

A Formal Treatment of Distributed Key Generation, and New Constructions

Chelsea Komlo, Ian Goldberg, Douglas Stebila

Abstract. In this work, we present a novel generic construction for a Distributed Key Generation (DKG) scheme. Our generic construction relies on three modular cryptographic building blocks. The first is an aggregatable Verifiable Secret Sharing (AgVSS) scheme, the second is a Non-Interactive Key Exchange (NIKE) scheme, and the third is a secure hash function. We give formal definitions for the AgVSS and NIKE schemes, as well as concrete constructions. The utility of this generic construction is *flexibility*; i.e., any aggregatable VSS and NIKE scheme can be employed, and the construction will remain secure.

To prove the security of our generic construction, we introduce formalized game-based notions of security for DKGs, building upon existing notions in the literature. However, these prior security notions either were presented informally, omitted important requirements, or assumed certain algebraic structure of the underlying scheme. Our security notions make no such assumption of underlying algebraic structure, and explicitly consider details such as participant consistency, communication patterns, and key validity. Further, our security notions imply simulatability with respect to a target key generation scheme without rewinding. Hence, any construction that is proven secure using our security notions additionally achieves UC security.

We then present STORM, a concrete instantiation of our generic construction that is secure in the discrete logarithm setting in the random oracle model. STORM is more efficient than related DKG schemes in the literature. Because of its simple design and composability, it is a practical choice for real-world settings and standardization efforts.

Table of Contents

1	Introduction	3
1.1	Our Contributions	3
1.2	Related Work	5
2	Preliminaries	6
2.1	Threshold Schemes	7
2.2	Non-Interactive Key Exchange (NIKE)	8
	A Concrete NIKE	9
2.3	Network Model	9
3	A Formalization of Verifiable Secret Sharing	10
3.1	Verifiable Secret Sharing (VSS)	10
	A Concrete VSS	12
3.2	Aggregatable Verifiable Secret Sharing (AgVSS)	13
	A Concrete AgVSS	16
4	Distributed Key Generation (DKG)	16
4.1	Robustness	18
	Strong Robustness	18
	Weak Robustness	18
4.2	Zero Knowledge	20
4.3	Indistinguishability	22
5	A Generic DKG Construction	23
	Key Generation Protocol	25
	Secret Recovery Sub-Protocol	25
	Recovery Algorithm	26
5.1	Security of the Generic Protocol	26
6	STORM, a Concrete DKG Construction	27
7	Conclusion	27
8	Acknowledgments	27
A	Proofs for Concrete NIKE	30
B	Proofs for Feldman's VSS	30
	B.1 Secrecy	30
	B.2 Uniqueness	31
C	Proofs for the Concrete AgVSS	31
	C.1 Correctness	31
	C.2 Aggregated Secrecy	32
	C.3 Uniqueness	34
D	Proofs for Generic Construction	35
	D.1 Robustness	35
	D.2 Zero-Knowledge	37

1 Introduction

Distributed Key Generation. Distributed Key Generation (DKG) schemes underlie many multi-party cryptographic primitives in use today, such as threshold signature schemes [5,34,37] and distributed randomness beacons [36]. A DKG allows a set of n participants to cooperate to generate a keypair, such that the public key represents the entire set of participants, and each participant holds a share of the corresponding secret key. Importantly, DKGs ensure that the public key is agreed upon by all (honest) parties at the end of the protocol, but no single party knows the corresponding secret key. However, DKGs ensure that this secret key can be recovered given a threshold t number of cooperating parties, where $t \leq n$.

Despite widespread interest in DKGs, existing constructions fall short in either their complexity or composability. Most glaringly is Pedersen-DKG [40]. While Pedersen-DKG is appealing because of its simplicity and efficiency, its security to date has only been proven in the context of specific applications [5,27,30], as it cannot be proven secure in a standalone fashion [28]. More specifically, it is possible that the use of Pedersen-DKG may be secure in some settings, but not in others. As such, the risk exists that Pedersen-DKG might be misused in contexts for which it cannot be proven secure. While many alternative DKG constructions have been presented in the literature [1,3,9,21,28–30,34,37,39], these constructions either are not proven secure with respect to a target key generation scheme, incur additional protocol complexity, or require pairings.

Conversely, growing interest in using threshold schemes in practice points towards the need for DKGs that are friendly to standardization efforts and securely composable across a range of use cases. For example, the FROST Schnorr threshold signature scheme [34], whose security was demonstrated at CRYPTO '22 [5], and for which an IETF draft exists [19], is written in a way to be agnostic to any particular key generation scheme. Ideally, a candidate DKG to be standardized should be applicable not only to threshold signature schemes like FROST that are compatible with single-party EdDSA verification [10,11], but also to other applications that require secret key material to be distributed among a set of trusted parties [22].

Proving the Security of DKGs. However, challenges arise when proving the security of a DKG in a standalone manner; i.e., that the DKG is secure in *any* setting in place of its target (single-party) key generation scheme. Although existing simulatability-based notions for DKGs exist [3,6,9,28,30,37], these notions prove to be either incomplete or insufficient. For example, these simulatability-based notions assume characteristics of the underlying algebraic structure of the scheme, by expecting a unique correspondence between secret and public keys. While such an assumption holds for schemes that reduce to discrete logarithm assumptions, this assumption is not universally held. For example, in lattice-based cryptosystems, public keys have a many-to-one relation with secret keys.

Further, many simulatability-based notions assume important details such as the consistency of honest participants' state or that adversarial players do not abort. However, such assumptions are *not* guaranteed in a distributed, multi-party setting. Additionally, critical security issues can arise when adversaries are allowed to abort at key times in the protocol. For example, allowing participants to abort the protocol during the last round can be employed as a vector for a key-bias attack [28]. Under the hood, many definitions require the protocol to be robust; i.e., that the protocol will always complete successfully in spite of misbehaving parties. However, doing so increases the complexity for the protocol by requiring additional complaint and voting rounds, which are difficult to implement correctly in practice.

1.1 Our Contributions

A Composable and Generic DKG. In this work, we present a novel generic DKG construction. Our construction relies on three simple cryptographic building blocks: a secure hash function, an *Aggregatable Verifiable Secret Sharing* (AgVSS) scheme, and a *Non-Interactive Key Exchange* (NIKE) scheme. The ability to aggregate Feldman's VSS [23] has been extensively employed in prior DKG constructions [28,30,34,37,40]; we simply formalize the notion of a VSS that is aggregatable, and present new security notions to reflect this property. We additionally provide formal definitions and notions of security for a NIKE, as well as concrete AgVSS and NIKE constructions.

	GJKR [28]	PedPop [5, 34]	STORM
Composable?	Yes	No	Yes
Rounds	3	2	3
Bandwidth Efficiency	$2nt + 3n$	$nt + 3n$	$nt + 5n$
Computational Efficiency	$2nt + 2n + t$	$nt + 2n$	$nt + 2n$
Security Model	DLP	KoE/AGM	CDH
Assumed Threshold	$n \geq 2t - 1$	$n \geq t$	$n \geq 2t - 1$

Table 1. Comparison between STORM and related synchronous constructions that are secure in the random oracle model, do not require pairings, and have formal proofs of security. Composable denotes if the construction has a standalone proof of security and can be used in any context that its target (single-party) key generation algorithm is used. The number of estimated rounds do not consider sub-protocols to identify and exclude cheating parties. Bandwidth efficiency denotes the number of group and field elements sent and received by each participant throughout the protocol. Computational efficiency denotes the number of exponentiations required (i.e., scalar multiplications in an elliptic curve group). n is the total number of participants in the protocol, t is the threshold number required to recover the secret key, and the number of corrupted parties is assumed to be *at most* $t - 1$. k is security parameters for the respective scheme [21]. DLP is the Discrete Logarithm Problem; KoE is the Knowledge of Exponent model; AGM is the Algebraic Group Model; CDH is the Computational Diffie-Hellman Problem.

Why a Generic Construction? The goal for presenting a generic DKG as opposed to a monolithic construction is to enable flexibility for implementations and future standardization efforts. For example, while we define concrete AgVSS and NIKE constructions, future implementations may instead wish to employ a standardized NIKE [4] instead. Further, as NIST is currently in the process of standardizing threshold primitives [11], it may be worthwhile to leverage future VSS standards, which may possibly diverge from the concrete VSS defined in this work. Or, implementations may wish to define a construction that employs publicly verifiable randomness, or derives secret values deterministically, such as from a seed phrase. The utility of a generic DKG is that so long as the chosen building blocks are secure, any instantiation of our generic DKG using these building blocks will remain secure.

A Concrete Instantiation: STORM. We present a concrete instantiation of our generic construction that we call STORM (Synchronous, distributed, and Optimized generation of key Material). STORM employs the concrete AgVSS and NIKE constructions defined in this paper. The security of STORM reduces to the Computational Diffie-Hellman (CDH) problem in the random oracle model. STORM is securely composable with existing threshold cryptosystems such as threshold Schnorr signatures [19, 34] and similar discrete-log based systems. We compare in Table 1 the practicality and security of STORM compared to related constructions in the literature that require only standard (i.e., non-pairing based) groups.

New Notions of Security for DKGs. Our goal is to prove the security of our generic construction in such a way that (i) demonstrates the DKG can be securely used in any setting where its target key generation scheme is used, (ii) captures implicit notions such as handling adversarial aborts and honest participant consistency, and (iii) assumes nothing about the algebraic structure of the scheme. Towards this end, we introduce formal game-based security notions for DKGs that generalize and extend prior notions in the literature. Our games make explicit important details such as communication patterns between honest and adversarial players, where prior notions have omitted some details. For example, our games explicitly model the adversary as *rushing*, in that it can always query for honest player’s contributions before producing its own. We then prove the security of our generic DKG with respect to these notions.

For a DKG to be secure, we require that the DKG be strongly or weakly robust, zero-knowledge, and indistinguishable. Strong versus weak robustness captures the distinction between a protocol that always

succeeds assuming some threshold of honest participants (strong robustness), and one which allows for non-fatal aborts, so long as all participants end in a consistent state (weak robustness). This distinction allows us to define a much simpler DKG construction. Similar to prior simulation-based notions, zero-knowledge requires that the DKG is simulatable with respect to a public key. Indistinguishability requires that the output of a DKG be indistinguishable from its target (single-party) key generation scheme.

By fulfilling the notion of indistinguishability, we can guarantee that the DKG is composable in any setting where the single-party key generation scheme could be used. Prior simulation-based notions of security for DKGs assumed that zero-knowledge implied indistinguishability, but such an assumption holds only under cryptosystems where secret and public keys have a unique correspondence. This is because simulatability-based experiments require the environment to simulate its participation with respect to a public challenge, and hence assert conditions *only* on public key material.

These notions could perhaps be captured by an alternative model such as Universal Composability [12]. However, we opt for a game-based approach, so that we can explicitly and independently define each notion. However, because our notions of security capture *perfect simulatability* with respect to a target key generation algorithm without rewinding, we can guarantee that DKGs that achieve our notions of security in fact achieve UC security [35].

What We Do Not Do. Our definitions and constructions assume a strictly synchronous setting; i.e., that every participant in the protocol terminates a round before beginning the next. While asynchronous DKGs exist [2, 21, 31, 33, 39], these constructions require the tradeoff of increased complexity. Instead, we target the setting where key generation is performed infrequently or in a more controlled environment, such as generating a long-lived keypair among a set of signing parties.

Additionally, while some prior DKGs build upon a publicly verifiable secret sharing scheme (PVSS) to reduce network round complexity [9, 29], we explicitly do not, due to the tradeoff in increased protocol complexity. For similar reasons, we do not consider the large-scale setting [15].

Open Questions. Our generic construction assumes a one-to-one correspondence between public and private keys. While STORM is secure assuming the hardness of the Computational Diffie-Hellman (CDH) problem, it is possible our generic construction can be instantiated by quantum-secure primitives that define a homomorphism between the secret and public domains. For example, performing rounds in a sequential manner may allow our generic construction to be instantiated by CSIDH-based primitives [17, 18]. However, because of the difficulty of instantiating a NIKE using lattice-based primitives, it is an open question as to how our generic construction can be extended to lattice-based primitives.

Summary of Our Contributions. In this work, we present the following contributions.

- We introduce the notion of an aggregated Verifiable Secret Sharing (AgVSS) scheme, building upon its implicit use in prior literature. We give concrete AgVSS and Non-Interactive Key Exchange (NIKE) constructions, and prove their security.
- We present formal game-based definitions for the security of a DKG, that makes explicit details such as communication patterns between honest and corrupted players, expectations of participant consistency and requirements of key validity. Further, our separation of the notion of robustness into a strong and weak variant allows for a simpler and more efficient DKG construction.
- We give a generic DKG construction, and prove its security using our new notions. Our generic construction requires a secure hash function, an AgVSS, and a NIKE.
- We introduce STORM, a concrete instantiation of our generic construction. The security of STORM reduces to the Computational Diffie-Hellman (CDH) problem in the random oracle model.

1.2 Related Work

Synchronous Constructions. Pedersen-DKG [41] is a two-round protocol among n parties, such that each party plays the role of the dealer in a Feldman Verifiable Secret Sharing (VSS) protocol [40]. Gennaro, Jarecki,

Krawczyk, and Rabin [28] describe an attack against Pedersen-DKG, which we refer to as the Key-Influence attack, that considers a *rushing adversary*, who is allowed to speak. In this attack, an adversary can influence the distribution of key material output by the protocol by adaptively choosing their contribution. While the authors later demonstrate unforgeability of a Schnorr threshold signature scheme with Pedersen-DKG employed for key generation [27], this proof of security does *not* extend to other applications. Gurkan, Jovanovic, Maller, Meiklejohn, Stern, and Tomescu [30] show that Pedersen-DKG can be used securely in any setting where the cryptographic scheme is “rekeyable”. In other words, if the scheme produces some cryptographic output o_1 (such as a ciphertext or signature) that is valid with respect to a secret key sk_1 , it is possible to obtain a second value $o_2 \leftarrow f_1(o_1, sk', \alpha)$, such that o_2 is valid with respect to a related secret key $sk_2 \leftarrow f_2(sk_1, sk', \alpha)$, where α is a side input and f_1, f_2 are efficient update functions. While some constructions such as BLS signatures [8] do support rekeyability, others such as Schnorr signatures do not.

Gennaro et al. [28] present a three-round construction that is secure against the Key-Influence attack. While their description of their construction implies that a round to broadcast failure messages is optional, their construction is insecure without this round. For this reason, our construction includes a mandatory “complaint or no complaint” round (see KeyGen_3 in Figure 11). Our construction employs similar techniques as Gennaro et al. [28], but offers improved performance.

Komlo and Goldberg [34] introduce a variant of Pedersen-DKG that we refer to as PedPop, where the dealer additionally sends a Schnorr proof of possession to prevent rogue-key attacks. Bellare, Crites, Komlo, Maller, Tessaro, and Zhu [5] present an alternative proof of security for PedPop, in the context of the FROST threshold signature, but this proof does not extend to the use of PedPop in other settings.

Lindell [37] introduced a DKG that builds upon Schnorr proofs of knowledge and online-extractable zero-knowledge proofs of knowledge. However, the protocol incurs computational costs due to the requirement of online extractors [24]. Further, the security of the scheme is only generally discussed without reference to a concrete notion of security.

Many efficient pairing-based DKGs are described in the literature [1, 29, 30]. Further, concrete DKGs have been presented in the ECDSA and RSA settings [13, 20, 25, 38]. This work focuses presenting a generic construction that can be instantiated without pairings, and describes STORM, a concrete construction that can be employed in the context of EdDSA. We leave extending our generic construction to the ECDSA or RSA setting for future work.

Definitions of DKG Security. Gurkan et al. [30] present a definition of DKG security, in which the DKG is secure if its use within some larger protocol does not weaken the security of the overall scheme from the centralized (single-party) setting. We build upon this notion by defining the security of a DKG with respect to a target key generation scheme. Bacho and Loss [3] introduce the security notions of consistency and oracle-aided algebraic simultatability for a DKG. Their notion of consistency is encompassed by our notions of strong and weak robustness. Because their notion of oracle-aided simultatability assumes a discrete logarithm oracle, it is not applicable to schemes without such algebraic structure. Boneh and Shoup [9] present a generic definition for the security of a distributed key generation protocol using a simulation-based approach, with respect to a target trusted secret sharing scheme. However, their definition implicitly assumes certain algebraic structure of the underlying scheme; whereas our notions makes no such assumptions.

2 Preliminaries

Notation. Let $\lambda \in \mathbb{Z}$ represent the security parameter in unary representation. We denote the assignment of an element y to the value x as $y \leftarrow x$, and sampling an element from some set S uniformly at random as $x \xleftarrow{\$} S$. For a randomized algorithm A , we write $x \xleftarrow{\$} A()$ to indicate the random variable x that is output from the execution of A . Let \mathbb{G} be a cyclic group of prime order q , and \mathbb{Z}_q be the field of integers modulo q . Let g be a generator of \mathbb{G} . Let \mathbb{F} denote a field. We use $[n]$ to represent the set $\{1, \dots, n\}$.

Polynomial Interpolation. A polynomial of degree $t - 1$ over a field \mathbb{F} can be interpolated by t (or more) points. Let η be the list of t distinct indices $\eta \subset [n]$ corresponding to the x-coordinates $x_i \in \mathbb{F}, i \in \eta$. Then,

$L_i(x)$ is the Lagrange polynomial defined by η , of the form

$$L_i(x) = \prod_{j \in \eta; j \neq i} \frac{x - x_j}{x_i - x_j}$$

where $L_i(x) = \sum_{k=0}^{t-1} L_{k,i} x^k$ has value 1 where $x = x_i$ and 0 otherwise. Given a set of t points $(x_i, f(x_i))$, any point $f(x_\ell)$ on the degree $t - 1$ polynomial f can be determined by Lagrange interpolation:

$$f(x_\ell) = \sum_{j \in \eta} f(x_j) \cdot L_j(x_\ell).$$

The polynomial f can also be described as a Vandermonde matrix:

$$\begin{pmatrix} 1 & x_0 & x_0^2 & \dots & x_0^{n-1} \\ 1 & x_1 & x_1^2 & \dots & x_1^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & x_n^2 & \dots & x_n^{n-1} \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_n \end{pmatrix} = \begin{pmatrix} f(x_0) \\ f(x_1) \\ \vdots \\ f(x_n) \end{pmatrix}$$

Given the set of t points $(x_i, f(x_i))$, each coefficient of the polynomial f can be solved for via

$$\begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_t \end{pmatrix} = \begin{pmatrix} 1 & x_1 & x_1^2 & \dots & x_1^{t-1} \\ 1 & x_2 & x_2^2 & \dots & x_2^{t-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_t & x_t^2 & \dots & x_t^{t-1} \end{pmatrix}^{-1} \begin{pmatrix} f(x_1) \\ f(x_2) \\ \vdots \\ f(x_t) \end{pmatrix} = \begin{pmatrix} L_{0,1} & L_{0,2} & \dots & L_{0,t} \\ L_{1,1} & L_{1,2} & \dots & L_{1,t} \\ \vdots & \vdots & \ddots & \vdots \\ L_{t-1,1} & L_{t-1,2} & \dots & L_{t-1,t} \end{pmatrix} \begin{pmatrix} f(x_1) \\ f(x_2) \\ \vdots \\ f(x_t) \end{pmatrix}$$

Where $L_{k,i}$ is the k^{th} coefficient of the Lagrange polynomial $L_i(x)$.

2.1 Threshold Schemes

A (t, n) threshold scheme is a multi-party protocol for a set of n parties. The scheme splits a secret s into n shares, which are distributed one to each of the n participants. To recover s , at least t participants are required to cooperate by pooling their shares. A common example of a threshold scheme is Shamir secret sharing [42],

Security Assumptions. The security of threshold schemes depend critically on the assumed number of corrupted parties with respect to the threshold t , as follows:

- *Honest Minority.* The adversary is assumed to control up to $t - 1$ parties, and at minimum one honest party is assumed.
- *Honest Majority.* Here, the adversary is still allowed to control up to $t - 1$ parties. However, at minimum, at least t honest parties are assumed.

Remark 1 (Adaptive Security). A threshold scheme can be *statically* secure or *adaptively* secure. Static security assumes the set of parties that an adversary corrupts is fixed, whereas adaptive security allows the adversary to corrupt any $t - 1$ parties. In this work, we model the adversary as having the power to choose n , t , the set of honest parties *honest*, and the set of corrupted parties *corrupt*. This modeling can easily be translated to the adaptive setting, where the adversary is given a corruption oracle instead of receiving state for the set of players *corrupt* at the beginning of the game. Consequently, our security definitions model the “single inconsistent party” notion of adaptive security [14], where the environment can handle adaptive

$\text{Exp}_{\text{NK}, \mathcal{A}}^{\text{rec}}(\lambda)$
1 : $(\text{sk}_1, \text{pk}_1) \xleftarrow{\$} \text{NK.KeyGen}(\lambda)$
2 : $(\text{sk}_2, \text{pk}_2) \xleftarrow{\$} \text{NK.KeyGen}(\lambda)$
3 : $\psi \leftarrow \text{NK.SharedKey}(\text{sk}_1, \text{pk}_2)$
4 : $\psi' \xleftarrow{\$} \mathcal{A}(\text{pk}_1, \text{pk}_2)$
5 : return 1 if $\psi' = \psi$
6 : return 0

Fig. 1. Unrecoverability experiment defining the advantage of an adversary \mathcal{A} against a non-interactive key exchange (NIKE) NK.

corruptions by the adversary, with the exception of a single “inconsistent” honest player. Hence, our security notions and proofs imply adaptive security, but with exponential tightness loss. For small choices of n and t , this tightness loss may be acceptable. Techniques to achieve tight proofs of adaptive security in stronger security models have recently been demonstrated in concurrent work [1, 3]; we expect these techniques to similarly apply to our constructions.

2.2 Non-Interactive Key Exchange (NIKE)

Our generic scheme additionally builds upon a Non-Interactive Key Exchange (NIKE) [7, 16, 26] scheme. We additionally define a verification algorithm that checks whether a tuple is in fact a valid keypair.

Definition 1. A **Non-Interactive Key Exchange (NIKE)** NK with verification is the tuple $\text{NK} = (\text{KeyGen}, \text{Verify}, \text{SharedKey})$, where

- $\text{KeyGen}(\lambda) \rightarrow (\text{sk}, \text{pk})$: Generate a secret key sk and corresponding public key pk with respect to the security parameter λ .
- $\text{Verify}(\text{sk}, \text{pk}) \rightarrow \{0, 1\}$: Verify that (sk, pk) is a valid output from NK.KeyGen .
- $\text{SharedKey}(\text{sk}_1, \text{pk}_2) \rightarrow \psi$: Output the shared key ψ generated by the combination of one party’s secret key sk_1 and the other party’s public key pk_2 .
- For correctness, we require that, for all $(\text{sk}_1, \text{pk}_1) \xleftarrow{\$} \text{NK.KeyGen}(\lambda), (\text{sk}_2, \text{pk}_2) \xleftarrow{\$} \text{NK.KeyGen}(\lambda)$, the following conditions hold:

$$\text{NK.Verify}(\text{sk}_i, \text{pk}_i) = 1, \text{ for } i \in \{1, 2\} \text{ and}$$

$$\text{NK.SharedKey}(\text{sk}_1, \text{pk}_2) = \text{NK.SharedKey}(\text{sk}_2, \text{pk}_1)$$

In addition to the requirement that a NIKE be *session-key unrecoverable*, we additionally require that a NIKE be *binding*. We discuss both properties next.

Session-Key Recovery. A NIKE that is *session-key unrecoverable* ensures that given the public keys of two parties pk_1, pk_2 , it should be hard for an adversary to obtain the shared secret. While prior security notions for non-interactive key exchange exist in the literature [16, 26], these notions allow an adversary to act as an active participant, and assume a distinguishing game. Our constructions require only the more restricted model where the adversary learns key material but does not contribute itself, and must compute the shared secret directly. We present this restricted notion in Figure 1.

The advantage of an adversary \mathcal{A} against NK in the unrecoverability experiment as defined in Figure 1 is

$$\text{Adv}_{\text{NK}, \mathcal{A}}^{\text{rec}}(\lambda) = \Pr[\text{Exp}_{\text{NK}, \mathcal{A}}^{\text{rec}}(\lambda) = 1].$$

Definition 2. A NIKE is **session-key unrecoverable** if for all probabilistic polynomial time adversaries \mathcal{A} , the function $\text{Adv}_{\text{NK}, \mathcal{A}}^{\text{rec}}(\lambda)$ is negligible.

$\text{Exp}_{\text{NK}, \mathcal{A}}^{\text{bind}}(\lambda)$	
1 :	$(\text{sk}_1^*, \text{pk}_1^*, \text{sk}_2^*, \text{pk}_2^*) \xleftarrow{\$} \mathcal{A}(\lambda)$
2 :	return 0 if $\text{NK.Verify}(\text{sk}_1^*, \text{pk}_1^*) \neq 1$
3 :	return 0 if $\text{NK.Verify}(\text{sk}_2^*, \text{pk}_2^*) \neq 1$
4 :	return 1 if $\text{NK.SharedKey}(\text{sk}_1^*, \text{pk}_2^*) \neq \text{NK.SharedKey}(\text{sk}_2^*, \text{pk}_1^*)$
5 :	return 0

Fig. 2. Binding experiment defining the advantage of an adversary \mathcal{A} against a non-interactive key exchange (NIKE) NK. Here, the adversary can perform all NK operations directly.

Binding. We introduce an additional security property that must hold when the adversary is allowed to generate both sides of the key exchange. Informally, a NIKE is *binding* when $\text{NK.SharedKey}(\text{sk}_1, \text{pk}_2)$ must equal $\text{NK.SharedKey}(\text{sk}_2, \text{pk}_1)$, even when the adversary is allowed to generate $(\text{sk}_1, \text{pk}_1), (\text{sk}_2, \text{pk}_2)$ itself. The only restriction to the adversary is that both keypairs must be valid. We formalize this requirement in Figure 2.

The advantage of an adversary \mathcal{A} against NK in the binding experiment as defined in Figure 2 is

$$\text{Adv}_{\text{NK}, \mathcal{A}}^{\text{bind}}(\lambda) = \Pr[\text{Exp}_{\text{NK}, \mathcal{A}}^{\text{bind}}(\lambda) = 1].$$

Definition 3. A NIKE is **binding** if for all probabilistic polynomial time adversaries \mathcal{A} , the function $\text{Adv}_{\text{NK}, \mathcal{A}}^{\text{bind}}(\lambda)$ is negligible.

A Concrete NIKE. We now describe a concrete NIKE which is a non-interactive Diffie–Hellman key exchange, with additionally a verification check to ensure that $\text{pk} = g^{\text{sk}}$.

- $\text{NK.KeyGen}(\lambda) \rightarrow (\text{sk}, \text{pk})$: Sample secret key $\text{sk} \xleftarrow{\$} \mathbb{Z}_q$; derive public key $\text{pk} \in \mathbb{G}$ as $\text{pk} = g^{\text{sk}}$. Output (sk, pk) .
- $\text{NK.Verify}(\text{sk}, \text{pk}) \rightarrow \{0, 1\}$: Output 1 if $\text{pk} = g^{\text{sk}}$; otherwise, output 0.
- $\text{NK.SharedKey}(\text{sk}_1, \text{pk}_2) \rightarrow \psi$: Derive the shared key $\psi \in \mathbb{G}$ as $\psi \leftarrow \text{pk}_2^{\text{sk}_1}$. Output ψ .

We demonstrate in Appendix A that the concrete scheme is unrecoverable assuming the Computational Diffie–Hellman (CDH) problem is hard, and is unconditionally binding.

2.3 Network Model

Underlying any DKG is the requirement that participants send and receive messages over a network channel. We now describe these channels and assumptions.

Synchronicity. In this work, we assume a DKG that operates purely in the synchronous model. In other words, we assume all honest parties wait to proceed to a subsequent round until they have received the expected inputs from all other participants from a prior round.

Broadcast channels. We assume an idealized broadcast channel. We represent messages sent and received over this idealized channel as inb_i , which is the set of broadcast messages that is received as input for the i^{th} protocol round. inb_i consists of broadcast messages $\text{bmsg}_1, \dots, \text{bmsg}_n$, sent by participants $1, \dots, n$. Participant k is the sender of a broadcast message bmsg_k , and all other participants are the receivers of this message.

Broadcast channels have the following properties:

1. *Consistent*. Each participant has the same view of the message sent over the channel.
2. *Authenticated*. Players know that the message was in fact sent by the claimed sender. In practice, this requirement is often fulfilled by a PKI.
3. *Reliable Delivery*. Player i knows that the message it sent was in fact received by the intended participants.
4. *Unordered*. The channel does *not* guarantee ordering of messages.

Peer to peer channels. We assume an idealized peer-to-peer channel. We represent messages sent and received over this channel as inp_i and outp_i , inp_i is the vector of peer-to-peer messages that is received as input for protocol round i , and outp_i is the vector of peer-to-peer messages that is output at the end of network round i . inp_i and outp_i consist of peer-to-peer messages which we denote as $\text{pmsg}_{i,j}$, where here, participant i is the sender and participant j is the receiver.

Peer-to-peer channels are authenticated, reliable, and unordered, per the definitions above. Additionally, peer-to-peer channels are *confidential*; i.e., only participants i and j are allowed to know the contents of $\text{pmsg}_{i,j}$.

Remark 2 (Rushing Adversaries). A **rushing adversary** is one that “speaks last;” i.e., it waits to receive inputs from honest participants before publishing its own. Importantly, a rushing adversary may be able to choose its input *adaptively* after observing inputs from all other participants. Our security definitions in Section 4 consider such an adversary, by allowing the adversary to query for honest participants’ contributions in a round before publishing its own.

3 A Formalization of Verifiable Secret Sharing

We now more formally introduce Verifiable Secret Sharing (VSS). We then extend the notion of a VSS to one that is aggregatable.

3.1 Verifiable Secret Sharing (VSS)

Verifiable Secret Sharing (VSS) allows a dealer to share a secret s in such a way that participants can ensure that combining their shares allows for recovery of a secret that corresponds to a public commitment to s , without learning s in the process. Participants verify that their shares can correctly recover s using this commitment. While prior definitions of VSS schemes [32] define algorithms to issue and verify shares, and recover the secret, we additionally define an algorithm to derive shares in a public domain. This notion has been employed implicitly in prior DKGs [34], and we formalize this construct here. Under the hood, our definition of a VSS encompasses both a secret-sharing primitive and a polynomial commitment primitive. It is possible breaking these notions into separate primitives may be useful for future work, as is the approach taken by Abraham et al. [1].

Definition 4. A **Verifiable Secret Sharing (VSS) scheme** S is the tuple (Share, Recover, Verify, GetPub) and parameterized by a one-way map $\text{SecretToPublic} : \hat{S} \rightarrow \hat{P}$ that maps elements in the secret domain \hat{S} to the public domain \hat{P} , where

- $\text{Share}(\lambda, s, n, t) \rightarrow (\{(1, w_1), \dots, (n, w_n)\}, D)$: A probabilistic algorithm performed by a dealer that accepts as input the security parameter λ , the secret $s \in \hat{S}$ from the secret domain \hat{S} , the number of participants n , and the threshold t , where n and t are positive integers, and $n \geq t$. Outputs the list of shares $\{(1, w_1), \dots, (n, w_n)\}$ and commitment D .
- $\text{Verify}(i, w_i, D) \rightarrow \{0, 1\}$: A deterministic algorithm performed by a recipient of a share that accepts an identifier i , a share w_i , and a commitment D . Outputs 1 if the share is valid with respect to D , otherwise outputs 0.
- $\text{Recover}(t, M) \rightarrow s/\text{fail}$: A deterministic algorithm that accepts a recovery set $M = \{(j, w_j)\}_{j \in C}$, where C is a set of participant identifiers such that $|C| \geq t$. Employ M to obtain s , or output fail in the case of failure.

$\text{Exp}_{\mathcal{S}, \mathcal{A}, \text{SimShare}}^{\text{sec}}(\lambda, s)$
1: $b \xleftarrow{\$} \{0, 1\}$
2: $(n, t, \text{corrupt}, \text{st}_A) \xleftarrow{\$} \mathcal{A}(\lambda)$
3: return 0 if $t > n$ or $\text{corrupt} \not\subseteq [n]$ or $ \text{corrupt} \geq t$
4: if $b = 0$
5: $(\{(j, w_j)\}_{j \in [n]}, D) \leftarrow \text{S.Share}(\lambda, s, n, t)$
6: else // $b = 1$ case
7: $\Delta \xleftarrow{\$} \text{S.SecretToPublic}(s)$
8: $(\{(j, w_j)\}_{j \in \text{corrupt}}, D) \xleftarrow{\$} \text{SimShare}(\Delta, n, t, \text{corrupt})$
9: for $j \in \text{corrupt}$ do
10: return 1 if $\text{S.Verify}(j, w_j, D) \neq 1$
11: // The simulation must produce $ \text{corrupt} $ valid shares
12: return 1 if $\text{S.GetPub}(0, D) \neq \Delta$
13: // The commitment must be valid with respect to Δ
14: $b' \xleftarrow{\$} \mathcal{A}(\text{st}_A, \{(j, w_j)\}_{j \in \text{corrupt}}, D)$
15: return 1 if $b' = b$
16: return 0

Fig. 3. Secrecy experiment defining the advantage of an adversary \mathcal{A} against a Verifiable Secret Sharing scheme (VSS) \mathcal{S} that is zero-indexed.

- $\text{GetPub}(i, D) \rightarrow W_i$: A deterministic algorithm that accepts an identifier i and a commitment D . Outputs a share in the public domain W_i associated with identifier i .

For correctness, for all security parameters λ , for all $s \in \hat{S}$, and n, t, C , such that $n \geq t$, $t \geq 1$, $C \subseteq [n]$, and $|C| \geq t$, the following must hold:

$$\begin{aligned}
& \text{S.Verify}(i, w_i, D) = 1 \text{ for all } i \in [n], \text{ and} \\
& \text{S.GetPub}(i, D) = \text{S.SecretToPublic}(w_i) \text{ for all } i \in [n], \text{ and} \\
& \text{S.Recover}(t, M) = s, \text{ where} \\
& \text{S.Share}(\lambda, s, n, t) \rightarrow (\{(i, w_i)\}_{i \in [n]}, D) \text{ and } M = \{(i, w_i)\}_{i \in C}
\end{aligned}$$

Zero-Indexing. For the purposes of our definitions, we additionally require that a VSS be *zero-indexed*, which we define next.

Definition 5. A **zero-indexed VSS** is one which satisfies the correctness condition augmented with the following line:

$$\text{S.Verify}(0, s, D) = 1 \text{ and } \text{S.GetPub}(0, D) = \text{S.SecretToPublic}(s). \quad (1)$$

A prototypical example of a zero-indexed VSS is Feldman’s secret sharing [23], which itself builds on Shamir secret sharing [42].

We now employ the notion of a zero-indexed VSS to define the security notions of secrecy and uniqueness for VSS schemes, building upon prior definitions in the literature [9, 32], but formalized in game-based notation.

Secrecy. A VSS scheme is *secret* if an adversary that is allowed to corrupt up to $t - 1$ players has a negligible probability of learning the secret s . This notion can be also expressed in terms of *simulatability*; i.e., an

$\text{Exp}_{\mathcal{S}, \mathcal{A}}^{\text{unq}}(\lambda)$
1: $(t^*, M_1^*, M_2^*, D^*) \xleftarrow{\$} \mathcal{A}(\lambda)$
2: // \mathcal{A} chooses two recovery sets and a commitment
3: // with respect to the same threshold
4: for $(i, w_i) \in M_1^*$ do
5: return 0 if $\text{S.Verify}(i, w_i, D^*) \neq 1$
6: for $(i', w'_i) \in M_2^*$ do
7: return 0 if $\text{S.Verify}(i', w'_i, D^*) \neq 1$
8: return 0 if $\text{S.Recover}(t^*, M_1^*) = \text{fail} \vee \text{S.Recover}(t^*, M_2^*) = \text{fail}$
9: return 1 if $\text{S.Recover}(t^*, M_1^*) \neq \text{S.Recover}(t^*, M_2^*)$
10: // \mathcal{A} wins if it can generate two different recovery sets for the same
11: // commitment, but which recover distinct secrets
12: return 0

Fig. 4. Uniqueness experiment defining the advantage of an adversary \mathcal{A} against a Verifiable Secret Sharing scheme (VSS) \mathcal{S} .

environment that can simulate secret sharing for an unknown challenge ensures that the adversary learns no additional information than is known by the environment. We formalize this notion in Figure 3, with respect to an algorithm SimShare that is defined in the context of the concrete construction. The experiment takes as input any secret s from the secret space \hat{S} . Importantly, we do not restrict secrets to be sampled uniformly at random; even degenerate cases should hold.

The advantage of an adversary \mathcal{A} against a VSS scheme \mathcal{S} in the secrecy experiment as defined in Figure 3 is

$$\text{Adv}_{\mathcal{S}, \mathcal{A}, \text{SimShare}}^{\text{sec}}(\lambda) = \max_{s \in \hat{S}} (|\Pr[\text{Exp}_{\mathcal{S}, \mathcal{A}, \text{SimShare}}^{\text{sec}}(\lambda, s) = 1] - 1/2|).$$

Definition 6. A VSS scheme \mathcal{S} is **secret** if there exists an algorithm SimShare such that for all probabilistic polynomial time adversaries \mathcal{A} , $\text{Adv}_{\mathcal{S}, \mathcal{A}, \text{SimShare}}^{\text{sec}}(\lambda)$ is a negligible function of λ .

Uniqueness. Informally, a VSS scheme \mathcal{S} that is *unique* is one where the public commitment D uniquely determines the output from S.Recover . We formalize this notion in Figure 4.

The advantage of an adversary \mathcal{A} against \mathcal{S} in the uniqueness experiment as defined in Figure 4 is

$$\text{Adv}_{\mathcal{S}, \mathcal{A}}^{\text{unq}}(\lambda) = \Pr[\text{Exp}_{\mathcal{S}, \mathcal{A}}^{\text{unq}}(\lambda) = 1].$$

Definition 7. A VSS scheme is **unique** if for all probabilistic polynomial time adversaries \mathcal{A} , $\text{Adv}_{\mathcal{S}, \mathcal{A}}^{\text{unq}}(\lambda)$ is a negligible function of λ .

A Concrete VSS. We now recall Feldman’s Verifiable Secret Sharing [23]. Feldman’s VSS itself is an augmentation to Shamir’s secret sharing [42]. Here, the map S.SecretToPublic is simply $a \mapsto g^a$, where g is the generator of the group \mathbb{G} , and $a \in \mathbb{Z}_q$.

- $\text{S.Share}(\lambda, s, n, t) \rightarrow (\{(1, w_1), \dots, (n, w_n)\}, D)$. Accepts as input the security parameter λ , a secret $s \in \mathbb{Z}_q$, the number of participants $n \in \mathbb{N}$, and the threshold $t \in \mathbb{N}$ such that $1 \leq t \leq n$. Perform the following steps:
 1. Sample $t - 1$ coefficients at random: $(a_1, \dots, a_{t-1}) \xleftarrow{\$} \mathbb{Z}_q^{t-1}$.
 2. Using s as the constant term and a_1, \dots, a_{t-1} as the remaining coefficients, define the $(t - 1)$ -degree polynomial $f(x) = s + \sum_{i=1}^{t-1} a_i x^i$.

3. Generate n shares $w_j \in \mathbb{Z}_q$ by deriving $w_j \leftarrow f(j)$, $j \in [n]$.
 4. Define $A_0 = g^s$ and $A_i = g^{a^i}$, for $i \in [t-1]$.
 5. Define the VSS commitment D as the vector of commitments to the coefficients of f : $D \leftarrow (A_0, \dots, A_{t-1})$.
 6. Output $(\{(i, w_i)\}_{i \in [n]}, D)$, where each i represents the identifier of the recipient of w_i .
- **S.Verify** $(i, w_i, D) \rightarrow \{0, 1\}$: Accepts as input a participant identifier $i \in [n]$, the share w_i that belongs to participant i , and a VSS commitment D that is a tuple of group elements, such that $|D| = t$. Verifies that (i, w_i) is a valid point on f as follows:
 1. Parse $\langle A_0, \dots, A_{t-1} \rangle \leftarrow D$.
 2. Output 1 if Equation 2 holds; otherwise, output 0.

$$g^{w_i} \stackrel{?}{=} \prod_{k=0}^{t-1} A_k^{i^k} \quad (2)$$

- **S.Recover** $(t, M) \rightarrow s/\text{fail}$: Accepts as input the threshold t , a recovery set $M = \{(j, w_j)\}_{j \in C}$ consisting of secret shares for a coalition $C \subseteq [n]$ where $|C| \geq t$. If $|C| < t$, output fail. Using the inputs, perform the following steps:
 1. Derive Lagrange coefficients for each identifier $\{L_j(0)\}_{j \in C}$.
 2. Derive $s = \sum_{j \in C} w_j \cdot L_j(0)$.
 3. Output s .
- **S.GetPub** $(i, D) \rightarrow W_i$: Accepts as input a participant identifier i and a VSS commitment. Outputs $W_i = g^{f(i)}$ by computing

$$W_i \leftarrow \prod_{k=0}^{|D|-1} A_k^{i^k} .$$

Correctness and security. For Feldman's VSS, correctness is straightforward to verify. For completeness, we discuss its security in Appendix B.

3.2 Aggregatable Verifiable Secret Sharing (AgVSS)

We now augment the notion of a verifiable secret sharing scheme as presented in Definition 4 to define an *aggregatable* verifiable secret sharing scheme (AgVSS) AV. The ability to aggregate Feldman VSS has been demonstrated before [30], and is employed implicitly in nearly every DKG that has been defined in the literature [27, 29, 37, 40] but we formalize the notion of an AgVSS more generally.

An AgVSS builds upon a base VSS scheme, and is defined with respect to the same map **SecretToPublic** : $\hat{S} \rightarrow \hat{P}$ as the base VSS that maps elements in the secret space \hat{S} to the public space \hat{P} . It is also defined with respect to a random oracle H , and an auxiliary distribution AX that has guessing entropy linear to the security parameter λ .

Definition 8. An **Aggregatable Verifiable Secret Sharing (AgVSS) scheme** $\text{AV}[H]$ is the tuple $(\text{Share}, \text{Recover}, \text{Verify}, \text{GetPub}, \text{AggPriv}, \text{AggPub}, \text{GetTweak})$, where

$\text{Share}, \text{Recover}, \text{Verify}, \text{GetPub}$ are identical to the base VSS scheme as in Definition 4, and where

- **GetTweak** $(O, \text{aux}) \rightarrow v$: A deterministic algorithm that accepts as input a public aggregation set O consisting of ℓ VSS commitments $\{D_1, \dots, D_\ell\}$, and an auxiliary string aux drawn from AX . Output a tweak $v \in \hat{S}$.
- **AggPriv** $(P, v) \rightarrow \hat{w}_i$: A deterministic algorithm that accepts as input a private aggregation set P consisting of ℓ shares $\{w_{ji}\}_{j \in [\ell]}$, and a tweak v . Output the aggregated share \hat{w}_i .
- **AggPub** $(O, v) \rightarrow C$: A deterministic algorithm that accepts as input a public aggregation set O consisting of ℓ VSS commitments $\{D_1, \dots, D_\ell\}$, and a tweak v . Output the aggregated commitment C .

Similarly to the VSS setting, we require that the AgVSS be zero-indexed.

$\text{Exp}_{\text{AV}[\text{H}]}^{\text{corr}}(\lambda, n, t, C, \ell)$
1: return 0 if $t \geq n \vee C \not\subseteq [n] \vee C < t$
2: $\text{aux} \xleftarrow{\$} \text{AV.AX}$
3: for $j \in [\ell]$ do
4: $s_j \xleftarrow{\$} \hat{S}; (\{(i, w_{ji})\}_{i \in [n]}, D_j) \leftarrow \text{AV}[\text{H}].\text{Share}(\lambda, s_j, n, t)$
5: $v \leftarrow \text{AV}[\text{H}].\text{GetTweak}(O, \text{aux})$
6: for $i \in [n]$ do
7: $\hat{w}_i \leftarrow \text{AV}[\text{H}].\text{AggPriv}(\{w_{ji}\}_{j \in [\ell]}, \{D_j\}_{j \in [\ell]}, v)$
8: $C \leftarrow \text{AV}[\text{H}].\text{AggPub}(\{D_j\}_{j \in [\ell]}, v)$
9: for $i \in [n]$ do
10: return 0 if $\text{AV}[\text{H}].\text{Verify}(i, \hat{w}_i, C) \neq 1$
11: $\hat{s} \leftarrow \text{AV}[\text{H}].\text{Recover}(t, \{(k, \hat{w}_k)\}_{k \in C}, C)$
12: return 0 if $\text{AV}[\text{H}].\text{Verify}(0, \hat{s}, C) \neq 1$
13: // This check assumes the AgVSS is zero-indexed
14: return 1

Fig. 5. Notion of correctness for an AgVSS.

Aggregation Correctness. An AgVSS must fulfill the notion of correctness for a VSS scheme. In addition, an AgVSS must fulfill the notion of correctness when the scheme is aggregatable; i.e., that for all n, t, ℓ, C , where $n, t, \ell \in \mathbb{N}$, $C \subseteq [n]$, and $|C| \geq t$,

$$\Pr[\text{Exp}_{\text{AV}[\text{H}]}^{\text{corr}}(\lambda, n, t, C, \ell) = 1] = 1$$

where $\text{Exp}_{\text{AV}[\text{H}]}^{\text{corr}}(\lambda, n, t, C, \ell)$ is presented in Figure 5.

An AgVSS must satisfy aggregated secrecy and uniqueness, as discussed next.

Aggregated Secrecy. Unlike in the plain VSS setting, an AgVSS allows for aggregating shares and commitments that could both be honestly as well as maliciously generated. Hence, we require that the AgVSS fulfills the notion of *aggregated secrecy*, which we formalize in Figure 6.

Remark 3 (Aggregated secrecy for an AgVSS is a generalization of VSS secrecy). The secrecy game for a VSS scheme as shown in Figure 3 only allows the adversary to receive a single set of shares and their corresponding commitment from the environment. Hence, the aggregated secrecy game in Figure 6 is a strict generalization of the VSS secrecy game, as the adversary in the secure aggregation game can not only query for many shares and commitments, but submit its own. As a consequence, any AgVSS that fulfills aggregated secrecy is also secret.

When executing the experiment shown in Figure 6, the environment simulates one secret sharing operation with respect to a challenge Δ when the random bit $b = 1$. The environment allows the adversary to query for corrupted parties' shares by $\mathcal{O}^{\text{GetShare}}$, and submit its own shares for honest parties by $\mathcal{O}^{\text{RecvShare}}$. The adversary is then provided with the auxiliary string aux . The adversary automatically wins if the aggregated commitment AV does not fulfill the check $\text{AV}[\text{H}].\text{GetPub}(0, C) \neq \Delta$. Otherwise, \mathcal{A} outputs a guess b' to guess if it is in the real or simulated environment.

The advantage of an adversary \mathcal{A} against AV in the aggregated secrecy experiment as defined in Figure 6 with respect to algorithms $\text{SimShare}, \text{SimTweak}$ is

$$\text{Adv}_{\text{AV}[\text{H}], \mathcal{A}, \text{SimShare}, \text{SimTweak}}^{\text{asec}}(\lambda) = \max_{s \in \hat{S}} |\Pr[\text{Exp}_{\text{AV}[\text{H}], \mathcal{A}, \text{SimShare}, \text{SimTweak}}^{\text{asec}}(\lambda, s) = 1] - 1/2|.$$

$\text{Exp}_{\text{AV}[\text{H}], \mathcal{A}, \text{SimShare}}^{\text{asec}}(\lambda, s)$	$\mathcal{O}^{\text{GetShare}}()$ // Performs secret sharing
1 : $b \xleftarrow{\$} \{0, 1\}$	1 : if issuedchal = 0
2 : $\Delta \xleftarrow{\$} \text{AV.SecretToPublic}(s)$	2 : issuedchal \leftarrow 1
3 : issuedchal \leftarrow 0	3 : if $b = 0$
4 : aux $\xleftarrow{\$} \text{AV.AX}$	4 : $(\{(j, w_j)\}_{j \in [n]}, D) \xleftarrow{\$} \text{AV}[\text{H}].\text{Share}(\lambda, s, n, t)$
5 : // Pick random auxiliary string	5 : $Q_2 \leftarrow Q_2 \cup \{\{(j, w_j)\}_{j \in \text{honest}, \cdot}, D\}$
6 : $\ell' \leftarrow 0$	6 : return $(\{(j, w_j)\}_{j \in \text{corrupt}}, D)$
7 : $Q_1 \leftarrow \emptyset, Q_2 \leftarrow \emptyset$	7 : else // $b = 1$ case
8 : $(n, t, \text{corrupt}, \text{st}_A) \xleftarrow{\$} \mathcal{A}(\lambda)$	8 : out $\xleftarrow{\$} \text{SimShare}(\lambda, \text{corrupt}, \Delta)$
9 : if $t > n \vee \text{corrupt} \not\subseteq [n]$	9 : $(\alpha^*, \{(j, w_j)\}_{j \in \text{corrupt}}, D^*) \leftarrow \text{out}$
10 : or if $ \text{corrupt} \geq t$	10 : $Q_1 \leftarrow (\alpha^*, D^*)$
11 : $z \xleftarrow{\$} \{0, 1\}$; return z	11 : return $(\{(j, w_j)\}_{j \in \text{corrupt}}, D^*)$
12 : honest $\leftarrow [n] \setminus \text{corrupt}$	12 : $\ell' \leftarrow \ell' + 1$
13 : $\text{st}'_A \xleftarrow{\$} \mathcal{A}^{\mathcal{O}^{\text{GetShare}}, \mathcal{O}^{\text{RecvShare}}, \mathcal{O}^{\text{GetTweak}}}(\text{st}_A)$	13 : $s_{\ell'} \xleftarrow{\$} \hat{S}$
14 : if issuedchal = 0	14 : $(\{(j, w_{\ell'_j})\}_{j \in [n]}, D_{\ell'}) \xleftarrow{\$} \text{AV}[\text{H}].\text{Share}(\lambda, s_{\ell'}, n, t)$
15 : $z \xleftarrow{\$} \{0, 1\}$; return z	15 : $Q_2 \leftarrow Q_2 \cup \{\{(j, w_{\ell'_j})\}_{j \in \text{honest}, \cdot}, D_{\ell'}\}$
16 : if $b = 1$	16 : return $(\{(j, w_{\ell'_j})\}_{j \in \text{corrupt}}, D_{\ell'})$
17 : in $\leftarrow ((\{D_j\}_{j \in [\ell']} \cup \{D^*\}), \text{aux})$	
18 : $C \leftarrow \text{AV}[\text{H}].\text{AggPub}(\text{in})$	$\mathcal{O}^{\text{RecvShare}}(\{(j, w_{\ell'_j})\}_{j \in \text{honest}}, D_{\ell'})$
19 : if $\text{AV}[\text{H}].\text{GetPub}(0, C) \neq \Delta$	1 : // Receives shares from \mathcal{A}
20 : return 1	2 : $\ell' \leftarrow \ell' + 1$
21 : $b' \xleftarrow{\$} \mathcal{A}^{\mathcal{O}^{\text{GetTweak}}}(\text{st}'_A, \text{aux})$	3 : for $j \in \text{honest}$
22 : // \mathcal{A} performs	4 : if $\text{AV}[\text{H}].\text{Verify}(j, w_{\ell'_j}, D_{\ell'}) \neq 1$
23 : // aggregation directly	5 : return \perp
24 : return 1 if $b' = b$	6 : $Q_2 \leftarrow Q_2 \cup \{\{(j, w_{\ell'_j})\}_{j \in \text{honest}}, D_{\ell'}\}$
25 : return 0	
	$\mathcal{O}^{\text{GetTweak}}(O_i, \text{aux}_i)$
	1 : if $b = 1$
	2 : $v \leftarrow \text{SimTweak}(O_i, \text{aux}_i, Q_1, Q_2)$
	3 : else
	4 : $v \leftarrow \text{AV.GetTweak}(O_i, \text{aux}_i)$
	5 : return v

Fig. 6. Aggregated secrecy experiment defining the advantage of an adversary \mathcal{A} against a zero-indexed Aggregatable Verifiable Secret Sharing scheme (AgVSS) AV.

Definition 9. An AgVSS scheme AV achieves **aggregated secrecy** if there exist algorithms SimShare, SimTweak such that for all probabilistic polynomial time adversaries \mathcal{A} , the function $\text{Adv}_{\text{AV}[\text{H}], \mathcal{A}, \text{SimShare}, \text{SimTweak}}^{\text{asec}}(\lambda)$ is a negligible function of λ .

Uniqueness. The uniqueness property for an AgVSS is the same as that of the base VSS, as presented in Figure 4. Notably, the adversary in this experiment can perform aggregation directly, and so the experiment does not change.

A Concrete AgVSS. We now present a concrete AgVSS scheme that builds upon Feldman’s VSS, and assumes an honest majority of participants. The algorithms `Share`, `Recover`, `Verify`, `GetPub` are identical to Feldman’s VSS, as presented in Section 3.1.

Let $H : \mathbb{G}^* \times \mathbb{Z}_q \rightarrow \mathbb{Z}_q$ be a cryptographically secure hash function, whose inputs are a tuple consisting of a vector of group elements in \mathbb{G} , and an element in \mathbb{Z}_q . Finally, let the distribution AX be the uniform distribution on \mathbb{Z}_q .

We now define the aggregation algorithms for this concrete AgVSS.

- `AV.GetTweak`(O, aux): Let $O = \{D_j\}_{j \in [\ell]}$ be the public aggregation set of the corresponding ℓ Feldman VSS commitments to P , such that $|O| = t$, and O is sorted in a canonical order. Generate $v \leftarrow H(O, \text{aux})$. Output v .
- `AV.AggPriv`(P_i, v): Let $P = \{w_{ji}\}_{j \in [\ell]}$ be a private aggregation set of ℓ Shamir secret shares, and $v \in \mathbb{Z}_q$ be an auxiliary string.
 1. Derive the aggregate share via Equation 3.

$$\hat{w}_i \leftarrow v + \sum_{j \in [\ell]} w_{ji} \quad (3)$$

2. Output \hat{w}_i
- `AV.AggPub`(O, v): Perform the following:
 1. Derive the commitment to the aggregated constant coefficient by Equation 4:

$$\hat{A}_0 = g^v \cdot \prod_{j \in [\ell]} D_j[0] \quad (4)$$

2. Derive the commitment to the $k \in \{1, \dots, t-1\}$ aggregated coefficients by Equation 5:

$$\hat{A}_k = \prod_{j \in [\ell]} D_j[k] \quad (5)$$

Let $C = \langle \hat{A}_0, \dots, \hat{A}_{t-1} \rangle$

3. Output C

We demonstrate the correctness and security properties of the concrete scheme in Appendix C.

4 Distributed Key Generation (DKG)

Informally, a *Distributed Key Generation* (DKG) scheme allows a set of n players to jointly generate a public key pk and a secret key sk , such that all parties learn the public key pk but no single party learns sk . Instead, each party id_i has a share sk_i of the secret key, such that a threshold t number of shares are required to recover sk , where $t \leq n$. A DKG generates secret and public keys that are compatible with some *target* key generation scheme which operates in the centralized (single-party) setting.

We now formalize the definition of a DKG. We begin by defining a target single-party key generation scheme TK that the DKG implements in a distributed manner.

Definition 10. A target key generation scheme TK is the tuple $\text{TK} = (\text{KeyGen}, \text{Verify})$, where:

- $\text{KeyGen}(\lambda) \rightarrow (\text{sk}, \text{pk})$: A probabilistic algorithm that accepts as input a security parameter λ and outputs a secret and public keypair (sk, pk) , where sk is an element of a secret domain \hat{S} and pk is an element of the public domain \hat{P} .
- $\text{Verify}(\text{sk}, \text{pk}) \rightarrow \{0, 1\}$: A deterministic algorithm that accepts as input a secret and public keypair (sk, pk) , and outputs 0 if the keypair is invalid; otherwise, outputs 1.

- For correctness, we require that for all $(\text{sk}, \text{pk}) \leftarrow \text{TK.KeyGen}(\lambda)$, we have

$$\text{TK.Verify}(\text{sk}, \text{pk}) = 1 .$$

Definition 11. A **distributed key generation scheme**, or **DKG**, is the tuple of the protocols $\text{D}[\text{H}] = (\text{KeyGen}, \text{Recover})$ parameterized by the number of rounds numrounds . A DKG is defined with respect to a target key generation scheme TK , and initialized with the hash function H , whose domain and codomain depend on D , where:

- $\text{KeyGen}(\lambda, n, t) \rightarrow \langle (\text{pk}, \text{qual}, \text{aux}), (\text{status}_i, \text{sk}_i)_{i \in [n]} \rangle$: An interactive protocol between n participants. Accepts as input the security parameter λ , and positive integers $n \geq t \geq 1$, where n is the number of participants and t is the threshold.

The protocol outputs a public value $(\text{pk}, \text{qual}, \text{aux})$ where:

- pk is the public key representing the group; it will be \perp if the protocol terminated unsuccessfully.
- $\text{qual} \subseteq [n]$ is the set of qualified players remaining at the end of the protocol; it will be \perp if the protocol terminated unsuccessfully.
- aux is any auxiliary information that may be required for a particular application, such as participant public keys,
- status_i indicates if participant i completed the protocol successfully, and can be either **abort**, **fail**, or **accept**. Here, **abort** indicates the protocol terminated early, **fail** indicates a critical fail state has been reached, and **accept** indicates that the protocol completed successfully.
- sk_i is participant i 's secret key share; it will be \perp if the protocol terminated unsuccessfully.

Internally to the KeyGen protocol, each participant performs the following:

- $\text{PerformRound}_0(\lambda, n, t, \text{id}_i) \rightarrow (\text{state}_i, \text{outp}, \text{bmsg}_i)$: The algorithm run by the participant to initiate the protocol. Accepts as input the security parameter λ , the number of participants n , the threshold t , and the participant identifier id_i . Outputs the state state_i for that participant, the vector outp of any outgoing peer-to-peer messages to other parties, and a broadcast message bmsg_i to all other parties. See Section 2.3 for definitions of these network channels.
 - $\text{PerformRound}_k(\text{state}_i, \text{inp}, \text{inb}) \rightarrow (\text{state}_i, \text{outp}, \text{bmsg}_i)$: The algorithm run for each intermediate round $k \in \{1, \dots, \text{numrounds} - 1\}$ in the protocol. Accepts as input the participant's internal state state_i , the vector inp of any peer-to-peer messages sent by other participants in the previous round, and the set inb of any broadcast messages sent by other participants in the previous round. Outputs the tuple $(\text{state}_i, \text{outp}, \text{bmsg}_i)$.
 - $\text{Finalize}(\text{state}_i, \text{inp}, \text{inb}) \rightarrow \langle (\text{pk}, \text{qual}, \text{aux}), (\text{status}_i, \text{sk}_i)_{i \in [n]} \rangle$: Accepts participant state state_i , peer-to-peer messages outp , and broadcast messages inb . Outputs a public output $(\text{pk}, \text{qual}, \text{aux})$ and a private output $(\text{status}_i, \text{sk}_i)$.
- $\text{Recover}(\text{pk}, t, \{(i, \text{sk}_i)\}_{i \in C}) \rightarrow \text{sk}/\text{fail}$: The algorithm run by the participant to finish its protocol execution. Accepts as input the public key pk , the threshold t , and a set of secret key shares $\{(i, \text{sk}_i)\}_{i \in C}$, where $|C| \geq t$. Derive sk from $\{\text{sk}_i\}_{i \in C}$. If $\text{TK.Verify}(\text{sk}, \text{pk}) \neq 1$, output **fail**. Otherwise, output sk .
 - For correctness, we require that, when all participants follow the protocol:

$$\begin{aligned} & \text{if } \text{D}[\text{H}].\text{KeyGen}(\lambda, n, t) \rightarrow ((\text{pk}, \text{qual}), \langle (\text{status}_i, \text{sk}_i) \rangle_{i \in [n]}) \\ & \text{then } \text{D}[\text{H}].\text{Recover}(\text{pk}, t, \{(i, \text{sk}_i)\}_{i \in C}) \rightarrow \text{sk} \text{ for all } C \subseteq [n], |C| \geq t, \\ & \text{and } \text{TK.Verify}(\text{sk}, \text{pk}) \rightarrow 1 \end{aligned}$$

A DKG must be correct, and fulfill the security notions of zero-knowledge and indistinguishability. Additionally, a DKG may fulfill the notion of either strong or weak robustness. We explain these concepts next.

4.1 Robustness

We now review the notion of robustness, building upon the formalism presented by Gurkan et al. [30]. In the literature, robustness today refers to the property that so long as the number of corrupted and honest participants meet some pre-defined threshold, then the protocol will complete successfully. The definition by Gurkan et al. deems a DKG to be robust if honest parties’ shares combine to output a secret key that is valid with respect to a public key. However, the notion by Gurkan et al. does not consider player’s ending states, and also assumes a one-to-one unique mapping between elements in the secret domain and elements in the public domain.

We present several clarifications to the notion of robustness using a game-based definition in Figure 7. First, our definition is generalizable beyond discrete-logarithm based constructions. Second, our definition encodes two different forms of robustness. *Strong robustness* defines the existing notion of robustness in the literature. We additionally introduce the notion of *weak robustness*, which allows the protocol to non-fatally abort, so long as all honest participants maintain the same status. In particular, participants can exit either with an **abort** status or a **fail** status, where **fail** indicates a critical failure. Prior notions in the literature considered only strong robustness; however, introducing a weaker variant allows us to define a more efficient construction that fulfills our other notions of secrecy and indistinguishability.

Remark 4 (Number of Honest Players). The robustness experiment in Figure 7 assumes an honest majority of players. If an adversary were allowed to control the majority of participants, it could split the view of honest participants, unless some additional structure is assumed, such as a public bulletin board.

Strong Robustness. We show the strong robustness attack experiment in Figure 7, where its differences with weak robustness are outlined in a box. The advantage of an adversary \mathcal{A} against D in the strong robustness experiment is

$$\text{Adv}_{D,\mathcal{A}}^{\text{str-rbst}}(\lambda) = \Pr[\text{Exp}_{D,\mathcal{A}}^{\text{str-rbst}}(\lambda) = 1] .$$

Definition 12. A DKG D is **strongly robust** if for all probabilistic polynomial time adversaries \mathcal{A} , the function $\text{Adv}_{D,\mathcal{A}}^{\text{str-rbst}}(\lambda)$ is negligible.

The strong robustness attack game allows \mathcal{A} to participate in key generation and control up to $t - 1$ players, the set of which we denote by **corrupt**. We denote the set of honest players as **honest**. The robustness attack game has oracles $\mathcal{O}^{\text{PerformRound}_r}$, $r \in [\text{numrounds} - 1]$, $\mathcal{O}^{\text{Finalize}}$, and **RO** in common with other games, so we show these oracles in Figure 8. The adversary participates in the protocol by exchanging peer-to-peer and broadcast messages with honest players in each round r via the oracle $\mathcal{O}^{\text{PerformRound}_r}$. The experiment is responsible for exchanging messages between parties and enforcing assumptions of the underlying network channels. Additionally, the experiment ensures that rounds must be queried in order, and that no round is repeated.

\mathcal{A} wins the strong robustness game under three conditions. First, if it causes any honest player to end in a non-accepting state. Second, if it causes any honest players’ view of **pk**, **qual**, or **aux** to be different from any other player’s. Finally, \mathcal{A} wins if it can cause the resulting keypair to be invalid with respect to the target key generation algorithm (recall that $D.\text{Recover}$ tests for validity of the keypair, and outputs **fail** if **sk** is invalid).

Weak Robustness. We similarly present the weak robustness attack experiment in Figure 7, where its differences with strong robustness are outlined in a dashed box. Intuitively, weak robustness differs from strong robustness in that it allows honest parties to end in an aborted state **abort**, so long as *all* honest parties end in the same state. Otherwise, if all honest parties end in an accepting state, it is identical to the notion of strong robustness.

The advantage of an adversary \mathcal{A} against the DKG D in the weak robustness experiment is

$$\text{Adv}_{D,\mathcal{A}}^{\text{wk-rbst}}(\lambda) = \Pr[\text{Exp}_{D,\mathcal{A}}^{\text{wk-rbst}}(\lambda) = 1] .$$

$\text{Exp}_{\mathcal{D}, \mathcal{A}}^{\text{str-rbst}}(\lambda)$	$\text{Exp}_{\mathcal{D}, \mathcal{A}}^{\text{wk-rbst}}(\lambda)$
1: $\text{rc} \leftarrow 0$ // Ensure rounds are queried in order	1: $\text{rc} \leftarrow 0$ // Ensure rounds are queried in order
2: $\text{inp}_{11} \leftarrow \emptyset, \dots, \text{inp}_{1n} \leftarrow \emptyset$ // Simulate peer-to-peer channels	2: $\text{inp}_{11} \leftarrow \emptyset, \dots, \text{inp}_{1n} \leftarrow \emptyset$ // Simulate peer-to-peer channels
3: $(n, t, \text{corrupt}, \text{st}_A) \leftarrow \mathcal{A}(\lambda)$; return 0 if $n < t \vee \text{corrupt} \not\subseteq [n] \vee \text{corrupt} \geq t$	3: $(n, t, \text{corrupt}, \text{st}_A) \leftarrow \mathcal{A}(\lambda)$; return 0 if $n < t \vee \text{corrupt} \not\subseteq [n] \vee \text{corrupt} \geq t$
4: $\text{honest} \leftarrow [n] \setminus \text{corrupt}$	4: $\text{honest} \leftarrow [n] \setminus \text{corrupt}$
5: return 0 if $ \text{honest} < t$	5: return 0 if $ \text{honest} < t$
6: for $i \in \text{honest}$ do	6: for $i \in \text{honest}$ do
7: $(\text{state}_i, \text{outp}_{0i}, \text{bmsg}_i) \leftarrow \text{PerformRound}_0[\text{H}](\lambda, n, t, i)$	7: $(\text{state}_i, \text{outp}_{0i}, \text{bmsg}_i) \leftarrow \text{PerformRound}_0[\text{H}](\lambda, n, t, i)$
8: $\text{inp}_{1k}[i] \leftarrow \text{outp}_{0i}[k], \forall k \in [n]$	8: $\text{inp}_{1k}[i] \leftarrow \text{outp}_{0i}[k], \forall k \in [n]$
9: // Set party i peer-to-peer messages for all other parties	9: // Set party i peer-to-peer messages for all other parties
10: $\mathcal{A}^{\mathcal{O}^{\text{PerformRound}_r}, \mathcal{O}^{\text{Finalize}}, \text{RO}}(\text{st}_A, \{\text{inp}_{1i}\}_{i \in \text{corrupt}}, \{\text{bmsg}_i\}_{i \in \text{honest}})$	10: $\mathcal{A}^{\mathcal{O}^{\text{PerformRound}_r}, \mathcal{O}^{\text{Finalize}}, \text{RO}}(\text{st}_A, \{\text{inp}_{1i}\}_{i \in \text{corrupt}}, \{\text{bmsg}_i\}_{i \in \text{honest}})$
11: return 0 if $\text{rc} \neq \text{numrounds}$ // Prevents trivial win by forcing coordination	11: return 0 if $\text{rc} \neq \text{numrounds}$ // Prevents trivial win by forcing coordination
12: if $\text{honest} \neq \{i : \text{status}_i = \text{accept}\}_{i \in \text{honest}}$	12: if $\text{honest} \neq \{i : \text{status}_i = \text{accept}\}_{i \in \text{honest}}$
13: return 1	13: return 1
14: for $i, j \in \text{honest}$ do	14: for $i, j \in \text{honest}$ do
15: return 1 if $\text{status}_i \neq \text{status}_j$	15: return 1 if $\text{status}_i \neq \text{status}_j$
16: // \mathcal{A} wins if any honest party ends in a disjoint state	16: // \mathcal{A} wins if any honest party ends in a disjoint state
17: if $\text{honest} = \{i : \text{status}_i = \text{abort}\}_{i \in \text{honest}}$	17: if $\text{honest} = \{i : \text{status}_i = \text{abort}\}_{i \in \text{honest}}$
18: return 0	18: return 0
19: // \mathcal{A} loses if all honest parties end in a consistent aborted state	19: // \mathcal{A} loses if all honest parties end in a consistent aborted state
20: for $i, j \in \text{honest}$ do	20: for $i, j \in \text{honest}$ do
21: return 1 if $\text{pk}_i \neq \text{pk}_j$	21: return 1 if $\text{pk}_i \neq \text{pk}_j$
22: return 1 if $\text{qual}_i \neq \text{qual}_j$	22: return 1 if $\text{qual}_i \neq \text{qual}_j$
23: return 1 if $\text{aux}_i \neq \text{aux}_j$	23: return 1 if $\text{aux}_i \neq \text{aux}_j$
24: $\text{pk} = \text{pk}_i, i \in \text{honest}$ // All parties at this point have the same view of pk	24: $\text{pk} = \text{pk}_i, i \in \text{honest}$ // All parties at this point have the same view of pk
25: $\text{sk} \leftarrow \text{D}[\text{H}].\text{Recover}(\text{pk}, \{(i, \text{sk}_i)\}_{i \in \text{honest}})$	25: $\text{sk} \leftarrow \text{D}[\text{H}].\text{Recover}(\text{pk}, \{(i, \text{sk}_i)\}_{i \in \text{honest}})$
26: return 1 if $\text{sk} = \text{fail}$	26: return 1 if $\text{sk} = \text{fail}$
27: // \mathcal{A} wins if honest parties' shares do not recover a valid secret	27: // \mathcal{A} wins if honest parties' shares do not recover a valid secret
28: return 0	28: return 0

Fig. 7. Game defining the advantage of an adversary \mathcal{A} to disrupt robustness for a DKG. We show the *strong robustness* experiment in a box, and the *weak robustness* experiment in a dashed box. The oracles $\mathcal{O}^{\text{PerformRound}_r}, r \in [\text{numrounds}-1]$, $\mathcal{O}^{\text{Finalize}}$, and RO are defined in Figure 8.

Definition 13. A DKG D is **weakly robust** if for all probabilistic polynomial time adversaries \mathcal{A} , the function $\text{Adv}_{\mathcal{D}, \mathcal{A}}^{\text{wk-rbst}}(\lambda)$ is negligible.

\mathcal{A} wins the weak robustness game under three conditions. First, \mathcal{A} wins if honest players do not all finish with the same status. Otherwise, if all honest players have status `accept`, then \mathcal{A} wins if the honest players' view of `pk` or `qual` are not all the same or if the resulting keypair is invalid with respect to the target key generation algorithm.

$\mathcal{O}^{\text{PerformRound}_r}(\{\{\text{pmsg}_{i,j}\}_{j \in \text{honest}}, \text{bmsg}_{r,i}\}_{i \in \text{corrupt}})$ 1 : // Executes round $r \in [\text{numrounds} - 1]$ 2 : return \perp if $r \neq \text{rc} + 1$ 3 : // Ensure all prior rounds 4 : // have been queried, 5 : // and that this round has 6 : // not yet been queried. 7 : $\text{rc} \leftarrow r$ 8 : for $i \in \text{corrupt}, j \in \text{honest}$ do 9 : $\text{inp}_{r,j}[i] \leftarrow \text{pmsg}_i[j]$ 10 : for $j \in \text{honest}$ do 11 : $\text{in} \leftarrow (\text{state}_j, \text{inp}_{r,j}, \{\text{bmsg}_{r,i}\}_{i \in [n]})$ 12 : $\text{out} \leftarrow \text{PerformRound}_r[\text{H}](\text{in})$ 13 : $\text{out} \leftarrow \text{SimRound}_r[\text{H}](\text{in})$ 14 : $(\text{state}_j, \text{outp}_{r+1,j}, \text{bmsg}_{r+1,j}) \leftarrow \text{out}$ 15 : for $k \in [n]$ do 16 : $\text{inp}_{r+1,k} \leftarrow \text{outp}_{r+1,j}[k]$ 17 : return $(\{\text{inp}_{r+1,i}\}_{i \in \text{corrupt}}, \{\text{bmsg}_j\}_{j \in \text{honest}})$	$\mathcal{O}^{\text{Finalize}}(\{\{\text{pmsg}_{i,j}\}_{j \in \text{honest}}, \text{bmsg}_{r,i}\}_{i \in \text{corrupt}})$ 1 : return \perp if $\text{rc} \neq \text{numrounds} - 1$ 2 : $\text{rc} \leftarrow \text{rc} + 1$ 3 : for $i \in \text{corrupt}, j \in \text{honest}$ do 4 : $\text{inp}_{r,j}[i] \leftarrow \text{pmsg}_i[j]$ 5 : for $j \in \text{honest}$ do 6 : $\text{in} \leftarrow (\text{state}_j, \text{inp}_{r,j}, \{\text{bmsg}_{r,i}\}_{i \in [n]})$ 7 : $\text{out} \leftarrow \text{Finalize}[\text{H}](\text{in})$ 8 : $\text{out} \leftarrow \text{SimFinalize}[\text{H}](\text{in})$ 9 : $(\text{pk}_j, \text{qual}_j, \text{aux}_j), (\text{status}_j, \text{sk}_j) \leftarrow \text{out}$ $\text{RO}(x)$ // Random oracle 1 : return $\text{H}(x)$
---	--

Fig. 8. Common oracles used within the games defining the security of a distributed key generation algorithm. Boxed algorithms are used by the robustness and indistinguishability games; dashed boxed algorithms are used by the zero-knowledge game.

4.2 Zero Knowledge

We next turn to defining zero-knowledge for a DKG. Zero knowledge requires that a probabilistic polynomial-time adversary \mathcal{A} allowed to control up to $t - 1$ participants does not gain any additional advantage in its effort to learn sk than it would against the target key generation algorithm. This notion is formalized by the requirement of *simulatability*: that the DKG can be simulated to an adversary with respect to some challenge $\hat{\Delta} \in \hat{P}$, where $\hat{\Delta}$ is a valid public key for the target key generation algorithm. We formalize this notion in Figure 9.

Our definition of zero-knowledge is distinct from the notion of secrecy introduced by Gennaro et al. [28] in several ways. First, we define the notion with respect to a target key generation algorithm, whereas the definition by Gennaro et al. is specifically within the discrete-logarithm setting. Second, Gennaro et al. assume that indistinguishability of the resulting keypair is an implicit consequence of the ability to simulate with respect to a public key. However, simulatability does not necessarily imply indistinguishability, because the zero-knowledge game *only* guarantees the indistinguishability of public key material between the real and simulated environments. If there is a one-to-one relation between public and secret keys, proving zero-knowledge implies indistinguishability. However, for cryptosystems where a one-to-many relation exists between public and private keys, such as in lattice-based cryptosystems, proving that a DKG is zero-knowledge does *not* imply indistinguishability from its target key generation scheme. For this reason, we instead present the notions of zero-knowledge and indistinguishability as separate notions.

We now define this notion of zero-knowledge more formally. The advantage of an adversary \mathcal{A} against a DKG D with respect to the simulation algorithms $(\text{SimRound}_i, i \in \{0, \dots, \text{numrounds} - 1\}, \text{SimFinalize})$ in the

```


$$\text{Exp}_{\mathcal{D}, \mathcal{A}, \text{SimRound}_i, \text{SimFinalize}}^{\text{zk}}(\lambda)$$

1:  $b \xleftarrow{\$} \{0, 1\}$ 
2:  $z \xleftarrow{\$} \{0, 1\}$  // Random coin for losing conditions
3:  $\text{rc} \leftarrow 0$  // Ensure rounds are queried in order
4:  $\text{inp}_{11} \leftarrow \emptyset, \dots, \text{inp}_{1n} \leftarrow \emptyset$  // Simulate peer-to-peer channels
5:  $(n, t, \text{corrupt}, \text{st}_A) \leftarrow \mathcal{A}(\lambda)$ ;  $\text{honest} \leftarrow [n] \setminus \text{corrupt}$ 
6: if  $n < t \vee \text{corrupt} \not\subseteq [n] \vee |\text{corrupt}| \geq t$ 
7:   return  $z$ 
8:  $(\cdot, \hat{\Delta}) \xleftarrow{\$} \text{TK.KeyGen}(\lambda)$ 
9: for  $i \in \text{honest}$  do
10:   if  $b = 0$  then  $(\text{state}_i, \text{outp}_i, \text{bmsg}_i) \leftarrow \text{PerformRound}_0[\text{H}](\lambda, n, t, i)$ 
11:   else  $(\text{state}_i, \text{outp}_i, \text{bmsg}_i) \leftarrow \text{SimRound}_0[\text{H}](\lambda, n, t, i, \hat{\Delta})$ 
12:    $\text{inp}_{1k}[i] \leftarrow \text{outp}_{0i}[k], \forall k \in [n]$ 
13:   // Set party  $i$  peer-to-peer messages for all other parties
14:  $b' \leftarrow \mathcal{A}^{\mathcal{O}^{\text{PerformRound}_r}, \mathcal{O}^{\text{Finalize}}, \text{RO}}(\text{st}_A, \{\text{inp}_{1j}\}_{j \in \text{corrupt}}, \{\text{bmsg}_i\}_{i \in \text{honest}})$ 
15: if  $\text{rc} \neq \text{numrounds}$  // Prevents trivial win by forcing coordination
16:   return  $z$ 
17: if  $\{j : \text{status}_j = \text{abort}\}_{j \in \text{honest}} = \text{honest}$ 
18:   return  $z$ 
19:   //  $\mathcal{A}$  cannot win if all honest parties non-fatally abort
20: if  $b = 1$ 
21:   return 1 if  $\exists i \in \text{honest} : \text{pk}_i \neq \hat{\Delta}$ 
22:   //  $\mathcal{A}$  wins when  $b = 1$  if the game is not simulated with respect to  $\hat{\Delta}$ 
23: return 1 if  $b \stackrel{?}{=} b'$ 
24: return 0

```

Fig. 9. Zero-knowledge experiment defining the advantage of an adversary \mathcal{A} that controls up to $t - 1$ players in a DKG. The oracles $\mathcal{O}^{\text{PerformRound}_r}$ and RO are defined in Figure 8.

zero-knowledge experiment as defined in Figure 9 is

$$\text{Adv}_{\mathcal{D}, \mathcal{A}, \text{SimRound}_i, \text{SimFinalize}}^{\text{zk}}(\lambda) = |\Pr[\text{Exp}_{\mathcal{D}, \mathcal{A}, \text{SimRound}_i, \text{SimFinalize}}^{\text{zk}}(\lambda) = 1] - 1/2|.$$

Definition 14. A DKG \mathcal{D} is **zero-knowledge** if there exist algorithms $(\text{SimRound}_i, i \in \{0, \dots, \text{numrounds} - 1\}, \text{SimFinalize})$ such that for all probabilistic polynomial time adversaries \mathcal{A} , $\text{Adv}_{\mathcal{D}, \mathcal{A}, \text{SimRound}_i, \text{SimFinalize}}^{\text{zk}}(\lambda)$ is negligible.

The zero-knowledge attack game in Figure 9 requires \mathcal{A} to successfully guess if the environment that it is in is real or simulated with respect to a challenge $\hat{\Delta}$, where $\hat{\Delta}$ is a public key generated by TK.KeyGen . The game accepts as input the security parameter λ , and internally to the experiment, samples a bit $b \in \{0, 1\}$ at random. In the simulated setting (when $b = 1$), the environment simulates the honest participants with respect to $\hat{\Delta}$ by employing the simulation algorithms $(\text{SimRound}_i, i \in \{0, \dots, \text{numrounds} - 1\}, \text{SimFinalize})$. The adversary wins by default if the environment cannot simulate the protocol in such a way that the resulting group public key pk does not equal $\hat{\Delta}$. Otherwise, the adversary wins if it successfully guesses whether it is playing against the real or simulated protocol.

```

ExpD,TK,A,εind(λ)
1:  $b \xleftarrow{\$} \{0, 1\}$ 
2:  $z \xleftarrow{\$} \{0, 1\}$  // Random coin for losing conditions
3:  $rc \leftarrow 0$  // Ensure rounds are queried in order
4:  $\text{inp}_{11} \leftarrow \emptyset, \dots, \text{inp}_{1n} \leftarrow \emptyset$  // Simulate peer-to-peer channels
5:  $(n, t, \text{corrupt}, \text{st}_A) \leftarrow \mathcal{A}(\lambda)$ 
6: if  $n < t \vee \text{corrupt} \not\subseteq [n] \vee |\text{corrupt}| \geq t$  then  $z \xleftarrow{\$} \{0, 1\}$ ; return  $z$ 
7:  $\text{honest} \leftarrow [n] \setminus \text{corrupt}$ 
8: for  $i \in \text{honest}$  do
9:    $(\text{state}_i, \text{outp}_{0i}, \text{bmsg}_i) \leftarrow \text{PerformRound}_0[\text{H}](\lambda, n, t, i)$ 
10:   $\text{inp}_{1k}[i] \leftarrow \text{outp}_{0i}[k], \forall k \in [n]$ 
11:  // Set party  $i$  peer-to-peer messages for all other parties
12:   $\mathcal{A}^{\mathcal{O}^{\text{PerformRound}_r}, \mathcal{O}^{\text{Finalize}_r}, \text{RO}}(\{\text{inp}_{1i}\}_{i \in \text{corrupt}}, \{\text{bmsg}_i\}_{i \in \text{honest}}, \text{st}_A)$ 
13: if  $rc \neq \text{numrounds}$  // Prevents trivial win by forcing coordination
14:   return  $z$ 
15: if  $\exists i \in \text{honest} : \text{status}_i \neq \text{accept}$ 
16:   return  $z$ 
17:  // Honest parties must complete successfully
18: for  $i, j \in \text{honest}$  do
19:   return 1 if  $\text{pk}_i \neq \text{pk}_j$ 
20:   return 1 if  $\text{qual}_i \neq \text{qual}_j$ 
21:  $\text{sk}_0 \leftarrow \text{D}[\text{H}].\text{Recover}(\text{pk}_0, t, \{(i, \text{sk}_i)\}_{i \in \text{honest}})$ 
22:  // All parties at this point have the same view of  $\text{pk}$  and  $\text{qual}$ 
23: return 1 if  $\text{sk}_0 = \text{fail}$ 
24:  // The output DKG keypair given to the distinguisher must be valid
25:  $(\text{sk}_1, \text{pk}_1) \xleftarrow{\$} \text{TK}.\text{KeyGen}(\lambda)$ 
26:  $b' \leftarrow \mathcal{E}((\text{pk}_b, \text{sk}_b), (\text{pk}_{1-b}, \text{sk}_{1-b}))$ 
27: return 1 if  $b' = b$ 
28: return 0

```

Fig. 10. Game defining the advantage of an adversary \mathcal{A} to compromise indistinguishability of a DKG. The distinguisher \mathcal{E} is challenged to distinguish between valid keypairs generated via the DKG and the target key generation algorithm. The oracles $\mathcal{O}^{\text{PerformRound}_r}$ and RO are defined in Figure 8.

4.3 Indistinguishability

The security of any key generation protocol requires that the secret key be hard to guess or learn from the public transcript. In the single-party setting, the adversary does not have any influence over the sampling of key material, and so guaranteeing this property is straightforward. However, because the adversary is an active participant in a distributed key generation protocol, we must ensure that a subset of at most $t - 1$ colluding players cannot bias key material.

To establish this property, we require that a DKG must generate keys that are *indistinguishable* from keys output by the target key generation algorithm. We formalize this requirement in the attack game in Figure 10.

The advantage of a distinguishing adversary \mathcal{E} against a DKG D in the indistinguishability experiment in Figure 10 with respect to the target key generation algorithm TK and participating adversary \mathcal{A} is

$$\text{Adv}_{D, \text{TK}, \mathcal{A}, \mathcal{E}}^{\text{ind}}(\lambda) = |\Pr[\text{Exp}_{D, \text{TK}, \mathcal{A}, \mathcal{E}}^{\text{ind}}(\lambda) = 1] - 1/2|.$$

Definition 15. A DKG D is **indistinguishable** from its target key generation algorithm if for all probabilistic polynomial time adversaries \mathcal{A} and computationally unbounded distinguishers \mathcal{E} , $\text{Adv}_{D, \text{TK}, \mathcal{A}, \mathcal{E}}^{\text{ind}}(\lambda)$ is negligible.

The attack game in Figure 10 allows for two distinct adversaries, \mathcal{A} and \mathcal{E} . \mathcal{A} participates in the protocol and controls up to $t - 1$ players. \mathcal{E} is simply given two sets of keypairs after the protocol completes, and is required to guess which keypair was generated by the distributed key generation protocol and which by the target key generation protocol. Note that \mathcal{A} and \mathcal{E} are assumed to be black box and act arbitrarily, but are not allowed to share state. Otherwise, the game could trivially be won simply by \mathcal{A} informing \mathcal{E} which public key was output from the distributed key generation protocol. To similarly prevent this trivial win condition, the keypair output from the distributed key generation execution must be valid, otherwise \mathcal{A} could force the protocol to output an invalid keypair, which \mathcal{E} could trivially distinguish from the output of TK.KeyGen.

5 A Generic DKG Construction

We now introduce a weakly robust generic construction that assumes an honest majority, shown in Figure 11. We present a concrete construction STORM of this generic construction that is secure in the discrete log setting in Section 6.

Our generic construction employs the following building blocks:

- A secure hash function H , where H_1 and H_2 are domain-separated instances of H , such that:
 - $H_1 : (\{D_j\}_{j \in [n]}, \text{aux}) \mapsto v$: Accepts as input a tuple, where the first element in the tuple is a set of n AgVSS commitments ordered canonically ordered, and the second element is an auxiliary element aux as output by H_2 below. The output is the tweak employed by the AgVSS aggregation algorithm.
 - $H_2 : \{\psi_j\}_{j \in [n]} \mapsto \text{aux}$: Accepts as input a set of n elements, and outputs an auxiliary element aux .
- An aggregatable verifiable secret sharing scheme $\text{AV}[H_1] = (\text{Share}, \text{Verify}, \text{Recover}, \text{GetPub}, \text{AggPriv}, \text{AggPub})$ that is zero-indexed.
- A non-interactive key exchange (NIKE) $\text{NK} = (\text{KeyGen}, \text{Verify}, \text{SharedKey})$.

The secret and public domains of the AgVSS and NIKE must be the same as that of the target key generation scheme, and the distribution over the secret domain for the NIKE must be the same as that of the target key generation scheme. Further, we assume that the NIKE outputs secret and public keys with strictly a one-to-one relation.

In addition to these building blocks, our construction requires that these primitives are algebraically composable. Intuitively, the secret key output from NK.KeyGen should be a valid secret that can be shared via AV.Share , and the AgVSS commitment should be verifiable with respect to the public NIKE key. Towards this latter point, we require Assumption 1 to hold.

Assumption 1 For any $(\text{sk}, \text{pk}) \xleftarrow{\$} \text{NK.KeyGen}(\lambda)$, and any $(n, t) \in \mathbb{N}$ such that $n \geq t$; and any $C \subseteq [n]$, $|C| \geq t$, Equation 6 holds:

$$\begin{aligned} & \text{NK.Verify}(\text{sk}, \text{pk}) = 1, \text{ where} \\ & (\{(i, w_i)\}_{i \in [n]}, D) \xleftarrow{\$} \text{AV}[H].\text{Share}(\lambda, \text{sk}, n, t) \text{ and} \\ & \text{pk} \leftarrow \text{AV}[H_1].\text{GetPub}(0, D) \text{ and} \\ & \text{sk} \leftarrow \text{AV}[H_1].\text{Recover}(t, \{(i, w_i)\}_{i \in C}) \end{aligned} \tag{6}$$

PerformRound ₀ [H](λ, n, t, i)	Finalize[H](state _i , \emptyset , inb ₃)
<pre> 1 : (α_i, A_i) $\xleftarrow{\\$}$ NK.KeyGen(λ) 2 : in \leftarrow (λ, α_i, n, t) 3 : ($\{(j, w_{ij})\}_{j \in [n]}, D_i$) $\xleftarrow{\\$}$ AV[H₁].Share(in) 4 : // The NIKE secret key α_i is 5 : // secret shared via the AgVSS. 6 : (β_i, B_i) $\xleftarrow{\\$}$ NK.KeyGen(λ) 7 : // Broadcast commitments 8 : bmsg_{0,i} \leftarrow (A_i, B_i, D_i) 9 : for $j \in [n], j \neq i$ do 10 : // Send shares via peer-to-peer channel 11 : pmsg_{0,ij} \leftarrow (j, w_{ij}) 12 : state_i \leftarrow ($\beta_i, w_{ii}, D_i, B_i$) 13 : outp \leftarrow {pmsg_{0,ij}}_{j ∈ [n], j ≠ i} 14 : return (state_i, outp, bmsg_{0,i}) </pre>	<pre> 1 : if state_i.status = abort 2 : return (\perp, \perp, \perp), (abort, \perp) 3 : parse {β_j}_{j ∈ [n], j ≠ i} \leftarrow inb₃ 4 : qual \leftarrow \emptyset; corrupt \leftarrow \emptyset 5 : for $j \in [n], j \neq i$ do 6 : if NK.Verify(β_j, B_j) = 1 7 : $\psi_j \leftarrow$ NK.SharedKey(β_j, A_j) 8 : // Qualified participants send 9 : // the correct opening β_j 10 : qual \leftarrow qual \cup {j} 11 : else corrupt \leftarrow corrupt \cup {j} 12 : for $j \in$ corrupt do 13 : // Perform secret recovery 14 : // sub-protocol for 15 : // misbehaving parties 16 : $M_j = \{(k, w_{jk})\}_{k \in \text{qual}}$ 17 : // All remaining players 18 : // combine their shares 19 : $\alpha_j \leftarrow$ AV[H₁].Recover(t, M_j) 20 : $\psi_j \leftarrow$ NK.SharedKey(α_j, B_j) 21 : $v \leftarrow$ AV[H₁].GetTweak($\{D_j\}_{j \in [n]}, \{\psi_j\}_{j \in [n]}$) 22 : sk_i \leftarrow AV[H₁].AggPriv($i, \{w_{ji}\}_{j \in [n]}, v$) 23 : $C \leftarrow$ AV[H₁].AggPub($\{D_j\}_{j \in [n]}, v$) 24 : pk \leftarrow AV[H₁].GetPub(0, C) 25 : // Derive public keys for all 26 : // remaining participants 27 : for $i \in$ qual do 28 : pk_i = AV[H₁].GetPub(i, C) 29 : aux \leftarrow {pk_i}_{i ∈ qual} 30 : return (pk, qual, aux), (accept, sk_i) </pre>
<pre> PerformRound₁[H](state_i, inp₁, inb₁) 1 : parse {(i, w_{ji})}_{j ∈ [n], j ≠ i} \leftarrow inp₁ 2 : parse {(A_j, B_j, D_j)}_{j ∈ [n], j ≠ i} \leftarrow inb₁ 3 : for $j \in [n], j \neq i$ do 4 : if AV[H₁].Verify(i, w_{ji}, D_j) \neq 1 5 : or if AV[H₁].GetPub(0, D_j) \neq A_j 6 : state_i.status = abort 7 : bmsg_{1,i} \leftarrow fail 8 : // If any verification check fails, abort 9 : return (state_i, \perp, bmsg_{1,i}) 10 : state_i \leftarrow state_i \cup {(j, w_{ji})}_{j ∈ [n]} 11 : bmsg_{1,i} \leftarrow accept 12 : return (state_i, \perp, bmsg_{1,i}) </pre>	
<pre> PerformRound₂[H](state_i, \emptyset, inb₂) 1 : if fail \in inb₂ 2 : state_i.status = abort 3 : // If any participant fails, abort 4 : return (state_i, \perp, \perp) 5 : bmsg_{2,i} \leftarrow β_i 6 : return (state_i, $\perp, \text{bmsg}_{2,i}$) </pre>	

Fig. 11. A generic DKG key generation protocol as defined in Section 4, initialized with a NIKE NK as defined in Section 2.2, an AgVSS AV as defined in Section 3.2, and a hash function H. Each PerformRound_r operation is defined with respect to participant i . The peer-to-peer channel inputs for round r are inp_r and outputs are outp_r, and are comprised of peer-to-peer messages pmsg_{r,j}. The broadcast channel inputs are inb_r and are comprised of broadcast messages bmsg_{r,j}, sent by participant $j, j \in [n]$.

Key Generation Protocol. As shown in Figure 11, the key generation protocol is completed in three network rounds.

First, in PerformRound_0 , each participant i performs NK.KeyGen twice; first to generate a secret key α_i and its commitment A_i , and second to generate a random value β_i and its commitment B_i . Then, each participant i performs an AgVSS secret sharing of α_i , outputting n secret shares $\{(j, w_{ij})\}_{j \in [n]}$, and an AgVSS commitment D_i . Finally, each participant i broadcasts its NIKE and AgVSS commitments to all other players via a broadcast message $\text{bmsg}_{0,i}$. Additionally, each participant i sends one secret share to every other player $j \neq i$ via a peer-to-peer message $\text{pmsg}_{0,ij}$, keeping one share for itself.

In PerformRound_1 , each participant i receives shares $\{(i, w_{ji})\}_{j \in [n], j \neq i}$ from all other players via the peer-to-peer channel input inp_1 and commitments $\{(A_j, B_j, D_j)\}_{j \in [n], j \neq i}$ from all other players via the broadcast channel input inb_1 . Each participant then verifies the correctness of their received share with respect to its corresponding commitment. Additionally, each participant i verifies that the secret committed to in D_i is the same secret committed to in A_i , by checking that $\text{AV.GetPub}(0, D_i) = A_i$. If any check fails, the participant will broadcast a fail message, and set their status to `abort`. Any participant that sets their status to `abort` results in that participant terminating (immediately exiting) the protocol. Else, it will broadcast an `accept` message.

In PerformRound_2 , participants receive as input status messages from all other participants via the broadcast channel input inb_2 . If any participant received a fail message, that participant immediately exits PerformRound_2 after setting their status to `abort`¹. Otherwise, each participant i broadcasts β_i to all other participants (which is the opening to the commitment B_i sent in PerformRound_0).

Finally, in Finalize , each participant i receives as input the openings $\{\beta_j\}_{j \in [n], j \neq i}$ from all other participants via the broadcast channel inb_3 . For the values they received from every other player $j \neq i$, they first verify the correctness of the opening β_j with respect to B_j . Depending on the output of this check, each participant i does the following:

- If the opening β_j is valid, participant i then derives the blinding factor ψ_j by performing $\psi_j \leftarrow \text{NK.SharedKey}(\beta_j, A_j)$.
- If the opening β_j is invalid, the set of participants `qual` that followed the protocol honestly cooperate to derive the correct ψ_j , excluding the cheating participant. They do so by pooling their respective shares in order to recover α_j . We discuss this secret recovery sub-protocol further below. After obtaining this set `qual` of at least t shares, each remaining participant then performs $\text{NK.SharedKey}(\alpha_j, B_j)$, obtaining ψ_j as the result. Due to the assumption that at least t honest participants exist, this step will always complete successfully.

Next, each participant derives $v \leftarrow \text{AV}[H_1].\text{GetTweak}(O, \{\psi_1, \dots, \psi_n\})$, where $O = \{D_j\}_{j \in [n]}$ is the set of all AgVSS commitments for each participant. Each participant i then performs AV.AggPriv and AV.AggPub , using as input the set of all their received shares $P = \{w_{ji}\}_{j \in [n]}$, the set of AgVSS commitments O , and v . Each participant i employs the output of AV.AggPriv and AV.AggPub to determine their final secret key share sk_i and the aggregated commitment C . The group public key pk is then derived using C . Similarly, the set of each participant's individual public key pk_i is derived using C , and then output as the auxiliary data.

Participants finally complete Finalize by setting their status to `accept`, updating their internal state, and exiting the protocol. The public output at the end of PerformRound_3 is $(\text{pk}, \text{qual}, \text{aux})$. The private output to each participant i is $(\text{status}, \text{sk}_i)$.

Secret Recovery Sub-Protocol. When performing the Finalize step in Figure 11, in case any party misbehaves and sends an incorrect opening β_j , all remaining honest participants broadcast the shares sent by the malicious party j in order to recover the secret s_j . We rely on existing definitions of this sub-protocol in prior related DKG constructions; see Gennaro et al. [28] for details of this sub-protocol, which is only performed in the event of a misbehaving party.

¹ Because the fail message is sent on a broadcast channel, if any participant receives a fail message, then all participants receive that message

Recovery Algorithm. While most practical DKG applications do not require the participants to actually come together to recover the secret key, it is possible that some applications may wish to do so. We define the recovery algorithm now.

- $\text{Recover}(\text{pk}, t, \{(i, \text{sk}_i)\}_{i \in C}) \rightarrow \text{sk}/\text{fail}$: Accepts the public key $\text{pk} \in \mathbb{G}$ and $|C| \geq t$ tuples $\{(i, \text{sk}_i)\}_{i \in C}$ consisting of participant i 's identifier and their secret key. Using these inputs, perform the following steps:
 1. Derive $\text{sk} = \text{AV.Recover}(t, \{(i, \text{sk}_i)\}_{i \in C})$.
 2. If $\text{TK.Verify}(\text{sk}, \text{pk}) \neq 1$, output fail.
 3. Output sk .

5.1 Security of the Generic Protocol

We next demonstrate that the generic protocol in Figure 11 is weakly robust, zero-knowledge, and indistinguishable.

Theorem 1. *The generic construction is weakly robust in the honest majority setting against an adversary \mathcal{A} controlling up to $t - 1$ participants in the random oracle model, a zero-indexed and secure AgVSS, a binding NIKE.*

Concretely, for every adversary \mathcal{A} that attacks the scheme, there exist adversaries $(\mathcal{B}_1, \mathcal{B}_2)$ that run in about the same time as \mathcal{A} such that

$$\text{Adv}_{\mathcal{D}, \mathcal{A}}^{\text{wk-rbst}}(\lambda) \leq \text{Adv}_{\text{NK}, \mathcal{B}_1}^{\text{bind}}(\lambda) + \text{Adv}_{\text{AV}[\mathbb{H}_1], \mathcal{B}_2}^{\text{unq}}(\lambda). \quad (7)$$

We prove Theorem 1 in Appendix D.1.

Theorem 2. *The generic construction is zero-knowledge in the honest majority setting in the random oracle model, against a PPT adversary \mathcal{A} controlling up to $t - 1$ participants, assuming the generic construction itself is weakly robust, the AgVSS achieves aggregated secrecy, and the NIKE is unrecoverable.*

Concretely, for every adversary \mathcal{A} that attacks the construction, there exist adversaries $(\mathcal{B}_1, \mathcal{B}_3, \mathcal{B}_4)$ that run in about the same time as \mathcal{A} , and algorithms $\text{SimRound}_i, i \in \{0, \dots, 2\}, \text{SimFinalize}$, such that

$$\begin{aligned} \text{Adv}_{\mathcal{D}, \mathcal{A}, \text{SimRound}_i, \text{SimFinalize}}^{\text{zk}}(\lambda) &\leq \text{Adv}_{\text{AV}[\mathbb{H}_1], \mathcal{B}_1}^{\text{sec}}(\lambda) + \text{Adv}_{\text{NK}, \mathcal{B}_3}^{\text{rec}}(\lambda) \\ &\quad + \text{Adv}_{\text{AV}[\mathbb{H}_1], \mathcal{B}_4, \text{SimShare}, \text{SimTweak}}^{\text{asec}}(\lambda). \end{aligned} \quad (8)$$

where $\text{SimShare}, \text{SimTweak}$ are as defined for aggregated secrecy of the AgVSS, as in Definition 9.

We prove Theorem 2 in Appendix D.2.

Theorem 3. *The generic construction is indistinguishable in the honest majority setting in the random oracle model, against a PPT adversary \mathcal{A} controlling up to $t - 1$ participants, assuming the generic construction itself is weakly robust, the AgVSS achieves aggregated secrecy, and the NIKE is unrecoverable.*

Concretely, for every adversary \mathcal{A} that attacks the construction, there exist adversaries $(\mathcal{B}_1, \mathcal{B}_3, \mathcal{B}_4)$ that run in about the same time as \mathcal{A} , such that

$$\begin{aligned} \text{Adv}_{\mathcal{D}, \text{TK}, \mathcal{A}, \varepsilon}^{\text{ind}}(\lambda) &\leq \text{Adv}_{\text{AV}[\mathbb{H}_1], \mathcal{B}_1}^{\text{sec}}(\lambda) + \text{Adv}_{\text{NK}, \mathcal{B}_3}^{\text{rec}}(\lambda) \\ &\quad + \text{Adv}_{\text{AV}[\mathbb{H}_1], \mathcal{B}_4, \text{SimShare}, \text{SimTweak}}^{\text{asec}}(\lambda). \end{aligned} \quad (9)$$

where $\text{SimShare}, \text{SimTweak}$ are as defined for aggregated secrecy of the AgVSS, as in Definition 9.

Because our generic construction requires a one-to-one mapping between secret and public keys, Theorem 3 follows from Theorem 2.

6 STORM, a Concrete DKG Construction

Our generic DKG presented in Section 5 must be instantiated with concrete building blocks to be usable. We now present one possible concrete construction, which we call STORM (Synchronous, distributed, and Optimized generation of key Material). We refer to Section 5 for the details of the DKG key generation and recovery steps, but describe the building blocks and security assumptions for STORM here.

STORM is secure under the Computational Diffie-Hellman (CDH) problem in the random oracle model, in the honest majority setting. STORM can be securely composed with any non-pairing based scheme that relies on these assumptions, as the distributed form of the following target key generation scheme.

Target Key Generation. The target key generation scheme TK for STORM is defined as follows. The secret domain \hat{S} is the field \mathbb{Z}_q and the public domain \hat{P} is the group \mathbb{G} generated by g .

- $\text{KeyGen}(\lambda) \rightarrow (\text{sk}, \text{pk})$: Sample $\text{sk} \xleftarrow{\$} \mathbb{Z}_q$. Generate $\text{pk} \leftarrow g^{\text{sk}}$. Output (sk, pk) .
- $\text{Verify}(\text{sk}, \text{pk}) \rightarrow \{0, 1\}$: Output 1 if $g^{\text{sk}} = \text{pk}$. Otherwise, output 0.

Building Blocks. STORM is instantiated with the following cryptographic building blocks:

1. the concrete NIKE scheme presented in Section 2.2,
2. the concrete AgVSS scheme presented in Section 3.2, and
3. a secure hash function H , where H_1 , and H_2 domain-separated instances of H , defined as follows:
 - $H_1 : (\mathbb{G}^t)^n \times \mathbb{Z}_q \rightarrow \mathbb{Z}_q$: Accepts as input a vector, where the first element is a tuple of n AgVSS commitments $D_j \in \mathbb{G}^t, j \in [n]$ which are each a tuple of group elements of size t , and the second element is an auxiliary element $\text{aux} \in \mathbb{Z}_q$.
 - $H_2 : (\mathbb{G})^n \rightarrow \mathbb{Z}_q$: Accepts as input a vector of n blinding factors such that $\psi_j \in \mathbb{G}, j \in [n]$, and outputs an auxiliary element $\text{aux} \in \mathbb{Z}_q$.

Correctness and Security. Because STORM is an instantiation of the generic construction presented in Section 5, its correctness and security automatically follow from the proofs for the generic scheme.

7 Conclusion

In this work, we present a generic construction for a DKG that can be securely employed in any setting in place of its target (single-party) key generation scheme. To prove its security, we require that the DKG be strongly or weakly robust, zero-knowledge, and indistinguishable from its target key generation scheme. We formalize these notions using a game-based approach, and use these notions to prove the security of our generic construction. We then introduce STORM, a concrete instantiation of our generic construction that is secure assuming the Computational Diffie-Hellman problem is hard, in the random oracle model.

8 Acknowledgments

We thank Dan Shumow, Greg Zaverucha, and Chris Wood for their helpful discussion, and Jack Grigg and Deirdre Connolly for their review and feedback. Many thanks to Olivier Sanders and Jonathan Katz for their discussion on simulation-based notions of security. Chelsea Komlo was supported in part by Microsoft Research in the process of completing this work. Douglas Stebila was supported in part by Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grant RGPIN-2022-0318. This research was undertaken, in part, thanks to funding from the Canada Research Chairs program.

References

1. I. Abraham, P. Jovanovic, M. Maller, S. Meiklejohn, and G. Stern. Bingo: Adaptively secure packed asynchronous verifiable secret sharing and asynchronous distributed key generation. Cryptology ePrint Archive, Paper 2022/1759, 2022. URL: <https://eprint.iacr.org/2022/1759>.
2. I. Abraham, P. Jovanovic, M. Maller, S. Meiklejohn, G. Stern, and A. Tomescu. Reaching consensus for asynchronous distributed key generation. In A. Miller, K. Censor-Hillel, and J. H. Korhonen, editors, *PODC '21: ACM Symposium on Principles of Distributed Computing*, pages 363–373. ACM, 2021.
3. R. Bacho and J. Loss. On the adaptive security of the threshold BLS signature scheme. In H. Yin, A. Stavrou, C. Cremers, and E. Shi, editors, *CCS 2022*, pages 193–207. ACM, 2022.
4. R. Barnes, K. Bhargavan, B. Lipp, and C. Wood. Hybrid public key encryption. <https://datatracker.ietf.org/doc/html/rfc9180>, 2022.
5. M. Bellare, E. C. Crites, C. Komlo, M. Maller, S. Tessaro, and C. Zhu. Better than advertised security for non-interactive threshold signatures. In Y. Dodis and T. Shrimpton, editors, *CRYPTO 2022*, volume 13510 of *LNCS*, pages 517–550. Springer, 2022.
6. R. Bendlin, S. Krehbiel, and C. Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In M. J. J. Jr., M. E. Locasto, P. Mohassel, and R. Safavi-Naini, editors, *Applied Cryptography and Network Security (ACNS) 2013*, volume 7954 of *LNCS*, pages 218–236. Springer, 2013.
7. D. J. Bernstein. Curve25519: New Diffie–Hellman speed records. In M. Yung, Y. Dodis, A. Kiayias, and T. Malkin, editors, *Public Key Cryptography - PKC 2006*, volume 3958 of *LNCS*, pages 207–228. Springer, 2006.
8. D. Boneh, B. Lynn, and H. Shacham. Short signatures from the weil pairing. *J. Cryptol.*, 17(4):297–319, 2004.
9. D. Boneh and V. Shoup. A graduate course in applied cryptography, 2023. URL: <http://toc.cryptobook.us/book.pdf>.
10. L. Brandão and M. Davidson. Notes on threshold eddsa/schnorr signatures. <https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8214B.ipd.pdf>, 2022.
11. L. Brandão and R. Peralta. Nist first call for multi-party threshold schemes. <https://nvlpubs.nist.gov/nistpubs/ir/2023/NIST.IR.8214C.ipd.pdf>, 2023.
12. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001*, pages 136–145. IEEE Computer Society, 2001.
13. R. Canetti, R. Gennaro, S. Goldfeder, N. Makriyannis, and U. Peled. UC non-interactive, proactive, threshold ECDSA with identifiable aborts. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, pages 1769–1787. ACM, 2020.
14. R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Adaptive security for threshold cryptosystems. In M. J. Wiener, editor, *CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999*, volume 1666 of *LNCS*, pages 98–115. Springer, 1999.
15. J. F. Canny and S. Sorkin. Practical large-scale distributed key generation. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 138–152. Springer, 2004.
16. D. Cash, E. Kiltz, and V. Shoup. The twin Diffie–Hellman problem and applications. In N. P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 127–145. Springer, 2008.
17. W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes. CSIDH: An efficient post-quantum commutative group action. Cryptology ePrint Archive, Paper 2018/383, 2018. URL: <https://eprint.iacr.org/2018/383>.
18. W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes. CSIDH: an efficient post-quantum commutative group action. In T. Peyrin and S. D. Galbraith, editors, *ASIACRYPT 2018*, volume 11274 of *LNCS*, pages 395–427. Springer, 2018.
19. D. Connolly, C. Komlo, I. Goldberg, and C. Wood. Two-round threshold Schnorr signatures with FROST, 2022. URL: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-frost/>.
20. I. Damgård and G. L. Mikkelsen. Efficient, robust and constant-round distributed RSA key generation. In D. Micciancio, editor, *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, volume 5978 of *Lecture Notes in Computer Science*, pages 183–200. Springer, 2010.
21. S. Das, T. Yurek, Z. Xiang, A. K. Miller, L. Kokoris-Kogias, and L. Ren. Practical asynchronous distributed key generation. In *43rd IEEE Symposium on Security and Privacy, SP 2022*, pages 2518–2534. IEEE, 2022.
22. A. Davidson, S. Sahib, P. Snyder, and C. Wood. Star: Distributed secret sharing for private threshold aggregation reporting, 2022. URL: <https://datatracker.ietf.org/doc/draft-dss-star>.

23. P. Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science*, pages 427–437. IEEE Computer Society, 1987.
24. M. Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In V. Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 152–168. Springer, 2005.
25. Y. Frankel, P. D. MacKenzie, and M. Yung. Robust efficient distributed rsa-key generation. In J. S. Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 663–672. ACM, 1998.
26. E. S. V. Freire, D. Hofheinz, E. Kiltz, and K. G. Paterson. Non-interactive key exchange. In K. Kurosawa and G. Hanaoka, editors, *Public-Key Cryptography - PKC 2013*, volume 7778 of *LNCS*, pages 254–271. Springer, 2013.
27. R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure applications of Pedersen’s distributed key generation protocol. In M. Joye, editor, *CT-RSA 2003*, volume 2612 of *LNCS*, pages 373–390. Springer, 2003.
28. R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure distributed key generation for discrete-log based cryptosystems. *J. Cryptol.*, 20(1):51–83, 2007.
29. J. Groth. Non-interactive distributed key generation and key resharing. Cryptology ePrint Archive, Paper 2021/339, 2021. URL: <https://eprint.iacr.org/2021/339>.
30. K. Gurkan, P. Jovanovic, M. Maller, S. Meiklejohn, G. Stern, and A. Tomescu. Aggregatable distributed key generation. In A. Canteaut and F. Standaert, editors, *EUROCRYPT 2021*, volume 12696 of *LNCS*, pages 147–176. Springer, 2021.
31. A. Kate, Y. Huang, and I. Goldberg. Distributed key generation in the wild. Cryptology ePrint Archive, Paper 2012/377, 2012. URL: <https://eprint.iacr.org/2012/377>.
32. J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.
33. E. Kokoris-Kogias, D. Malkhi, and A. Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *CCS ’20*, pages 1751–1767. ACM, 2020.
34. C. Komlo and I. Goldberg. FROST: flexible round-optimized Schnorr threshold signatures. In O. Dunkelman, M. J. J. Jr., and C. O’Flynn, editors, *Selected Areas in Cryptography - SAC 2020*, volume 12804 of *LNCS*, pages 34–65. Springer, 2020.
35. E. Kushilevitz, Y. Lindell, and T. Rabin. Information-theoretically secure protocols and security under composition. *SIAM J. Comput.*, 39(5):2090–2112, 2010.
36. League of Entropy. Threshold partially-oblivious PRFs with applications to key management. <https://docs.leep.network/random-beacon/dkg.html>, 2019.
37. Y. Lindell. Simple three-round multiparty Schnorr signing with full simulatability. Cryptology ePrint Archive, Report 2022/374, 2022. <https://ia.cr/2022/374>.
38. Y. Lindell and A. Nof. Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. In D. Lie, M. Mannan, M. Backes, and X. Wang, editors, *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 1837–1854. ACM, 2018.
39. W. Neji, K. B. Sinaoui, and N. B. Rajeb. Distributed key generation protocol with a new complaint management strategy. *Secur. Commun. Networks*, 9(17):4585–4595, 2016.
40. T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In J. Feigenbaum, editor, *CRYPTO ’91*, volume 576 of *LNCS*, pages 129–140. Springer, 1991.
41. T. P. Pedersen. A threshold cryptosystem without a trusted party (extended abstract). In D. W. Davies, editor, *EUROCRYPT 1991*, volume 547 of *LNCS*, pages 522–526. Springer, 1991.
42. A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.

SimShare($\Delta, n, t, \text{corrupt}$)
1 : $A_0 = \Delta$
2 : for $j \in \text{corrupt}$ do
3 : $w_j \xleftarrow{\$} \mathbb{Z}_q$
4 : for $k \in \{1, \dots, t-1\}$ do
5 : $A_k = A_0^{L'_{0,k}} \cdot \prod_{j \in \text{corrupt}} g^{L'_{j,k} \cdot w_j}$
6 : $D = \langle A_0, \dots, A_{(t-1)} \rangle$
7 : return $(\{(j, w_j)\}_{j \in \text{corrupt}}, D)$

Fig. 12. Algorithm to simulate secret sharing in a Feldman VSS to an adversary controlling $t - 1$ participants denoted by the set corrupt with respect to a challenge Δ . Here, $L'_{i,k}$ denotes the k^{th} coefficient of the Lagrange polynomial $L_i(x)$, which is computed with respect to the set $\{0\} \cup \text{corrupt}$.

A Proofs for Concrete NIKE

In Section 2.2, we introduce a concrete NIKE scheme. We now demonstrate that it is unrecoverable and binding.

Theorem 4. *The concrete NIKE of Section 2.2 is unrecoverable against a probabilistic polynomial-time adversary \mathcal{A} , assuming the Computational Diffie–Hellman (CDH) assumption holds in \mathbb{G} . More specifically, for any PPT adversary \mathcal{A} that attacks NK, there exists an adversary \mathcal{B} with approximately the same running time as \mathcal{A} such that*

$$\text{Adv}_{\text{NK}, \mathcal{A}}^{\text{rec}}(\lambda) \leq \text{Adv}_{\mathbb{B}}^{\text{cdh}}(\lambda)$$

Because breaking session key recovery in our concrete scheme simply requires an adversary to guess a valid Diffie–Hellman shared secret knowing only two public keys, we omit the proof for Theorem 4 as it is a direct reduction to CDH.

Theorem 5. *The concrete NIKE of Section 2.2 is binding against a computationally unbounded adversary \mathcal{A} .*

Proof. We now show that any adversary \mathcal{A} cannot win the binding experiment Exp^{bind} in Figure 2 for the concrete NIKE of Section 2.2. \mathcal{A} wins when it can produce some $(\text{sk}_1^*, \text{pk}_1^*, \text{sk}_2^*, \text{pk}_2^*)$ such that:

1. $\text{NK.Verify}(\text{sk}_1^*, \text{pk}_1^*)$ outputs 1, and
2. $\text{NK.Verify}(\text{sk}_2^*, \text{pk}_2^*)$ outputs 1, but
3. $\text{NK.SharedKey}(\text{sk}_1^*, \text{pk}_2^*) \neq \text{NK.SharedKey}(\text{sk}_2^*, \text{pk}_1^*)$

However, in the setting of the concrete scheme, this means that $\text{pk}_1^* = g^{\text{sk}_1^*}$ and $\text{pk}_2^* = g^{\text{sk}_2^*}$, yet $(\text{pk}_2^*)^{\text{sk}_1^*} \neq (\text{pk}_1^*)^{\text{sk}_2^*}$, which is impossible.

B Proofs for Feldman’s VSS

In Section 3.1, we define Feldman’s VSS. We now discuss its security.

B.1 Secrecy

Theorem 6. *Feldman’s VSS is perfectly secret.*

Proof. We prove Theorem 6 by showing that the adversary cannot distinguish Game 0 (which corresponds to the VSS secrecy experiment from Figure 3 when $b = 0$ and the protocol is honestly executed) from Game 1 (Figure 3 when $b = 1$ and the protocol is simulated).

Game 0. This is simply the secrecy experiment as shown in Figure 3, instantiated with Feldman’s VSS, and with $b = 0$.

Game 1. This is the game shown in Figure 3, when $b = 1$. Recall that when $b = 1$, the environment simulates secret sharing for $\Delta \in \mathbb{G}$ whose discrete logarithm is unknown, employing an algorithm `SimShare`.

We show `SimShare` for our concrete construction in Figure 12. To begin, the algorithm picks $t - 1$ shares w_j at random. This number of shares is sufficient for the simulation, as we assume $|\text{corrupt}| = t - 1$, without loss of generality. `SimShare` then performs polynomial interpolation “in the exponent” using Δ and the $t - 1$ shares, to derive $A_k, k \in [t - 1]$. The commitment to f is then set as $D = \langle A_0, \dots, A_{t-1} \rangle$, where $A_0 = \Delta$.

Difference between Game 0 and Game 1. We now show that \mathcal{A} has no additional advantage in distinguishing the games.

The simulation is perfect because D in Game 1 simply defines a random polynomial f' “in the exponent”, whose constant term $f'(0)$ (the secret) is the discrete logarithm of Δ . Hence, the output of `SimShare` is indistinguishable to the adversary from a real output of `S.Share`.

Finishing the proof. Because \mathcal{A} has identical advantage in Game 0 and Game 1 (i.e., the simulation is perfect), the advantage that \mathcal{A} wins the VSS secrecy game is then zero. This concludes the proof. \square

B.2 Uniqueness

We now give a proof sketch for the uniqueness of Feldman’s VSS; we omit the full proof as it is a well-known result in the literature [32].

Feldman’s VSS is information-theoretically unique because Shamir’s secret sharing is information-theoretically unique. Put differently, a polynomial f is uniquely defined by t points. Hence, given some threshold $t^* \geq 1$ chosen by the adversary playing against `Expunq` in Figure 4, the adversary cannot pick two different recovery sets M_1^*, M_2^* and a single D^* where each share in M_1^* and M_2^* is valid with respect to t^* and D^* , but `S.Recover`(t^*, M_1^*) \neq `S.Recover`(t^*, M_2^*).

C Proofs for the Concrete AgVSS

In Section 3.2, we define a concrete AgVSS which builds upon Feldman’s VSS. We demonstrate the correctness, aggregated secrecy, and uniqueness of that scheme now.

C.1 Correctness

The concrete AgVSS scheme of Section 3.2 fulfills the notion of correctness for the base VSS scheme, as the algorithms `Share`, `Recover`, `Verify`, `GetPub` are identical to Feldman’s VSS, as is the map `SecretToPublic`. We now demonstrate that the concrete AgVSS scheme fulfills the additional notion of correctness for an AgVSS scheme.

Theorem 7. *The concrete AgVSS scheme of Section 3.2 is correct in the sense of Figure 5.*

Proof. We analyze each case where a return value of 0 could occur in Figure 5, and demonstrate that that event will not occur.

Case 1. We first show that for each $i \in [n]$, `AV.Verify`(i, \hat{w}_i, C) = 1. Recall that \hat{w}_i is simply

$$\hat{w}_i = v + \sum_{j \in [l]} w_{ji}$$

by Equation 3. Because each w_{ji} was generated by the `AV.Share` algorithm, then $w_{ji} = f_j(i)$ for some polynomials f_1, \dots, f_ℓ each of degree $t - 1$. Let $f' = v + \sum_{i \in [\ell]} f_i$. Then (i, \hat{w}_i) is a valid point on f' , because

$$\hat{w}_i = v + \sum_{j \in [\ell]} w_{ji} = v + \sum_{j \in [\ell]} f_j(i) = f'(i)$$

Because C is a commitment to the same f' as shown in Equations 4 and 5, then `AV.Verify`(i, \hat{w}_i, C) will output 1. Note also that `AV.Recover` will output $\hat{s} = f'(0) = v + \sum_{j \in [\ell]} s_j$.

Case 2. We next prove that `AV.Verify`($0, \hat{s}, C$) = 1. As above, after performing `AV.AggPub`, the resulting aggregated commitment C can be represented as a commitment to the polynomial f' . `AV.Verify`($0, \hat{s}, C$) checks that $g^{\hat{s}} = \hat{A}_0$ (from Eq. 2 with $i = 0$). From Eq. 4, $\hat{A}_0 = g^v \cdot \prod_{j \in [\ell]} D_j[0] = g^v \cdot \prod_{j \in [\ell]} g^{f_j(0)} = g^v \cdot \prod_{j \in [\ell]} g^{s_j}$. Since $\hat{s} = v + \sum_{j \in [\ell]} s_j$ as above, `AV.Verify`($0, \hat{s}, C$) will output 1.

This concludes the proof. □

C.2 Aggregated Secrecy

We now show that the concrete scheme of Section 3.2 fulfills the notion of aggregated secrecy from Figure 6. To do so, we first show how the environment simulates `SimShare` and `H`. Without loss of generality, we assume $|\text{corrupt}| = t - 1$.

Lemma 1. *When playing the secure aggregation game as in Figure 6 against the concrete scheme, the output of $\mathcal{O}^{\text{GetShare}}$ when $b = 0$ (i.e., when the environment honestly performs `AV.Share`) is indistinguishable to an adversary from the output when $b = 1$ (i.e., when the environment performs `SimShare` as shown in Figure 13), assuming the adversary controls no more than $|\text{corrupt}| = t - 1$ participants.*

Proof. `SimShare` accepts a challenge $\Delta \in \mathbb{G}$, and outputs a randomly sampled secret α , $|\text{corrupt}|$ shares, and an `AgVSS` commitment. To begin, it picks $\alpha \in \mathbb{Z}_q$ at random, and then generates a blinded commitment A_0 to the (unknown) discrete logarithm of Δ using α as the blinding factor, by Equation 10.

$$A_0 = \Delta \cdot g^\alpha \tag{10}$$

`SimShare` then picks $t - 1$ shares for the corrupted parties at random, resulting in the list of shares $\langle w_j \rangle_{j \in \text{corrupt}} \xleftarrow{\$} \mathbb{Z}_q^{(t-1)}$ and then generates commitments to coefficients $A_1, \dots, A_{(t-1)}$ with respect to these shares and A_0 , by Equation 11.

$$A_k = A_0^{L'_{0,k}} \cdot \prod_{j \in \text{corrupt}} g^{L'_{j,k} \cdot w_j} \tag{11}$$

Here, $L'_{i,k}$ denotes the k^{th} coefficient of the Lagrange polynomial $L'_i(x)$, which is computed with respect to the set $\{0\} \cup \text{corrupt}$. The commitment D to the $|\text{corrupt}|$ shares is then $D = \langle A_0, \dots, A_{(t-1)} \rangle$. The simulation is perfect because D simply is committing to a random polynomial f “in the exponent”, whose constant term is the discrete logarithm of A_0 , (where A_0 is a blinded commitment to the discrete logarithm of Δ) and with $t - 1$ random points that are the shares distributed to the adversary. Each additional coefficient $a_k, k \in [t - 1]$ of f is committed to “in the exponent” via A_k , which can be determined via Equation 11. Hence, so long as the adversary controls fewer than $t - 1$ players, `AV.Verify`(i, w_i, D) for all $i \in \text{corrupt}$ will output 1. □

Lemma 2. *The adversary has negligible advantage in distinguishing the environment’s simulation of `GetTweak` in the concrete scheme of Section 3.2 in the (programmable) random oracle model, assuming an honest majority.*

Proof. Recall that `GetTweak` is instantiated by `H` in the concrete scheme. For any query (O', aux') where $\text{aux}'_i \neq \text{aux}$, the environment simply programs `H` to return a random value. However, when it receives a query for (O, aux) , it programs `H` to return v as in Equation 12. Recall that aux is provided to the adversary

SimShare(λ , corrupt, Δ)	H(O_i , aux $_i$)
1 : $\alpha \xleftarrow{\$} \mathbb{Z}_q$; $A_0 = \Delta \cdot g^\alpha$	1 : // Q_1, Q_2, Q_3 are defined
2 : for $j \in \text{corrupt}$ do	2 : // by the experiment
3 : $w_j \xleftarrow{\$} \mathbb{Z}_q$	3 : if $Q_3[O_i, \text{aux}_i] \neq \perp$
4 : for $k \in \{1, \dots, t-1\}$ do	4 : return $Q_3[O_i, \text{aux}_i]$
5 : $A_k = A_0^{L'_{0,k}} \cdot \prod_{j \in \text{corrupt}} g^{L'_{j,k} \cdot w_j}$	5 : $D^* \leftarrow Q_1$
6 : $D = \langle A_0, \dots, A_{(t-1)} \rangle$	6 : $\{D_j\}_{j \in [\ell']} \leftarrow Q_2$
7 : return $(\alpha, \{(j, w_j)\}_{j \in \text{corrupt}}, D)$	7 : $O = \{D_j\}_{j \in [\ell']} \cup D^*$
	8 : if aux $_i \neq \text{aux}$ or $O_i \neq O$
	9 : // Respond honestly if this is
	10 : // not the challenge query
	11 : $Q_3[O_i, \text{aux}_i] \xleftarrow{\$} \mathbb{Z}_q$
	12 : return $Q_3[O_i, \text{aux}_i]$
	13 : $\alpha \leftarrow Q_1$
	14 : for $\{(j, w_{kj})\}_{j \in \text{honest}, k \in [\ell']} \in Q_2$
	15 : in $\leftarrow (t, \{(j, w_{kj})\}_{j \in \text{honest}})$
	16 : $s_k \leftarrow \text{AV}[\text{H}].\text{Recover}(\text{in})$
	17 : $v \leftarrow (-1) \cdot (\alpha + \sum_{i \in [j]} s_k)$
	18 : $Q_3[O_i, \text{aux}_i] \leftarrow v$
	19 : return v

Fig. 13. Algorithms to simulate secret sharing and tweak generation in an AgVSS to an adversary controlling $t - 1$ participants denoted by the set **corrupt** with respect to a challenge Δ and at least t honest players denoted by the set **honest**, where **corrupt** \cup **honest** = $[n]$ and **corrupt** \cap **honest** = \emptyset . Here, $L'_{i,k}$ denotes the k^{th} coefficient of the Lagrange polynomial $L'_i(x)$, which is computed with respect to the set $\{0\} \cup \text{corrupt}$.

only after it has finished querying $\mathcal{O}^{\text{GetShare}}$ and $\mathcal{O}^{\text{RecvShare}}$, and so similarly O is fixed at the time that the adversary can query with **aux** as input (i.e, the environment will never have to guess which query to program).

$$v \leftarrow -\alpha - \sum_{j \in [\ell']} s_j \quad (12)$$

There are three requirements that must be satisfied for the environment to program **H** correctly.

1. The programmed output v must be indistinguishable from a random value.
2. The environment must be able to derive $\{s_j\}_{j \in [\ell']}$ without the involvement or detection of \mathcal{A} .
3. The environment can correctly program **H** strictly before \mathcal{A} can query on inputs (O, aux) .

We first demonstrate that the output v from Equation 12 is indistinguishable from any other output from **H**. Recall that α is chosen at random when **SimShare** is performed, and published only within the commitment A_0 as shown in Equation 10, serving as a blinding factor for Δ . Because Δ is chosen at uniformly at random from the public domain, and because α is chosen at random, the v value in Equation 12 is indistinguishable from a random element to \mathcal{A} . Moreover, even if the adversary submitted inputs to $\mathcal{O}^{\text{RecvShare}}$ in the attempt to cancel out the outputs from $\mathcal{O}^{\text{GetShare}}$, the adversary would be unable to do so, because the adversary sees only the blinded commitment A_0 as shown in Equation 10, and because Δ is chosen uniformly at random. Therefore, programming **H** at exactly one point with v is indistinguishable to any other point that is honestly programmed with a random value chosen from \mathbb{Z}_q .

Second, the environment can derive $\{s_j\}_{j \in [\ell']}$ without the involvement or detection of \mathcal{A} , because we operate in the honest majority setting, and assume that the environment simulates at least t honest players. The environment can therefore reconstruct the adversary's s_j values (from Q_2), and it knows its own (from both Q_1 and Q_2).

Finally, the adversary does not learn \mathbf{aux} until *after* it has completed its queries to $\mathcal{O}^{\text{GetShare}}$ and $\mathcal{O}^{\text{RecvShare}}$. Before returning \mathbf{aux} to \mathcal{A} , the environment programs \mathbf{H} with v on inputs (O, \mathbf{aux}) . Because $\mathbf{AX} = \mathbb{Z}_q$ and \mathbf{aux} is some uniform element from \mathbf{AX} , then \mathcal{A} had negligible probability in learning the output \mathbf{H} with input \mathbf{aux} until the environment has programmed it. The environment is able to correctly program on these inputs using v because it is guaranteed that \mathcal{A} cannot further influence O . In other words, there is no danger that the environment must make a decision on which input to program using v .

This completes the proof. \square

Theorem 8. *The concrete scheme of Section 3.2 fulfills aggregated secrecy against an adversary \mathcal{A} as in Figure 6, assuming an honest majority.*

Proof. Recall that in the aggregated secrecy game, \mathcal{A} has three opportunities to distinguish the $b = 0$ game from the $b = 1$ game:

1. by how the environment responds to $\mathcal{O}^{\text{GetShare}}$,
2. by how the environment responds to $\mathcal{O}^{\text{GetTweak}}$, and
3. whether or not $\text{AV.GetPub}(0, C)$ evaluates to the challenge Δ when $b = 1$.

For (1), by Lemma 1, $\mathcal{O}^{\text{GetShare}}$ is indistinguishable to the adversary when $b = 0$ (i.e., shares are honestly generated) and when $b = 1$ (i.e., SimShare is performed). For (2), by Lemma 2, the environment can simulate \mathbf{H} in a way that is indistinguishable to the adversary in the $b = 0$ case (the output is chosen at random) and the $b = 1$ case (the output is programmed). We complete the proof by showing that the adversary has zero advantage for (3).

Recall that the concrete AgVSS is zero-indexed, per Definition 5. By Lemma 1, the environment can embed Δ in the first query of the adversary to $\mathcal{O}^{\text{GetShare}}$ in a way that is indistinguishable to \mathcal{A} . \mathcal{A} must query $\mathcal{O}^{\text{GetShare}}$ at least once (i.e., if $\text{issuedchal} = 0$ at the end of the experiment, then the experiment returns 0). Hence, it is guaranteed that the public recovery set $\{D_j\}_{j \in [\ell]}$ that is input to AV.AggPub will include a contribution with respect to the discrete logarithm of Δ .

Equation 13 demonstrates that the commitment to the constant term of $C = \{\hat{A}_0, \dots, \hat{A}_{t-1}\}$ evaluates to Δ ; and so, because the concrete AgVSS is zero-indexed, then $\text{AV.GetPub}(0, C)$ evaluates to Δ .

$$\begin{aligned}
\hat{A}_0 &= g^v \cdot \prod_{j \in [\ell]} D_j[0] \\
&= g^v \cdot \Delta \cdot g^\alpha \cdot g^{\sum_{i \in [\ell']} s_i} \\
&= g^{-\alpha - \sum_{i \in [\ell']} s_i} \cdot \Delta \cdot g^\alpha \cdot g^{\sum_{i \in [\ell']} s_i} \\
&= \Delta
\end{aligned} \tag{13}$$

Hence, the adversary gains no advantage in its guess of b' , regardless of whether it is operating in the real ($b = 0$) or simulated ($b = 1$) setting. This concludes the proof. \square

C.3 Uniqueness

The uniqueness of the concrete AgVSS also follows from that of Feldman's VSS, as we now show.

Theorem 9. *The concrete AgVSS scheme is unique.*

Proof. We now show that the additional aggregation algorithm in our AgVSS construction gives the adversary no additional power in the uniqueness game than in the plain Feldman’s VSS setting.

Recall the VSS uniqueness experiment in Figure 4. The adversary is required to output a valid tuple (M_1^*, M_2^*, D^*) , for which $S.Recover(M_1^*) \neq S.Recover(M_2^*)$. As above, we know that the adversary cannot win the uniqueness game in the Feldman’s VSS setting, no matter what strategy it employs, including using the AgVSS aggregation algorithm, since the aggregation algorithm uses no information hidden from the adversary. Therefore the adversary cannot win the uniqueness game in the AgVSS setting either. \square

D Proofs for Generic Construction

D.1 Robustness

Proof. We prove Theorem 1 by a sequence of games.

Game 0. This is the DKG weak robustness game as in Figure 7, applied to construction in Figure 11.

Let W_0 be the event that \mathcal{A} wins in Game 0. Then,

$$\text{Adv}_{\mathcal{A}, \mathcal{D}}^{\text{wk-rbst}}(\lambda) = \Pr[W_0]$$

Let \mathcal{C} be the algorithm that simulates the robustness game to \mathcal{A} . We now describe how \mathcal{C} simulates the game.

Setup. To begin, \mathcal{C} initializes $Q_1 \leftarrow \emptyset$ to simulate H_1 and $Q_2 \leftarrow \emptyset$ to simulate H_2 , in addition to initializing each honest party as shown in Figure 7. \mathcal{C} handles \mathcal{A} ’s random oracle queries by lazy sampling, as follow:

H_1 : When the adversary queries H_1 on inputs (O_i, aux_i) , the environment checks if it is in Q_1 , and if so, returns the corresponding v_i . Otherwise, the environment randomly samples v_i from its respective domain, sets $Q_1[(O_i, \text{aux}_i)] = v_i$, and then returns v_i .

H_2 : When the adversary queries H_2 on inputs $\{\psi_1, \dots, \psi_n\}$, the environment checks if it is in Q_2 , and if so, returns the corresponding aux_i . Otherwise, the environment randomly samples aux_i from its respective domain, sets $Q_2[\{\psi_1, \dots, \psi_n\}] = \text{aux}_i$, and then returns aux_i .

\mathcal{C} performs all signing oracle queries honestly.

Game 1. The only difference in Game 1 is if \mathcal{A} outputs a tuple $(B_j, \text{pk}_j, D_j, \beta_j)$, $j \in \text{corrupt}$ where (B_j, pk_j, D_j) are output in PerformRound_0 and β_j is output in PerformRound_2 such that Equation 14 holds, then \mathcal{C} aborts.

$$\begin{aligned} \text{AV}[H_1].\text{GetPub}(0, D_j) &= \text{pk}, \text{ and} \\ \text{NK}.\text{Verify}(\beta_j, B_j) &\rightarrow 1, \text{ but} \\ \text{NK}.\text{SharedKey}(\beta_j, \text{pk}_j) &\neq \text{NK}.\text{SharedKey}(\alpha_j, B_j) \end{aligned} \tag{14}$$

However, if the NIKE is binding and Assumption 1 holds, then Game 1 and Game 0 are indistinguishable to \mathcal{A} .

Reduction to NIKE Binding. Let \mathcal{B}_1 be an adversary playing against the NIKE binding game. \mathcal{B}_1 simulates Game 1 to \mathcal{A} in exactly the same way as Game 0. However, in Finalize , if \mathcal{A} outputs some β_j with respect to the commitments (pk_j, B_j, D_j) output in PerformRound_0 such that Equation 14 holds, then \mathcal{B}_1 recovers α_j , and outputs $(\alpha_j, \text{pk}_j, \beta_j, B_j)$ as its output to the NIKE binding game.

\mathcal{B}_1 recovers α_j by performing $\alpha_j \leftarrow \text{AV}[H_1].\text{Recover}(t, M_j)$, where $M_j = \{(k, w_{jk})\}_{k \in \text{honest}}$. \mathcal{B}_1 can recover s_j because the checks in PerformRound_1 completed successfully (else the honest participants would have aborted the protocol in PerformRound_1 and PerformRound_2), then each share in M_j is a valid share of the secret α_j committed to by D_j and pk_j . Further, because \mathcal{B}_1 simulates at least t honest players, it has a sufficient number of shares to perform this recovery step.

Hence, when \mathcal{A} distinguishes Game 1 from Game 0 (i.e., by causing Game 1 to abort), then \mathcal{B}_1 wins.

Difference between Game 0 and Game 1. Let W_1 be the event that \mathcal{A} wins in Game 1. Then,

$$|\Pr[W_1] - \Pr[W_0]| \leq \text{Adv}_{\text{NK}, \mathcal{B}_1}^{\text{bind}}(\lambda) \quad (15)$$

Game 2. The only difference in Game 2 is that in `Finalize`, if Equation 16 holds, then \mathcal{C} aborts.

$$\begin{aligned} \text{sk} &\leftarrow \text{AV.Recover}(t, \{(i, \text{sk}_i)\}_{i \in \text{honest}}), \text{ but} \\ \text{NK.Verify}(\text{sk}, \text{pk}) &\rightarrow 0 \end{aligned} \quad (16)$$

However, if the AgVSS is unique, then Game 1 and Game 2 are indistinguishable.

Reduction to AgVSS Uniqueness. Let \mathcal{B}_2 be an adversary playing against the AgVSS uniqueness game.

\mathcal{B}_2 follows the protocol honestly for all parties $i \in \text{honest}$. However, in `PerformRound3`, \mathcal{B}_2 recovers sk and tests to see if Equation 16 holds. \mathcal{B}_2 can recover sk without the involvement of \mathcal{A} because it simulates at least t honest parties.

Recall that $\text{sk}_i \leftarrow \text{AV}[\text{H}_1].\text{AggPriv}(i, \{w_{i,j}\}_{j \in [n]}, \{D_j\}_{j \in [n]}, \text{aux})$, and $C \leftarrow \text{AV}[\text{H}_1].\text{AggPub}(\{D_j\}_{j \in [n]}, \text{aux})$. The only values that differ between each honest party when performing this aggregation step are the sets $\{w_{i,j}\}_{j \in [n]}$.

If Equation 16 holds, then this means that at least one honest party $k \in \text{honest}$ received a share $w_{j,k}$ from a corrupted party $j \in \text{corrupt}$ such that $\text{AV}[\text{H}_1].\text{Verify}(k, w_{j,k}, D_j) = 1$ but where Equation 17 holds.

$$\begin{aligned} \text{AV}[\text{H}_1].\text{Recover}(t, \{(\ell, w_{j,\ell})\}_{\ell \in \text{honest}}) &\neq \\ \text{AV}[\text{H}_1].\text{Recover}(t, \{(\ell, w_{j,\ell})\}_{\ell \in \text{honest}, \ell \neq k}) & \end{aligned} \quad (17)$$

Let $M_1 = \{(\ell, w_{j,\ell})\}_{\ell \in \text{honest}}$ and $M_2 = \{(\ell, w_{j,\ell})\}_{\ell \in \text{honest}, \ell \neq k}$. When this event has occurred, \mathcal{B}_2 submits (t, M_1, M_2, D_j) as its output to the AgVSS uniqueness game. When \mathcal{A} can distinguish Game 2 from Game 1, then \mathcal{B}_2 wins the AgVSS uniqueness game.

Difference between Game 1 and Game 2. Let W_2 be the event that \mathcal{A} wins in Game 2. Then,

$$|\Pr[W_2] - \Pr[W_1]| \leq \text{Adv}_{\text{AV}[\text{H}_1], \mathcal{B}_2}^{\text{unq}}(\lambda) \quad (18)$$

Finishing the Proof. Recall that in the weak robustness game, the adversary has four opportunities to win, assuming all honest parties have not aborted the protocol. First, if any honest party has a disjoint view of pk , second, if any honest party has a disjoint view of qual , and third, if any honest party has a disjoint view of aux . Finally, the adversary wins if the sk that is recovered from the honest parties' shares is invalid.

Because we assume a broadcast channel, if any honest party aborts in either `PerformRound1` or `PerformRound2`, then all honest parties will abort in `PerformRound2`. And so if any honest party proceeds to `Finalize`, then all honest parties will proceed to `Finalize`, and all honest parties will set their status to `accept`. While honest parties might abort (meaning that strong robustness cannot be fulfilled), honest parties will consistently either end in an aborted or accepting status, meaning that the weak robustness game can be fulfilled. And so the only opportunity for the adversary to win against the generic construction in the weak robustness game is to disrupt consistency of pk , qual , aux , or to make $\text{D}[\text{H}].\text{Recover}$ output fail.

Again, because we assume a broadcast channel, all parties will maintain the same view of $\{(\beta_j, B_j, \text{pk}_j, D_j)\}$. And so all parties will determine the set qual in `Finalize` using the same information. Hence, the adversary has no advantage in forcing a disjoint qual .

From Game 1, we know that all honest parties derive the same view of $\{\psi_j\}$ such that $\text{NK.SharedKey}(\beta_j, \text{pk}_j) \rightarrow \psi_j$. Because the tweak value v is derived from H_2 using the values $\{\psi_j\}_{j \in [n]}$, then all parties will obtain the same v . Because pk is derived using v and the set $\{D_j\}_{j \in [n]}$, then the adversary has no advantage in forcing a disjoint pk . Similarly, the adversary has no advantage in forcing a disjoint $\text{aux} = \{\text{pk}_i\}_{i \in \text{qual}}$, which are also derived from v and the set $\{D_j\}_{j \in [n]}$.

Finally, from Game 2, we know that $\text{D}[\text{H}].\text{Recover}$ will output a sk that is valid with respect to pk .

The adversary in Game 2 then has no advantage in winning the DKG weak robustness game. Concretely, combining the advantages in Equations 15 and 18 gives Equation 7. This completes the proof. \square

D.2 Zero-Knowledge

Proof. We now complete the proof of Theorem 2 via a sequence of games.

Game 0. This is the DKG zero-knowledge game when $b = 0$ as in Figure 9, applied to the construction in Figure 11.

Let W_0 be the event that \mathcal{A} outputs 1 in Game 0. Then,

$$\text{Adv}_{\mathcal{D}, \mathcal{A}, \text{SimRound}_i, \text{SimFinalize}}^{\text{zk}}(\lambda) = \Pr[W_0]$$

Let \mathcal{C} be the algorithm that simulates the DKG zero-knowledge game to \mathcal{A} . We now describe in further detail how \mathcal{C} simulates the game.

Setup. To begin, \mathcal{C} initializes $Q_1 \leftarrow \emptyset$ to simulate H_1 and $Q_2 \leftarrow \emptyset$ to simulate H_2 , in addition to initializing each honest party as shown in Figure 9. \mathcal{C} handles \mathcal{A} 's random oracle queries by lazy sampling, as follows:

H_1 : When the adversary queries H_1 on inputs (O_i, aux_i) , the environment checks if it is in Q_1 , and if so, returns the corresponding v_i . Otherwise, the environment randomly samples v_i from its respective domain, sets $Q_1[(O_i, \text{aux}_i)] = v_i$, and then returns v_i .

H_2 : When the adversary queries H_2 on inputs $\{\psi_1, \dots, \psi_n\}$, the environment checks if it is in Q_2 , and if so, returns the corresponding aux_i . Otherwise, the environment randomly samples aux_i from its respective domain, sets $Q_2[\{\psi_1, \dots, \psi_n\}] = \text{aux}_i$, and then returns aux_i .

\mathcal{C} performs all signing oracle queries honestly.

Game 1. The only difference between Game 0 and Game 1 is that in Game 1, \mathcal{C} performs $\text{SimRound}_0, \dots, \text{SimRound}_{\text{numrounds}-1}, \text{SimFinalize}$ for participant $\tau \stackrel{\$}{\leftarrow} \text{honest}$ with respect to the challenge $\hat{\Delta}$. More specifically, when performing SimRound_0 , \mathcal{C} employs $(\{w_{\tau,j}^*\}_{j \in \text{corrupt}}, D_\tau^*) \leftarrow \text{AV}[H_1].\text{SimShare}(\lambda, n, t, \tau, \hat{\Delta})$ for an honest participant τ sampled from the set of honest participants at random, such that Equation 19 holds:

$$\text{S.GetPub}(0, D_\tau^*) = \hat{\Delta} \quad (19)$$

\mathcal{C} performs $\text{SimRound}_1, \text{SimRound}_2, \text{SimFinalize}$ as in the real protocol. However, one consequence of \mathcal{C} simulating secret sharing with respect to $\hat{\Delta}$ is that \mathcal{C} cannot correctly derive each honest players' sk_i , for all $i \in \text{honest}$ when performing SimFinalize . This however does not impact the ability of \mathcal{C} to perfectly simulate all honest players, because \mathcal{C} is never required to perform operations with respect to $\hat{\Delta}$ directly.

If the AgVSS is secret, then Game 0 and Game 1 are indistinguishable to \mathcal{A} .

Reduction to AgVSS (Non-Aggregated) Secrecy. We will construct a reduction \mathcal{B}_1 which is an adversary against the AgVSS (non-aggregated) secrecy game as defined in Definition 6, such that \mathcal{B}_1 simulates either Game 0 or Game 1 to \mathcal{A} . When the AgVSS hidden bit $b = 0$, \mathcal{B}_1 ends up simulating Game 0; when the AgVSS hidden bit $b = 1$, \mathcal{B}_1 ends up simulating Game 1.

\mathcal{B}_1 is constructed similar to Game 0, with the following exception. \mathcal{B}_1 simply performs PerformRound_1 for each honest non-simulated party $i \in \text{honest}, i \neq \tau$. However, for the simulated honest party τ , instead of following the protocol honestly, \mathcal{B}_1 instead employs the challenge received as input $(\{(j, w_j)\}_{j \in \text{corrupt}}, D)$ from the AgVSS secrecy experiment. In particular, it sets $D_\tau \leftarrow D$, and $(\{(j, w_{\tau,j})\}_{j \in \text{corrupt}} \leftarrow (\{(j, w_j)\}_{j \in \text{corrupt}})$. It then follows the remainder of the protocol honestly. While it will not be able to derive the honest player's secret keys in PerformRound_3 , this will not have an impact on its ability to correctly simulate.

The simulation of Game 0 by \mathcal{B}_1 is identical to when $b = 0$ in the AgVSS secrecy game. Hence, the only distinction between Game 0 and Game 1 is when $b = 1$ in the AgVSS secrecy game.

When \mathcal{A} outputs b' as its guess, then \mathcal{B}_1 outputs b' . Hence, when \mathcal{A} distinguishes Game 0 from Game 1, \mathcal{B}_1 wins the AgVSS secrecy game.

Difference between Game 0 and Game 1. If the AgVSS is secret, then the difference between Game 0 and Game 1 is negligible. Let W_1 be the event that \mathcal{A} wins in Game 1. Then

$$|\Pr[W_1] - \Pr[W_0]| \leq \text{Adv}_{\text{AV}, \mathcal{B}_1}^{\text{sec}}(\lambda) \quad (20)$$

```

PerformRound2(statei, inp, inb)
1: for  $j \in [n], j \neq i$  do
2:   if AV[H].Verify( $i, w_{ji}, D_j$ )  $\neq 1$ 
3:   or if AV[H].GetPub( $0, D_j$ )  $\neq \text{pk}_j$ 
4:     statei.status = abort
5:   return (statei,  $\perp$ , fail)
6: statei  $\leftarrow$  statei  $\cup \{(i, w_{ji})\}_{j \in [n]}$ 
7:  $M_k = \{(j, w_{kj})\}_{k \in \text{corrupt}, j \in \text{honest}}$ 
8:  $\alpha_k \leftarrow$  AV.Recover[H]( $t, M_k, D_k$ )
9: return  $\perp$  if  $\alpha_k = \text{fail}$ 
10:  $\psi_k \leftarrow$  NK.SharedKey[H]( $\alpha_k, B_k$ ),  $\forall i \in \text{corrupt}$ 
11:  $O = \{D_j\}_{j \in [n]}$ 
12:  $\text{aux} \leftarrow \{\psi_j\}_{j \in [n]}$ 
13: return (statei,  $\perp$ , accept)

```

Fig. 14. Game 2: Code movement, highlighting lines moved from PerformRound₃ to PerformRound₂

Game 2. The difference between Game 1 and Game 2 is code movement from Finalize to PerformRound₁, as shown in Figure 14. Rather than learning corrupted players' blinding factors as input to PerformRound₃, \mathcal{C} instead derives blinding factors for all corrupted parties in PerformRound₂, (importantly, before it reveals the honest players' blinding factors to \mathcal{A}) using the shares from all honest players. If \mathcal{C} is unable to recover these blinding factors, then it returns \perp .

However, because \mathcal{C} simulates at minimum t honest players (honest majority model), it can perform AV.Recover without the involvement of corrupted parties.

Difference between Game 1 and Game 2. Let W_2 be the event that \mathcal{A} wins in Game 2. Then,

$$|\Pr[W_2]| = |\Pr[W_1]| \quad (21)$$

Game 3. The only difference between Game 3 and Game 2 is that if the adversary queries the random oracle H_1 with respect to any of the honest parties' blinding values $\{\psi_i\}_{i \in \text{honest}}$ before the adversary is provided with the witness values $\{\beta_i\}_{i \in \text{honest}}$ as input to PerformRound₃, then the game aborts.

If the NIKE is session-key unrecoverable, the additional advantage to \mathcal{A} is negligible.

Reduction to NIKE Unrecoverability. We will construct a reduction \mathcal{B}_3 which is an adversary against the NIKE unrecoverability game, such that \mathcal{B}_3 simulates Game 3 to \mathcal{A} .

\mathcal{B}_3 begins by receiving $(\text{pk}_1, \text{pk}_2)$ in the NIKE unrecoverability game. It then simulates Game 3 to \mathcal{A} in exactly the same manner as in Game 2, with the following exception. \mathcal{B}_3 simply performs PerformRound₀ honestly for each honest party $i \in \text{honest}, i \neq \tau$. However, for honest party τ , \mathcal{B}_4 instead simulates the protocol via SimRound; it sets $\text{pk}_\tau \leftarrow \text{pk}_1$ and $B_\tau \leftarrow \text{pk}_2$. \mathcal{B}_3 then simulates secret sharing via SimShare with respect to pk_τ . From Game 1, we know that doing so is indistinguishable to \mathcal{A} . \mathcal{B}_3 then follows the remainder of PerformRound₁ as for all the other honest parties.

When simulating $\mathcal{O}_2^{\text{PerformRound}_\tau}$, \mathcal{B}_3 acts honestly for all honest parties $i \in \text{honest}, i \neq \tau$. However, for participant τ , \mathcal{B}_4 does the following. Because it does not know the corresponding β_c to B_c , it instead must guess. It looks at all of the queries made to H_1 that \mathcal{A} has made up to that point, and randomly selects one of the inputs. Let this randomly selected input be $(\{\psi_j\}_{j \in [n]}^*, \text{aux}^*)$.

\mathcal{B}_3 then picks the simulated party's element $\psi_\tau \leftarrow \{\psi_j\}_{j \in [n]}^*$, and outputs $\psi' = \psi_\tau$ as its guess for the NIKE unrecoverability game. If \mathcal{A} queries H_1 with ψ_τ before \mathcal{B}_3 has completed $\mathcal{O}_2^{\text{PerformRound}_r}$, then \mathcal{B}_3 wins the NIKE game with probability $\frac{1}{q}$.

Difference between Game 2 and Game 3. Let W_3 be the event that \mathcal{A} wins in Game 3. The additional ability for \mathcal{A} to win Game 3 is then bounded by its ability to win the NIKE unrecoverability game. Hence,

$$|\Pr[W_3] - \Pr[W_2]| \leq \text{Adv}_{\text{NK}, \mathcal{B}_3}^{\text{unpred}}(\lambda) \cdot \frac{1}{q} \quad (22)$$

Game 4. The only difference between Game 3 and Game 4 is that in Game 4, \mathcal{C} simulates secret sharing for participant $\tau \stackrel{s}{\leftarrow}$ honest with respect to the challenge $\hat{\Delta}$. From Game 1, we know that \mathcal{C} can do so in such a way that is indistinguishable to \mathcal{A} . However, in Game 4, it additionally does so in such a way that Equation 23 holds after performing Finalize.

$$\mathcal{C}[H_1].\text{GetPub}(0, \mathcal{C}) = \hat{\Delta} \quad (23)$$

More specifically, in $\mathcal{O}_0^{\text{PerformRound}_r}$, \mathcal{C} performs

$$(\{(j, w_{\tau j})\}_{j \in \text{corrupt}}, D_\tau) \stackrel{s}{\leftarrow} \text{AV}[H_1].\text{SimShare}(\lambda, \text{corrupt}, \hat{\Delta}).$$

It then follows the remainder of PerformRound_0 as in Game 3.

Then, in $\mathcal{O}_3^{\text{PerformRound}_r}$, it derives v' as $v' \stackrel{s}{\leftarrow} \text{AV}[H_1].\text{SimTweak}(O, \text{state}_\tau)$. It then programs $H_2(O, \text{aux})$ with v' . From Game 3, we know that \mathcal{C} can derive aux in $\mathcal{O}_3^{\text{PerformRound}_r}$, but \mathcal{A} has only negligible probability of doing so.

We will see that, if the AgVSS is securely aggregatable, then the additional advantage to the adversary in Game 4 is negligible.

Reduction to AgVSS Aggregated Secrecy. We will construct a reduction \mathcal{B}_4 which is an adversary against the AVSS aggregated secrecy game, such that \mathcal{B}_4 simulates either Game 3 or Game 4 to \mathcal{A} . When the hidden bit b in the AVSS aggregated secrecy game is $b = 0$, \mathcal{B}_4 ends up simulating Game 3; when $b = 1$, \mathcal{B}_4 ends up simulating Game 4.

\mathcal{B}_4 is constructed similar to Game 3, except as follows. When \mathcal{A} queries $\mathcal{O}_1^{\text{PerformRound}_r}$, for each honest player $k \in \text{honest}$, \mathcal{B}_4 queries the AVSS oracle $\mathcal{O}^{\text{GetShare}}()$, receiving $(\{w_j\}_{j \in \text{honest}}, D^*)$ in return. \mathcal{B}_4 then uses these outputs to simulate participant τ for $\mathcal{O}_1^{\text{PerformRound}_r}$. \mathcal{B}_4 follows the protocol honestly for all other players, and submits all $(\{w_{kj}\}_{j \in \text{honest}}, D_k), k \in [n] \setminus \tau$ to the AVSS oracle $\mathcal{O}^{\text{RecvShare}}$.

Then, \mathcal{B}_4 receives aux' , and then queries the AgVSS aggregated secrecy oracle $\mathcal{O}^{\text{GetTweak}}$ on inputs (O, aux') , receiving v' in return. \mathcal{B}_4 then programs its own $H_1(O, \text{aux}) = v'$. From Game 3, \mathcal{A} can query for its output with negligible probability before \mathcal{B}_4 programs it. \mathcal{B}_4 then follows the remainder of Game 3 as before.

When \mathcal{A} outputs b' , then \mathcal{B}_4 outputs b' . Hence, when \mathcal{A} distinguishes Game 3 from Game 4, \mathcal{B}_4 wins the AgVSS aggregated secrecy game.

Difference between Game 3 and Game 4. Let W_4 be the event that \mathcal{A} wins in Game 4. The difference between W_3 and W_4 is bounded by the advantage of \mathcal{A} in the AgVSS aggregated security game.

$$|\Pr[W_4] - \Pr[W_3]| \leq \text{Adv}_{\text{AV}, \mathcal{B}_4, \text{SimShare}}^{\text{asec}}(\lambda) \quad (24)$$

Finishing the Proof. The resulting public key in Game 4 is identical to the challenge $\hat{\Delta}$, and hence is identical to the DKG zero-knowledge game when $b = 1$. The adversary then has at most negligible advantage in winning the DKG zero-knowledge game. Concretely, combining the advantage of the adversary in (20)-(24) proves Equation 9. This concludes the proof. \square