Repudiate*, *VerRepud*), then we additionally define C-RS.Repudiate and C-RS.VerRepud to simply run the Repudiate and VerRepud algorithms of RS, as follows.

C-RS.Repudiate($R, sk, \sigma = (\sigma_{\text{RS}}, c), m$)

1. Output $\text{RS.Repudiate}(\text{RS}(R), sk, \sigma_{\text{RS}}, m)$.

C-RS.VerRepud($R, vk, \sigma = (\sigma_{\text{RS}}, c), m, \xi$)

1. Output $\text{RS.VerRepud}(\text{RS}(R), sk, \sigma_{\text{RS}}, m, \xi)$.

Theorem 5.5. C-RS is a claimable ring signature scheme (Theorem 5.6). Moreover, if RS is a repudiable ring signature scheme, then C-RS is repudiable-and-claimable (Theorem 5.11).

Proof. Follows from Theorems 5.6 and 5.11. □

Theorem 5.6 (Claimability of C-RS). C-RS is a claimable ring signature scheme.

Proof. Follows from Lemmata 5.7–5.10, which establish the properties of correctness, claimability, unforgeability, and anonymity, respectively. □

Lemma 5.7 (Correctness of C-RS). C-RS satisfies correctness (Definition 2.2).

Correctness is immediate, so we omit the proof.

Lemma 5.8 (Claimability of C-RS). C-RS is claimable (Definition 3.9).

Proof. We show that C-RS satisfies each of the three conditions of Definition 3.9. The first condition is immediate by the correctness of the signature scheme Σ , since the use of the PRF ensures that the values $(r'_{\Sigma}, r'_{\text{Com}}, \sigma'_{\Sigma})$ computed in Claim are the same as the corresponding values computed in Sign and that the commitments c, c' match.

For the second condition, assume for contradiction that there exists some PPT malicious claiming algorithm $\mathcal{A}_{\text{Claim}} = (\mathcal{A}_1, \mathcal{A}_2)$ that is able to claim a signature produced by a different party.

That is, that the following probability is non-negligible:

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{C-RS.Gen}(1^k) \\ (R', m) \leftarrow \mathcal{A}_1^{\mathcal{O}}(vk) \\ \sigma \leftarrow \text{C-RS.Sign}(R' \cup \{vk\}, sk, m) \\ (\zeta, vk') \leftarrow \mathcal{A}_2^{\mathcal{O}}(R' \cup \{vk\}, \sigma) \\ b \leftarrow \text{C-RS.VerClaim}(R' \cup \{vk\}, vk', \sigma, \zeta) \\ b' \leftarrow \text{C-RS.Verify}(R' \cup \{vk\}, \sigma, m) \end{array} : \begin{array}{l} b = 1 \wedge b' = 1 \\ \wedge vk' \neq vk \end{array} \right], \quad (24)$$

where $\mathcal{O} = \{\text{OSign}, \text{OClaim}_{vk,sk}\}$. We will produce an adversary \mathcal{B} that breaks the binding property of the commitment scheme. Let \mathcal{B} first run the experiment in Equation 24, invoking the malicious claiming algorithm $\mathcal{A}_{\text{Claim}}$ and using its knowledge of the secret key sk to answer the oracle queries of $\mathcal{A}_{\text{Claim}}$. Let \mathcal{B} then compute $\zeta' = \text{C-RS.Claim}(R' \cup \{vk\}, sk, \sigma)$ and $b'' = \text{C-RS.VerClaim}(R' \cup \{vk\}, vk, \sigma, \zeta')$. Since the signature verifies correctly with non-negligible probability, the check on line 7 of C-RS.Sign must pass, and so we must have that $b'' = 1$ whenever $b' = 1$. Consequently with non-negligible probability we have that $b = b' = b'' = 1$ and $vk \neq vk'$. Conditioning on this event, we have that $\sigma = (\sigma_{\text{RS}}, c)$, $\zeta = (r_{\text{Com}}, \sigma_{\Sigma})$, and $\zeta' = (r'_{\text{Com}}, \sigma'_{\Sigma})$ are such that $c = \text{Com}((vk, \sigma_{\Sigma}), r_{\text{Com}}) = \text{Com}((vk', \sigma'_{\Sigma}), r'_{\text{Com}})$. But $vk \neq vk'$, so with non-negligible probability \mathcal{B} generates two different openings to the same commitment, breaking the binding property of the commitment. This concludes the proof of the second property of Definition 3.9.

For the third condition, assume for contradiction that there exists some PPT malicious signing-and-claiming algorithm $\mathcal{A}_{\text{S\&C}}$ that produces a signature and claims it on behalf of a different party. That is, assume that the following probability is non-negligible:

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{C-RS.Gen}(1^k) \\ (R', m, \sigma, \zeta) \leftarrow \mathcal{A}_{\text{S\&C}}^{\mathcal{O}, \text{OClaim}_{(vk,sk)}}(vk) \\ b \leftarrow \text{C-RS.VerClaim}(R' \cup \{vk\}, vk, \sigma, \zeta) \\ b' \leftarrow \text{C-RS.Verify}(R' \cup \{vk\}, \sigma, m) \end{array} : \begin{array}{l} b = 1 \wedge b' = 1 \\ \wedge Q \cap \{(\cdot, \sigma)\} = \emptyset \end{array} \right], \quad (25)$$

where $\mathcal{O} = \{\text{OSign}\}$ and Q is the set of queries made to the oracle $\text{OClaim}_{(vk,sk)}$. We will construct an adversary \mathcal{B} that breaks the unforgeability property of the signature scheme Σ . \mathcal{B} first requests a verification key vk_{Σ}^* from the challenger for Σ . It samples keys $(vk, sk) \leftarrow \text{C-RS.Gen}(1^k)$ as in the beginning of the experiment in Equation 25, but replaces vk_{Σ} with vk_{Σ}^* and sk_{Σ} with \perp in vk and sk , respectively. It then invokes the adversary $\mathcal{A}_{\text{S\&C}}$ on inputs $(1^k, vk)$ to produce values (R', m, σ, ζ) , using the challenger for the signature scheme Σ to respond to the oracle queries of $\mathcal{A}_{\text{S\&C}}$ as follows:

- When $\mathcal{A}_{\text{S\&C}}$ queries oracle OSign on input (m, R) , algorithm \mathcal{B} first computes σ_{RS} as in C-RS.Sign . It then invokes the challenger for Σ on message (vk, σ_{RS}) , receiving signature σ_{Σ} . (If the challenger has already been queried on this message (vk, σ_{RS}) in a previous invocation of OSign , then instead of querying the challenger again, use the same value σ_{Σ} sent by the challenger in the previous invocation.) It then proceeds as in steps 4–8 of algorithm C-RS.Sign computing value r_{Com} and commitment c , testing whether C-RS.VerClaim succeeds (where in the inner invocation of C-RS.Claim we set σ'_{Σ} to be the challenger-produced signature σ_{Σ} instead of running $\Sigma.\text{Sign}$), and returning $\sigma = (\sigma_{\text{RS}}, c)$.
- When $\mathcal{A}_{\text{S\&C}}$ queries oracle $\text{OClaim}_{(vk,sk)}$ on input (R, σ) , algorithm \mathcal{B} first tests whether σ was an output returned by a previous invocation of the oracle OSign . If not, it immediately returns

\perp . If it is, then let σ_Σ be the signature returned by the challenger for Σ on that invocation of OSign . Compute value r_{Com} as in algorithm C-RS.Claim , and output $\zeta = (r_{\text{Com}}, \sigma_\Sigma)$.

Finally, algorithm \mathcal{B} parses $\zeta = (r'_{\text{Com}}, \sigma'_\Sigma)$ and outputs σ'_Σ as its forgery.

We now outline a hybrid argument to show that this adversary \mathcal{B} breaks the unforgeability property of the signature scheme.

Hybrid 1. The experiment with adversary \mathcal{B} as just described.

Hybrid 2. Instead of the signature scheme challenger using actual randomness to generate the signatures in the oracle queries to OSign , use pseudorandomness obtained by a PRF invocation $r_\Sigma = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{\text{RS}}, 0))$

Hybrid 3. Instead of responding to both types of oracle queries as above, use actual invocations of C-RS.Sign and C-RS.Claim , using the secret key sk_Σ known to the challenger for the unforgeability game.

Indistinguishability of Hybrids 1 and 2 follows from the pseudorandomness of the PRF. For Hybrids 2 and 3, note first that the two experiments behave identically on invocations of the oracle OSign . It remains to consider invocations of the oracle OClaim . By the binding property of the commitment scheme and the pseudorandomness of the PRF, with all but negligible probability adversary $\mathcal{A}_{\text{S\&C}}$ in Hybrid 3 will be unable to find an input to the oracle OClaim that does not yield output \perp , except for inputs that were previously produced as output to the oracle OSign . Consequently, except with negligible probability, the oracle OClaim will behave identically in Hybrids 2 and 3. But Hybrid 3 is exactly the experiment in Equation 25, so with non-negligible probability adversary $\mathcal{A}_{\text{S\&C}}$ in this experiment produces a signature $\sigma = (\sigma_{\text{RS}}, c)$ and claim $\zeta = (r'_{\text{Com}}, \sigma'_\Sigma)$ such that σ was not a query to oracle $\text{OClaim}_{(vk, sk)}$ and $\Sigma.\text{Verify}(vk_\Sigma, \sigma'_\Sigma, (vk, \sigma_{\text{RS}}))$ outputs 1, i.e. σ'_Σ verifies as a valid signature of the message (vk, σ_{RS}) under key vk_Σ . By the hybrid argument, it follows that with non-negligible probability, \mathcal{B} successfully produces a valid signature for message (vk, σ_{RS}) .

It remains to argue that this signature is a valid forgery, i.e. that its message is distinct from each of the signatures produced by the challenger. To achieve this, we will make a small modification to the adversary \mathcal{B} ; call the new adversary \mathcal{B}' . Let $L = L(k)$ be a (polynomial) upper bound on the number of queries made by \mathcal{B} to the oracle OSign , and choose at random an index $i^* \in [L]$. On the i^* th query to OSign , rather than invoking the ring signature challenger to obtain σ_Σ and computing commitment c , instead choose a random string σ^* and compute the commitment with respect to that. By the hiding property of the commitment scheme, as long as the oracle OClaim is not queried on this signature, the output of the modified adversary \mathcal{B}' is indistinguishable from that of \mathcal{B} . We already have by assumption that with non-negligible probability $\mathcal{A}_{\text{S\&C}}$ wins the experiment in Equation 25, but moreover we have with non-negligible probability that in addition either the signature σ produced by $\mathcal{A}_{\text{S\&C}}$ was never the output of oracle OSign (and that i^* is greater than the number of queries made to oracle OSign) or that it was the output of the i^* th query to OSign (and no earlier query). In the latter event, since $\mathcal{A}_{\text{S\&C}}$ wins the experiment in Equation 25, this signature cannot have been a query to oracle OClaim , and so in either case, the behavior of \mathcal{B}' is indistinguishable from \mathcal{B} . But in either case, the signature σ in the experiment with \mathcal{B}' was not produced by invoking the challenger to the signature scheme, so it follows that (except with negligible probability), the signature σ'_Σ from the claim ζ was not produced by the signature scheme challenger. Putting everything together, we have that with non-negligible probability, the

adversary $\mathcal{A}_{S\&C}$ in the experiment with \mathcal{B}' produces a valid signature for a message distinct from any message signed by the signature scheme challenger. This violates the unforgeability property of the signature scheme and yields a contradiction, and so the third property of Definition 3.9 is also satisfied. \square

Lemma 5.9 (Unforgeability of C-RS). *C-RS is unforgeable (in the sense of Definition 3.11).*

Proof. The proof is by reduction to the unforgeability of RS.²⁰ Suppose, for contradiction, that there is a PPT adversary \mathcal{A} that violates unforgeability of C-RS (Definition 3.11). Then we construct another adversary \mathcal{B} that violates unforgeability of RS *without having access to a OClaim oracle*. On input (vk_1, \dots, vk_N) which are verification keys of RS, \mathcal{B} behaves as follows.

1. For each $i \in [N]$:
 - Sample $(vk_\Sigma^i, sk_\Sigma^i) \leftarrow \Sigma.\text{Gen}(1^k)$.
 - Sample $sk_{\text{PRF}}^i \leftarrow \text{PRF}.\text{Gen}(1^k)$.
 - Let $vk_i^* = (vk_i, vk_\Sigma^i)$ and $sk_i^* = (sk_i, sk_\Sigma^i, sk_{\text{PRF}}^i)$.
2. Run \mathcal{A} on input (vk_1^*, \dots, vk_N^*) , answering \mathcal{A} 's oracle queries as follows.
 - (a) For each query (i, m, R) to C-RS.OSign:
 - Query RS.OSign on (i, m, R) and receive response σ_{RS} .
 - Let $r_\Sigma = \text{PRF}.\text{Eval}(sk_{\text{PRF}}^i, (vk_i^*, \sigma_{\text{RS}}, 0))$.
 - Let $\sigma_\Sigma = \Sigma.\text{Sign}(sk_\Sigma^i, (vk_i^*, \sigma_{\text{RS}}); r_\Sigma)$.
 - Let $r_{\text{Com}} = \text{PRF}.\text{Eval}(sk_{\text{PRF}}^i, (vk_i^*, \sigma_{\text{RS}}, 1))$.
 - Let $c = \text{Com}(\sigma_\Sigma; r_{\text{Com}})$.
 - Output $\sigma = (\sigma_{\text{RS}}, c)$.
 - (b) For each query (i, R, σ) to C-RS.OClaim:
 - Parse σ as (σ_{RS}, c) .
 - Let $r'_\Sigma = \text{PRF}.\text{Eval}(sk_{\text{PRF}}^i, (vk_i^*, \sigma_{\text{RS}}, 0))$.
 - Let $r'_{\text{Com}} = \text{PRF}.\text{Eval}(sk_{\text{PRF}}^i, (vk_i^*, \sigma_{\text{RS}}, 1))$.
 - Let $\sigma'_\Sigma = \Sigma.\text{Sign}(sk_\Sigma^i, (vk_\Sigma^i, \sigma_{\text{RS}}); r'_\Sigma)$.
 - Output $\zeta = (r'_{\text{Com}}, \sigma'_\Sigma)$.
3. Upon receiving an output (R', m', σ') from \mathcal{A} : parse σ' as $(\sigma'_{\text{RS}}, c')$, define $R'' = \{vk_i : vk_i^* \in R'\}$, and output $(R'', m', \sigma'_{\text{RS}})$.

By construction of C-RS and \mathcal{B} , whenever \mathcal{A} successfully forges with respect to C-RS, \mathcal{B} successfully forges with respect to RS. Moreover, \mathcal{B} 's responses to \mathcal{A} 's oracle queries are, by construction, distributed identically to the oracle responses in the unforgeability experiment for C-RS. Therefore, \mathcal{A} 's probability of successful forgery is the same when \mathcal{B} runs \mathcal{A} in the above reduction, as in the unforgeability experiment. By our supposition, \mathcal{A} 's forging probability in the unforgeability experiment is some non-negligible ϵ , so it follows that \mathcal{A} 's forging probability in the above reduction is also ϵ , and therefore \mathcal{B} 's forging probability with respect to RS is in turn ϵ . This contradicts the unforgeability of RS. The lemma follows. \square

²⁰Note that the unforgeability guarantee we have on RS is standard unforgeability of ring signatures (Definition 2.8), which does not give the adversary a OClaim oracle. If we had the stronger guarantee that RS were unforgeable in the presence of a OClaim oracle, then the unforgeability of C-RS would follow immediately, since signatures of C-RS contain signatures of RS.

Lemma 5.10 (Anonymity of C-RS). *If RS satisfies anonymity (resp., adaptive anonymity) against adversarially chosen keys (Definition 2.5), then C-RS satisfies anonymity (resp., adaptive anonymity) against adversarially chosen keys (Definition 3.10).*

Proof. We give the proof that C-RS satisfies *adaptive* anonymity whenever RS satisfies adaptive anonymity. The non-adaptive version of the statement has a slightly simpler proof: the proof follows the same structure, but certain steps of the proof become unnecessary. In the rest of the proof, we write “anonymity” to mean “adaptive anonymity against adversarially chosen keys.”

We begin with a hybrid argument. Recall the anonymity experiment (Definition 2.3) for anonymity against adversarially chosen keys (Definition 3.10), shown below as “Hybrid 0.”

Anonymity experiment (Hybrid 0)

$$\begin{aligned}
& (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\
& ((m^*, i_0^*, i_1^*, R^*), \mathfrak{s}) \leftarrow \mathcal{A}_1^{\text{OSign, OClaim, Corr}}(vk_1, \dots, vk_N) \\
& b \leftarrow \{0, 1\} \\
& \sigma \leftarrow \text{Sign}(R^* \cup \{vk_{i_0^*}, vk_{i_1^*}\}, sk_{i_b^*}, m^*) \\
& b' \leftarrow \mathcal{A}_2^{\text{OSign, OClaim}^{(\sigma)}, \text{Corr}}(\mathfrak{s}, \sigma)
\end{aligned}$$

We now define two signing algorithms Sign_1 and Sign_2 which are slight variants of C-RS.Sign . For $\iota \in \{1, 2\}$, we define Hybrid ι to be the same as Hybrid 0 except that the invocation of Sign in the fourth line of the experiment is replaced by an invocation of Sign_ι . In the descriptions of Sign_1 and Sign_2 below, changes from the preceding hybrid are marked in blue, and steps which are entirely removed are “crossed out” and shown in red.

$\text{Sign}_1(R, sk, m)$

1. Let $\sigma_{\text{RS}} \leftarrow \text{RS.Sign}(\text{RS}(R), sk_{\text{RS}}, m)$.
2. Let $r_\Sigma = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{\text{RS}}, 0))$.
3. Let $\sigma_\Sigma = \Sigma.\text{Sign}(sk_\Sigma, (vk, \sigma_{\text{RS}}); r_\Sigma)$.
4. Let $r_{\text{Com}} \leftarrow \{0, 1\}^\nu$.
5. Let $c = \text{Com}(\sigma_\Sigma; r_{\text{Com}})$.
6. Output $\sigma = (\sigma_{\text{RS}}, c)$.

$\text{Sign}_2(R, sk, m)$

1. Let $\sigma_{\text{RS}} \leftarrow \text{RS.Sign}(\text{RS}(R), sk_{\text{RS}}, m)$.
2. ~~Let $r_\Sigma = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{\text{RS}}, 0))$.~~
3. ~~Let $\sigma_\Sigma = \Sigma.\text{Sign}(sk_\Sigma, (vk, \sigma_{\text{RS}}); r_\Sigma)$.~~
4. Let $r_{\text{Com}} \leftarrow \{0, 1\}^\nu$.
5. Let $c = \text{Com}(0; r_{\text{Com}})$.
6. Output $\sigma = (\sigma_{\text{RS}}, c)$.

Hybrid 1 is indistinguishable from Hybrid 0. This follows from PRF security as long as there are no other variables in \mathcal{A} 's view that are correlated with the PRF output in Hybrid 0, namely, $r_{\text{Com}} = \text{PRF.Eval}(sk_{\text{PRF}}, (vk, \sigma_{\text{RS}}, 1))$. Since PRF security guarantees that PRF outputs on different inputs are computationally indistinguishable from uniform and independent strings, it suffices to establish that nowhere else in the anonymity experiment is the PRF evaluated on the specific input $(vk, \sigma_{\text{RS}}, 1)$. The only PRF evaluations during the experiment are to compute the r_Σ and r_{Com} values used by the OSign oracle when responding to oracle queries. The PRF inputs used to compute r_Σ values are distinct, by construction, from those used to compute r_{Com} values (since the former end in 0 and the latter end in 1). The PRF inputs used to compute r_{Com} values for OSign queries are also, with overwhelming probability, distinct from the challenge

input $(vk, \sigma_{\text{RS}}, 1)$. Since each such PRF input is of the format $(vk', \sigma'_{\text{RS}}, 1)$ where σ'_{RS} is an honestly generated signature under RS, this follows from the following two observations.

1. The anonymity of RS implies that multiple honestly generated signatures under the *same* key pair must be *distinct* with overwhelming probability. (Otherwise, the adversary could break anonymity by querying a signature under every key in the ring on many messages, and then checking the challenge signature for equality with any of the preceding signatures.)
2. The unforgeability of RS implies that honestly generated signatures under an honestly generated key pair (vk, sk) are *distinct* from honestly generated signatures under any other — possibly adversarially generated²¹ — key pair.

Hybrid 2 is indistinguishable from Hybrid 1. This follows from the hiding property of the commitment, since all that has changed from Hybrid 1 is the value committed to by c .

The rest of the proof gives a reduction between Hybrid 2 and the anonymity of RS. Note that the anonymity guarantee we have on RS is anonymity of standard ring signatures (Definition 2.8), which does not give the adversary a OClaim oracle.

Suppose, for contradiction, that there is a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that violates anonymity of C-RS (Definition 3.10). Then we construct another adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ that violates anonymity of RS *without having access to a OClaim oracle*. On input (vk_1, \dots, vk_N) which are verification keys of RS, \mathcal{B}_1 behaves as follows.

1. For each $i \in [N]$, construct $vk_i^* = (vk_i, vk_\Sigma^i)$ and $sk_i^* = (sk_i, sk_\Sigma^i, sk_{\text{PRF}}^i)$ exactly as described in Step 1 in the proof of Lemma 5.9.
2. Run \mathcal{A} on input (vk_1^*, \dots, vk_N^*) , answering \mathcal{A} 's oracle queries exactly as described in Step 2 in the proof of Lemma 5.9.
3. Upon receiving an output $((m', i'_0, i'_1, R'), \mathfrak{s}')$ from \mathcal{A} : define $R'' = \{vk_i : vk_i^* \in R'\}$, let \mathfrak{s}'' be all of \mathcal{B} 's internal state, and output $((m', i'_0, i'_1, R''), \mathfrak{s}'')$.

Then, on input $(\mathfrak{s}'', \sigma_{\text{RS}})$, \mathcal{B}_2 behaves as follows.

1. Let $c = \text{Com}(0; r)$ for truly random r .
2. Let $\sigma = (\sigma_{\text{RS}}, c)$.
3. Run \mathcal{A}_2 to obtain $b' \leftarrow \mathcal{A}_2(\mathfrak{s}', \sigma)$.
4. Output b' .

\mathcal{A} 's view between the reduction run by \mathcal{B} is identically distributed to \mathcal{A} 's view in the experiment of Hybrid 2. Moreover, by construction, \mathcal{B}_2 's guess is correct exactly when \mathcal{A}_2 guesses b' correctly. Therefore, \mathcal{B} 's success probability in the anonymity experiment of RS is negligibly close to \mathcal{A} 's success probability in the anonymity experiment of C-RS. By supposition, the latter probability is non-negligibly greater than $1/2$. It follows that \mathcal{B} violates the anonymity of RS, which is contradiction. The lemma follows. \square

Theorem 5.11 (Repudiability-and-claimability of C-RS). *If RS is repudiable, then C-RS is a repudiable-and-claimable ring signature scheme. (Definition A.1).*

Proof sketch. Suppose RS is repudiable. We need to prove that C-RS satisfies the definitions (which are given in Appendix A) of repudiability, claimability, anonymity, and unforgeability of repudiable-and-claimable ring signature schemes.

²¹With knowledge of vk but not sk .

The proofs of anonymity and unforgeability are essentially identical to the proofs of Lemmata 5.10 and 5.9, respectively. Those proofs reduce the anonymity/unforgeability of C-RS to the anonymity/unforgeability of RS (as defined in Section 2), respectively, under the assumption that RS is a *standard* (i.e., not necessarily repudiable) ring signature scheme. The same proof structure suffices to argue that in the case that RS is a repudiable ring signature scheme, that the anonymity/unforgeability of C-RS reduces to the anonymity/unforgeability notions for *repudiable* ring signature schemes (as defined in Section 3.1), which by assumption are satisfied by RS.

It remains to prove repudiability and claimability of C-RS, according to Definitions A.2 and A.3, respectively.

Repudiability We reduce the repudiability of C-RS to the repudiability of RS. Suppose, for contradiction, that C-RS did not satisfy Definition A.2. Then the repudiability of RS could be violated by an adversary \mathcal{A} that instantiates its own commitment and signature schemes and PRF, and runs the adversary \mathcal{B} that breaks the repudiability of C-RS, while:

- augmenting each signature under RS with a corresponding commitment so that it appears indistinguishable from a signature under C-RS in \mathcal{B} 's view; and
- remembering the commitment randomness for any such commitments; and
- answering oracle calls to OClaim by producing the appropriate decommitments (using the remembered commitment randomness); and
- answering oracle calls to ORpd by passing them to its own ORpd oracle for RS.

Claimability The proof structure for claimability is very similar to that of Lemma 5.8, which argues that C-RS satisfies each of the three conditions of Definition 3.9 in turn.

- The argument for the first condition (i.e., that honest claims are validated by VerClaim) goes through unchanged.
- The argument for the second condition (i.e., that non-signers cannot successfully claim) also goes through unchanged: the adversary \mathcal{B} constructed in the proof of Lemma 5.8 has knowledge of the secret key sk , which it can use to answer ORpd queries using the honest repudiation algorithm.
- The argument for the third condition needs to be augmented with a description of how \mathcal{B} responds to oracle queries to ORpd that are made by \mathcal{A} . Just as above, \mathcal{B} answers these oracle queries by running the honest repudiation algorithm. \square

6 Unclaimable construction

In this section we show how to construct unclaimable ring signatures from lattice assumptions. The scheme is exactly the SIS-based ring signature scheme of Brakerski and Kalai [BK10], augmented with an additional algorithm ExtractRandomness.

We first give a very brief summary of necessary background on lattice trapdoors; see [GPV08] and Appendix D for details.

6.1 Lattice trapdoor sampling

Let $q \in \mathbb{N}$, $m' \in \mathbb{N}$, and $\beta \in \mathbb{Z}$ be functions of security parameter n . The (inhomogeneous, average-case) *short integer solution* ($\text{SIS}_{q,m,\beta}$) assumption states that given $A \leftarrow \mathbb{Z}_q^{n \times m'}$, $v \leftarrow \mathbb{Z}_q^n$, it is

computationally hard to find $x \in \mathbb{Z}_q^{m'}$ such that $Ax = v$ and $\|x\| \leq \beta$. For polynomial m', β and prime $q \geq \beta \cdot \omega(\sqrt{n \log n})$, the SIS problem is known to be as hard as approximating worst-case lattice problems, in particular the Shortest Independent Vectors Problem (SIVP), to within a factor of $\beta \cdot \tilde{O}(\sqrt{n})$ [MR07; GPV08].

Let $D_{\Lambda, s, c}$ denote the discrete Gaussian distribution over n -dimensional lattice Λ , centered at $c \in \mathbb{R}^n$ and with parameter s . We note the existence of the following algorithms, described in [GPV08]²²:

- There is an algorithm `TrapdoorSamp` that on input a security parameter 1^n produces a matrix $A \in \mathbb{Z}_q^n$ and a trapdoor T , where A is statistically close to uniform and T is a short basis for the lattice $\Lambda^\perp(A)$.
- There is an algorithm `SampleDist` sampling from the discrete Gaussian distribution $D_{\mathbb{Z}^{m'}, s, 0}$.
- There is an algorithm `SampleCond` that on input a matrix A , trapdoor T , parameter s and vector u , produces a sample x distributed statistically close to the discrete Gaussian distribution $D_{\mathbb{Z}^{m'}, s, 0}$ conditioned on $Ax = u$. We have that $\|x\|_2 \leq s\sqrt{n}$ with probability 1.

We will also require additional algorithms that given output values of the algorithms `SampleDist` and `SampleCond`, respectively, sample randomness under which the algorithm produces the desired output.

- There is an algorithm `ExplainDist` that on input an image vector x and parameter s , samples randomness ρ that yields output x under algorithm `SampleDist`, i.e. samples from the distribution $\{\rho \mid \text{SampleDist}(s; \rho) = x\}$.
- There is an algorithm `ExplainCond` that on input matrix A , trapdoor T , parameter s , vector u and image vector x , samples randomness ρ that yields output x under algorithm `SampleCond` with inputs (A, T, s, u) , i.e. samples from the distribution $\{\rho \mid \text{SampleCond}(A, T, s, u; \rho) = x\}$.

We describe the algorithms `ExplainDist` and `ExplainCond` in Appendix D. We will use a slight modification of the `SampleCond` algorithm of [GPV08] that uses the basis randomization technique of [CHKP10]. We need the following lemma.

Lemma 6.1. *Let (A_1, T_1) and (A_2, T_2) be sampled from `TrapdoorSamp`, let $y \in \mathbb{Z}_q^n$, and let $s \geq \max(\|\tilde{T}_1\|, \|\tilde{T}_2\|) \cdot \omega(\sqrt{\log n})$, where the tilde denotes Gram-Schmidt orthogonalization. Sample vectors x_1 and x'_2 from `SampleDist`. Let $x_2 \leftarrow \text{SampleCond}(A_2, T_2, s, y - A_1 x_1)$, and let $x'_1 \leftarrow \text{SampleCond}(A_1, T_1, s, y - A_2 x'_2)$. Then the distributions $(A_1, T_1, A_2, T_2, x_1, x_2)$ and $(A_1, T_1, A_2, T_2, x'_1, x'_2)$ are statistically close.*

Intuitively, this lemma says that the sampled vectors are distributed the same independently of which of the two trapdoors was used. This follows immediately from Lemma 3.3 of [CHKP10].

6.2 The basic construction of [BK10]

We now describe the construction of [BK10].²³ Brakerski and Kalai first construct a base version of their scheme that satisfies a weaker security notion, and augment it to fully secure ring signatures in a series of steps.

²²These are given by the algorithms `TrapGen`, `SampleD` and `SampleSIS` in [GPV08].

²³The original presentation of [BK10] introduces an abstraction they call *ring trapdoor functions*, and instantiates this abstraction from bilinear group assumptions as well as the SIS assumption. We present their SIS-based construction more explicitly, without the additional layer of abstraction, in order to make the role of the SIS trapdoor more apparent.

Let the message space be $\{0, 1\}^\ell$, and let $X = \{x \in \mathbb{Z}_q^{m'} : \|x\|_2 \leq s\sqrt{m'}\}$ for some $s = \omega(\sqrt{n \log n \log q})$ be the set of “short” vectors.

The key generation algorithm samples a matrix with an SIS trapdoor, and an additional set of 2ℓ matrices, two corresponding to each bit of the message. It additionally samples a target vector y , and outputs the matrices and target vector as the verification key and the trapdoor as the signing key.

BK-RS.Gen(1^k)

1. Let $(A, T) \leftarrow \text{TrapdoorSamp}(1^k)$.
2. For $(i, b) \in [\ell] \times \{0, 1\}$, let $A_{i,b} \leftarrow \mathbb{Z}_q^{n \times m'}$.
3. Let $y \leftarrow \mathbb{Z}_q^n$.
4. Output $vk = (A, (A_{j,b})_{(j,b) \in [\ell] \times \{0,1\}}, y)$ and $sk = (vk, T)$.

The signing algorithm proceeds as follows. A target vector y is selected from the lexicographically first verification key. For each identity in the ring, short vectors are sampled for matrices corresponding to each bit of the message to be signed, as well as the additional matrix. Finally, the trapdoor in the signing key is used to obtain a short vector, which is sampled from the same distribution conditioned on having a particular product with the matrix A_{i^*} corresponding to the signer, i.e., conditioned on Equation 26 being satisfied. The signature consists of the list of short vectors for each identity and each index of the message.

BK-RS.Sign($R, sk, m; \rho$)

1. Parse $R = (vk_1, \dots, vk_N)$ and $sk = (vk, T)$.
2. For $i \in [N]$, parse $vk_i = (A_i, (A_{j,b}^{(i)})_{(j,b) \in [\ell] \times \{0,1\}}, y_i)$.
3. Let $y = y_i$, where $i \in [N]$ is the index for which vk_i is lexicographically first.
4. If $vk \notin R$, output \perp and halt.
5. Define $i^* \in [N]$ be such that $vk_{i^*} = vk$.
6. Using the trapdoor T_A for A_{i^*} , we can sample $(x_j^{(i)})_{i \in [N], j \in \{0\} \cup [\ell]}$ such that

$$\sum_{i \in [N]} A_i x_0^{(i)} + \sum_{\substack{i \in [N] \\ j \in [\ell]}} A_{j,m_j}^{(i)} x_j^{(i)} = y. \quad (26)$$

That is, for $(i, j) \in [N] \times \{0\} \cup [\ell]$ other than the pair $(i^*, 0)$, we invoke algorithm **SampleDist** to sample $x_j^{(i)} \in \mathcal{X}$ independently from the discrete Gaussian distribution \mathcal{X} . Finally, we invoke algorithm **SampleCond** use the trapdoor T for A_{i^*} to sample $x_0^{(i^*)}$ from a distribution statistically close to the distribution \mathcal{X} conditioned on Equation 26 being satisfied.

7. Output $\sigma = (x_j^{(i)})_{i \in [N], j \in \{0\} \cup [\ell]}$.

The verification procedure simply checks that each vector in the signature has short entries and that Equation 26 is satisfied.

BK-RS.Verify(R, σ, m)

1. Parse $R = (vk_1, \dots, vk_N)$.
2. For $i \in [N]$, parse $vk_i = (A_i, (A_{j,b}^{(i)})_{(j,b) \in [\ell] \times \{0,1\}}, y_i)$.

3. Parse $\sigma = (x_j^{(i)})_{i \in [N], j \in \{0\} \cup [\ell]}$.
4. For each $x_j^{(i)}$ for $i \in [N], j \in \{0\} \cup [\ell]$, if $x_j^{(i)} \notin X$ then immediately reject.
5. Let $y = y_i$, where $i \in [N]$ is the index for which A_{i^*} is lexicographically first.
6. Accept if Equation 26 above is satisfied, and otherwise reject.

Thus far we have simply described the basic ring signature scheme of [BK10]. We augment this scheme by providing an ExtractRandomness algorithm. In order to do so, we must produce “explaining randomness” that maps to the desired output vector under the algorithms SampleDist and SampleCond. We do this using algorithms ExplainDist and ExplainCond, as described in Appendix D.

BK-RS.ExtractRandomness(R, sk, σ, m)

1. Parse $R = (vk_1, \dots, vk_N)$ and $sk = (vk, T)$.
2. For $i \in [N]$, parse $vk_i = (A_i, (A_{j,b}^{(i)})_{(j,b) \in [\ell] \times \{0,1\}}, y_i)$.
3. Parse $\sigma = (x_j^{(i)})_{i \in [N], j \in \{0\} \cup [\ell]}$.
4. If $vk \notin R$, output \perp and halt.
5. Define $i^* \in [N]$ be such that $vk_{i^*} = vk$.
6. For $(i, j) \in [N] \times \{0\} \cup [\ell]$ other than the pair $(i^*, 0)$, invoke algorithm ExplainDist to sample random coins $\rho_j^{(i)}$ that produce output $x_j^{(i)}$ under the discrete Gaussian sampling algorithm.
7. Invoke algorithm ExplainCond to sample random coins $\rho_0^{(i^*)}$ that produce output $x_0^{(i^*)}$ under the conditional random sampling algorithm using trapdoor T .
8. Output $(\rho_j^{(i)})$.

Theorem 6.2. *Under the $\text{SIS}_{q,m',\beta}$ assumption, BK-RS is a unclaimable ring signature scheme satisfying a weak notion of unforgeability in which the challenge is sampled at random at the beginning of the experiment.*

Proof sketch. Completeness, anonymity and unforgeability are proven in [BK10]. It remains to show that the scheme is unclaimable. Consider the experiment described in Definition 3.12. The components of the signature corresponding to matrices with no trapdoor are distributed identically on the two sides of the experiment. By the correctness of algorithm ExplainDist, the components of ρ_1 and ρ_2 corresponding to identities other than vk_1, vk_2 are distributed statistically close, jointly with S and the corresponding components of the signature. It remains to consider the portions of the signature corresponding to identities vk_1 and vk_2 . But Lemma 6.1 implies that the distribution of vectors $(x_0^{(1)}, x_0^{(2)})$ is statistically close, regardless of which trapdoor was used to sample. By the correctness of algorithm ExplainCond, the corresponding components of ρ_1 and ρ_2 are also statistically close, even conditioned on the other values in the experiment. The conclusion follows. \square

6.3 Unclaimability for the full ring signature scheme of [BK10]

The ring signature scheme just described satisfies a weak notion of unforgeability, in which the message on which a signature must be forged is sampled at random by the challenger and sent to the forger in the beginning of the experiment. To achieve full unforgeability, [BK10] proceed

through a sequence of four reductions to construct schemes satisfying successively stronger notions of unforgeability. We now provide a brief overview of these reductions and describe how to modify the `ExtractRandomness` algorithm for each scheme.

The first modified scheme appends a description of the ring to the message to be signed. This only affects that message, so the `ExtractRandomness` algorithm is unchanged and is simply invoked on a different message.

The second modification is the most complicated, and introduces a variant of chameleon hash functions. A chameleon hash function h is sampled during `Gen` and is included as part of the verification key vk . During the `Sign` algorithm, randomness r is sampled from some distribution,²⁴ and a value $y = h(m, r)$ is computed, where m is the message to be signed and h is the hash function corresponding to the lexicographically first identity in the ring. The previous signature scheme is invoked on message y , and the signature is augmented to include the randomness r as well. Observe that the only randomness to explain is the choice of r and the randomness used in the invocation of the previous signature scheme. Consequently the only modification to `ExtractRandomness` is that it now must also provide random coins resulting in a particular choice of the vector r , which is straightforward.

The third modification simply computes a signature under the previous scheme of every prefix of the message to be signed, and outputs a list of these $|m|$ signatures as its signature. We can invoke the previous `ExtractRandomness` algorithm for the previous scheme on each of these $|m|$ messages. The final modification produces a random pad α as part of the `Gen` algorithm, and computes the signature on the exclusive or of the original message with the pad corresponding to the lexicographically first identity in the ring. This is the full ring signature scheme of [BK10]. As above, this only affects the message to be signed, and so the `ExtractRandomness` algorithm is simply invoked on a different message.

Given the `ExtractRandomness` algorithm for the weakly-unforgeable ring signature scheme in the previous section, the modifications we have just described yield a `ExtractRandomness` algorithm for the fully-unforgeable ring signature scheme of [BK10]. It is not difficult to see that this scheme satisfies Definition 3.12 and is an unclaimable ring signature scheme. Consequently, we obtain the following theorem.

Theorem 6.3. *Under the $\text{SIS}_{q,m',\beta}$ assumption, the ring signature scheme of [BK10] combined with the `ExtractRandomness` algorithm described above is an unclaimable ring signature scheme.*

Acknowledgements. We thank Yael Tauman Kalai for advice on an earlier draft, Hao Lin and Mingqiang Wang for pointing out an error in an earlier version, and anonymous reviewers for their helpful comments. Both authors' research was supported by the following grants: NSF MACS (CNS-1413920), DARPA IBM (W911NF-15-C-0236), Simons Investigator award agreement dated June 5th, 2012, and the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370. Sunoo Park was additionally supported by the MIT Media Lab's Digital Currency Initiative. Adam Sealfon was additionally supported by a DOE CSGF fellowship, DARPA/NJIT Palisade 491512803, Sloan/NJIT 996698, and MIT/IBM W1771646.

²⁴The distribution over which r is generated is uniform over the set of vectors with bounded ℓ_2 norm, i.e., $\{r \in \mathbb{Z}_q^{m'} : 0 < \|r\|_2 \leq \beta\}$ for some m', β .

References

- [Ajt99] Miklós Ajtai. “Generating hard instances of the short basis problem”. In: *ICALP*. 1999.
- [BLO18] Carsten Baum, Huang Lin, and Sabine Oechsner. *Towards Practical Lattice-Based One-Time Linkable Ring Signatures*. Cryptology ePrint Archive 2018/107. 2018.
- [BKM09] Adam Bender, Jonathan Katz, and Ruggero Morselli. “Ring Signatures: Stronger Definitions, and Constructions without Random Oracles”. In: *J. Cryptology* 22.1 (2009), pp. 114–138.
- [Bit17] Nir Bitansky. “Verifiable Random Functions from Non-interactive Witness-Indistinguishable Proofs”. In: *TCC*. 2017.
- [Boo+15] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, Jens Groth, and Christophe Petit. “Short Accountable Ring Signatures Based on DDH”. In: *ESORICS*. 2015.
- [BK10] Zvika Brakerski and Yael Tauman Kalai. “A Framework for Efficient Signatures, Ring Signatures and Identity Based Encryption in the Standard Model.” In: *IACR Cryptology ePrint Archive 2010/086* (2010).
- [CDNO97] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. “Deniable encryption”. In: *CRYPTO*. 1997.
- [CPP18] Ran Canetti, Sunoo Park, and Oxana Poburinnaya. “Fully Bideniable Interactive Encryption”. In: *IACR Cryptology ePrint Archive 2018* (2018), p. 1244. URL: <https://eprint.iacr.org/2018/1244>.
- [CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. “Bonsai Trees, or How to Delegate a Lattice Basis”. In: *EUROCRYPT*. 2010.
- [CH91] David Chaum and Eugène van Heyst. “Group Signatures”. In: *EUROCRYPT*. 1991.
- [DN07] Cynthia Dwork and Moni Naor. “Zaps and Their Applications”. In: *SIAM J. Comput.* 36.6 (2007).
- [FS07] Eiichiro Fujisaki and Koutarou Suzuki. “Traceable Ring Signature”. In: *PKC*. 2007.
- [GPV07] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. “Trapdoors for Hard Lattices and New Cryptographic Constructions”. In: *IACR Cryptology ePrint Archive 2007/432* (2007).
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions”. In: *STOC*. 2008.
- [GO92] Shafi Goldwasser and Rafail Ostrovsky. “Invariant Signatures and Non-Interactive Zero-Knowledge Proofs are Equivalent (Extended Abstract)”. In: *CRYPTO*. 1992.
- [GHKW17] Rishab Goyal, Susan Hohenberger, Venkata Koppula, and Brent Waters. “A Generic Approach to Constructing and Proving Verifiable Random Functions”. In: *TCC*. 2017.
- [Ish+16] Ai Ishida, Keita Emura, Goichiro Hanaoka, Yusuke Sakai, and Keisuke Tanaka. “Group Signature with Deniability: How to Disavow a Signature”. In: *CANS*. 2016.
- [KL14] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014. ISBN: 9781466570269.

- [Lin19] Hao Lin. *Personal communication*. Oct. 2019.
- [LW19] Hao Lin and Mingqiang Wang. “Repudiable Ring Signature: Stronger Security and Logarithmic-Size”. In: *IACR Cryptology ePrint Archive 2019* (2019), p. 1269. URL: <https://eprint.iacr.org/2019/1269>.
- [LNWX17] San Ling, Khoa Nguyen, Huaxiong Wang, and Yanhong Xu. “Lattice-Based Group Signatures: Achieving Full Dynamicity with Ease”. In: *ACNS*. 2017.
- [LSW06] Joseph K. Liu, Willy Susilo, and Duncan S. Wong. “Ring Signature with Designated Linkability”. In: *IWSEC*. 2006.
- [LWW04] Joseph K. Liu, Victor K. Wei, and Duncan S. Wong. “Linkable Spontaneous Anonymous Group Signature for Ad Hoc Groups (Extended Abstract)”. In: *ACISP*. 2004.
- [Man00] R. Mankiewicz. *The story of mathematics*. The story of mathematics. Princeton University Press, 2000. URL: <https://books.google.com/books?id=JXxrpwAACAAJ>.
- [Mel+13] Carlos Aguilar Melchor, Slim Bettaleb, Xavier Boyen, Laurent Fousse, and Philippe Gaborit. “Adapting Lyubashevsky’s Signature Schemes to the Ring Signature Setting”. In: *AFRICACRYPT*. 2013.
- [MRV99] Silvio Micali, Michael O. Rabin, and Salil P. Vadhan. “Verifiable Random Functions”. In: *FOCS*. 1999.
- [MP12] Daniele Micciancio and Chris Peikert. “Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller”. In: *EUROCRYPT*. 2012.
- [MR07] Daniele Micciancio and Oded Regev. “Worst-Case to Average-Case Reductions Based on Gaussian Measures”. In: *SIAM J. Comput.* 37.1 (2007), pp. 267–302.
- [Mon] Monero. *Monero: Private Digital Currency*. <https://www.getmonero.org>.
- [Ngu05] Lan Nguyen. “Accumulators from Bilinear Pairings and Applications”. In: *Topics in Cryptology - CT-RSA*. 2005.
- [RST01] Ronald L. Rivest, Adi Shamir, and Yael Tauman. “How to Leak a Secret”. In: *ASIACRYPT*. 2001.
- [SW14] Amit Sahai and Brent Waters. “How to use indistinguishability obfuscation: deniable encryption, and more”. In: *STOC*. 2014, pp. 475–484.
- [SP14] Vishal Saraswat and Sumit Kumar Pandey. “How to Leak a Secret and Reap the Rewards Too”. In: *Progress in Cryptology - LATINCRYPT 2014 - Third International Conference on Cryptology and Information Security in Latin America, Florianópolis, Brazil, September 17-19, 2014, Revised Selected Papers*. Ed. by Diego F. Aranha and Alfred Menezes. Vol. 8895. Lecture Notes in Computer Science. Springer, 2014, pp. 348–367. DOI: [10.1007/978-3-319-16295-9_19](https://doi.org/10.1007/978-3-319-16295-9_19). URL: https://doi.org/10.1007/978-3-319-16295-9_19.
- [SS10] Sven Schäge and Jörg Schwenk. “A CDH-Based Ring Signature Scheme with Short Signatures and Public Keys”. In: *FC*. 2010.
- [XY04] Shouhuai Xu and Moti Yung. “Accountable Ring Signatures: A Smart Card Approach”. In: *CARDIS*. 2004, pp. 271–286.

A Definitions for repudiable-and-claimable ring signatures

Definition A.1 (Repudiable-and-claimable ring signature). A *repudiable-and-claimable ring signature scheme* is a ring signature scheme with an additional quadruple of algorithms

$$(\text{Repudiate}, \text{VerRepud}, \text{Claim}, \text{VerClaim}) ,$$

satisfying the five properties of *correctness* (Definition 2.2), *repudiability* (Definition A.2) *claimability* (Definition A.3), *anonymity* (Definition A.4), and *unforgeability* (Definition A.5).

The syntax of $(\text{Repudiate}, \text{VerRepud})$ and $(\text{Claim}, \text{VerClaim})$ are as defined in Definitions 3.3 and 3.9 respectively.

Definition A.2 (Repudiability of repudiable-and-claimable ring signatures). A ring signature scheme $\Sigma = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies *repudiability* if equipped with algorithms

$$(\text{Repudiate}, \text{VerRepud}, \text{Claim}, \text{VerClaim})$$

such that for any (possibly adversarial) PPT signing algorithm $\mathcal{A}_{\text{Sign}}$, there exists a negligible function ε such that (3) and (4) (from Definition 3.3) are satisfied when $\mathcal{O} = \{\text{OSign}, \text{ORpd}, \text{OClaim}\}$.

Definition A.3 (Claimability of repudiable-and-claimable ring signatures). A ring signature scheme $\Sigma = (\text{Gen}, \text{Sign}, \text{Verify})$ satisfies *claimability* if equipped with algorithms

$$(\text{Repudiate}, \text{VerRepud}, \text{Claim}, \text{VerClaim})$$

such that conditions 1, 2, and 3 of Definition 3.9 hold when $\mathcal{O} = \{\text{OSign}, \text{ORpd}\}$.

Definition A.4 (Anonymity of repudiable-and-claimable ring signatures). A repudiable-and-claimable ring signature scheme

$$(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Repudiate}, \text{VerRepud}, \text{Claim}, \text{VerClaim}))$$

satisfies anonymity against

$$\{\text{adversarially chosen keys, attribution attacks, full key exposure}\}$$

if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $(\{\text{OSign}, \text{ORpd}, \text{OClaim}\}, \emptyset, \alpha)$ -anonymous (Definition 2.3) for, respectively,

$$\alpha \in \{2, 1, 0\} .$$

Moreover, the claimable ring signature satisfies the *adaptive* variants of the above anonymity definitions if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $(\{\text{OSign}, \text{ORpd}, \text{OClaim}\}, \{\text{OSign}, \text{ORpd}^{(R^*, m^*)}, \text{OClaim}\}, \alpha)$ -anonymous for $\alpha \in \{2, 1, 0\}$ respectively.

Definition A.5 (Unforgeability of repudiable-and-claimable ring signatures). A repudiable-and-claimable ring signature scheme

$$(\text{Gen}, \text{Sign}, \text{Verify}, (\text{Repudiate}, \text{VerRepud}, \text{Claim}, \text{VerClaim}))$$

is *unforgeable* if $(\text{Gen}, \text{Sign}, \text{Verify})$ is $\{\text{ORpd}, \text{OClaim}\}$ -unforgeable (Definition 2.7).

B Completing the proof of Lemma 4.12

Proof (continued). It remains to show that R-RS satisfies (4). Suppose not, for the sake of contradiction. Then there is a PPT algorithm $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$ such that the following probability is non-negligible:

$$\Pr \left[\begin{array}{l} (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\ (\sigma, R', m, \{\xi_{vk}\}_{vk \in R' \setminus R}) \leftarrow \mathcal{A}_{\mathcal{S}\&\mathcal{R}}^{\mathcal{O}}(R) \\ \forall vk \in R' \setminus R, b_{vk} \leftarrow \text{VerRepud}(R', vk, \sigma, m, \xi_{vk}) \\ b' \leftarrow \text{Verify}(R', \sigma, m) \end{array} \quad : \quad \begin{array}{l} R' \cap R \neq \emptyset \wedge \bigwedge_{vk \in R' \setminus R} b_{vk} = 1 \\ \wedge b' = 1 \wedge Q \cap \{(\cdot, m, R')\} = \emptyset \end{array} \right] \quad (27)$$

where $R = \{vk_1, \dots, vk_N\}$, $\mathcal{O} = \{\text{OSign}, \text{ORpd}\}$, and Q is the set of queries made to oracle OSign .

Based on $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$, we construct another adversary \mathcal{A}' to Parallel VRF Game for $2N$ keys, as follows. \mathcal{A}' first samples

$$(vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k)$$

and parses each vk_i as

$$vk_i = (\vec{vk}_{\text{VRF}}^i = (vk_{\text{VRF}}^{i,1}, vk_{\text{VRF}}^{i,2}, vk_{\text{VRF}}^{i,3}, vk_{\text{VRF}}^{i,4}), \rho_i, \vec{\alpha}_i).$$

Next, \mathcal{A}' obtains $2N$ verification keys from the VRF challenger. For notational convenience in the rest of the proof, let these $2N$ keys be denoted by $\{vk_{\text{VRF}}^{i,3,*}, vk_{\text{VRF}}^{i,4,*}\}_{i \in [N]}$. Let the corresponding $2N$ secret keys be denoted by $\{sk_{\text{VRF}}^{i,3,*}, sk_{\text{VRF}}^{i,4,*}\}_{i \in [N]}$.²⁵

Define vk_i^* as

$$vk_i^* = (\vec{vk}_{\text{VRF}}^i = (vk_{\text{VRF}}^{i,1}, vk_{\text{VRF}}^{i,2}, vk_{\text{VRF}}^{i,3,*}, vk_{\text{VRF}}^{i,4,*}), \rho_i, \vec{\alpha}_i).$$

That is, vk_i^* is identical to vk_i except that $vk_{\text{VRF}}^{i,3}$ and $vk_{\text{VRF}}^{i,4}$ are replaced by $vk_{\text{VRF}}^{i,3,*}$ and $vk_{\text{VRF}}^{i,4,*}$.

Let $R^* = \{vk_1^*, \dots, vk_N^*\}$. \mathcal{A}' then runs $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$ on input R^* , answering its oracle queries as follows.

- On query (R'', i'', m'') to OSign : \mathcal{A}' runs the honest signing algorithm R-RS.Sign on input $(R'' \cup \{vk_{i''}^*\}, sk_{i''}, m'')$, with the following modification: in step 6, instead of using $sk_{i''}$ to generate y_3, τ_3 , and y_4, τ_4 , \mathcal{A}' invokes the corresponding VRF oracle for keys $vk_{\text{VRF}}^{i'',3,*}, vk_{\text{VRF}}^{i'',4,*}$.
- On query $(R'', i'', \sigma'', m'')$ to ORpd : \mathcal{A}' runs the honest repudiation algorithm R-RS.Repudiate on input $(R'' \cup \{vk_{i''}^*\}, sk_{i''}, \sigma'', m'')$. (As noted above, sk_{VRF}^3 and sk_{VRF}^4 are not used in algorithm R-RS.Repudiate , so \mathcal{A}' does not need to invoke the VRF oracle here.)

Let (σ, ξ, m, R') be the output of $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$. \mathcal{A}' parses $\sigma = (\vec{\pi}, \vec{y}, \varphi)$ and $\vec{y} = (y_1, \dots, y_4)$, then submits (R', m, φ) to the VRF challenger and, for each $i \in [N]$, receives responses $y_{3,i}, y_{4,i}$. If there exists any index $i \in [N]$ such that $y_{3,i} = y_3$ or $y_{4,i} = y_4$, \mathcal{A}' outputs 0. Otherwise, \mathcal{A}' outputs a random bit. Let us now consider the behavior of \mathcal{A}' in the two cases where the VRF challenger's bit b is equal to 0 and equal to 1.

Case $b = 0$. In this case, the view of $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$ is identical to the view in (27), so by assumption $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$ wins the game described in (27) with non-negligible probability. Note that whenever $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$ wins the game, the condition $Q \cap \{(\cdot, m, R')\} \neq \emptyset$ in (27) implies that \mathcal{A} has not previously made an oracle query on the VRF challenge message (R', m, φ) during the query phase. Let us suppose that $\mathcal{A}_{\mathcal{S}\&\mathcal{R}}$ wins the game described in (27) and consider the implications.

²⁵Note that $\{sk_{\text{VRF}}^{i,3,*}, sk_{\text{VRF}}^{i,4,*}\}_{i \in [N]}$ are generated by the VRF challenger and not accessible by \mathcal{A}' .

By definition, if R-RS.Verify accepts (with non-negligible probability) on input (R', σ, m) , then

$$\forall i \in [|R'|], \text{ZAP.Verify}_L(\rho_i, \pi_i, (R', m, \varphi, \vec{y})) = 1 . \quad (28)$$

$R' \cap R \neq \emptyset$ from our assumption that $\mathcal{A}_{S\&R}$ wins the game described in (27), and therefore we have that at least one $\rho_{i'}$ corresponding to some $vk_{i'} \in R' \cap R$ is honestly generated. Thus, the soundness of the ZAP holds w.r.t. this $\rho_{i'}$, and (28) implies that $(R', m, \varphi, \vec{y}) \in L$. Moreover, by the definition of L ,

$$(R', m, \varphi, \vec{y}) \in L \Rightarrow (b_3 \vee b_4), \quad (29)$$

where b_3, b_4 are as defined in Definition 4.6. Expanding the definitions of b_3, b_4 , we have that the right-hand side of (29) implies:

$$\begin{aligned} \exists \eta \in \{3, 4\}, i^* \in [|R'|], \tau_1, \dots, \tau_4, \gamma \text{ s.t. } \forall i' \in [|R'|], j' \in [M], \\ \text{VRF.Verify}(vk_{\text{VRF}}^{i^*, \eta, *}, (R', m, \varphi), y_\eta, \tau_\eta; \alpha_{j'}^{i'} \oplus \gamma) = 1 . \end{aligned} \quad (30)$$

Then applying Corollary 4.11 (again setting the algorithm \mathcal{V} to be VRF.Verify): (30) implies *either*

$$\begin{aligned} \exists \eta \in \{3, 4\} \text{ and } \tau \text{ s.t.} \\ \Pr \left[\text{VRF.Verify}(vk_{\text{VRF}}^{i^*, \eta, *}, (R', m, \varphi), y_\eta, \tau) = 1 \right] \text{ is overwhelming,} \end{aligned} \quad (31)$$

or a negligible probability event occurred. By the complete and unique provability of the VRF, (31) implies that

$$y_3 = \text{VRF.Eval}(sk_{\text{VRF}}^{i^*, 3, *}, (R', m, \varphi)) \text{ or } y_4 = \text{VRF.Eval}(sk_{\text{VRF}}^{i^*, 4, *}, (R', m, \varphi)) . \quad (32)$$

Chaining together the implications, we obtain that (32) holds with all but negligible probability conditioned on the non-negligible-probability event of $\mathcal{A}_{S\&R}$ winning the game described in (27).

Finally, by definition of Parallel VRF Game, when $b = 0$, for all $i \in [N]$,

$$y_{3,i} = \text{VRF.Eval}(sk_{\text{VRF}}^{i, 3, *}, (R', m, \varphi)) \text{ and } y_{4,i} = \text{VRF.Eval}(sk_{\text{VRF}}^{i, 4, *}, (R', m, \varphi)) .$$

It follows that conditioned on $b = 0$, there is a non-negligible probability that

$$\exists i^* \in [N] \text{ s.t. } y_3 = y_{3,i^*} \text{ or } y_4 = y_{4,i^*} . \quad (33)$$

Recall that (33) is the trigger condition for \mathcal{A}' to output 0. Therefore, when $b = 0$, \mathcal{A}' outputs 0 with non-negligible probability (and outputs a random bit the rest of the time).

Case $b = 1$. In this case, y'_3 and y'_4 are uniformly random and independent of the rest of the experiment, so with overwhelming probability, they will be distinct from $y_{3,i}$ and $y_{4,i}$ for every $i \in [N]$. Consequently, in this case, with all but negligible probability \mathcal{A} outputs a random bit.

Thus, \mathcal{A}' wins Parallel VRF Game with non-negligible probability. This contradicts the security of the VRF. We therefore conclude that R-RS satisfies (4). The lemma follows. \square

C Anonymity of R-RS

Lemma 4.14. R-RS satisfies adaptive anonymity against adversarially chosen keys (Definition 3.4).

Proof. Suppose that this is not the case. Then there exists some $N = \text{poly}(k)$ and PPT $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that the following probability is non-negligibly greater than $1/2$, where I is the set of queries to the corruption oracle, $\mathcal{O}_1 = \{\text{OSign}, \text{ORpd}\}$, and $\mathcal{O}_2 = \{\text{OSign}, \text{ORpd}^{(R^*, m^*)}\}$:

$$\Pr \left[\begin{array}{l} (vk_1, sk_1), \dots, (vk_N, sk_N) \leftarrow \text{Gen}(1^k) \\ ((m^*, i_0^*, i_1^*, R^*), \mathfrak{s}) \leftarrow \mathcal{A}_1^{\mathcal{O}_1, \text{Corr}}(vk_1, \dots, vk_N) \\ b \leftarrow \{0, 1\} \\ \sigma \leftarrow \text{Sign}(R^* \cup \{vk_{i_0^*}, vk_{i_1^*}\}, sk_{i_b^*}, m^*) \\ b' \leftarrow \mathcal{A}_2^{\mathcal{O}_2}(\mathfrak{s}, \sigma) \end{array} : b' = b \wedge \{i_0^*, i_1^*\} \cap I = \emptyset \right]. \quad (34)$$

Note that unlike (2), the second part of the adversary in (34) does not have a corruption oracle. This is because when $\alpha = 0$ (where α is defined as in Definition 2.3), as here, all corruptions can be performed by \mathcal{A}_1 without loss of generality. We proceed via a sequence of hybrids.

Hybrid 1. The honest experiment (as described in (34)), with $b = 0$.

Hybrid 2. Identical to the above, but in the generation of signature $\sigma = (\vec{\pi}, (y_1, \dots, y_4), \varphi)$, y_2 and y_4 are generated at random while y_1 and y_3 are still generated using the VRFs for party i_0^* .

To prove that Hybrid 1 is indistinguishable from Hybrid 2, we want to rely on the security of the VRF, but we can't invoke VRF security immediately because the OSign and ORpd oracles contain the VRF secret key. In order to use VRF security and show indistinguishability, it suffices to prove that with overwhelming probability, these oracles cannot be used to learn VRF outputs (for VRFs corresponding to identities i_0^*, i_1^*) on the challenge value (R^*, m^*, φ) . Any oracle queries that reveal *other* VRF outputs can be submitted to the VRF challenger in the security reduction: that is, supposing an adversary that distinguishes Hybrids 1 and 2, we can build an adversary that wins the Parallel VRF game while making only legal VRF queries to the challenger.

The following three points complete the proof of indistinguishability, addressing each category of oracle access in turn: (1) queries to OSign, (2) queries by \mathcal{A}_1 to ORpd, and (3) queries by \mathcal{A}_2 to ORpd.

1. Each time the oracle OSign is queried it will choose a random salt, so with overwhelming probability it will also never output the same salt φ as the challenge.
2. The salt φ is not sampled until after the first part \mathcal{A}_1 of the adversary completes, so since \mathcal{A}_1 can make only polynomially many oracle queries and the salt is sampled from $\{0, 1\}^k$, \mathcal{A}_1 makes no oracle queries to ORpd with the same salt value φ as the challenge, except with negligible probability.
3. It remains to consider the oracle queries made by \mathcal{A}_2 to the punctured repudiation oracle $\text{ORpd}^{(R^*, m^*)}$. But this oracle is punctured at the challenge ring and message, so it cannot be used to query the VRF at the challenge value.

Consequently, even though the oracles contain the VRF secret key, with overwhelming probability they cannot be used to learn the VRF output on the challenge value. Thus, the residual pseudorandomness property of VRF security ensures pseudorandomness of the VRF output on the challenge value even in the presence of the oracles.

Hybrid 3. Identical to the above, but in the generation of signature $\sigma = (\vec{\pi}, (y_1, \dots, y_4), \varphi)$, y_2 and y_4 are generated using the VRFs for party i_1^* rather than i_0^* . That is, parsing $sk_{i_1^*} = ((sk_{\text{VRF}}^{i_1^*,1}, \dots, sk_{\text{VRF}}^{i_1^*,4}), vk_{i_1^*})$, we set

$$y_2 = \text{VRF.Eval}(sk_{\text{VRF}}^{i_1^*,2}, (R, m, \varphi)) \quad \text{and} \quad y_4 = \text{VRF.Eval}(sk_{\text{VRF}}^{i_1^*,4}, (R, m, \varphi)) .$$

Hybrid 3 is indistinguishable from Hybrid 2 by the security of the VRF, using an essentially identical argument to the previous pair of hybrids (i.e., arguing residual pseudorandomness for the challenge value even in the presence of oracles OSign and ORpd).

Hybrid 4. Identical to the above, but in the generation of signature σ , the ZAPs in the Sign procedure are proven with respect to the witnesses for party i_1^* rather than i_0^* . That is, parsing $sk_{i_1^*} = ((sk_{\text{VRF}}^{i_1^*,1}, \dots, sk_{\text{VRF}}^{i_1^*,4}), vk_{i_1^*})$, in the invocation of Sign to generate σ , let $\tau_2' = \text{VRF.Prove}(sk_{\text{VRF}}^{i_1^*,2}, (R, m, \varphi))$ and $\tau_4' = \text{VRF.Prove}(sk_{\text{VRF}}^{i_1^*,4}, (R, m, \varphi))$. For each $i \in [N]$, in step 7 of Sign , generate π_i by invoking

$$\text{ZAP.Prove}(\rho_i, \text{stmt}, (i_1^*, \perp, \tau_2', \perp, \tau_4', \gamma))$$

instead of using the witness for i_0^* .

By the witness indistinguishability of the ZAP, Hybrid 4 is indistinguishable from Hybrid 3.

Hybrid 5. Identical to the above, but in the generation of signature $\sigma = (\vec{\pi}, (y_1, \dots, y_4), \varphi)$, y_1 and y_3 are generated at random while y_2 and y_4 are still generated using the VRFs for party i_1^* .

Hybrid 5 is indistinguishable from Hybrid 4 by the security of the VRF, using an essentially identical argument to that used to argue indistinguishability of Hybrids 1 and 2.

Hybrid 6. Identical to the above, but in the generation of signature $\sigma = (\vec{\pi}, (y_1, \dots, y_4), \varphi)$, y_1 and y_3 (as well as y_2 and y_4) are now generated using the VRFs for party i_1^* .

Hybrid 6 is indistinguishable from Hybrid 5, by the security of the VRF, using an essentially identical argument to that used to argue indistinguishability of Hybrids 1 and 2.

Hybrid 7. The honest experiment, with $b = 1$.

By the witness indistinguishability of the ZAP, Hybrid 7 is indistinguishable from Hybrid 6. \square

D Explaining randomness of discrete Gaussian samples

Our construction requires “explaining” algorithms ExplainDist and ExplainCond . ExplainDist must, on input (x, y) , output ρ distributed according to the conditional distribution

$$\left\{ \rho \mid \text{SampleDist}(x; \rho) = y \right\} .$$

Similarly, ExplainCond must, on input (x, y) , output ρ distributed according to

$$\left\{ \rho \mid \text{SampleCond}(x; \rho) = y \right\} .$$

In this appendix, we outline how this “explaining randomness” is able to be efficiently computed according to the required distribution.

In our underlying instantiation, as in the original construction of [BK10]:

- `SampleDist` is instantiated by the function `SampleD`(B, s, c) defined in [GPV07, §4.2].²⁶
- `SampleCond` is instantiated by the function `SampleSIS`(A, T, s, u) defined in [GPV07, §5.3.2].

We first recall the basis randomization technique of [CHKP10]. In the following, tildes denote Gram-Schmidt orthogonalization.

Lemma D.1 (Lemma 3.3 in [CHKP10]). *There exists a probabilistic polynomial-time algorithm `RandBasis`(T, s) that takes as input a basis T of an m' -dimensional integer lattice and a parameter $s \geq \|\tilde{T}\| \cdot \omega(\sqrt{\log n})$, and outputs a basis T' for the same lattice, such that:*

1. *With overwhelming probability, $\|\text{RandBasis}(T, s)\| \leq s \cdot \sqrt{m'}$*
2. *For any pair of bases T_1, T_2 for the same lattice and any $s \geq \max(\|\tilde{T}_1\|, \|\tilde{T}_2\|) \cdot \omega(\sqrt{\log n})$, the outputs of `RandBasis`(T_1, s) and `RandBasis`(T_2, s) have negligible statistical distance.*

We now describe the relatively simple algorithms `SampleD` and `SampleSIS` in order to outline how one can efficiently compute randomness to explain any particular output. `SampleSIS` invokes `SampleD`, and `SampleD` in turn invokes a simpler algorithm `SampleZ` which is also defined in [GPV07]. We describe these three algorithms in the order they are listed in the preceding sentence.

Definition D.2. `SampleSIS` takes as input (A, T, s, u) where A and T are matrices, s is a Gaussian parameter, and u is a vector.

1. Let $T' \leftarrow \text{RandBasis}(T, s)$.
2. By Gaussian elimination, choose an arbitrary t such that $At = u \pmod{q}$.
3. Sample $v \leftarrow \text{SampleD}(T', s, -t)$.
4. Output $e = t + v$.

Claim D.3. There is an efficient algorithm for `ExplainCond` if there is an efficient algorithm for `ExplainDist`.

Proof. Essentially, the claim follows from the fact that Step 3 is the only randomized step in `SampleSIS`. Recall that `ExplainCond` needs to correctly sample the randomness distribution of `SampleCond` (i.e., `SampleSIS`) conditioned on a particular input and output. Given the input, t (Step 2) can be computed deterministically. Finally, given the output e and t , it remains only to explain the randomness of `SampleD` conditioned on input $(T, s, -t)$ and output $e - t$. \square

Definition D.4. `SampleD` takes as input (B, s, c) where $B = (b_1, \dots, b_n)$ is a matrix describing a lattice basis, s is a Gaussian parameter, and c is a vector.

1. Let $v_n := 0$ (the zero vector) and $c_n := c$. For $i = n, \dots, 1$:
 - (a) Let $c'_i := \langle c_i, \tilde{b}_i \rangle / \langle \tilde{b}_i, \tilde{b}_i \rangle \in \mathbb{R}$ and $s'_i = s / \|\tilde{b}_i\|_2 > 0$.
 - (b) Sample $z_i \leftarrow \text{SampleZ}(s'_i, c'_i)$.
 - (c) Let $c_{i-1} := c_i - z_i b_i$ and let $v_{i-1} := v_i + z_i b_i$.
2. Output v_0 .

Claim D.5. There is an efficient algorithm for `ExplainDist` if there is an efficient algorithm to “explain `SampleZ`” — i.e., an efficient algorithm that, on input (s, c, x) , samples from the following distribution:

$$\left\{ \rho \mid \text{SampleZ}(s, c; \rho) = x \right\}. \quad (35)$$

²⁶We cite the full version here for the detailed description of the function; the conference version is [GPV08].

Proof (sketch). Each output of `SampleDist` on a given input induces a unique set of z_i s — in other words, as noted in [GPV07], there is a bijective correspondence between the random choices of the z_i s and the output. After recovering these z_i s, it remains only to explain the randomness of `SampleZ` conditioned on inputs (s'_i, c'_i) and outputs z_i for each i . \square

Definition D.6. `SampleZ` takes input (s, c) where s is a Gaussian parameter and $c \in \mathbb{R}$. In the following, $t(n) \in \omega(\sqrt{\log(n)})$ is some fixed function, say, $t(n) = \log(n)$; $\text{Ber}(p)$ denotes the Bernoulli distribution with probability p of outputting 1; and $\rho_s(\cdot)$ is the Gaussian measure, defined as $\rho_s(x) = \exp(-\pi\|x\|_2^2/s^2)$.

1. Let $X = \mathbb{Z} \cap [c - s \cdot t(n), c + s \cdot t(n)]$ and sample uniformly $x \leftarrow X$.
2. Sample $b \leftarrow \text{Ber}(\rho_s(x - c))$.
3. If $b = 1$, output x . Else, go back to step 1.

Claim D.7. There is an efficient algorithm that, on input (s, c, x) , samples from the distribution described in (35).

Proof (sketch). `SampleZ` consists of rejection sampling based on a biased coinflip, where the bias depends on the present sample x . The randomness ρ to explain a particular output of `SampleZ` can be thought to consist of a series of “rejected” samples $x_1, \dots, x_{\ell-1}$ followed by the “accepted” sample x_ℓ (which must be equal to x). The length ℓ of this sequence follows a geometric distribution parametrized by the *expected* bias β over the entire domain X . That is,

$$\beta = \frac{1}{|X|} \sum_{x \in X} \rho_s(x - c) .$$

Once a value of ℓ is sampled from the appropriately parametrized geometric distribution, it remains only to sample the “failed attempts” $x_1, \dots, x_{\ell-1}$. This can be done by the following procedure. For each $i \in [\ell - 1]$:

1. Sample uniformly $x' \leftarrow X$.
2. Sample $b \leftarrow \text{Ber}(\rho_s(x' - c))$.
3. If $b = 0$, then set $x_i := x'$ and continue. Else, go back to Step 1.

\square