Aggregatable Distributed Key Generation

Kobi Gurkan [⋆], Philipp Jovanovic [⋆⋆], Mary Maller [⋆⋆⋆], Sarah Meiklejohn [†], Gilad Stern [‡], Alin Tomescu [§]

Abstract. In this paper, we introduce a distributed key generation (DKG) protocol with aggregatable and publicly-verifiable transcripts. Compared with prior *publicly-veriable* approaches, our DKG reduces the size of the final transcript and the time to verify it from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \log n)$, where n denotes the number of parties. As compared with prior non-publicly-verifiable approaches, our DKG leverages *gossip* rather than *all-to-all communication* to reduce verification and communication complexity. We also revisit existing DKG security definitions, which are quite strong, and propose new and natural relaxations. As a result, we can prove the security of our aggregatable DKG as well as that of several existing DKGs, including the popular Pedersen variant. We show that, under these new definitions, these existing DKGs can be used to yield secure threshold variants of popular cryptosystems such as El-Gamal encryption and BLS signatures. We also prove that our DKG can be securely combined with a new efficient verifiable unpredictable function (VUF), whose security we prove in the random oracle model. Finally, we experimentally evaluate our DKG and show that the per-party overheads scale linearly and are practical. For 64 parties, it takes 71 ms to share and 359 ms to verify the overall transcript, while for 8192 parties, it takes 8 s and 42.2 s respectively.

1 Introduction

System designers who strive to remove single points of failure often rely on tools provided by threshold cryptography [22, 60] and secure multi-party computation [20, 35]. In this paper, we study *distributed key generation (DKG)* [32, 54], a method from threshold cryptography that often plays an essential role during the setup of distributed systems, including Byzantine consensus [6, 66], time-stamping services [14, 63], public randomness beacons [30, 61], and data archive systems [46, 65]. A DKG enables a set of parties to generate a keypair such that any sufficiently large subset can perform an action that requires the secret key while any smaller subset cannot. To achieve this, a DKG essentially turns each party into a dealer for a verifiable secret sharing (VSS) scheme [19, 26, 55]. This

[⋆] cLabs, Ethereum Foundation. Email: me@kobi.one
[⋆⋆] University College London. Email: p.jovanovic@ucl.ac.uk
[⋆⋆⋆] Ethereum Foundation. Email: mary.maller@ethereum.org
[†] University College London, Google. Email: s.meiklejohn@ucl.ac.uk
[‡] Hebrew University. Email: gilad.stern@mail.huji.ac.il
[§] VMware Research. Email: alint@vmware.com

process yields a single collective public key, generated in a distributed manner, with each party keeping a share of the secret key for themselves.

Current DKGs [28, 32, 33, 54] commonly require that all n parties broadcast $\mathcal{O}(n)$-sized messages that are then used by each party to verify the shares they received from their peers. This results in each party communicating $\mathcal{O}(n^2)$ sized messages via broadcast. While some DKGs have $\mathcal{O}(n)$ communication and verification per party, they rely on constant-sized polynomial commitment schemes that require trusted setup [44, 43]. In this work, we show how to reduce the size of the final DKG transcript to $\mathcal{O}(n)$ by making the parties' contributions *aggregatable*. This enables us to relay (partial) transcripts in an efficient and resilient manner, *e.g.*, over gossip networks, ensuring that transcripts do not grow in size since aggregation can be done in a continuous manner. Aggregatability also enables us to refresh the transcript if and when shares get compromised.

Our DKG transcripts contain the information needed for parties to decrypt their secret shares. During aggregation it is therefore essential to ensure that only valid (partial) transcripts are aggregated. We achieve this by making our transcripts *publicly-verifiable* so that anybody receiving and aggregating transcripts can verify their correctness. Making the transcripts publicly-verifiable has several other advantages: It ensures that all parties can obtain their secret shares, even if they go offline momentarily, and also enables us to remove the "complaint rounds" that are used in previous DKGs to expose misbehaving parties. This improves overall latency, since fewer communication rounds are required, and reduces the protocol's complexity from an implementation perspective.

A consequence of our approach is that the DKG secret key and its shares are group elements rather than field elements. While this prevents us from using it for many well-known cryptosystems, we demonstrate its applicability by introducing a *verifiable unpredictable function (VUF)* [24, 49] whose secret key is a group element, and prove its security in the random oracle model. Threshold VUFs are useful in the construction of verifiable random beacons, which themselves are invaluable in building proof-of-stake-based cryptocurrencies [34, 45]. To the best of our knowledge, our VUF is the first that takes a group element as the secret key, and its performance is also reasonable: our VUF output consists of 6 source group elements and can be verified using 10 pairings. We also provide further optimizations enabling us to reduce the VUF contributions in the threshold protocol to just 2 source group elements that can be verified with 3 pairings.

We also revisit the definitions for DKGs in the hope of reducing complexities and inefficiencies. In particular, previous definitions [32] required *secrecy*, in the sense that the output of the DKG must be indistinguishable from random. While this notion has the benefit of making the DKG modular (one can replace key generation with a DKG in any context), it also is difficult to realise. Indeed, Gennaro et al. [32] demonstrated that the popular Pedersen DKG does not have secrecy, and introduced an alternative and considerably less efficient protocol to achieve secrecy. Additionally, no one-round DKG can achieve secrecy because a rushing adversary (an adversary that plays last) can always influence the final distribution. Instead, we look to prove that a DKG is *robust* (see Definition 6)

and *security-preserving* (see Definition 8) in the sense that any adversary that breaks the security of a threshold version of the scheme (i.e., one using a DKG) also breaks the original security property. These security notions are new to this paper.

Gennaro et al. [33] previously observed that the Pedersen DKG suffices to construct threshold Schnorr signatures [59]. Recently, Benhamouda et al. [10] found an attack on this approach when the adversary is *concurrent*. (Gennaro et al. had not considered concurrent adversaries.) Komlo and Goldberg [48] show that it is possible to avoid the attack but, in doing so, they lose robustness (*e.g.*, if a single party goes offline a signature will not be produced). This raises questions as to whether it is still okay to use the Pedersen DKG with respect to other signature schemes such as BLS. In this paper, we provide a positive answer in the form of a security proof that holds concurrently and does not rely on rewinding the adversary. Specifically, we show that the Pedersen DKG is security-preserving with respect to any *rekeyable* encryption scheme, signature scheme, or VUF scheme where the sharing algorithm is the same as encryption or signing (see Definition 5).

Our contributions. In Section 5, we construct an *aggregatable* and *publicly-verifiable* distributed key generation protocol. The aggregation can be completed by any party (there are no additional secrets) and can also be done incrementally. The cost of verifying our transcripts is $\mathcal{O}(n \log n)$ whereas prior approaches were $\mathcal{O}(n^2)$ [28]. If any user temporarily goes offline, they can still recover their secret shares. Dealing DKG shares takes $\mathcal{O}(n \log n)$ time and aggregation costs are $\mathcal{O}(n)$.

We prove security of our DKG using a natural definition (see Section 3.6), which roughly states that, if it is possible to break a cryptosystem's security game with a DKG swapped in, then it is possible to break that cryptosystem's original security game that did not use the DKG for key generation. We further demonstrate that, counter-intuitively, it is possible to prove that a DKG realises this definition without needing a separate proof for each cryptosystem. Indeed, we show that any encryption scheme, signature scheme, or VUF that are *rekeyable*, and where the sharing algorithm is the same as encryption or signing, can be securely instantiated using a *key-expressable* DKG (see Definition 7). This includes El-Gamal encryption, BLS signatures, a new VUF we introduce in Section 7, and, we suspect, many others.

We further demonstrate the applicability of our techniques by showing that all three of the Pedersen DKG [54], the Fouque-Stern DKG [28], and our aggregatable DKG are key-expressable and thus can be used securely with rekeyable encryption schemes, signature schemes, and VUFs whose decryption/ signing algorithms are the same as the algorithms to generate decryption/ signature shares. Our proof allows for rushing adversaries and holds concurrently (*i.e.*, with respect to an adversary that can open many sessions at the same time). We cannot cover Schnorr signatures, however, because their threshold variants do not appear to be rekeyable.

Our final contribution, in Section 8, is a Rust implementation of our aggregatable DKG to demonstrate its practicality by showing that its overheads are indeed linear. For example, the evaluation of our implementation shows that for 64 / 128 / 8192 nodes it takes 71 ms / 137 ms / 8,000 ms to share one secret and 359 ms / 747 ms / 42,600 ms to verify the corresponding transcript.

2 Related Work

DKG	**Broadcasts**		**P2P**	**PV**	**Complaints**	**Rounds**		**Prover**	**Verifier**	
	$\mathcal{O}(n)$	$\mathcal{O}(1)$				*Broadcast*	*Gossip*		*Local*	*Global*
Pedersen	n	–	n	no	yes	3	–	$n \lg n$	n^2	–
Kate	–	n	n	no	yes	3	–	n^2	n	–
AMT	–	n	$n \lg n$	no	yes	3	–	$n \lg n$	$n \lg n$	–
Fouque-Stern	n	–	n^2	yes	no	1	–	$n \lg n$	n^2	n^2
Our work	$\lg n$	n	$n \lg^2 n$	yes	no	2	$\lg n$	$n \lg n$	$n \lg^2 n$	$n \lg n$

Table 1. Complexities of prior DKG protocols with n parties, per party. In the "Broadcast" column, we count the number of broadcasts by size (either $O(n)$ or $O(1)$-sized). "P2P" means the total size of the messages sent over public and private communication channels (excludes broadcast messages). "PV" means publicly-verifiable. "Verifier local" indicates the per-party time spent verifying their shares from other parties, while "global" indicates the time to verify the final DKG transcript.

We provide an asymptotic overview of the state-of-the-art for DKGs in Table 2. Here we assume that the threshold t is linear in n. Our comparisons consider the optimistic case where there is no more than a constant number of complaints for protocols where these are relavant (recall in our protocol there are no complaints).

Pedersen introduced the first efficient DKG protocol for discrete log-based cryptosystems [55], building on top of Feldman's VSS [26]. Gennaro et al. [33] showed Pedersen's DKG does not generate uniformly distributed secrets, and proposed a protocol that does but at the cost of lower efficiency. They also fix problems with the complaint phase in Pedersen's DKG. Neji et al. [51] gave a more efficient protocol that ensures uniformity in Pedersen's DKG. Kate [43] reduced the broadcast overhead per DKG party from $\mathcal{O}(n)$ to $\mathcal{O}(1)$ using their constant-sized polynomial commitment scheme [44]. However, there scheme depends on a trusted setup algorithm, the costs of which are not considered in Table 2. Trusted setup algorithms have a round complexity of t, and each of these rounds require users to broadcast $\mathcal{O}(n)$ sized messages [15, 37]. Unlike our protocol, all of these protocols rely on complaints rounds, are not publicly verifiable and have $\mathcal{O}(n^2)$ communication complexity.

Fouque and Stern present a one-round, publicly-verifiable DKG that uses only public channels. However, their final transcript size is $\mathcal{O}(n^2)$ whereas ours is $\mathcal{O}(n)$ because we can aggregate. Furthermore, their security proof does not allow for rushing adversaries. While they do not measure performance, their use of Paillier encryption [53] is likely to make their DKG slow and have high communication costs. Nonetheless, unlike our DKG, theirs has the advantage of outputting secrets that are field, rather than group, elements.

Other works tackle the DKG problem from different angles. Canetti et al's DKG [16] has adaptive security, while ours is secure only against static adversaries that fix the set of corrupted parties before the protocol starts. Canny and Sorkin [17] study DKG protocols with poly-logarithmic communication and computation cost per-party, but their protocol relies on a trusted dealer that permutes the parties before the protocol starts. Kate et al. [42, 43] and Kokoris-Kogias et al. [47] study DKG protocols in the *asynchronous* setting, unlike our work and most previous work. Schindler et al. [58] use the Ethereum blockchain to instantiate the synchronous broadcast channel all DKG protocols mentioned so far assume, including ours. Tomescu et al. [62] lower the computational cost of dealing in Kate et al.'s DKG [44], at a logarithmic increase in communication. Lastly, several works implement and benchmark synchronous, statically-secure DKG protocols for discrete log-based cryptosystems [58, 52, 40, 23, 39].

Abe et al. [1] observed that any fully structure preserving signature scheme that depends solely on algebraic operations cannot be used as a VUF or VRF. Unlike our VUF (which is not algebraic), this rules out a number of structure preserving signatures from being candidates for building VUFs [4, 2, 64, 3].

3 Definitions

3.1 Preliminaries

If x is a binary string then $|x|$ denotes its bit length. If S is a finite set then $|S|$ denotes its size and $x \xleftarrow{\$} S$ denotes sampling a member uniformly from S and assigning it to x. We use $\lambda \in \mathbb{N}$ to denote the security parameter and 1^λ to denote its unary representation. Algorithms are randomized unless explicitly noted otherwise. "PPT" stands for "probabilistic polynomial time." We use $\vec{y} \leftarrow A(\vec{x}; r)$ to denote running algorithm A on inputs $\vec{x}$ and randomness r and assigning its output to $\vec{y}$. We use $\vec{y} \xleftarrow{\$} A(\vec{x})$ to denote $y \leftarrow A(x; r)$ for uniformly random r. We use $[A(\vec{x})]$ to denote the set of values that have non-zero probability of being output by A on input $\vec{x}$. For two functions $f, g : \mathbb{N} \to [0, 1]$, we use $f(\lambda) \approx g(\lambda)$ to denote $|f(\lambda) - g(\lambda)| = \lambda^{-\omega(1)}$. We use code-based games in security definitions and proofs [9]. A game $\mathsf{Sec}_{\mathcal{A}}(\lambda)$, played with respect to a security notion Sec and adversary $\mathcal{A}$, has a MAIN procedure whose output is the output of the game. The notation $\mathsf{Pr}[\mathsf{Sec}_{\mathcal{A}}(\lambda)]$ denotes the probability that this output is 1.

We formalize bilinear groups via a *bilinear group sampler*, which is an efficient *deterministic* algorithm $\mathsf{GroupGen}$ that given a security parameter 1^λ

(represented in unary), outputs a tuple $\mathsf{bp} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1)$ where $\mathbb{G}_1$, $\mathbb{G}_2$, $\mathbb{G}_T$ are groups with order divisible by the prime $p \in \mathbb{N}$, g_1 generates $\mathbb{G}_1$, $\hat{h}_1$ generates $\mathbb{G}_2$, and $e\colon \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is a (non-degenerate) bilinear map. Galbraith et al. distinguish between three types of bilinear group samplers [29]. Type I groups have $\mathbb{G}_1 = \mathbb{G}_2$ and are known as *symmetric* bilinear groups. Types II and III are *asymmetric* bilinear groups, where $\mathbb{G}_1 \neq \mathbb{G}_2$. Type II groups have an efficiently computable homomorphism $\psi : \mathbb{G}_2 \to \mathbb{G}_1$, while Type III groups do not have an efficiently computable homomorphism in either direction. Certain assumptions are provably false with respect to certain group types (*e.g.*, SXDH only holds for Type III groups), and we work only with Type III groups.

3.2 Communication and threat models

In this section we discuss our communication and threat models.

Synchrony: We assume perfect synchrony. There is a strict time bound between rounds. All messages (honest and adversarial) within a round will be seen by all parties by the end of the round.

Communication channel: We assume the existence of a broadcast channel for sending messages. If a non-faulty party broadcasts a message then it will be seen by everyone by the end of the round. It is not possible to forge messages from non-faulty parties.

Adversarial threshold: We denote by t the adversarial threshold; i.e., the number of parties that the adversary can corrupt. The total number of parties is denoted by n. We set no specific bounds on the adversarial threshold because a rational adversary might prefer to attack the secrecy of the DKG over blocking the communication channels [5, 31].

Assumptions on the adversary: Our security proofs are given with respect to static adversaries, meaning the adversary must state at the start of the security game all of the parties that it has corrupted. We allow the adversary to control the ordering of messages within a round, and in particular the adversary can wait to receive all messages within a round before it broadcasts its message (this is called a *rushing* adversary). The adversary can also choose not to participate at all.

Byzantine adversary A byzantine adversary is a malicious entity that may differ arbitrarily from the protocol.

Crashed party A crashed party is a party that has gone offline e.g. due to a faulty internet connection. After a party has crashed they will not send any more messages.

3.3 Assumptions

Our security proofs are provided in the random oracle model; i.e., there exists a simulator that can program the output of a hash function provided that their chosen outputs are indistinguishable from random.

We rely on the SXDH assumption, which is an extension of the DDH assumption to Type III bilinear groups. Informally, it states that given g_1^α and g_1^β it is hard to distinguish $g_1^{\alpha\beta}$ from random.

Assumption 1 (Symmetric External Diffie Hellman SXDH [7, 8]) *For an adversary $\mathcal{A}$, define*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{A}}(\lambda) = |\Pr[\mathsf{bp} \leftarrow \mathsf{GroupGen}(1^\lambda); \alpha, \beta \xleftarrow{\$} \mathbb{F} \ : \ \mathcal{A}(\mathsf{bp}, g_1^\alpha, g_1^\beta, g_1^{\alpha\beta}) = 1] \\ - \Pr[\mathsf{bp} \leftarrow \mathsf{GroupGen}(1^\lambda); \alpha, \beta, \gamma \xleftarrow{\$} \mathbb{F} \ : \ \mathcal{A}(\mathsf{bp}, g_1^\alpha, g_1^\beta, g_1^\gamma) = 1]|.\end{aligned}$$

We say that the SXDH assumption holds if $\mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{A}}(\lambda) < \mathsf{negl}(\lambda)$ for all PPT adversaries.

The BDH assumption is an extension of the CDH assumption to Type III bilinear groups. Informally, it states that given $g_1^\alpha, g_1^\beta, \hat{h}_1^\gamma, \hat{h}_1^{\alpha\gamma}$ it is hard to compute $e(g_1, \hat{h}_1)^{\alpha\beta}$.

Assumption 2 (Computational Bilinear Diffie Hellman BDH [12]) *Define an adversary $\mathcal{A}$'s advantage $\mathsf{Adv}^{\mathsf{BDH}}_{\mathcal{A}}(\lambda)$ against BDH by*

$$\Pr[\mathsf{bp} \leftarrow \mathsf{GroupGen}(1^\lambda); \alpha, \beta, \gamma \xleftarrow{\$} \mathbb{F} \ : \ \mathcal{A}(\mathsf{bp}, g_1^\alpha, g_1^\beta, \hat{h}_1^\gamma, \hat{h}_1^{\alpha\gamma}) = e(g_1, \hat{h}_1)^{\alpha\beta}].$$

We say that the BDH assumption holds if $\mathsf{Adv}^{\mathsf{BDH}}_{\mathcal{A}}(\lambda) < \mathsf{negl}(\lambda)$ for all PPT adversaries.

3.4 Verifiable unpredictable functions (VUFs)

A VUF allows a party with a secret key to compute a deterministic (keyed) function and prove to an external verifier that the result is correct. The notion is related to signatures, with the extra requirement that the output of the signer must be unique, even to a party that can choose the secret key. We have made the following changes to prior definitions [24] in order to better suit our setting: (1) we include a global setup algorithm to generate a common reference string; (2) we include a derive algorithm to map the prover's output onto the unique function output.

Definition 1 (Verifiable Unpredictable Function). *Let Π = (VUF.Setup, VUF.Gen, VUF.Eval, VUF.Sign, VUF.Derive, VUF.Ver) be the following set of efficient algorithms:*

$\mathsf{crs}_{\mathsf{vuf}} \leftarrow \mathsf{VUF.Setup}(1^\lambda)$ *: a DPT algorithm that takes as input the security parameter and outputs a common reference string.*

$(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{VUF.Gen}(\mathsf{crs}_{\mathsf{vuf}})$*: a PPT algorithm that takes as input a common reference string and returns a public key and a secret key.*

$\mathsf{out} \leftarrow \mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)$*: a DPT algorithm that takes as input a common reference string, secret key, and message $m \in \{0,1\}^\lambda$ and returns $\mathsf{out} \in \{0,1\}^\lambda$.*

$\sigma \xleftarrow{\$} \mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)$*: a PPT algorithm that takes as input a common reference string, secret key, and message, and returns a signature* σ*.*

$\mathsf{out} \leftarrow \mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)$*: a DPT algorithm that takes as input a common reference string, public key, message and signature and returns* $\mathsf{out} \in \{0,1\}^\lambda$*.*

$0/1 \leftarrow \mathsf{VUF.Ver}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)$*: a DPT algorithm that takes as input a common reference string, public key, message and signature and returns* 1 *to indicate acceptance and* 0 *to indicate rejection.*

We say that Π *is a* verifiable unpredictable function (VUF) *if it satisfies correctness, uniqueness, and unpredictability (defined below).*

A VUF is correct if an honest signer always convinces an honest verifier and always outputs a seed such that the derive function outputs the correct value.

Definition 2 (Correctness). *A VUF is correct if for all* $\lambda \in \mathbb{N}$ *and* $m \in \{0,1\}^\lambda$ *we have that*

$$\Pr\left[\begin{array}{l|l} \mathsf{crs}_{\mathsf{vuf}} \leftarrow \mathsf{Setup}(1^\lambda), & \mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma) = \\ (\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{VUF.Gen}(\mathsf{crs}_{\mathsf{vuf}}), & \mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m) \\ \sigma \xleftarrow{\$} \mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m) & \wedge\ \mathsf{VUF.Ver}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma) = 1 \end{array}\right] = 1.$$

A VUF is unique if an adversary (even one that chooses the secret key) cannot output a verifying signature such that the derive function outputs the wrong value.

Definition 3 (Uniqueness). *For a VUF* Π *and an adversary* $\mathcal{A}$*, let* $\mathsf{Adv}^{\mathsf{unique}}_{\mathcal{A}}(\lambda) = \Pr[\mathsf{Game}^{\mathsf{unique}}_{\mathcal{A}}(\lambda)]$*, where* $\mathsf{Game}^{\mathsf{unique}}_{\mathcal{A}}(\lambda)$ *is defined as follows:*

$\underline{\text{MAIN } \mathsf{Game}^{\mathsf{unique}}_{\mathcal{A}}(\lambda)}$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow \mathsf{VUF.Setup}(1^\lambda)$
$(\mathsf{pk}, m, \sigma_1, \sigma_2) \xleftarrow{\$} \mathcal{A}(\mathsf{crs}_{\mathsf{vuf}})$
$y_1 \leftarrow \mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma_1)$
$y_2 \leftarrow \mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma_2)$
return $(y_1 \neq y_2) \ \wedge\ \mathsf{VUF.Ver}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma_1) \ \wedge\ \mathsf{VUF.Ver}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma_2)$

We say that Π *is unique if for all PPT adversaries* $\mathcal{A}$ *we have that* $\mathsf{Adv}^{\mathsf{unique}}_{\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

Finally, a VUF is unpredictable if an adversary cannot predict the output of the function $\mathsf{VUF.Eval}$ on a message for which it has not seen any valid signatures.

Definition 4 (Unpredictability). *For a VUF* Π *and an adversary* $\mathcal{A}$*, let* $\mathsf{Adv}^{\mathsf{predict}}_{\mathcal{A}}(\lambda) = \Pr[\mathsf{Game}^{\mathsf{predict}}_{\mathcal{A}}(\lambda)]$ *where* $\mathsf{Game}^{\mathsf{predict}}_{\mathcal{A}}(\lambda)$ *is defined as follows:*

$\underline{\text{MAIN } \mathsf{Game}^{\mathsf{predict}}_{\mathcal{A}}(\lambda)}$
$H \leftarrow \emptyset$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow \mathsf{VUF.Setup}(1^\lambda)$
$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{VUF.Gen}(\mathsf{crs}_{\mathsf{vuf}})$
$(m, y) \xleftarrow{\$} \mathcal{A}^{\mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, \cdot)}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk})$
return $(\mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m) = y) \ \wedge \ (m \notin H)$

$\underline{\text{ORACLE } \mathcal{O}^{\mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)}}$
add m *to query set* H
return $\mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)$

We say that Π *is unpredictable if for all PPT adversaries* $\mathcal{A}$ *we have that* $\mathsf{Adv}^{\mathsf{predict}}_{\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

3.5 Rekeyability

To show that existing cryptographic primitives can be instantiated with our DKG, and other DKGs in the literature, we rely on a property called *rekeyability*. Intuitively, rekeyability says that it is possible to transform an object (e.g., a ciphertext or signature) that was formed using one cryptographic key into an object formed with a related key. As one concrete example, in the BLS signature scheme, in which a signature on a message m is of the form $\sigma = H(m)^{\mathsf{sk}_1}$, it is possible to transform this into a signature under the key $\alpha\mathsf{sk}_1 + \mathsf{sk}_2$ by computing $\sigma^\alpha \cdot H(m)^{\mathsf{sk}_2}$. This means that BLS can be efficiently rekeyed with respect to the secret key. While this notion is related to the idea of re-randomizability [36, 57, 27], we are not aware of any formalizations in the literature and it may be of independent interest.

Definition 5 (Rekeyability). *For a public-key primitive* $\Pi = (\mathsf{KeyGen}, \Pi_1, \ldots, \Pi_n)$ *and functions* $f_{\mathsf{k}}(\alpha, \mathsf{k}_1, \mathsf{k}_2)$ *that outputs* $\alpha\mathsf{k}_1 \oplus \mathsf{k}_2$ *for some binary operator* $\oplus$ *(typically* $+$ *or* $\times$*), we define rekeyability as follows for all* $\alpha \in \mathbb{N}$ *and* $(\mathsf{pk}_1, \mathsf{sk}_1), (\mathsf{pk}_2, \mathsf{sk}_2) \in [\mathsf{KeyGen}(1^\lambda)]$*:*

- *We say that an algorithm* Π_i *is* rekeyable *with respect to the secret key if there exists an efficient function* rekey_i *such that*

$$\mathsf{rekey}_i(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, x, \Pi_i(\mathsf{sk}_1, x; r)) = \Pi_i(f_{\mathsf{sk}}(\alpha, \mathsf{sk}_1, \mathsf{sk}_2), x; r)$$

  for all $x \in \mathsf{Domain}(\Pi_i)$ *and randomness* r*. Likewise, we say that an algorithm* Π_j *is* rekeyable *with respect to the public key if there exists an efficient function* rekey_j *such that*

$$\mathsf{rekey}_j(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, \Pi_j(\mathsf{pk}_1, x; r)) = \Pi_j(f_{\mathsf{pk}}(\alpha, \mathsf{pk}_1, \mathsf{pk}_2), x; r)$$

  for all $x \in \mathsf{Domain}(\Pi_i)$ *and randomness* r*.*
- *We say that* (Π_i, Π_j) *is* rekeyable *with respect to the secret key if (1)* Π_i *is rekeyable with respect to the secret key and (2)*

$$\Pi_j(\mathsf{pk}_1, y) = \Pi_j(f_{\mathsf{pk}}(\alpha, \mathsf{pk}_1, \mathsf{pk}_2), \mathsf{rekey}_i(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, y)).$$

  Likewise we say that (Π_i, Π_j) *is* rekeyable *with respect to the public key if (1)* Π_i *is rekeyable with respect to the public key and (2)*

$$\Pi_j(\mathsf{sk}_1, y) = \Pi_j(f_{\mathsf{sk}}(\alpha, \mathsf{sk}_1, \mathsf{sk}_2), \mathsf{rekey}_i(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, y)).$$

For encryption, we would want that $(\mathsf{Encrypt}, \mathsf{Decrypt})$ is rekeyable with respect to the public key, meaning new key material can be folded into ciphertexts without affecting the ability to decrypt. For signing, we would want that $(\mathsf{Sign}, \mathsf{Verify})$ is rekeyable with respect to the secret key, meaning that if signatures verify then so do their rekeyed counterparts.

3.6 Distributed key generation (DKG)

We define a *distributed key generation* (DKG) as an interactive protocol that is used to generate a keypair $(\mathsf{pk}, \mathsf{sk})$. We define this as $(\mathsf{transcript}, \mathsf{pk}) \xleftarrow{\$} \mathtt{DKG}(I, n)$, where n is the number of participants in the DKG, I is the indices of the adversarial participants (so $|I| \leq t$), pk is the resulting public key, and $\mathsf{transcript}$ is some representation of the messages that have been exchanged.

We additionally consider an algorithm $\mathsf{Reconstruct}$ that, given $\mathsf{transcript}$ and the shares submitted by $t+1$ honest parties, outputs the secret key sk corresponding to pk. With this in place, we can define an omniscient interactive protocol $(\mathsf{transcript}, (\mathsf{pk}, \mathsf{sk}), \mathsf{state}_{\mathcal{A}}) \xleftarrow{\$} \mathtt{OmniDKG}(I, n)$ that is aware of the internal state of each participant and thus can output sk (by running the $\mathsf{Reconstruct}$ algorithm) and $\mathsf{state}_{\mathcal{A}}$; i.e., the internal state of the adversary.

The $\mathsf{Reconstruct}$ algorithm is useful not only in defining this extra interactive protocol, but also in defining a notion of *robustness* for DKGs (initially called correctness by Gennaro et al. [32]). We define this as follows:

Definition 6 (Robustness). *A DKG protocol is* robust *if the following properties hold:*

- *A DKG transcript* dkg *determines a public key* pk *that all honest parties agree on.*
- *There is an efficient algorithm* $\mathsf{Reconstruct}$

$$\mathsf{sk} \leftarrow \mathsf{Reconstruct}(\mathsf{dkg}, \mathsf{sk}_1, \ldots, \mathsf{sk}_\ell) \text{ for } t+1 \leq \ell \leq n$$

  that takes as input a set of secret key shares where at least $t+1$ *are from honest parties and verifies them against the public transcript produced by the DKG protocol. It outputs the unique value* sk *such that* $\mathsf{pk} \leftarrow \mathsf{KeyGen}(1^\lambda; \mathsf{sk})$.

Beyond robustness, we also want a DKG to *preserve security* of the underlying primitive for which it is run. Previous related definitions of *secrecy* for DKGs required there to exist a simulator that could fix the output of the DKG; i.e., given an input y, could output $(\mathsf{transcript}, y)$ that the adversary could not distinguish from a real $(\mathsf{transcript}, \mathsf{pk})$ output by the DKG run with t adversarial participants. While general, this definition is strong and required previous constructions to have more rounds or constraints than would otherwise be necessary; *e.g.*, there seem to be significant barriers to satisfying this definition in any DKG where the adversary is allowed to go last, as they the know the entire transcript and can bias the final result.

In defining what it means for a DKG to preserve security, we first weaken this previous definition. Rather than require a simulator given pk_1 to have the DKG output exactly pk_1, we consider that it can instead fix the output public key to have a known relation with its input public key. In particular, a simulator given pk_1 can fix the output of the DKG to be pk, where the simulator knows $(\alpha, \mathsf{pk}_2, \mathsf{sk}_2)$ such that $\mathsf{pk} = f(\alpha, \mathsf{pk}_1, \mathsf{pk}_2)$ for $\alpha \neq 0$ and f as defined in the rekeyability definition (see Definition 5). We call this property *key expressability*.

Definition 7 (Key expressability). *For a simulator* Sim*, define as* $(\mathsf{transcript}, \mathsf{pk}, \alpha, \mathsf{pk}_2, \mathsf{sk}_2) \xleftarrow{\$} \mathtt{SimDKG}(\mathsf{Sim}, I, n)$ *a run of the DKG protocol in which all honest participants are controlled by* Sim*, which takes as input a public key* pk_1 *and has private outputs* α*,* pk_2*, and* sk_2*. We say that a DKG is* key-expressable *if there exists such a simulator* Sim *such that (1)* $(\mathsf{transcript}, \mathsf{pk})$ *is distributed identically to the output of* $\mathtt{DKG}(I, n)$*, (2)* $(\mathsf{pk}_2, \mathsf{sk}_2)$ *is a valid keypair, and (3)* $\mathsf{pk} = f(\alpha, \mathsf{pk}_1, \mathsf{pk}_2) = \alpha\mathsf{pk}_1 \oplus \mathsf{pk}_2$.

To now define a security-preserving DKG, we intuitively consider a DKG being run in the context of a security game. To keep our definition as general as possible, our only requirements are that (1) the security game contains a line of the form $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\lambda)$ (it also works if KeyGen takes a common reference string as additional input), and (2) pk is then later given as input to the adversary. We then say that the DKG preserves security if it is not possible for an adversary participating in the DKG to do better than it would have done in the original security game, in which it was given pk directly. Formally, we have the following definition.

Definition 8 (Security-preserving). *Define* Game *as any security game containing the line* $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\lambda)$*, denoted* $\mathsf{line}_{\mathsf{pk}}$*, and where* pk *is later input to an adversary* $\mathcal{A}$ *(in addition to other possible inputs). Define* $\mathsf{Game}'(\mathsf{line}, x)$*, parameterized by a starting line* line *and some value* x*, as* Game *but with* $\mathsf{line}_{\mathsf{pk}}$ *replaced by* line *and* $\mathcal{A}$ *given* x *as input rather than* pk*. It is clear that* $\mathsf{Game} = \mathsf{Game}'(\mathsf{line}_{\mathsf{pk}}, \mathsf{pk})$.

Define $\mathsf{line}_{\mathsf{dkg}}$ *as the line* $(\mathsf{transcript}, (\mathsf{pk}, \mathsf{sk}), \mathsf{state}_{\mathcal{A}}) \xleftarrow{\$} \mathtt{OmniDKG}(I, n)$*, and define* $\mathsf{DKG\text{-}Game} \leftarrow \mathsf{Game}'(\mathsf{line}_{\mathsf{dkg}}, \mathsf{state}_{\mathcal{A}})$*. We say the DKG* preserves security *for* Game *if*

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{DKG\text{-}Game}}(\lambda) \leq \mathsf{Adv}_{\mathcal{A}}^{\mathsf{Game}}(\lambda) + \mathsf{negl}(\lambda)$$

for all PPT adversaries $\mathcal{A}$*.*

We do not view our requirements for the original security game as restrictive, given the number of security games that satisfy them. For signature unforgeability, for example, our definition says that an adversary that participates in the DKG, and can carry its state from that into the rest of the game (including all of the messages it saw), cannot achieve better advantage than when it is just given the public key (as in the standard EUF-CMA game).

While the relationship between key expressability and security-preserving DKGs is not obvious, we show in Section C that it is typically the case that

when key-expressable DKGs are used for rekeyable primitives, they preserve the security of that primitive's underlying security game.

4 Our Enhanced Scrape PVSS

A *secret sharing scheme* allows a *dealer* to deal out n secret *shares* so that any subset of $t+1$ shares suffices to reconstruct the secret, but subsets of size $\leq t$ shares do not. A *publicly verifiable secret sharing (PVSS)* scheme is a secret sharing scheme in which any third party can verify that the dealer has behaved honestly. Importantly, PVSS obviates the need for a complaint round in VSS protocols, which simplifies designing PVSS-based DKG protocols [28]. Cascudo and David designed an elegant PVSS scheme called *Scrape* [18] with $O(n)$ verification costs. In this section, we describe a slightly-modified variant of Scrape that supports *aggregation* and uses Type III pairings and an additional element $\hat{u}_2 \in \mathbb{G}_2$ that will help our DKG security proofs later on. We rely on Type III pairings, not only for efficiency, but also because the SXDH assumption does not hold in symmetric groups. We give a formal description in Fig. 1. In Section 5, we use this slightly-modified variant of Scrape to construct our DKG.

Common reference string (CRS). All parties use the same CRS consisting of (1) a bilinear group description bp (which fixes $g_1 \in \mathbb{G}_1$ and $\hat{h}_1 \in \mathbb{G}_2$), (2) a group element $\hat{u}_1 \in \mathbb{G}_2$ and (3) *encryption keys* $\mathsf{ek}_i \in \mathbb{G}_2$ for every party P_i with corresponding *decryption keys* $\mathsf{dk}_i \in \mathbb{F}$ known only to P_i such that $\mathsf{ek}_i = \hat{h}_1^{\mathsf{dk}_i}$.

Dealing. Scrape resembles other Shamir-based [60] secret sharing schemes. The Scrape dealer will share a secret $\hat{h}_1^{a_0} \in \mathbb{G}_2$, whose corresponding $a_0 \in \mathbb{F}$ the dealer knows. (This is different than other VSS schemes, which typically share a secret in $\mathbb{F}$ rather than in $\mathbb{G}_2$.) The dealer picks a random, degree-t polynomial $f(X) = (a_0, a_1, \ldots, a_t)$, where $f(0) = a_0$, and commits to it via Feldman [26] as $F_i = g_1^{a_i}, \forall i \in [0, t]$. Party P_i's share will be $\hat{h}_1^{f(\omega_i)}$. The dealer then computes Feldman commitments $A_i = g_1^{f(\omega_i)}$ and encryptions $\mathsf{ek}_i^{f(\omega_i)}$ of each share. (The term "encryption" here is slightly abused since these are not IND-CPA-secure ciphertexts.) The *PVSS transcript* will consist of the Feldman commitments to $f(X)$ and to the shares, plus the encryptions of the shares. Additionally, we augment the transcript with $\hat{u}_2 = \hat{u}_1^{a_0}$, which helps our DKG security proofs.

Verifying. Each party P_i can verify that the PVSS transcript is a correct sharing of $\hat{h}_1^{a_0}$. For this, P_i checks the Feldman commitments A_i to the shares $f(\omega_i)$ are consistent with the Feldman commitment to $f(X)$ via Lagrange interpolation in the exponent (see Fig. 1). Then, each P_i checks their encryption of $f(\omega_i)$ against A_i. Altogether, this guarantees that the encrypted shares are indeed the evaluations of the committed polynomial f.

Aggregating transcripts. One of our key contributions is an algorithm for aggregating two Scrape PVSS transcripts pvss_1 and pvss_2 for polynomials f_1 and f_2 into a single transcript for their sum $f_1 + f_2$. This is a key ingredient of our DKG from Section 5. Our aggregation leverages the homomorphism of

$\mathsf{Scrape.Deal}(\mathsf{bp}, \mathbf{ek}, \hat{u}_1, a_0) \to \mathsf{pvss}$

$(a_1, \ldots, a_t) \xleftarrow{\$} \mathbb{F}^t$, $f(X) \leftarrow \sum_{i=0}^{t} a_i X^i$
$F_0, \ldots, F_t \leftarrow g_1^{a_0}, \ldots, g_1^{a_t}$
$\hat{u}_2 \leftarrow \hat{u}_1^{a_0}$
$A_1, \ldots, A_n \leftarrow g_1^{f(\omega_1)}, \ldots, g_1^{f(\omega_n)}$
$\hat{Y}_1 \ldots, \hat{Y}_n \leftarrow \mathsf{ek}_1^{f(\omega_1)}, \ldots, \mathsf{ek}_n^{f(\omega_n)}$
return $\boldsymbol{F}, \hat{u}_2, \boldsymbol{A}, \hat{\boldsymbol{Y}}$

$\mathsf{Scrape.Verify}(\mathsf{bp}, \mathbf{ek}, \hat{u}_1, \hat{u}_2, \mathsf{pvss}) \to 0/1$

$\boldsymbol{F}, \hat{u}_2, \boldsymbol{A}, \hat{\boldsymbol{Y}} \leftarrow \mathsf{parse}(\mathsf{pvss})$
$\alpha \xleftarrow{\$} \mathbb{F}$
check $\prod_{j=1}^{n} A_j^{\ell_j(\alpha)} = \prod_{j=0}^{t} F_j^{\alpha^j}$
check $e(F_0, \hat{u}_1) = e(g_1, \hat{u}_2)$
check $e(g_1, \hat{Y}_j) = e(A_j, \mathsf{ek}_j)$ for $1 \le j \le n$
return 1 if all checks pass, else return 0

$\mathsf{Scrape.Aggregate}(\mathsf{bp}, \mathsf{pvss}_1, \mathsf{pvss}_2) \to \mathsf{pvss}$

$((F_{1,0}, \ldots, F_{1,t}), \hat{u}_{1,2}, (A_{1,1}, \ldots, A_{1,n}), (\hat{Y}_{1,1}, \ldots, \hat{Y}_{1,n})) \leftarrow \mathsf{parse}(\mathsf{pvss}_1)$
$((F_{2,0}, \ldots, F_{2,t}), \hat{u}_{2,2}, (A_{2,1}, \ldots, A_{2,n}), (\hat{Y}_{2,1}, \ldots, \hat{Y}_{2,n})) \leftarrow \mathsf{parse}(\mathsf{pvss}_2)$

for $0 \le i \le t$:
  $F_i \leftarrow F_{1,i} F_{2,i}$

for $1 \le i \le n$:
  $A_i \leftarrow A_{1,i} A_{2,i}$, $\hat{Y}_i \leftarrow \hat{Y}_{1,i} \hat{Y}_{2,i}$

$\hat{u}_2 \leftarrow \hat{u}_{1,2} \hat{u}_{2,2}$
return $\boldsymbol{F}, \hat{u}_2, \boldsymbol{A}, \hat{\boldsymbol{Y}}$

Fig. 1. Dealing, verification and aggregation algorithms for the Scrape PVSS. Here, $\mathbf{ek}, \boldsymbol{F}, \boldsymbol{A}, \hat{\boldsymbol{Y}}$ denote vectors of ek_i's, F_i's, A_i's and $\hat{Y}_i$'s. The polynomial $\ell_j(X)$ denotes the Lagrange polynomial equal to 1 at ω_j and 0 at $\omega_i \neq \omega_j$. The ω_i's are public predetermined values which, for efficiency purposes, should be chosen as roots of unity of degree n. For more details, see Appendix A.

Feldman commitments and of the encryption scheme. Indeed, suppose we have Feldman commitments to f_b consisting of $F_{b,i} = g_1^{a_{b,i}}, \forall i \in [0, t]$, where $a_{b,i}$'s are the coefficients of f_b, for $b \in \{1, 2\}$. Then, $F_i = F_{1,i} F_{2,i} = g^{a_{1,i} + a_{2,i}}, \forall i \in [0, t]$ will be a Feldman commitment to $f_1 + f_2$. Similarly, we can aggregate the share commitments $A_{b,i} = g^{f_b(\omega_i)}$ as $A_i = A_{1,i} A_{2,i} = g^{(f_1 + f_2)(\omega_i)}, \forall i \in [n]$. Lastly, the encryptions $\mathsf{ek}_i^{f_b(\omega_i)}$ can be aggregated as $\mathsf{ek}_i^{(f_1+f_2)(\omega_i)} = \mathsf{ek}_i^{f_1(\omega_i)} \mathsf{ek}_i^{f_2(\omega_i)}$. We summarize this aggregation algorithm in Fig. 1.

Reconstructing the secret. At the end of the PVSS protocol, each party P_i can decrypt their share as $\hat{A}_i = \hat{Y}_i^{\mathsf{dk}_i^{-1}} = (\mathsf{ek}_i^{f(\omega_i)})^{\mathsf{dk}_i^{-1}} = \hat{h}_1^{f(\omega_i)}$. Recall that the degree t polynomial $f(X)$ encodes the secret $f(0) = a_0$. Thus, any set S of $\ge t + 1$ honest parties can reconstruct $\mathsf{sk} = \hat{h}_1^{f(0)}$ as follows:

1. For each share $\hat{A}_i$ provided, check that $e(A_i, \hat{h}_1) = e(g_1, \hat{A}_i)$, where $A_i = g_1^{f(\omega_i)}$ is part of the PVSS transcript. If this check fails, or if P_i does not provide a share, then remove P_i from S.
2. Return, $\mathsf{sk} = \prod_{i \in S} \hat{A}_i^{\ell_{S,i}(0)}$ where $\ell_{S,i}(X)$ is a Lagrange polynomial equal to 0 at $\omega_j \in S$ for $i \neq j$, and 1 at ω_i.

5 Distributed Key Generation

In this section, we describe our distributed key generation (DKG) protocol for generating a key-pair $(\mathsf{pk}, \mathsf{sk})$ of the form

$$\mathsf{pk} = (g_1^a, \hat{u}_1^a) \in \mathbb{G}_1 \times \mathbb{G}_2 \quad \text{and} \quad \mathsf{sk} = \hat{h}_1^a \in \mathbb{G}_2, \text{ where } a \in \mathbb{F}$$

We often refer to $a \in \mathbb{F}$ as the *DKG secret*. All parties P_i use the same Scrape CRS (see Section 4) but augmented with *verification keys* vk_i (defined later).

At a high level, our DKG protocol resembles previous protocols based on verifiable secret sharing: each party P_i deals a secret $\hat{h}_1^{c_i}$ to all other parties using the Scrape PVSS from Section 4. Additionally, each party P_i includes a proof-of-knowledge of their secret c_i. At this point, each party P_j would have to verify the PVSS transcript of every other party P_i, resulting in $O(n^2)$ work. Then, the final secret would be $\mathsf{sk} = \hat{h}_1^a$ with $a = \sum_{i \in Q} c_i$, where Q is the set of all parties who dealt honestly (i.e., whose PVSS transcript verified). Note that since PVSS transcripts are publicly-verifiable, all parties P_i agree on Q and there is no need for a complaint round. We often refer to an honest party P_i as having *contributed* to the final secret key and to c_i as its *contribution*.

Gossip and aggregate. To avoid the $O(n^2)$ verification work per party, we leverage aggregation of Scrape PVSS transcripts. We observe that a party who verified several transcripts can aggregate them into a single one and forward it to another party, who can now verify this aggregated transcript faster. By carefully aggregating and *gossiping* transcripts in this manner, we decrease verification time per party from $O(n^2)$ to $O(n \log^2 n)$. One caveat is that, due to the randomized nature of gossiping, a party's contribution c_i might be incorporated multiple times, say w_i times, into the final secret $\mathsf{sk} = \hat{h}_1^a$. As a result, the final $a = \sum_{i \in Q} w_i c_i$, where w_i is called the *weight* of each c_i.

Signatures-of-knowledge of contributions. Similar to previous DKGs [33], our DKG requires each party P_i to prove knowledge of its contribution c_i to the final DKG secret. However, since our DKG transcripts must be publicly-verifiable, we also require each party to sign their contributions. We achieve both of these goals using a *signature-of-knowledge (SoK)*. Specifically, P_i signs $C_i = g_1^{c_i}$ using its secret key sk_i, with corresponding *verification key* $\mathsf{vk}_i = g_1^{\mathsf{sk}_i}$:

$$\sigma_i = (\sigma_{i,1}, \sigma_{i,2}) = (\mathsf{Hash}_{\mathbb{G}_2}(C_i)^{c_i}, \mathsf{Hash}_{\mathbb{G}_2}(\mathsf{vk}_i, C_i)^{\mathsf{sk}_i})$$

where $\mathsf{Hash}_{\mathbb{G}_2}$ is a hash function that maps to $\mathbb{G}_2$. Any verifier with vk_i can verify the signature-of-knowledge σ_i of c_i as:

$$e(C_i, \mathsf{Hash}_{\mathbb{G}_2}(C_i)) = e(g_1, \sigma_{i,1}) \ \wedge \ e(\mathsf{vk}_i, \mathsf{Hash}_{\mathbb{G}_2}(\mathsf{vk}_i, C_i)) = e(g_1, \sigma_{i,2})$$

Our signatures of knowledge are simulation-sound and thus cannot be compressed or combined. However, since they are constant-sized, this is not problematic. We refer to the signing algorithm as $\mathsf{SoK.Sign}(C_i, \mathsf{sk}_i, c_i) \to \sigma_i$ and the verification algorithm as $\mathsf{SoK.Verify}(\mathsf{vk}_i, C_i, \sigma_i) \to 0/1$.

DKG transcripts. To maintain their public-verifiability, aggregated PVSS transcripts must keep track of the weights w_i of each party's contribution c_i and of the σ_i's. This gives rise to a new notion of a *DKG transcript* defined as:

$$\mathsf{transcript} = ((C_1, \ldots, C_n), (w_1, \ldots, w_n), (\sigma_1, \ldots, \sigma_n), \mathsf{pvss}), \tag{1}$$

where $C_i = g_1^{c_i}$ is a commitment to the contribution c_i of party P_i, w_i is its weight, σ_i is the SoK of c_i and pvss is an (aggregated) PVSS transcript for secret $a = \sum_{i \in [n]} w_i c_i$.

Recall from Fig. 1 that pvss stores a Feldman commitment $\boldsymbol{F}$ to a polynomial $f(X)$ with $f(0) = a$ and that $F_0 = g_1^a$. In our protocol, each party P_i initializes their DKG transcript by picking $c_i \xleftarrow{\$} \mathbb{F}$ and setting $\mathsf{pvss} \leftarrow \mathsf{Scrape.Deal}(\mathsf{bp}, \mathbf{ek}, \hat{u}_1, c_i)$, $C_i \leftarrow g_1^{c_i}$, $w_i \leftarrow 1$ and $\sigma_i \leftarrow \mathsf{SoK.Sign}(C_i, \mathsf{sk}_i, c_i)$. For $j \neq i$, P_i sets $C_j \leftarrow \bot$, $w_j \leftarrow 0$ and $\sigma_j \leftarrow \bot$. Importantly, in our protocol, each party will broadcast the C_i commitment to their contribution and gossip the rest of their DKG transcript to a subset of the other parties (we discuss this in more detail later on).

Verifying DKG transcripts. To verify the DKG $\mathsf{transcript}$ from Eq. (1), one first checks that its inner pvss transcript verifies. Second, for all *non-trivial contributions* with $w_i \neq 0$, one first checks if their signature of knowledge σ_i verifies. Finally, one checks that the contributions correctly combine to the commitment F_0 to the zero coefficient of $f(X)$ shared in pvss; i.e., that $C_1^{w_1} \cdots C_n^{w_n} = F_0$. If $\mathsf{transcript}$ passes these checks, then one can be sure that the players P_i which have $w_i \neq 0$ in $\mathsf{transcript}$ have contributed to its corresponding DKG secret. See Fig. 2 for a full description.

Aggregating DKG transcripts. Given two input DKG transcripts

$$(C_{b,1}, \ldots, C_{b,n}), (w_{b,1}, \ldots, w_{b,n}), (\sigma_{b,1}, \ldots, \sigma_{b,n}), \mathsf{pvss}_b, \text{ for } b \in \{1, 2\}$$

we can easily aggregate them into a single DKG transcript

$$(C_1, \ldots, C_n), (w_1, \ldots, w_n), (\sigma_1, \ldots, \sigma_n), \mathsf{pvss}$$

We first aggregate the pvss_b transcripts into pvss via $\mathsf{Scrape.Aggregate}$ (see Fig. 1). Second, we aggregate the weights, which are field elements, as $w_i = w_{1,i} + w_{2,i}, \forall i \in [n]$. Third, if P_i contributed in one of the input transcripts, then P_i's contribution should also be reflected in the aggregated transcript. In other words, for any $C_{b,i} \neq \bot$ and valid $\sigma_{b,i}$, we simply set $C_i = C_{b,i}$ and $\sigma_i = \sigma_{b,i}$. The choice of $C_{b,i}$ does not matter when they are both $\neq \bot$ since they were both obtained from the broadcast channel, so they must be equal. As a result, their corresponding $\sigma_{b,i}$'s will also be equal since our signatures of knowledge are unique.

Reconstructing the secret. As explained in the beginning of this section, the final key-pair will be $\mathsf{pk} = (g_1^{f(0)}, \hat{u}_2) = (g_1^{f(0)}, \hat{u}_1^{f(0)})$ and $\mathsf{sk} = \hat{h}_1^{f(0)}$. Since the final DKG transcript is just an augmented Scrape PVSS transcript, reconstruction of sk works as explained in Section 4.

$\underline{\mathsf{DKG.Aggregate}(\mathsf{bp}, \mathsf{transcript}_1, \mathsf{transcript}_2) \to \mathsf{transcript}}$

$((C_{1,1}, \ldots, C_{1,n}), (w_{1,1}, \ldots, w_{1,n}), (\sigma_{1,1}, \ldots, \sigma_{1,n}), \mathsf{pvss}_1) \leftarrow \mathsf{parse}(\mathsf{transcript}_1)$
$((C_{2,1}, \ldots, C_{2,n}), (w_{2,1}, \ldots, w_{2,n}), (\sigma_{2,1}, \ldots, \sigma_{2,n}), \mathsf{pvss}_2) \leftarrow \mathsf{parse}(\mathsf{transcript}_2)$

for $1 \leq i \leq n$:
  $w_i \leftarrow w_{1,i} + w_{2,i}$
  if $\sigma_{1,i} \neq \perp$: $\sigma_i \leftarrow \sigma_{1,i}$, else: $\sigma_i \leftarrow \sigma_{2,i}$
  if $C_{1,i} \neq \perp$: $C_i \leftarrow C_{1,i}$, else: $C_i \leftarrow C_{2,i}$

$\mathsf{pvss} \leftarrow \mathsf{Scrape.Aggregate}(\mathsf{bp}, \mathsf{pvss}_1, \mathsf{pvss}_2)$
return $(C_1, \ldots, C_n), (w_1, \ldots, w_n), (\sigma_1, \ldots, \sigma_n), \mathsf{pvss}$

$\underline{\mathsf{DKG.Verify}(\mathsf{bp}, (\mathsf{ek}_i, \mathsf{vk}_i)_{i \in [n]}, \hat{u}_1, \mathsf{transcript}) \to 0/1}$

$((C_1, \ldots, C_n), (w_1, \ldots, w_n), (\sigma_1, \ldots, \sigma_n), \mathsf{pvss}) \leftarrow \mathsf{parse}(\mathsf{transcript})$
$((F_0, \ldots, F_t), \hat{u}_2, (A_1, \ldots, A_n), (\hat{Y}_1, \ldots, \hat{Y}_n)) \leftarrow \mathsf{parse}(\mathsf{pvss})$
check $\mathsf{Scrape.Verify}(\mathsf{bp}, (\mathsf{ek}_1, \ldots, \mathsf{ek}_n), \hat{u}_1, \hat{u}_2, \mathsf{pvss}) = 1$

for $1 \leq i \leq n$:
  if $w_i \neq 0$: check $\mathsf{SoK.Verify}(\mathsf{vk}_i, C_i, \sigma_i) = 1$

check $C_1^{w_1} \cdots C_n^{w_n} = F_0$
return 1 if all checks pass, else return 0

Fig. 2. Aggregation algorithm for the distributed key generation protocol.

5.1 A gossip protocol

In Step 4 of our DKG, we rely on a gossip protocol to communicate the $\mathcal{O}(n)$-sized DKG transcripts. By using gossip, we avoid both the need to broadcast these larger messages, which is expensive, and the need for a central aggregator. We detail our protocol in Appendix B, but provide some insight here into how it works.

We take an optimistic approach and provide robustness for up to $t_r < n/2 - \log n$ crashed parties but only up to $\log n$ Byzantine adversaries. We believe this approach is often reasonable in practice because if a Byzantine adversary attacks the robustness of a DKG, the only outcome is that the computation required to output the DKG is higher. Furthermore, Byzantine attacks on robustness are detectable, so any faulty party can be manually removed from the system. This is in contrast to an attack on the security preservation of the DKG, which could have far more serious consequences. If we want a security threshold of t_s, then we have to assume that t_s parties respond. A direct implication is that $n - t_r$ must be at least t_s, showing an inherent tradeoff between the security and robustness thresholds. In our scheme we can set t_s to be exactly equal to $n - t_r$.

The gossip protocol has each party send its currently aggregated DKG transcript to $\mathcal{O}(c \log n)$ parties in expectation in each round, and terminate when it has agreed on a *"full" transcript*; i.e., a valid transcript with at least $t_s + 1$ contributions. Here c is a small success parameter such that $c \geq 4$. However,

Our aggregatable DKG protocol

Common reference string: Scrape CRS consiting of $\mathsf{bp} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1)$, encryption and verficiation keys $(\mathsf{ek}_i, \mathsf{vk}_i)_{i\in[n]}$, nth roots of unity $(\omega_i)_{i\in[n]}$ in $\mathbb{F}$, random $\hat{u}_1 \in \mathbb{G}_2$ such that nobody knows $\log_{\hat{h}_1}(\hat{u}_1)$.

Party P_i's private input: Decryption key dk_i for ek_i and secret key sk_i for vk_i.

1. Each P_i picks random $c_i \in \mathbb{F}$, computes $C_i = g_1^{c_i}$ and broadcasts C_i.
2. Each P_i picks random polynomial $f_i(X) \in \mathbb{F}[X]$ of degree at most t

$$f_i(X) = a_{i,0} + a_{i,1}X + \cdots + a_{i,t}X^t$$

such that $a_{i,0} = c_i$. They compute $f_i(\omega_j)$ for $j \in [n]$. Each party *gossips* (see Section 5.1) their DKG transcript consisting of (1) $F_{i,k} = g_1^{a_{i,k}}$ for $k \in [0, t]$; (2) $\hat{u}_{i,2} = \hat{u}_1^{c_i}$; (3) a vector $\vec{w}_i$ such that $w_{i,j} = 1$ if $i = j$ and 0 otherwise; (4) $A_{i,j} = g_1^{f_i(\omega_j)}$ for $j \in [n]$; (5) $Y_{i,j} = \mathsf{ek}_j^{f_i(\omega_j)}$ for $j \in [n]$; and (6) a vector $\vec{\sigma}_i$ such that $\sigma_{i,j} = (\mathsf{Hash}_{\mathbb{G}_2}(C_i)^{c_i}, \mathsf{Hash}_{\mathbb{G}_2}(\mathsf{vk}_i, C_i)^{\mathsf{sk}_i})$ if $i = j$ and $\perp$ otherwise.
3. During the gossip phase, each P_i verifies the transcripts it receives using $\mathsf{DKG.Verify}$ (see Fig. 2). If two transcripts verify, it aggregates them using $\mathsf{DKG.Aggregate}$ (also in Fig. 2), and gossips the aggregated transcript. The aggregated transcripts contain a list of weights $(w_1, \ldots, w_n) \in \mathbb{F}^n$ indicating how many times each party has contributed to the current transcript. When a party receives a *"full" transcript* with $\geq t + 1$ non-zero weights, it broadcasts this as a candidate final transcript.
4. Parties terminate in the round where they first broadcast a "full" transcript. If several candidate "full" transcripts were broadcast, the one whose pk has the lowest bit count is chosen as the final one. The final $\mathsf{pk} = (\prod_{i=1}^n C_i^{w_i}, \prod_{i=1}^n \hat{u}_{i,2}^{w_i})$. Each party computes their secret key share as $Y_i^{\mathsf{dk}_i^{-1}}$ such that they can reconstruct.

Fig. 3. Our DKG with reconstruction threshold $t + 1$ run by parties $P_1, \ldots, P_n$.

deciding when to terminate is non-trivial, because the aggregated "full" transcripts may all be different. We thus still rely on broadcast to agree on which transcript to use, but our goal is to minimize the number of total broadcasts. We do this by having each party with a full transcript broadcast it with probability $2/n$ in a given round. We argue that this makes the protocol likely to terminate within $\mathcal{O}(c \log n)$ rounds. Parties agree to use the transcript whose public key has a binary representation with the smallest bit-count (but any other publicly-verifiable convention works too). In terms of complexity, our gossip protocol requires $\mathcal{O}(cn^2 \log n)$ total words to be communicated in private messages and $\mathcal{O}(c \log^2 n)$ broadcasts.

5.2 Security analysis

Robustness Our DKG is robust in the sense that all honest parties agree on the final public key, and in the sense that any set S containing at least $t+1$ honest parties can reconstruct the secret key.

Theorem 1 (DKG is robust). *The scheme in Figure 3 is robust for any primitive with keys of the form* $\mathsf{pk} = (g_1^a, \hat{u}_1^a) \in \mathbb{G}_1 \times \mathbb{G}_2$.

Proof. First we show that all honest parties have the same value pk. By perfect synchrony we have that in each round all honest parties agree on a completing set of broadcasts. From the broadcast messages that complete and verify, one must have the most sparse binary decomposition. This message defines a public key pk that all parties agree on.

We show that reconstruction always succeeds on input of n shares where at least $t+1$ are input by non-faulty parties. First observe that if the DKG transcript verifies, then for some random value α we have that

$$A_1^{\ell_1(\alpha)} \cdots A_n^{\ell_n(\alpha)} = F_0 F_1^{\alpha} \cdots F_t^{\alpha^t}$$

By the Schwartz-Zippel Lemma this implies that with overwhelming probability

$$f(X) = f_0 + f_1 X + \cdots + f_t X^t = a_1 \ell_1(X) + \cdots + a_n \ell_n(X)$$

and $a_i = f(\omega_i)$. Second observe that $e(A_i, \hat{h}_1) = e(g_1, \hat{A}_i)$ if and only if $\hat{A}_i = \hat{h}_1^{f(\omega_i)}$. Where at least $t+1$ parties are honest the reconstruction algorithm receives at least $t+1$ verifying shares. With $t+1$ verifying shares the reconstruction algorithm always succeeds because f has degree t.

Security preserving We now prove that our DKG satisfies key expressability; i.e., we construct a simulator that is able to fix the output to be a value $\alpha \mathsf{pk}_1 + \mathsf{pk}_2$, where pk_1 is given as input and $\alpha \neq 0$. This does not directly prove that the DKG preserves security, but in Appendix C we detail how combining a key-expressable DKG with rekeyable encryption schemes, signature schemes, and VUFs implies that the DKG also preserves security of these primitives. We cover

these three due to their popularity (and our VUF construction in Section 7), but envisage that there are many other primitives that are rekeyable and thus similarly preserve their security when combined with key-expressable DKGs.

Theorem 2 (DKG). *The scheme in Figure 3 is key-expressable as per Definition 7 in the random oracle model for any primitive with keys of the form* $\mathsf{pk} = (g_1^a, \hat{u}_1^a) \in \mathbb{G}_1 \times \mathbb{G}_2$ *and* $\mathsf{sk} = \hat{h}_1^a \in \mathbb{G}_2$.

Proof. We design an adversary $\mathcal{B}$ that takes as input pk_1 such that whenever the DKG outputs pk, $\mathcal{B}$ outputs $\alpha, \mathsf{pk}_2, \mathsf{sk}_2$ such that $\mathsf{pk} = \alpha\mathsf{pk}_1 + \mathsf{pk}_2$. Suppose $\mathcal{B}$ receives input $\mathsf{pk}_1 = (g_2, \hat{v}_2)$.

First $\mathcal{B}$ runs the DKG with $\mathcal{A}$. Let $\mathbb{I}_B \subset [1, n]$ be the set of corrupted (i.e. "bad") parties and $\mathbb{I}_G \subset [1, n]$ be the set of uncorrupted ("good") parties. For good parties P_k, $\mathcal{B}$ *simulates* the adversarial view of this party's output, so that public view $C_k, \hat{u}_{k,2}$ sent by P_k is equal to $(g_2^{a_k}, \hat{v}_2^{a_k})$.

In the course of this simulation, $\mathcal{B}$ answers $\mathcal{A}$'s queries to the oracle $\mathsf{Hash}_{\mathbb{G}_2}$ by selecting $r \xleftarrow{\$} \mathbb{F}$ at random, and returning $\hat{h}_1^r$.

In the registration round, when $\mathcal{A}$ queries $\mathcal{B}$ on the k-th honest value, $\mathcal{B}$ chooses $\mu_k, \kappa_k \xleftarrow{\$} \mathbb{F}$ randomly from the field and returns the public key $(\mathsf{ek}_k, \mathsf{vk}_k) = (\hat{u}_1^{\mu_k}, g_2^{\kappa_k})$.

In the broadcast round, $\mathcal{B}$ chooses $a_k \xleftarrow{\$} \mathbb{F}$ randomly for each honest party and computes $C_k = g_2^{a_k}$. It then samples $\chi_k, \psi_k \xleftarrow{\$} \mathbb{F}$ and programs $\mathsf{Hash}_{\mathbb{G}_2}$ to return $\hat{u}_1^{\chi_k}$ and $\hat{u}_1^{\psi_k}$ on input C_k and (vk_k, C_k) respectively. Finally it broadcasts C_k. With overwhelming probability, $\mathcal{A}$ is yet to query the randomised value C_k.

In the share creation round, when queried on P_k, $\mathcal{B}$ is required to output

$$(\boldsymbol{F}_k, \hat{u}_{k,2}, \hat{\sigma}_k, \boldsymbol{A}_k, \hat{\boldsymbol{Y}}_k)$$

that are indistinguishable from a valid output. Assume without loss of generality that $|\mathbb{I}_B| = t$. It then behaves as follows

1. Choose random $\bar{x}_{k,j} \xleftarrow{\$} \mathbb{F}$ for each $j \in \mathbb{I}_B$ and interpolate in the exponent to find $(F_{k,0}, \ldots, F_{k,t})$ such that $F_{k,i} = g_1^{c_i}$, where $\sum_{i=0}^{t} c_i X^i$ evaluates to $\bar{x}_{k,j}$ at ω_j for $j \in \mathbb{I}_B$ and $a_k \log_{g_1}(g_2)$ at 0. These c_i values are unknown to $\mathcal{B}$.
2. Set $\hat{u}_{k,2} = \hat{v}_2^{a_k}$.
3. Set $\sigma_k = (\hat{v}_2^{a_k\chi_k}, \hat{v}_2^{\kappa_k\psi_k})$.
4. To compute $A_{k,1}, \ldots, A_{k,n}$, set $A_{k,j} = \prod_{i=0}^{t} F_{k,i}^{\omega_j^i}$.
5. To compute $\hat{Y}_{k,j}$ for $j \in \mathbb{I}_B$, return $\mathsf{ek}_j^{\bar{x}_{k,j}}$. To compute $\hat{Y}_{k,j}$ for $j \in \mathbb{I}_G$, interpolate in the exponent to find $\hat{u}_1^{c_0}, \ldots, u_1^{c_t}$ for $c_0, \ldots, c_{t-1}$ as in Step 1 (recall that $\mathcal{B}$ knows $\hat{u}_1^{\log_{g_1}(g_2)}$). Return $\hat{Y}_{k,j} = \prod_{i=0}^{t} \hat{u}_1^{c_i\mu_j\omega_j^i}$.

This simulation is perfect. Indeed $c_0 = \log_{g_1}(C_k)$ and $c_1, \ldots, c_t$ are randomly distributed. We have that $\hat{u}_{k,2} = \hat{v}_2^{a_k(\nu+1)} = \hat{u}_1^{\log_{g_1}(C_k)}$. Also, $\sigma_{k,1} = \mathsf{Hash}_{\mathbb{G}_2}(C_k)^{\log_{g_1}(C_k)}$ and $\sigma_{k,2} = \mathsf{Hash}_{\mathbb{G}_2}(\mathsf{vk}_k, C_k)^{\log_{g_1}(\mathsf{vk}_k)}$. The values $A_{k,1}, \ldots, A_{k,n}$ are computed honestly and are the unique encryptions that satisfy the verifier.

Suppose that the DKG terminates with transcript

$$((C'_1, \ldots, C'_n), (w_1, \ldots, w_n), (\sigma'_1, \ldots, \sigma'_n), \mathsf{pvss}) \ .$$

The public key is given by

$$C = C'_1 \cdots C'_n, \hat{u}_2 = \hat{u}_{1,2} \cdots \hat{u}_{n,2}$$

For each adversarial contribution C'_i, $\mathcal{B}$ looks up r such that $\mathsf{Hash}_{\mathbb{G}_2}(C'_j) = \hat{h}_1^r$. Here, i can be any index, as $\mathcal{A}$ might have forged one of $\mathcal{B}$'s contributions. If the adversary has not queried $\mathsf{Hash}_{\mathbb{G}_2}$ on C' then the probability of them returning a verifying signature σ is negligible. To get the secret key share, $\mathcal{B}$ extracts $\hat{C}_i = \hat{\sigma}^{\frac{1}{r}}$ such that $\hat{C}_i = \hat{h}_1^{\log_{g_1}(C_i)}$.

If $\mathcal{A}$ has included at least one of $\mathcal{B}$'s contributions, then $\mathcal{B}$ computes $z = \sum_{k \in S} w_k$ for S the set of honest participants whose contribution is included in the transcript. Additionally, $\mathcal{B}$ computes $\mathsf{pk}_2 = (\prod_{i \notin S} C'_i, \prod_{i \notin S} \hat{u}_{i,2})$ and $\mathsf{sk}_2 = \prod_{i \notin S} \hat{C}_i$. Then, we have that $\mathsf{pk} = \alpha \mathsf{pk}_1 + \mathsf{pk}_2$ for $\alpha \neq 0$ and sk_2 is a key for pk_2. Thus $\mathcal{B}$ returns (α, sk_2).

If $\mathcal{A}$ has not included any contributions from $\mathcal{B}$, then that $\mathcal{A}$ has forged a signature σ'_k with respect to some $\mathsf{vk}_k = g_2^{\kappa_k}$ and contribution C'_k. Using the oracle queries, $\mathcal{B}$ looks up r such that $\mathsf{Hash}_{\mathbb{G}_2}(\mathsf{vk}_k, C'_k) = \hat{h}_1^r$. Since $\sigma'_k = (\sigma'_{k,1}, \sigma'_{k,2})$ verifies, we have that $\sigma'_{k,2} = \hat{h}_1^{r\kappa_k \log_{g_1}(g_2)}$. Thus, $\mathcal{B}$ computes $\mathsf{sk}_1 = (\sigma'_{k,2})^{\frac{1}{r\kappa_k}}$. Additionally, $\mathcal{B}$ computes $\mathsf{pk}_2 = (g_2^{-1} \prod_i C'_i, \hat{v}_2^{-1} \prod_i \hat{u}_{i,2})$ and $\mathsf{sk}_2 = \mathsf{sk}_1^{-1} \prod_i \hat{C}_i$. Then, we have that $\mathsf{pk} = \mathsf{pk}_1 + \mathsf{pk}_2$ and sk_2 is a key for pk_2 and $\mathcal{B}$ returns $(1, \mathsf{sk}_2)$.

6 Alternative DKGs Have Provable Security

In this section we demonstrate that two popular DKGs, the Pedersen DKG and the Fouque-Stern DKG, are also key-expressable. As a direct consequence, they can be used to securely instantiate a DKG for both El-Gamal encryption and BLS signatures, as we prove in Appendix D. Our results generalise to other rekeyable constructions that have public keys in $\mathbb{G}$ and secret keys in $\mathbb{F}$. In addition to justifying the applicability of our security definitions and proof techniques, we hope this also fills a gap in the literature as we are unaware of other works that provide correct proofs for these DKGs.

6.1 Pedersen DKG from Feldman's VSS

We prove that key expressability holds for Pedersen's DKG provided the threshold of adversarial participants is less than $n/2$. It is our belief that this bound on the number of adversarial participants can be removed provided that one gives signatures of knowledge of the individual contributions. Pedersen's DKG can be seen as n parallel instantiations of the Feldman VSS [26]. We remind

the reader that key expressability does not imply secrecy (invalidating the attack of Gennaro et al. [32]) but does allow us to prove the security preservation of certain rekeyable schemes. A proof of the following theorem is provided in Appendix E.2.

Theorem 3. *The scheme in Figure 6 is a key-expressable DKG against static adversaries with adversarial threshold $t < n/2$ for any scheme whose key generation outputs values $\mathsf{pk} = g_1^a \in \mathbb{G}_1$, $\mathsf{sk} = a \in \mathbb{F}$.*

6.2 The Fouque-Stern publicly verifiable DKG

We now show the key expressability of the publicly verifiable Fouque Stern DKG [28]. This DKG has the benefit of outputting field elements as secret keys, but the total communication and verification costs are of order $\mathcal{O}(n^2)$. Unlike Fouque and Stern's original argument, we allow for the existence of rushing adversaries. Indeed Fouque and Stern rely in their reduction on an honest party playing last. Instantiating such an assumption would require the use of a trusted third party and therefore negate the benefits of distributing the key generation. A proof of the following theorem is provided in Appendix E.3.

Theorem 4. *The scheme in Figure 7 is a key-expressable DKG in the random oracle model against static adversaries under the decisional composite residuosity assumption for any scheme whose key generation outputs values $\mathsf{pk} = g_1^a \in \mathbb{G}_1$, $\mathsf{sk} = a \in \mathbb{F}$.*

6.3 El-Gamal and BLS

In Appendix D, we observe that El-Gamal encryption and BLS signatures are both rekeyable (and both have field elements as secret keys). We thus obtain the following two corollaries:

Corollary 1. *The El-Gamal encryption scheme is IND-CPA-secure when instantiated with the Pedersen DKG or the Fouque-Stern DKG.*

Corollary 2. *The BLS signature scheme is EUF-CMA-secure when instantiated with the Pedersen DKG or the Fouque-Stern DKG.*

7 A Structure-Preserving VUF

In this section, we introduce a verifiable unpredictable function (VUF), secure in the random oracle model, that has group elements as the secret key. We can thus securely instantiate our VUF using our DKG.

As one application, VUFs can be used to create randomness beacons, where unlike in, *e.g.*, BLS multi-signatures [11], if a threshold of signers is reached, then the same signature is always produced. By hashing the outcome of this VUF with a random oracle we can obtain a verifiable random function (VRF). Abe et al. [1] proved that it is impossible to construct an algebraic VUF with a secret key as a group element. Since we are using a hash function, however, we are not fully algebraic and therefore sidestep this impossibility result.

7.1 Our construction

Our VUF scheme is given in Figure 4. The techniques were inspired by a combination of BLS signatures [13] and Escala-Groth NIZKs [25] (which are an improvement of Groth-Sahai proofs [38]). Unlike BLS signatures our secret keys are group elements and unlike Escala-Groth NIZKs our VUFs are non-malleable.

Given an input $m \in \mathbb{F}$ under public key $g_1^a, \hat{u}_1^a$ and secret key $\hat{h}_1^a$, the unique output given by $\mathsf{VUF.Eval}(\mathsf{sk}, m)$ is $e(\mathsf{Hash}_{\mathbb{G}_1}(m), \hat{h}_1^a)$. Given $g_1^a \in \mathbb{G}_1$ and $\mathsf{Hash}_{\mathbb{G}_1}(m) \in \mathbb{G}_1$, it is hard for an adversary to compute $e(\mathsf{Hash}_{\mathbb{G}_1}(m), \hat{h}_1)^a \in \mathbb{G}_T$. We formally prove in Theorem 5 and 6 that our VUF satisfies uniqueness (see Definition 3) and unpredictability (see Definition 4) under the SXDH and BDH assumptions.

$\underline{\mathsf{VUF.Setup}(\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1})}$
$\hat{u}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4 \xleftarrow{\$} \mathbb{G}_2$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{h}_2, \hat{h}_3, \hat{h}_4)$
return $\mathsf{crs}_{\mathsf{vuf}}$

$\underline{\mathsf{VUF.Gen}(\mathsf{crs}_{\mathsf{vuf}})}$
$a \xleftarrow{\$} \mathbb{F}$, $\mathsf{pk} \leftarrow g_1^a, \hat{u}_1^a \in (\mathbb{G}_1 \times \mathbb{G}_2)$
$\mathsf{sk} \leftarrow \hat{h}_1^a \in \mathbb{G}_2$
return $(\mathsf{pk}, \mathsf{sk})$

$\underline{\mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)}$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
return $e(Z, \mathsf{sk})$

$\underline{\mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)}$
$(\pi_1, \pi_2, \pi_3, \pi_4 \in \mathbb{G}_1^4, \hat{\pi}_1, \hat{\pi}_2 \in \mathbb{G}_2^2) \leftarrow \mathsf{parse}(\sigma)$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
return $e(Z, \hat{\pi}_2) e(\pi_2, \hat{h}_3) e(\pi_4, \hat{h}_4)$

$\underline{\mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)}$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
$\alpha, \beta \xleftarrow{\$} \mathbb{F}$
$\pi_1, \pi_2, \pi_3, \pi_4 \leftarrow g_1^\alpha, Z^\alpha, g_1^\beta, Z^\beta$
$\hat{\pi}_1, \hat{\pi}_2 \leftarrow \hat{h}_1^{-\alpha} \hat{h}_2^{-\beta},\ \hat{h}_3^{-\alpha} \hat{h}_4^{-\beta} \cdot \mathsf{sk}$
return $(\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)$

$\underline{\mathsf{VUF.Ver}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)}$
$(A, \hat{u}_2) \leftarrow \mathsf{parse}(\mathsf{pk})$
$(\pi_1, \pi_2, \pi_3, \pi_4 \in \mathbb{G}_1^4, \hat{\pi}_1, \hat{\pi}_2 \in \mathbb{G}_2^2) \leftarrow \mathsf{parse}(\sigma)$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
check $1 = e(g_1, \hat{\pi}_1) e(\pi_1, \hat{h}_1) e(\pi_3, \hat{h}_2)$
check $1 = e(Z, \hat{\pi}_1) e(\pi_2, \hat{h}_1) e(\pi_4, \hat{h}_2)$
check $e(A, \hat{h}_1) = e(g_1, \hat{\pi}_2) e(\pi_1, \hat{h}_3) e(\pi_3, \hat{h}_4)$
return 1 if all checks pass, else return 0

Fig. 4. Verifiable unpredictable function with group elements as the secret key.

Setup: The setup algorithm is a transparent algorithm that takes as input the bilinear group $\mathsf{bp} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1)$ and returns four group elements in the second source group: $\hat{u}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4 \in \mathbb{G}_2^4$.

KeyGen: The VUF.Gen algorithm takes as input the common reference string. It samples a random field element $a \xleftarrow{\$} \mathbb{F}$. The public key $\mathsf{pk} \in \mathbb{G}_1 \times \mathbb{G}_2$ and the secret key $\mathsf{sk} \in \mathbb{G}_2$ are given as $\mathsf{pk} = (g_1^a, \hat{u}_1^a)$ and $\mathsf{sk} = \hat{h}_1^a$.

Sign: The VUF.Sign algorithm first hashes the message m to obtain $Z \in \mathbb{G}_1$ as $Z = \mathsf{Hash}_{\mathbb{G}_1}(m)$. The signer generates a commitment to sk by sampling random

elements $\alpha, \beta \in \mathbb{F}$ and computing

$$(\hat{\pi}_1, \hat{\pi}_2) = (\hat{h}_1^{-\alpha}\hat{h}_2^{-\beta}, \mathsf{sk} \cdot \hat{h}_3^{-\alpha}\hat{h}_4^{-\beta}).$$

If $\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4$ are randomly distributed, this commitment is perfectly hiding. However, if $\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4$ form an SXDH challenge, then there exists some ξ such that $\hat{h}_3 = \hat{h}_1^{\xi}$ and $\hat{h}_4 = \hat{h}_2^{\xi}$, meaning that the commitment forms an El-Gamal encryption of sk. In this case, we say that the commitment is perfectly binding.

Having generated $(\hat{\pi}_1, \hat{\pi}_2)$, the signer now generates $(\pi_1, \pi_2, \pi_3, \pi_4) \in \mathbb{G}_1^4$ such that

$$(\pi_1, \pi_2, \pi_3, \pi_4) = (g_1^{\alpha}, Z^{\alpha}, g_1^{\beta}, Z^{\beta})$$

These signature elements have been designed such that the random blinders α, β are canceled out in the verifier's equations.

The signer returns the output $\sigma = (\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)$.

Derive: The $\mathsf{VUF.Derive}$ computes $Z = \mathsf{Hash}_{\mathbb{G}_1}(m)$ and then returns

$$T = e(Z, \hat{\pi}_2)e(\pi_2, \hat{h}_3)e(\pi_4, \hat{h}_4)$$

as the unique and unpredictable component. If the signer is honest then $T = e(Z, \mathsf{sk}) = \mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)$.

Verify: The $\mathsf{VUF.Ver}$ algorithm parses the signature to check that $(\pi_1, \pi_2, \pi_3, \pi_4)$ is in $\mathbb{G}_1^4$, and $(\hat{\pi}_1, \hat{\pi}_2)$ is in $\mathbb{G}_2^2$. The verifier computes Z identically to the signer, *i.e.*, $Z = \mathsf{Hash}_{\mathbb{G}_1}(m)$. The verifier then checks that three pairing equations are satisfied in order to be convinced that there exist α, β such that

$$(\pi_2, \pi_4, \hat{\pi}_2) = (Z^{\alpha}, Z^{\beta}, \hat{h}_3^{-\alpha}\hat{h}_4^{-\beta} \cdot \mathsf{sk})$$

Specifically, they check that:

$$1 = e(g_1, \hat{\pi}_1)e(\pi_1, \hat{h}_1)e(\pi_3, \hat{h}_2) \tag{2}$$

$$1 = e(Z, \hat{\pi}_1)e(\pi_2, \hat{h}_1)e(\pi_4, \hat{h}_2) \tag{3}$$

$$e(\mathsf{pk}, \hat{h}_1) = e(g_1, \hat{\pi}_2)e(\pi_1, \hat{h}_3)e(\pi_3, \hat{h}_4) \tag{4}$$

They return 1 if all these checks pass and 0 otherwise.

Given a signature that satisfies these equations, an extractor that knows a trapdoor SXDH relation between the CRS elements can output a valid witness sk. However, there also exists a simulated CRS indistinguishable from random such that we can simulate signatures without knowing sk.

Threshold VUF Scheme We discuss how to transform our VUF into a threshold VUF. The individual VUF shares can be made shorter using our optimisation in Appendix G. Suppose that there are n parties $P_1, \ldots, P_n$ and we want that any $t+1$ of them can jointly sign a message, but that t of them cannot. We

use Shamir's secret sharing scheme and choose a degree t polynomial $f(X)$. Let $\omega_1, \ldots, \omega_n$ denote unique evaluation points and $\ell_{S,1}(X), \ldots, \ell_{S,t+1}(X)$ denote the Lagrange polynomials such that for all $\omega_j \in S$ we have that $\ell_{S,i}(\omega_j)$ is equal to 1 if $i = j$ and 0 otherwise.

The threshold setup algorithm runs identically to the non-threshold version to return $\mathsf{crs}_{\mathsf{vuf}}$. The key generation outputs a public key and n secret key shares of the form

$$\mathsf{pk} = (g_1^{f(0)}, \hat{u}_1^{f(0)}), \mathsf{sk}_1 = \hat{h}_1^{f(\omega_1)}, \ldots, \mathsf{sk}_n = \hat{h}_1^{f(\omega_n)}.$$

To compute their share of the threshold signature on m party P_i outputs

$$\sigma_i = (\pi_{i,1}, \pi_{i,2}, \pi_{i,3}, \pi_{i,4}, \hat{\pi}_{i,1}, \hat{\pi}_{i,2}) \xleftarrow{\$} \mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}_i, m)$$

To aggregate t signature shares on m from parties $\{P_i\}_{i \in S}$ compute

$$\sigma = \left(\prod_{i \in S} \pi_{i,1}^{\ell_{S,i}(0)}, \prod_{i \in S} \pi_{i,2}^{\ell_{S,i}(0)}, \prod_{i \in S} \pi_{i,3}^{\ell_{S,i}(0)}, \prod_{i \in S} \pi_{i,4}^{\ell_{S,i}(0)}, \prod_{i \in S} \hat{\pi}_{i,1}^{\ell_{S,i}(0)}, \prod_{i \in S} \hat{\pi}_{i,2}^{\ell_{S,i}(0)} \right)$$

The verification and derive algorithms run identically to their non-threshold counterparts on the input $(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)$

We briefly show that σ is correct. Set $Z = \mathsf{Hash}_{\mathbb{G}_1}(m)$ and see that $\sigma = (\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)$ is given by

$$\begin{aligned}
\pi_1 &= \textstyle\prod_{i \in S} \pi_{i,1}^{\ell_{S,i}(0)} &&= g_1^{\sum_{i \in S} \alpha_i \ell_{S,i}(0)} \\
\pi_2 &= \textstyle\prod_{i \in S} \pi_{i,2}^{\ell_{S,i}(0)} &&= Z^{\sum_{i \in S} \alpha_i \ell_{S,i}(0)} \\
\pi_3 &= \textstyle\prod_{i \in S} \pi_{i,3}^{\ell_{S,i}(0)} &&= g_1^{\sum_{i \in S} \beta_i \ell_{S,i}(0)} \\
\pi_4 &= \textstyle\prod_{i \in S} \pi_{i,4}^{\ell_{S,i}(0)} &&= Z^{\sum_{i \in S} \beta_i \ell_{S,i}(0)} \\
\hat{\pi}_1 &= \textstyle\prod_{i \in S} \hat{\pi}_{i,1}^{\ell_{S,i}(0)} &&= \hat{h}_1^{-\sum_{i \in S} \alpha_i \ell_{S,i}(0)} \hat{h}_2^{-\sum_{i \in S} \beta_i \ell_{S,i}(0)} \\
\hat{\pi}_2 &= \textstyle\prod_{i \in S} \hat{\pi}_{i,2}^{\ell_{S,i}(0)} &&= \hat{h}_3^{-\sum_{i \in S} \alpha_i \ell_{S,i}(0)} \hat{h}_4^{-\sum_{i \in S} \beta_i \ell_{S,i}(0)} \textstyle\prod_{i \in S} \mathsf{sk}_i^{\ell_{S,i}(0)} \\
&&&= \hat{h}_3^{-\sum_{i \in S} \alpha_i \ell_{S,i}(0)} \hat{h}_4^{-\sum_{i \in S} \beta_i \ell_{S,i}(0)} \hat{h}_1^{f(\omega_i) \ell_{S,i}(0)}
\end{aligned}$$

Since f has degree t we have that $f(\omega_i)\ell_{S,i}(0) = f(0)$. Denote $\alpha = \sum_{i \in S} \alpha_i \ell_{S,i}(0)$ and $\beta = \sum_{i \in S} \beta_i \ell_{S,i}(0)$ in the above equation to get that

$$(\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2) = (g_1^{\alpha}, Z^{\alpha}, g_1^{\beta}, Z^{\beta}, h_1^{-\alpha} h_2^{-\beta}, h_3^{-\alpha} h_4^{-\beta} \hat{h}_1^{f(0)}) \ .$$

Thus the threshold signature is distributed identically to the non-threshold counterpart and the verifier and deriver output 1 and $e(Z, \hat{h}_1)^{f(0)}$, respectively.

Aggregatable Signature Scheme It is also possible to use our VUF to instantiate an aggregatable signature scheme with secret keys as group elements. For aggregating, one simply takes the product of the public key elements output by $\mathsf{VUF.Gen}$ and the signature elements output by $\mathsf{VUF.Sign}$. Similar to the BLS scheme, this aggregatable signature scheme would be susceptible to *rogue key attacks* [50]. It is thus important to provide simulation-extractable proofs of knowledge of secret keys as part of a public key infrastructure.

7.2 Security analysis

To prove that our VUF is secure, we need to prove that it satisfies uniqueness and unpredictability.

Theorem 5. *The VUF in Figure 4 satisfies uniqueness (Definition 3) under the* SXDH *assumption in the random oracle model.*

Theorem 6. *The VUF in Figure 4 satisfies unpredictability (Definition 4) under the* SXDH *and the* BDH *assumption in the random oracle model.*

We provide formal proofs of these theorems in Appendix F. Intuitively, uniqueness relies on the fact that it would be statistically impossible to satisfy the verifiers equations for a wrong evaluation if the CRS was made up of an SXDH instance. Since VUF.Eval is deterministic there can only be one correct evaluation. Thus if an adversary could break uniqueness in the general case, then we could use them as a subroutine to determine SXDH instances from random.

Our unpredictability proof uses an adversary who predicts the VUF to compute a BDH output. To do this we embed one component of the BDH challenge into the public key being targeted, and the other into the adversaries random oracle queries. However, we also need to simulate responses to the adversaries signature requests, and to do this (after jumping to a hybrid game with a structured CRS) we need to program the oracle such that we know a discrete log. This could present a collision as the adversary may have already queried that point. To counteract, we take a random guess as to which oracle query the adversary will output their prediction for, and if we guess wrong we abort. Thus our reduction is not tight, but does provide us with a polynomial chance of success whenever the adversary succeeds.

After observing that our VUF is rekeyable, we prove the following corollary in Appendix D.

Corollary 3. *The VUF in Figure 4 is unique and unpredictable when instantiated with the DKG in Figure 3.*

8 Implementation

We implement our DKG and VUF and summarise the performance of our schemes in Tables 2 and 3. Our implementation is written in Rust on top of the `libzexe` library, which performs efficient finite field arithmetic, elliptic curve arithmetic, and finite field FFTs. We evaluate our DKG and VUF on a desktop machine with an *i*7-8700k CPU at 3.7GHz and 32GB of DDR4 RAM. We use the BLS12-381 curve. For hashing to groups, we use the try-and-reject method by instantiating a ChaCha20 RNG with a Blake2s hash of the input message, sampling field elements and checking if they are valid x-coordinates, deriving the corresponding point if so. Our implementation is not constant-time. Upon publication, we plan to release our implementation as open-source software.

Parties	DKG.Deal (ms)	Scrape.Verify (ms)	DKG.Verify (ms)	Transcript size (kB)
64	72	96	376	25
128	124	178	704	50
256	271	346	1305	99
8192	8000	9900	42 600	3146

Table 2. The performance of our DKG, averaged across 10 samples of each operation. For n parties, we use a threshold of $t = 2n/3$.

	Our VUF	Our optimised VUF	BLS [13]
Key prove (ms)	-	2.89	-
Public key (bytes)	48	336	96
Key verify (ms)	-	4.00	-
Sign (ms)	3.47	0.58	0.44
Signature size (bytes)	384	96	48
Verify (ms)	4.73	2.39	2.15
Derive (ms)	2.37	2.37	-

Table 3. The performance of our VUF (Section 7), our optimised VUF (Appendix G), and the BLS signature scheme. These numbers were averaged across four distinct runs, with 100 samples of each operation per run.

We utilise a few optimization techniques throughout the implementation. First, when verifying multiple pairing equations, we instead compute a randomised check of a single pairing equation so as to amortise the cost of the final exponentiations. We then compute the pairing product efficiently using the underlying `libzexe` implementation. In the same vein, when verifying pairing equations where two pairings are computed with respect to the same source group element, we combine the two into a randomised check. For large multi-exponentiations we use the `libzexe` implementation of Pippenger's algorithm. For large polynomial evaluations we use FFTs. We additionally utilise batch normalization of projective points.

We evaluate our DKG with respect to 64, 128, 256, and 8192 parties. We see that the time taken to compute, verify, and aggregate a transcript all increase linearly in the number of parties. Verifying a transcript with 256 parties takes a little more than a second.

In addition to our VUF presented in Section 7, we also evaluate an optimised VUF that we present in Appendix G. We compare the performance of our VUF and our optimised VUF with BLS [13], which is the state of the art in the random oracle model. We do not give the derivation time for BLS because this is the identity function. It can be seen that signing and verifying our optimised VUF

is only fractionally more expensive than BLS, but that verifying our full VUF is approximately twice as expensive.

Acknowledgements

Thank you to Ittai Abraham for helpful discussions and feedback. Sarah Meiklejohn was supported in part by EPSRC Grant EP/N028104/1. Gilad Stern was supported by the HUJI Federmann Cyber Security Research Center in conjunction with the Israel National Cyber Directorate (INCD) in the Prime Minister's Office.

References

[1] M. Abe, J. Camenisch, R. Dowsley, and M. Dubovitskaya. "On the Impossibility of Structure-Preserving Deterministic Primitives". In: *J. Cryptology* 32.1 (2019), pp. 239–264.

[2] M. Abe, M. Chase, B. David, M. Kohlweiss, R. Nishimaki, and M. Ohkubo. "Constant-Size Structure-Preserving Signatures: Generic Constructions and Simple Assumptions". In: *J. Cryptol.* 29.4 (2016), pp. 833–878.

[3] M. Abe, J. Groth, M. Kohlweiss, M. Ohkubo, and M. Tibouchi. "Efficient Fully Structure-Preserving Signatures and Shrinking Commitments". In: *J. Cryptol.* 32.3 (2019), pp. 973–1025.

[4] M. Abe, J. Groth, M. Ohkubo, and M. Tibouchi. "Unified, Minimal and Selectively Randomizable Structure-Preserving Signatures". In: *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*. 2014, pp. 688–712.

[5] I. Abraham, D. Dolev, R. Gonen, and J. Y. Halpern. "Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation". In: *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing, PODC 2006, Denver, CO, USA, July 23-26, 2006*. 2006, pp. 53–62.

[6] I. Abraham, D. Malkhi, and A. Spiegelman. "Asymptotically Optimal Validated Asynchronous Byzantine Agreement". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. PODC '19. 2019.

[7] G. Ateniese, J. Camenisch, and B. de Medeiros. "Untraceable RFID tags via insubvertible encryption". In: *Proceedings of the 12th ACM Conference on Computer and Communications Security, CCS 2005, Alexandria, VA, USA, November 7-11, 2005*. 2005, pp. 92–101.

[8] L. Ballard, M. Green, B. de Medeiros, and F. Monrose. *Correlation-Resistant Storage via Keyword-Searchable Encryption*. IACR Cryptol. ePrint Arch. 2005/417. 2005.

[9] M. Bellare and P. Rogaway. "The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs". In: *Advances in Cryptology - EUROCRYPT 2006*. Springer Berlin Heidelberg, 2006, pp. 409–426.

[10] F. Benhamouda, T. Lepoint, M. Orrù, and M. Raykova. *On the (in)security of ROS*. Cryptology ePrint Archive, Report 2020/945. 2020.

[11] A. Boldyreva. "Threshold Signatures, Multisignatures and Blind Signatures Based on the Gap-Diffie-Hellman-Group Signature Scheme". In: *PKC 2003*. Ed. by Y. G. Desmedt. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 31–46.

[12] D. Boneh and X. Boyen. "Efficient Selective Identity-Based Encryption Without Random Oracles". In: *J. Cryptology* 24.4 (2011), pp. 659–693.

[13] D. Boneh, B. Lynn, and H. Shacham. "Short Signatures from the Weil Pairing". In: *Advances in Cryptology - ASIACRYPT 2001*. 2001, pp. 514–532.

[14] A. Bonnecaze and P. Trebuchet. "Threshold Signature for Distributed Time Stamping Scheme". In: *Ann. Telecommun.* 62 (2007), pp. 1353–1364.

[15] S. Bowe, A. Gabizon, and I. Miers. "Scalable Multi-party Computation for zk-SNARK Parameters in the Random Beacon Model". In: *IACR Cryptol. ePrint Arch.* 2017 (2017), p. 1050. URL: http://eprint.iacr.org/2017/1050.

[16] R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. "Adaptive Security for Threshold Cryptosystems". In: *Advances in Cryptology — CRYPTO' 99*. Ed. by M. Wiener. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 98–116. ISBN: 978-3-540-48405-9.

[17] J. Canny and S. Sorkin. "Practical Large-Scale Distributed Key Generation". In: *Advances in Cryptology - EUROCRYPT 2004*. Ed. by C. Cachin and J. L. Camenisch. Springer Berlin Heidelberg, 2004, pp. 138–152.

[18] I. Cascudo and B. David. "SCRAPE: Scalable Randomness Attested by Public Entities". In: *Applied Cryptography and Network Security - 15th International Conference, ACNS 2017, Kanazawa, Japan, July 10-12, 2017, Proceedings*. 2017, pp. 537–556.

[19] B. Chor, S. Goldwasser, S. Micali, and B. Awerbuch. "Verifiable secret sharing and achieving simultaneity in the presence of faults". In: *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*. 1985, pp. 383–395.

[20] A. De Santis, Y. Desmedt, Y. Frankel, and M. Yung. "How to Share a Function Securely". In: *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*. STOC '94. 1994, pp. 522–533.

[21] A. Demers et al. "Epidemic Algorithms for Replicated Database Maintenance". In: *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing*. PODC '87. 1987, pp. 1–12.

[22] Y. Desmedt and Y. Frankel. "Threshold cryptosystems". In: *Advances in Cryptology — CRYPTO'89*. 1990, pp. 307–315.

[23] DFINITY. *Distributed Key Generation in JS*.

[24] Y. Dodis and A. Yampolskiy. "A Verifiable Random Function with Short Proofs and Keys". In: *Public Key Cryptography - PKC 2005*. Ed. by S. Vaudenay. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 416–431. ISBN: 978-3-540-30580-4.

[25] A. Escala and J. Groth. "Fine-Tuning Groth-Sahai Proofs". In: *Public-Key Cryptography - PKC 2014 - 17th International Conference on Practice and Theory in Public-Key Cryptography, Buenos Aires, Argentina, March 26-28, 2014. Proceedings*. 2014, pp. 630–649.

[26] P. Feldman. "A Practical Scheme for Non-interactive Verifiable Secret Sharing". In: *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*. SFCS '87. IEEE Computer Society, 1987, pp. 427–438.

[27] N. Fleischhacker, J. Krupp, G. Malavolta, J. Schneider, D. Schröder, and M. Simkin. "Efficient Unlinkable Sanitizable Signatures from Signatures with Re-Randomizable Keys". In: *Proceedings of PKC 2016*. 2016.

[28] P. Fouque and J. Stern. "One Round Threshold Discrete-Log Key Generation without Private Channels". In: *Public Key Cryptography, 4th International Workshop on Practice and Theory in Public Key Cryptography, PKC 2001, Cheju Island, Korea, February 13-15, 2001, Proceedings*. 2001, pp. 300–316.

[29] S. D. Galbraith, K. G. Paterson, and N. P. Smart. "Pairings for cryptographers". In: *Discrete Applied Mathematics* 156.16 (2008). Applications of Algebra to Cryptography, pp. 3113 –3121.

[30] D. Galindo, J. Liu, M. Ordean, and J.-M. Wong. *Fully Distributed Verifiable Random Functions and their Application to Decentralised Random Beacons*. Cryptology ePrint Archive, Report 2020/096. 2020.

[31] J. A. Garay, J. Katz, U. Maurer, B. Tackmann, and V. Zikas. "Rational Protocol Design: Cryptography against Incentive-Driven Adversaries". In: *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*. 2013, pp. 648–657.

[32] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. "Secure Distributed Key Generation for Discrete-Log Based Cryptosystems". In: *Journal of Cryptology* 20 (2007), pp. 51–83.

[33] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. "Secure Applications of Pedersen's Distributed Key Generation Protocol". In: *Topics in Cryptology - CT-RSA 2003, The Cryptographers' Track at the RSA Conference 2003, San Francisco, CA, USA, April 13-17, 2003, Proceedings*. 2003, pp. 373–390.

[34] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. "Algorand: Scaling Byzantine Agreements for Cryptocurrencies". In: *Proceedings of the 26th Symposium on Operating Systems Principles*. SOSP '17. 2017.

[35] O. Goldreich, S. Micali, and A. Wigderson. "How to Play ANY Mental Game". In: *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*. STOC '87. Association for Computing Machinery, 1987, pp. 218–229.

[36] J. Groth. "Rerandomizable and replayable adaptive chosen ciphertext attack secure cryptosystems". In: *TCC 2004*. Vol. 2951. LNCS. 2004, pp. 152–170.

[37] J. Groth, M. Kohlweiss, M. Maller, S. Meiklejohn, and I. Miers. "Updatable and Universal Common Reference Strings with Applications to zk-SNARKs". In: *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*. 2018, pp. 698–728.

[38] J. Groth and A. Sahai. "Efficient Noninteractive Proof Systems for Bilinear Groups". In: *SIAM J. Comput.* 41.5 (2012), pp. 1193–1232.

[39] Heiko Stamer. *Distributed Privacy Guard*.

[40] GNOSIS. *Distributed Key Generation*.

[41] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. "Randomized rumor spreading". In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. 2000, pp. 565–574.

[42] A. Kate and I. Goldberg. "Distributed Key Generation for the Internet". In: *29th IEEE International Conference on Distributed Computing Systems*. 2009, pp. 119–128.

[43] A. Kate. "Distributed Key Generation and Its Applications". PhD thesis. Waterloo, Ontario, Canada, 2010.

[44] A. Kate, G. M. Zaverucha, and I. Goldberg. "Constant-Size Commitments to Polynomials and Their Applications". In: *ASIACRYPT'10*. 2010. ISBN: 978-3-642-17373-8.

[45] A. Kiayias, A. Russell, B. David, and R. Oliynykov. "Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol". In: *Advances in Cryptology – CRYPTO 2017*. 2017.

[46] E. Kokoris-Kogias, E. C. Alp, L. Gasser, P. Jovanovic, E. Syta, and B. Ford. *Verifiable Management of Private Data under Byzantine Failures*. Cryptology ePrint Archive, Report 2018/209. 2018.

[47] E. Kokoris-Kogias, D. Malkhi, and A. Spiegelman. *Asynchronous Distributed Key Generation for Computationally-Secure Randomness, Consensus, and Threshold Signatures.* Cryptology ePrint Archive, Report 2019/1015. 2019.

[48] C. Komlo and I. Goldberg. "FROST: Flexible Round-Optimized Schnorr Threshold Signatures". In: *IACR Cryptol. ePrint Arch.* 2020 (2020), p. 852.

[49] S. Micali, M. Rabin, and S. Vadhan. "Verifiable Random Functions". In: *40th Annual Symposium on Foundations of Computer Science*. 1999, pp. 120–130.

[50] S. Micali, K. Ohta, and L. Reyzin. "Accountable-Subgroup Multisignatures: Extended Abstract". In: *Proceedings of the 8th ACM Conference on Computer and Communications Security*. CCS '01. Philadelphia, PA, USA: Association for Computing Machinery, 2001, 245–254.

[51] W. Neji, K. Blibech, and N. Ben Rajeb. "Distributed key generation protocol with a new complaint management strategy". In: *Security and Communication Networks* 9.17 (2016), pp. 4585–4595.

[52] Orbs Network. *Orbs Network: DKG for BLS threshold signature scheme on the EVM using solidity*. 2018.

[53] P. Paillier. "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes". In: *Advances in Cryptology — EUROCRYPT '99*. Ed. by J. Stern. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 223–238.

[54] T. P. Pedersen. "A Threshold Cryptosystem without a Trusted Party". In: *Advances in Cryptology – EUROCRYPT '91*. Ed. by D. W. Davies. Springer Berlin Heidelberg, 1991, pp. 522–526.

[55] T. P. Pedersen. "Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing". In: *Proceedings on Advances in Cryptology*. Ed. by J. Feigenbaum. CRYPTO '91. Berlin, Heidelberg: Springer Berlin Heidelberg, 1992, pp. 129–140.

[56] B. Pittel. "On Spreading a Rumor". In: *SIAM Journal on Applied Mathematics* 47.1 (1987), pp. 213–223. ISSN: 00361399.

[57] M. Prabhakaran and M. Rosulek. "Rerandomizable RCCA encryption". In: *Advances in Cryptology - CRYPTO 2007*. Vol. 4622. LNCS. 2007, pp. 517–534.

[58] P. Schindler, A. Judmayer, N. Stifter, and E. Weippl. *ETHDKG: Distributed Key Generation with Ethereum Smart Contracts*. Cryptology ePrint Archive, Report 2019/985. 2019.

[59] C. P. Schnorr. "Efficient Identification and Signatures for Smart Cards". In: *Advances in Cryptology — CRYPTO' 89 Proceedings*. Ed. by G. Brassard. New York, NY: Springer New York, 1990, pp. 239–252. ISBN: 978-0-387-34805-6.

[60] A. Shamir. "How to Share a Secret". In: *Commun. ACM* 22.11 (1979).

[61] E. Syta et al. "Scalable Bias-Resistant Distributed Randomness". In: *38th IEEE Symposium on Security and Privacy*. San Jose, CA, May 2017.

[62] A. Tomescu et al. "Towards Scalable Threshold Cryptosystems". In: *IEEE S&P'20*. 2020.

[63] D. Tulone. "A Scalable and Intrusion-tolerant Digital Time-stamping System". In: *2006 IEEE International Conference on Communications*. Vol. 5. 2006, pp. 2357–2363.

[64] Y. Wang, Z. Zhang, T. Matsuda, G. Hanaoka, and K. Tanaka. "How to Obtain Fully Structure-Preserving (Automorphic) Signatures from Structure-Preserving Ones". In:

[65] T. M. Wong, C. Wang, and J. M. Wing. "Verifiable Secret Redistribution for Archive Systems". In: *First International IEEE Security in Storage Workshop*. 2002, pp. 94–105.

[66] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham. "HotStuff: BFT Consensus with Linearity and Responsiveness". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. PODC '19. 2019.

A The Scrape PVSS

We describe in more detail the operation of Scrape.Deal and Scrape.Verify. For the reconstruction algorithm, which is unused in this work, we refer the reader back to the original paper [18].

Dealer

Commit to the coefficients: The dealer takes as input $(\hat{F}_0, a_0) \in \mathbb{G} \times \mathbb{F}$ where $\hat{F}_0 = \hat{h}_1^{a_0}$ is the secret being shared. They first samples a random polynomial

$$f(X) = a_0 + a_1 X + \cdots + a_t X^t$$

of degree at most t with a_0 as the constant coefficient. They then commit to the coefficients by computing $(F_0, \ldots, F_t) \in \mathbb{G}_1^{t+1}$ as

$$(F_0, \ldots, F_t) = (g_1^{a_0}, \ldots, g_1^{a_t})$$

Compute additional element: The dealer sets $\hat{u}_2 = \hat{u}_1^{a_0}$.

Commit to the evaluation points: Suppose that $\omega_1, \ldots, \omega_n$ form an order n multiplicative subgroup in $\mathbb{F}$ (so that FFTs are efficient to compute). The dealer computes

$$(A_0, \ldots, A_{n-1}) = (g_1^{f(\omega_1)}, \ldots, g_1^{f(\omega_n)}).$$

Note that the verifier could compute these points for itself but by having the prover send these values we can save on verification costs.

Encrypt the evaluation points: To encrypt the evaluation points, the dealer takes as input the encryption keys $(\mathsf{ek}_1, \ldots, \mathsf{ek}_n) \in \mathbb{G}_2^n$ as input and returns the encryptions $(\hat{Y}_1, \ldots, \hat{Y}_n) \in \mathbb{G}_2^n$ such that

$$(\hat{Y}_1, \ldots, \hat{Y}_n) = (\mathsf{ek}_1^{f(\omega_1)}, \ldots, \mathsf{ek}_n^{f(\omega_n)}).$$

Decrypting these values will yield group elements as opposed to field elements, hence our final secret is a group element in $\mathbb{G}_2$.

Return the shares: The dealer returns a commitment to coefficients $\boldsymbol{F} \in \mathbb{G}_1^{t+1}$, a commitment to the evaluation points $\boldsymbol{A} \in \mathbb{G}_1^n$, and the encryptions of evaluation points $\hat{\boldsymbol{Y}} \in \mathbb{G}_2^n$.

Verifier

Compute the committed evaluations: The verifier ensures that the committed evaluations $\boldsymbol{A} \in \mathbb{G}_1^n$ are consistent with the committed polynomial in $\boldsymbol{F}$

with a randomised check. They consider the Lagrange polynomials with respect to the set of distinct points $\omega_1, \ldots, \omega_n \in \mathbb{F}^n$ given by

$$\ell_j(X) = \prod_{i=1, i\neq j}^{n} \frac{X - \omega_i}{\omega_j - \omega_i}$$

such that $\ell_j(X)$ is equal to 0 on $\omega_i \neq \omega_j$ and $\ell_j(\omega_j) = 1$. The verifier samples a random point $\alpha \in \mathbb{F}$. They check that

$$A_1^{\ell_1(\alpha)} \cdots A_n^{\ell_n(\alpha)} = F_0 F_1^{\alpha} \cdots F_t^{\alpha^t}.$$

If this equation verifies then with overwhelming probability the evaluations in $A_1, \ldots, A_n$ are correct.

Check the additional value: The verifier checks that the additional value $\hat{u}_2 \in \mathbb{G}_2$ has been computed correctly i.e. that

$$e(F_0, \hat{u}_1) = e(g_1, \hat{u}_2)$$

Check the encrypted evaluations are correct: The verifier checks that the encryptions $\hat{\boldsymbol{Y}} \in \mathbb{G}_2^n$ contain the same evaluations as $\boldsymbol{A}$ using n pairing equations:

$$e(g_1, \hat{Y}_1) = e(A_1, \mathsf{ek}_1), \ldots, e(g_1, \hat{Y}_n) = e(A_n, \mathsf{ek}_n)$$

If all checks pass the verifier returns 1 to indicate acceptance. Else it returns 0 to indicate rejection.

Observe that if this check passes then the secret share $\hat{h}_1^{f(\omega_i)}$ can be decrypted as $\hat{Y}_i^{\frac{1}{\mathsf{dk}}}$.

B Gossip in more detail

Here we describe our gossip protocol in more detail. We assume that up to $\frac{n}{2} - \log n$ parties may be crashed, and $\log n$ parties may be Byzantine. Denote $t_r < \frac{n}{2}$ to be the total number of faulty parties, crashed and Byzantine alike. In practice, we might set the threshold of faulty parties to be lower to allow for a higher threshold in the security preservation. We describe our protocol with respect to some success parameter $c \geq 4$ which will usually be chosen to be very small, and never greater than $\frac{n}{\ln n}$. A higher success parameter would result in larger round and communication complexity, but also result in a higher probability of the protocol succeeding in the given number of rounds. More precisely, the probability of succeeding in the given number of rounds is in the order of $n^{-c+O(1)}$.

Party P_i takes as input the public key infrastructure, the round number $r \in \mathbb{N}$, their secret key sk_i and their DKG transcript $\mathsf{transcript}_i$. They behave as follows.

for $1 \leq j \leq n$:
  send $(y_i, \pi_i, \mathsf{transcript}_i)$ to P_j with probability $\frac{\lceil 2c \ln(n) \rceil}{n}$

for $1 \leq j \leq n$ such that a send request is received from P_j:
  check $\mathsf{DKG.Verify}(\mathsf{bp}, (\mathsf{ek}_i, \mathsf{vk}_i)_{i \in [n]}, \hat{u}_1, \mathsf{transcript}_j) = 1$
  if yes set $\mathsf{transcript}_i = \mathsf{DKG.Aggregate}(\mathsf{bp}, \mathsf{transcript}_i, \mathsf{transcript}_j)$

if $r \geq \ln(n)$:
  if $\mathsf{transcript}_i$ contains $t + 1$ contributions:
    broadcast $\mathsf{transcript}_i$ with probability $\frac{2}{n}$

Terminate if broadcast completes with verifying transcript of $t + 1$ contributions.

Fig. 5. Gossip protocol for Party P_i

B.1 Construction

In Figure B.1 we explain the actions that each party takes in the gossip protocol. Intuitively, in each round, every party uniformly picks an expected $2c \ln n$ parties to send its current aggregated DKG transcript. Then, after $c \ln n$ rounds, every party that has aggregated a DKG transcript with enough contributions, broadcasts it with probability $\frac{2}{n}$, resulting in $O(1)$ expected broadcasts in each such round.

B.2 Complexity

The protocol requires $\mathcal{O}(c \log n)$ rounds with all but a polynomially small probability (see Lemma 2). In each round, every nonfaulty party sends an expected $\mathcal{O}(c \log n)$ messages of size $\mathcal{O}(n)$ words (see Lemma 1). In addition, each party reads an expected $\mathcal{O}(c \log n)$ messages containing full DKG transcripts of size $\mathcal{O}(n)$ words each. This includes messages from all Byzantine parties as well. Due to standard measure-concentration arguments, this asymptotic complexity holds w.h.p. as well. Summing up all of the messages, we find that each party reads and sends $\mathcal{O}(c^2 n \log^2 n)$ words in private channels throughout the protocol w.h.p.

In addition, in every round an expected $\mathcal{O}(1)$ broadcasts are sent by nonfaulty parties, totalling in $\mathcal{O}(c \log n)$ broadcasts. Due to standard measure-concentration arguments, the number of broadcasts sent per round is $\mathcal{O}(\log n)$ w.h.p., including possible broadcasts from faulty parties. This means that the total number of broadcasts sent throughout the protocol is $\mathcal{O}(c \log^2(n))$ w.h.p.

B.3 Expected number of rounds before parties have enough contributions

Lemma 1. *In $\mathcal{O}(c \log n)$ rounds, all nonfaulty parties have a DKG transcript with $t+1$ contributions with probability $1 - \frac{1}{p(n)}$ or greater for some polynomial p.*

Proof. First, note that at the time the adversary chooses which parties to corrupt, it does not know which parties will communicate with each other throughout the protocol. This means that the adversary's choice of parties to corrupt is entirely independent of the parties to whom the nonfaulty parties send messages.

Once a nonfaulty party j receives a message with a contribution that verifies from party i, every message it sends will contain a contribution from party i, and its transcript will verify as well. This means that we can essentially see the protocol as n simultaneous runs of a gossip protocol with different sources for each rumor. Observe some party i. In each round r, the probability that i does not send a message to any nonfaulty party is:

$$\left(1 - \frac{\lceil 2c \ln n \rceil}{n}\right)^{n-t_r} \leq \left(1 - \frac{2c \ln n}{n}\right)^{\frac{n}{2}} \leq e^{-c \ln n} = n^{-c}$$

Using the union bound, the probability that there exists some nonfaulty party that doesn't send a message to some nonfaulty party in any given round is no greater than $n \cdot n^{-c} = n^{-c+1}$. Note that clearly the probability that any specific nonfaulty party is chosen is entirely symmetric to the probability that any other party is chosen, so at least one nonfaulty party is uniformly chosen by each other nonfaulty party with probability n^{-c+1} or greater. Observe only the communication between nonfaulty parties. Using well-known results [21, 41, 56], we know that if in each round all nonfaulty parties communicate with at least one nonfaulty party, then with all but a polynomially small probability of failure $\frac{1}{q(n)}$, all parties receive a contribution from party i in $\mathcal{O}(\log n)$ rounds. We can also make sure that $\frac{1}{q(n)} = \mathcal{O}(n^{-c+4})$ by running the gossip protocol for $\mathcal{O}(c \log n)$ rounds. Let d be the constant in the $\mathcal{O}(c \log n)$ number of rounds required. Now define a failure event in which either there exists some nonfaulty party that doesn't send a message to any nonfaulty party in the first $d \cdot c \log n$ rounds, or that the gossip initiated by some nonfaulty party requires more than $d \cdot c \log n$ rounds. Again, using the union bound and the fact that $\frac{n}{\ln n} \geq c$, the probability that this event takes place is no greater than:

$$d \cdot c \log n \cdot \frac{n}{2} \cdot n^{-c+1} + \frac{n}{2} \cdot \frac{1}{q(n)} = \mathcal{O}\left(n^{-c+3}\right)$$

which is polynomially small.

B.4 Total expected number of rounds

Lemma 2. *Some nonfaulty party sends a broadcast of a verifying DKG transcript in $\mathcal{O}(c \log n)$ rounds with all but a polynomially small probability, if no such broadcast has previously been sent.*

Proof. As shown in Lemma 1, after $\mathcal{O}(c \log n)$ rounds all nonfaulty parties have aggregated a DKG transcript with contributions from all nonfaulty parties with all but a polynomially small probability, $\frac{1}{p(n)}$. Since $n > 2t$, this means that at that time they all have at least $t + 1$ contributions in their aggregated DKG transcripts. From that point on, every party has a $\frac{2}{n}$ probability of broadcasting its aggregated DKG transcript in each round. Now observe the next $\lceil c \ln n \rceil$ rounds. The probability that none of the nonfaulty parties send a broadcast in a single one of those rounds is:

$$\left(1 - \frac{2}{n}\right)^{n-t_r} \leq \left(1 - \frac{2}{n}\right)^{\frac{n}{2}} \leq e^{-1}$$

Therefore, the probability that none of the nonfaulty parties broadcast some message in any of those $\lceil c \ln n \rceil$ rounds is no greater than:

$$\left(e^{-1}\right)^{\lceil c \ln n \rceil} \leq \left(e^{-1}\right)^{c \ln n} = n^{-c}$$

Now define a failure event in which either there exists a nonfaulty party whose aggregated DKG transcript after $\mathcal{O}(c \log n)$ rounds doesn't consist of at least $t+1$ contributions, or no nonfaulty party broadcasts a message in the $\lceil c \ln n \rceil$ rounds after all of them have $t + 1$ contributions in their aggregated DKG transcript. Using the union bound, this failure event occurs with probability no greater than $\frac{1}{p(n)} + n^{-c}$. Note that if neither of those events occur, some nonfaulty party broadcasts a message in $\mathcal{O}(c \log n)$ rounds, completing the proof.

C Rekeyability Implies Security Preservation

In this section, we prove that using a key-expressable DKG for several rekeyable primitives implies that the DKG preserves security as well. In particular, we show that provided that the signing/encryption shares are identical to signatures/encryptions, IND-CPA is preserved for rekeyable encryption schemes, EUF-CMA is preserved for rekeyable signature schemes, and both uniqueness and unpredictability are preserved for rekeyable VUFs. We then show in Appendix D that common constructions such as El-Gamal encryption and BLS signatures are rekeyable, as is our VUF from Section 7.

C.1 Encryption

We show that any IND-CPA secure and rekeyable encryption scheme (KeyGen, $\mathsf{Encrypt}, \mathsf{Decrypt}$) such that $\mathsf{Encrypt}_{\mathsf{pk}}$ and $\mathsf{EncryptShare}_{\mathsf{pk}}$ are identical is IND-CPA secure under a key-expressable DKG. As a reminder, the IND-CPA security game is defined as follows:

$$\begin{array}{ll}
\underline{\textsc{main}\ \mathsf{Game}_{\mathcal{A}}(1^\lambda)} & \underline{\mathcal{O}^{\mathsf{Encrypt}}_{\mathsf{pk}}(m_0, m_1)} \\
(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\lambda) & \text{return } \mathsf{Encrypt}(\mathsf{pk}, m_b) \\
b \xleftarrow{\$} \{0,1\} & \\
b' \xleftarrow{\$} \mathcal{A}^{\mathsf{Encrypt}}(\mathsf{pk}) & \\
\text{return } b = b' &
\end{array}$$

Lemma 3. *If* $\mathsf{Encrypt}$ *is rekeyable with respect to the public key, and if* $\mathsf{Encrypt}_{\mathsf{pk}} = \mathsf{EncryptShare}_{\mathsf{pk}}$*, then any key-expressable DKG is also security-preserving for IND-CPA security.*

Proof. Let $\mathcal{A}$ be an adversary playing the IND-CPA security game with a key-expressable DKG. We design $\mathcal{B}$ such that

$$\mathsf{Adv}^{\text{IND-CPA},\mathsf{DKG\text{-}Game}}_{\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\text{IND-CPA}}_{\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda).$$

To start, $\mathcal{B}$ takes as input a public key pk_1. It runs $(\mathsf{transcript}, \mathsf{pk}, \alpha, \mathsf{pk}_2, \mathsf{sk}_2) \xleftarrow{\$} \mathtt{SimDKG}(\mathsf{Sim}, I, N)$, acting as Sim to interact with $\mathcal{A}$; recall that by the definition of key expressability we now have that $\mathsf{pk} = f(\alpha, \mathsf{pk}_1, \mathsf{pk}_2)$. When $\mathcal{A}$ queries $\mathsf{Encrypt}$ on (m_0, m_1), $\mathcal{B}$ queries its own oracle $\mathsf{Encrypt}$ on (m_0, m_1) to get a ciphertext c. It then returns $\mathsf{rekey}(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, c)$. When $\mathcal{A}$ returns b' then $\mathcal{B}$ also returns b'.

By the rekeyability of $\mathsf{Encrypt}$,

$$\mathsf{rekey}(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, \mathsf{Encrypt}(\mathsf{pk}_1, m_b; r)) = \mathsf{Encrypt}(f_{\mathsf{pk}}(\alpha, \mathsf{pk}_1, \mathsf{pk}_2), m_b; r).$$

$\mathcal{B}$ thus perfectly simulates the $\mathsf{Encrypt}$ oracle that $\mathcal{A}$ expects, and key expressability implies that it also perfectly simulates the DKG. It thus wins whenever $\mathcal{A}$ does.

C.2 Signatures

We show that any existentially unforgeable and rekeyable signature scheme ($\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify}$) with $\mathsf{Sign}_{\mathsf{sk}} = \mathsf{SignShare}_{\mathsf{sk}}$ is existentially unforgeable under a key-expressable DKG. As a reminder, the EUF-CMA security game is defined as follows:

$$\begin{array}{ll}
\underline{\textsc{main}\ \mathsf{Game}_{\mathcal{A}}(1^\lambda)} & \underline{\mathcal{O}^{\mathsf{Sign}}_{\mathsf{sk}}(m)} \\
(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\lambda) & \sigma \xleftarrow{\$} \mathsf{Sign}(\mathsf{sk}, m) \\
Q \leftarrow \emptyset & Q \leftarrow Q \cup \{m\} \\
(m, \sigma) \xleftarrow{\$} \mathcal{A}^{\mathsf{Sign}}(\mathsf{pk}) & \text{return } \sigma \\
\text{return } m \notin Q \wedge \mathsf{Verify}(\mathsf{pk}, m, \sigma) &
\end{array}$$

Lemma 4. *If* $(\mathsf{Sign}, \mathsf{Verify})$ *is rekeyable with respect to the secret key and if* $\mathsf{Sign}_{\mathsf{sk}} = \mathsf{SignShare}_{\mathsf{sk}}$*, then any key-expressable DKG is also security-preserving for EUF-CMA security.*

Proof. Let $\mathcal{A}$ be an adversary playing the EUF-CMA security game with a key-expressable DKG. We design $\mathcal{B}$ such that

$$\mathsf{Adv}^{\text{EUF-CMA},\mathsf{DKGGame}}_{\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\text{EUF-CMA}}_{\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda).$$

To start, $\mathcal{B}$ takes as input a public key pk_1. It then runs $(\mathsf{transcript}, \mathsf{pk}, \alpha, \mathsf{pk}_2, \mathsf{sk}_2) \xleftarrow{\$} \mathtt{SimDKG}(\mathsf{Sim}, I, N)$, acting as Sim to interact with $\mathcal{A}$; by the definition of key expressability we now have that $\mathsf{pk} = f_{\mathsf{pk}}(\alpha, \mathsf{pk}_1, \mathsf{pk}_2)$. When $\mathcal{A}$ queries Sign on m, $\mathcal{B}$ queries its own oracle Sign on m to get a signature σ. It then returns $\mathsf{rekey}(\alpha, \mathsf{pk}_1, -\mathsf{sk}_2, m, \sigma)$. When $\mathcal{A}$ returns (m, σ) then $\mathcal{B}$ returns $y \leftarrow \mathsf{rekey}(\frac{1}{\alpha}, \mathsf{pk}, \frac{-\mathsf{sk}_2}{\alpha}, m, \sigma)$.

By the rekeyability of Sign,

$$\mathsf{rekey}(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, m, \mathsf{Sign}(\mathsf{sk}_1, m; r)) = \mathsf{Sign}(f_{\mathsf{sk}}(\alpha, \mathsf{sk}_1, \mathsf{sk}_2), m; r).$$

$\mathcal{B}$ thus perfectly simulates the Sign oracle that $\mathcal{A}$ expects, and key expressability implies that it also perfectly simulates the DKG. If $\mathcal{A}$ has not queried on m then neither has $\mathcal{B}$, and by the rekeyability of $(\mathsf{Sign}, \mathsf{Verify})$ we have that

$$\begin{aligned}
\mathsf{Verify}(\mathsf{pk}, m, \sigma) &= \mathsf{Verify}\left(f_{\mathsf{pk}}(\frac{1}{\alpha}, \mathsf{pk}, -\frac{\mathsf{pk}_2}{\alpha}), m, \mathsf{rekey}(\frac{1}{\alpha}, \mathsf{pk}, \frac{-\mathsf{sk}_2}{\alpha}, m, \sigma)\right) \\
&= \mathsf{Verify}(\frac{1}{\alpha}(\alpha\mathsf{pk}_1 \oplus \mathsf{pk}_2) \oplus \frac{-\mathsf{pk}_2}{\alpha}, m, y) \\
&= \mathsf{Verify}(\mathsf{pk}_1 \oplus \frac{\mathsf{pk}_2}{\alpha} \oplus \frac{-\mathsf{pk}_2}{\alpha}, m, y) \\
&= \mathsf{Verify}(\mathsf{pk}_1, m, y).
\end{aligned}$$

The output y of $\mathcal{B}$ thus verifies under pk_1 whenever $\mathcal{A}$'s output (m, σ) verifies under pk, so $\mathcal{B}$ wins whenever $\mathcal{A}$ wins.

C.3 VUFs

We show that any unique, unpredictable, and rekeyable VUF scheme with $\mathsf{VUF.Sign}_{\mathsf{sk}} = \mathsf{VUF.SignShare}_{\mathsf{sk}}$ is unique and unpredictable under a key-expressable DKG.

Lemma 5. *If* $(\mathsf{VUF.Sign}, \mathsf{VUF.Ver})$ *is rekeyable with respect to the secret key,* $\mathsf{VUF.Sign}_{\mathsf{sk}} = \mathsf{VUF.SignShare}_{\mathsf{sk}}$*, and* $\mathsf{VUF.Eval}$ *is rekeyable with respect to the secret key, then any key-expressable DKG is also security-preserving for the uniqueness and unpredictability games for a VUF.*

Proof. To prove that uniqueness is preserved, we observe that the adversary picks the public key pk in the uniqueness game. This renders the DKG setup irrelevant, meaning uniqueness holds regardless.

To prove that unpredictability is preserved, let $\mathcal{A}$ be an adversary playing the VUF unpredictability game with a key-expressable DKG. We design $\mathcal{B}$ such that

$$\mathsf{Adv}^{\mathsf{predict},\mathsf{DKGGame}}_{\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\mathsf{predict}}_{\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda).$$

To start, $\mathcal{B}$ takes as input a public key pk_1. It then runs $(\mathsf{transcript}, \mathsf{pk}, \alpha, \mathsf{pk}_2, \mathsf{sk}_2) \xleftarrow{\$} \mathtt{SimDKG}(\mathsf{Sim}, I, N)$, acting as Sim to interact with $\mathcal{A}$; by the definition of key expressability we now have that $\mathsf{pk} = f_{\mathsf{pk}}(\alpha, \mathsf{pk}_1, \mathsf{pk}_2)$. When $\mathcal{A}$ queries $\mathsf{VUF.Sign}$ on m, $\mathcal{B}$ queries its own oracle $\mathsf{VUF.Sign}$ on m to get vuf. It then returns $\mathsf{rekey}_{\mathsf{VUF.Sign}}(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, m, \mathsf{vuf})$. When $\mathcal{A}$ returns (m, y) then $\mathcal{B}$ returns $z \leftarrow \mathsf{rekey}_{\mathsf{VUF.Eval}}(\frac{1}{\alpha}, \mathsf{pk}, \frac{-\mathsf{sk}_2}{\alpha}, (m, y))$.

By the rekeyability of $\mathsf{VUF.Sign}$, we have that

$$\begin{aligned}&\mathsf{rekey}(\alpha, \mathsf{pk}_1, \mathsf{sk}_2, m, \mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}_1, m; r))\\ &\qquad = \mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, f_{\mathsf{sk}}(\alpha, \mathsf{sk}_1, \mathsf{sk}_2), m; r),\end{aligned}$$

so $\mathcal{B}$ perfectly simulates $\mathcal{A}$'s queries. If $\mathcal{A}$ has not queried on m then neither has $\mathcal{B}$. If we define the other winning condition of the unpredictability game as an auxiliary function $0/1 \leftarrow \mathsf{VUF.Ver}'(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m, y)$ that outputs 1 if $\mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m) = y$ and 0 otherwise, then by the rekeyability of $(\mathsf{VUF.Eval}, \mathsf{VUF.Ver}')$ we have that

$$\begin{aligned}&\mathsf{VUF.Ver}'(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m, y)\\ &= \mathsf{VUF.Ver}'\left(\mathsf{crs}_{\mathsf{vuf}}, f_{\mathsf{sk}}(\tfrac{1}{\alpha}, \mathsf{sk}, -\tfrac{\mathsf{sk}_2}{\alpha}), \mathsf{rekey}_{\mathsf{VUF.Eval}}(\tfrac{1}{\alpha}, \mathsf{pk}, \tfrac{-\mathsf{sk}_2}{\alpha}, m, y)\right)\\ &= \mathsf{VUF.Ver}'(\mathsf{crs}_{\mathsf{vuf}}, \tfrac{1}{\alpha}(\alpha\mathsf{sk}_1 \oplus \mathsf{sk}_2) \oplus \tfrac{-\mathsf{sk}_2}{\alpha}, z)\\ &= \mathsf{VUF.Ver}'(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}_1 \oplus \tfrac{\mathsf{sk}_2}{\alpha} \oplus \tfrac{-\mathsf{sk}_2}{\alpha}, z)\\ &= \mathsf{VUF.Ver}'(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}_1, z).\end{aligned}$$

Thus $\mathcal{B}$ wins whenever $\mathcal{A}$ wins.

D Rekeyable Cryptographic Primitives

In this section, we show that several cryptographic constructions are rekeyable, and thus can be securely instantiated using a key-expressable DKG, as we showed in Appendix C. Specifically we cover the El-Gamal encryption scheme, the BLS signature scheme, and our VUF scheme. The first two of these have field elements as secret keys so can use the Pedersen DKG or Fouque-Stern DKG, while the latter has group elements as secret keys so can use our DKG. In general we can cover schemes that have some degree of malleability but not those that do not. For example, we can cover the BLS signature scheme that is susceptible to rogue key attacks, but we do not know how to extend our methods to Schnorr signatures. This we consider to be somewhat reassuring because there is a known attack against applying the Pedersen DKG to Schnorr signatures assuming a concurrent adversary [10].

D.1 El-Gamal Encryption

El-Gamal encryption consists of the following three algorithms:

- $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(\mathbb{G}_1)$: Choose $x \xleftarrow{\$} \mathbb{F}$. Return $\mathsf{pk} = g_1^x$ and $\mathsf{sk} = x$.
- $(c_1, c_2) \xleftarrow{\$} \mathsf{Encrypt}(\mathsf{pk}, m)$: Choose $r \xleftarrow{\$} \mathbb{F}$. Return $c_1 = g_1^r$ and $c_2 = m\mathsf{pk}^r$.
- $m \leftarrow \mathsf{Decrypt}(\mathsf{sk}, (c_1, c_2))$: Return $m = c_2 c_1^{-\mathsf{sk}}$.

We now prove Corollary 1, which states that El-Gamal preserves IND-CPA security when instantiated with the Pedersen or Fouque-Stern DKG.

Proof. This follows directly from Lemma 3, Theorems 3 and 4, and the fact that the El-Gamal encryption scheme is rekeyable, using $\alpha = 1$ (as this is its value for both DKGs), $f_{\mathsf{pk}}(\mathsf{pk}_1, \mathsf{pk}_2) = \mathsf{pk}_1 \cdot \mathsf{pk}_2$, and $f_{\mathsf{sk}}(\mathsf{sk}_1, \mathsf{sk}_2) = \mathsf{sk}_1 + \mathsf{sk}_2$.

To see, this we define rekey for $\mathsf{Encrypt}$ as follows:

$$\mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, (c_1, c_2)) = (c_1, c_2 c_1^{\mathsf{sk}_2})$$

Observe that

$$\mathsf{Decrypt}(\mathsf{sk}_1, (c_1, c_2)) = c_2 c_1^{-\mathsf{sk}_1}$$
$$\mathsf{Decrypt}(\mathsf{sk}_1 + \mathsf{sk}_2, \mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, (c_1, c_2))) = \mathsf{Decrypt}(\mathsf{sk}_1 + \mathsf{sk}_2, (c_1, c_2 c_1^{\mathsf{sk}_2})) = c_2 c_1^{-\mathsf{sk}_1},$$

so $(\mathsf{Encrypt}, \mathsf{Decrypt})$ is rekeyable with respect to the public key. Observe also that

$$\mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, \mathsf{Encrypt}(\mathsf{pk}_1, m; r)) = \mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, (g_1^r, m \cdot \mathsf{pk}_1^r)) = (g_1^r, m \cdot (\mathsf{pk}_1^r \mathsf{pk}_2^r))$$

and is thus equal to $\mathsf{Encrypt}(\mathsf{pk}_1 \cdot \mathsf{pk}_2, m; r)$.

D.2 BLS Signatures

A BLS signature consists of the following three algorithms:

- $(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(\mathsf{bp})$: Choose $x \xleftarrow{\$} \mathbb{F}$. Set $\mathsf{pk} = \hat{h}_1^x$ and $\mathsf{sk} = x$. Return $(\mathsf{pk}, \mathsf{sk})$.
- $\sigma \xleftarrow{\$} \mathsf{Sign}(m, \mathsf{sk}:)$ Compute $Z = \mathsf{Hash}_{\mathbb{G}_1}(m)$. Return Z^x.
- $0/1 \leftarrow \mathsf{Verify}(\mathsf{pk}, m, \sigma)$: Compute $Z = \mathsf{Hash}_{\mathbb{G}_1}(m)$. Check $e(Z, \mathsf{pk}) = e(\sigma, \hat{h}_1)$. Return 1 if check passes, else return 0.

We now prove Corollary 2, which states that BLS preserves EUF-CMA security when instantiated with the Pedersen or Fouque-Stern DKG.

Proof. This follows directly from Lemma 4, Theorems 3 and 4, and the fact that the BLS signature scheme is rekeyable, using $\alpha = 1$ (as this is its value for both DKGs), $f_{\mathsf{pk}}(\mathsf{pk}_1, \mathsf{pk}_2) = \mathsf{pk}_1 \cdot \mathsf{pk}_2$, and $f_{\mathsf{sk}}(\mathsf{sk}_1, \mathsf{sk}_2) = \mathsf{sk}_1 + \mathsf{sk}_2$.

To see this, we define rekey for Sign as follows:

$$\mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, m, \sigma) = \sigma \cdot \mathsf{Hash}(m)^{-\mathsf{sk}_2}$$

Observe that

$$\begin{array}{ll} \mathsf{Verify}(\mathsf{pk}_1 \cdot \mathsf{pk}_2, m, \sigma) = 1 & \Leftrightarrow e(\mathsf{pk}_1 \cdot \mathsf{pk}_2, \mathsf{Hash}(m)) = e(g_1, \sigma) \\ \mathsf{Verify}(\mathsf{pk}_1, m, \mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, m, \sigma)) = 1 & \Leftrightarrow e(\mathsf{pk}_1, \mathsf{Hash}(m)) = e(g_1, \sigma \mathsf{Hash}(m)^{-\mathsf{sk}_2}) \\ & \Leftrightarrow e(\mathsf{pk}_1 \cdot \mathsf{pk}_2, \mathsf{Hash}(m)) = e(g_1, \sigma), \end{array}$$

so $(\mathsf{Sign}, \mathsf{Verify})$ is rekeyable with respect to the secret key. Observe also that

$$\mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, \mathsf{Sign}(\mathsf{sk}_1 + \mathsf{sk}_2, m)) = \mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, \mathsf{Hash}(m)^{\mathsf{sk}_1 + \mathsf{sk}_2}) = \mathsf{Hash}(m)^{\mathsf{sk}_1}$$

and is thus equal to $\mathsf{Sign}(\mathsf{sk}_1, m)$.

D.3 Our VUF

We now prove Corollary 3, which says that our VUF from Section 7 preserves its security when using our DKG.

Proof. This follows directly from Lemma 5, Theorem 2, and the fact that the VUF scheme is rekeyable.

To see this, we define $\mathsf{rekey}_{\mathsf{VUF.Sign}}$ as follows:

$$\begin{aligned} &\mathsf{rekey}_{\mathsf{VUF.Sign}}(\mathsf{pk}_1, \mathsf{sk}_2, m, (\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)) \\ &\quad = (\pi_1 g_1^{\alpha}, \pi_2 \mathsf{Hash}(m)^{\alpha}, \pi_3 g_1^{\beta}, \pi_4 \mathsf{Hash}(m)^{\beta}, \hat{\pi}_1 \hat{h}_1^{-\alpha} \hat{h}_2^{-\beta}, \mathsf{sk}_2^{-1} \hat{\pi}_2 \hat{h}_3^{-\alpha} \hat{h}_4^{-\beta}) \end{aligned}$$

Observe for $Z = \mathsf{Hash}(m)$ that

$$\begin{array}{l} e(g_1, \hat{\pi}_1 \hat{h}_1^{-\alpha} \hat{h}_2^{-\beta}) e(\pi_1 g_1^{\alpha}, \hat{h}_1) e(\pi_3 g_1^{\beta}, \hat{h}_2) = e(g_1, \hat{\pi}_1) e(\pi_1, \hat{h}_1) e(\pi_3, \hat{h}_2) \\ e(Z, \hat{\pi}_1 \hat{h}_1^{-\alpha} \hat{h}_2^{-\beta}) e(\pi_2 Z^{\alpha}, \hat{h}_1) e(\pi_4 Z^{\beta}, \hat{h}_2) = e(Z, \hat{\pi}_1) e(\pi_2, \hat{h}_1) e(\pi_4, \hat{h}_2) \\ e(g_1, \hat{\pi}_2 \mathsf{sk}_2^{-1} \hat{h}_3^{-\alpha} \hat{h}_4^{-\beta}) e(\pi_1 g_1^{\alpha}, \hat{h}_3) e(\pi_3 g_1^{\beta}, \hat{h}_2) = e(\mathsf{pk}_2^{-1}, \hat{h}_1) e(g_1, \hat{\pi}_2) e(\pi_1, \hat{h}_3) e(\pi_3, \hat{h}_2) \end{array}$$

Thus m and $(\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)$ verify under $\mathsf{pk}_1 + \mathsf{pk}_2$ if and only if

$$m, \mathsf{rekey}_{\mathsf{VUF.Sign}}(\mathsf{pk}_1, \mathsf{sk}_2, m, (\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2))$$

verifies under pk_1. Furthermore, $\mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, \mathsf{VUF.Sign}(\mathsf{bp}, \mathsf{sk}_1 + \mathsf{sk}_2, m))$ is identical to $\mathsf{VUF.Sign}(\mathsf{bp}, \mathsf{sk}_1, m)$.

We define $\mathsf{rekey}_{\mathsf{VUF.Eval}}$ as follows:

$$\mathsf{rekey}_{\mathsf{VUF.Eval}}(\mathsf{pk}_1, \mathsf{sk}_2, m, T) = T \cdot e(\mathsf{Hash}(m), \mathsf{sk}_2^{-1}).$$

Observe that

$$\begin{array}{l} e(\mathsf{Hash}(m), \mathsf{sk}_1) = \mathsf{VUF.Eval}(\mathsf{sk}_1, m) \\ e(\mathsf{Hash}(m), \mathsf{sk}_1) = \mathsf{rekey}(\mathsf{pk}_1, \mathsf{sk}_2, m, \mathsf{VUF.Eval}(\mathsf{sk}_1 + \mathsf{sk}_2, m)) \end{array}$$

so the two outputs are identical.

The Pedersen DKG

1. Each party P_i chooses a random polynomial $f_i(z)$ over $\mathbb{F}$ of degree t

$$f_i(z) = a_{i,0} + a_{i,1}X + \cdots + a_{i,t}X^t$$

and broadcasts $X_{i,k} = g^{a_{i,k}}$ for $k = 0, \ldots, t$. Denote $a_{i,0}$ by x_i and $X_{i,0}$ by y_i. Each P_i computes $\bar{x}_{i,j} = f_i(\omega_j)$ for $j = 1, \ldots, n$ and sends $\bar{x}_{i,j}$ secretly to party P_j.

2. Each party P_j verifies the shares they received from the other parties by checking that

$$g^{\bar{x}_{i,j}} = \prod_{k=0}^{t} (X_{i,k})^{\omega_j^k}$$

If the check fails for an index i then P_j broadcasts a complaint against P_i.

3. If more than t parties complain against a party P_i then that party is disqualified. Otherwise P_i reveals the share $\bar{x}_{i,j}$ for each complaining parties P_j. If any of the revealed shares fails the equation, P_i is disqualified. The share of a disqualified party P_i is set to 1.
4. The public value y is computed as $y = \prod_{i=1}^{n} y_i$.

Fig. 6. Pedersen's Distributed Key Generation protocol.

E Specification and Security of Alternative DKGs

E.1 Background: Alternative DKGs constructions

We provide an overview of the Pedersen DKG construction in Figure 6 and the Fouque-Stern DKG construction in Figure 7.

E.2 Security proof for Theorem 3

We design an adversary $\mathcal{B}$ that takes as input pk_1 such that whenever the DKG outputs pk, $\mathcal{B}$ ouputs a, sk_2 such that the associated public key pk_2 satisfies $\mathsf{pk} = \alpha\mathsf{pk}_1 + \mathsf{pk}_2$. Suppose $\mathcal{B}$ receives input $\mathsf{pk}_1 = g_2$.

First $\mathcal{B}$ runs the DKG with $\mathcal{A}$. Let $\mathbb{I}_B \subset [1, n]$ be the set of corrupted (i.e. "bad") parties and $\mathbb{I}_G \subset [1, n]$ be the set of uncorrupted ("good") parties. Assume without loss of generality that party P_k is uncorrupted i.e. $k \in \mathbb{I}_G$. The adversary $\mathcal{B}$ follows the DKG protocol on behalf of the parties $P_i \in \mathbb{I}_G/\{k\}$ as prescribed, but for party P_k $\mathcal{B}$ *simulates* the adversarial view of this party's output, so that public view y_k broadcasted by P_k is equal to g_2.

When queried on P_i for $i \in \mathbb{I}_G/\{k\}$ the simulator samples $a_i \xleftarrow{\$} \mathbb{F}$ and computes $y_i = g_1^{a_i}$. They continue to return the public values $X_{i,j}$ and the private values $\bar{x}_{i,j}$ honestly.

The Fouque Stern DKG

The Fouque Stern DKG relies on

- A public key infrastructure where the public keys $\mathsf{pk}_1, \ldots, \mathsf{pk}_n$ consist of (N_j, G_j) for $1 \leq j \leq n$ where N_j is an RSA modulus and G_j has order N_j modulo N_j^2.
- The Paillier encryption scheme [53], where messages $m \in \mathbb{F}$ are encrypted as $\mathsf{Encrypt}(\mathsf{pk}_j, m) = G_j^m u^{N_j} \mod N_j^2$, where u us sampled randomly from $\mathbb{Z}_{N_j}^*$.
- A proving system with public parameters $\mathbb{G}_1$ to show that $(y, Y) \in \mathbb{G} \times \mathbb{Z}_{N_j^2}$ is such that $y = g_1^m$ and Y is a Paillier encryption of m under public key pk_j. The prover computes ciphertexts as follows: (1) choose random $r \xleftarrow{\$} \mathbb{F}$ and $s \xleftarrow{\$} \mathbb{Z}_{N_j}^*$; (2) set $\tau = (g_1^r, G_j^r s^{N_j} \mod N_j^2)$; (3) compute $e = \mathsf{Hash}(g, G_j, y, Y, \tau)$; (4) set $z = r + ex$ and $w = su^e \mod N_j$; (5) return (e, z, w). The verifier checks that $e = \mathsf{Hash}(g, G_j, y, Y, g^z y^{-e}, G_j^z w^{N_j} Y^{-e})$ and returns 1 if it is.

The Fouque Stern DKG works as follows.

1. Each party P_i chooses a random polynomial $f_i(z)$ over $\mathbb{F}$ of degree t

$$f_i(z) = a_{i,0} + a_{i,1}X + \cdots + a_{i,t}X^t$$

   and computes $\bar{x}_{i,j} = f_i(\omega_j)$ for $j = 1, \ldots, n$. Denote $a_{i,0}$ by x_i and $X_{i,0}$ by y_i. Each party broadcasts y_i, $X_{i,k} = g^{a_{i,k}}$ for $k = 1, \ldots, t$, $A_{i,j} = g^{\bar{x}_{i,j}}$ for $j = 1, \ldots, n$, $Y_{i,j} = \mathsf{Encrypt}(\mathsf{pk}_j, \bar{x}_{i,j})$ for $j = 1, \ldots, n$, and $\pi_{i,j} = \mathsf{Prove}(\mathbb{G}_1, (\mathsf{pk}_j, A_{i,j}, Y_{i,j}), (\bar{x}_{i,j}, u_{i,j}))$.
2. Each party verifies the shares by checking that

$$\mathsf{Verify}(\mathbb{G}_1, (\mathsf{pk}_j, A_{i,j}, Y_{i,j}), \pi_{i,j}) = 1$$

   for all i, j. They further choose $\alpha \xleftarrow{\$} \mathbb{F}$ and verify that

$$\prod_{k=0}^{t} X_{i,k}^{\alpha^k} = \prod_{k=1}^{n} A_{i,k}^{\ell_k(\alpha)}$$

   for all i. If the check fails for an index i then P_i is disqualified and their share is set to 1.
3. The public value y is computed as $y = \prod_{i=1}^{n} y_i$.

Fig. 7. Fouque and Stern's Distributed Key Generation protocol.

When queried on P_k the adversary $\mathcal{B}$ is required to output

$$g_2, X_{k,j}$$

that are indistinguishable from a valid output as well as t values $\bar{x}_{k,j}$. Assume without loss of generality that $1 \notin \mathbb{I}_B$ and that $|\mathbb{I}_B| = t$. Then $\mathcal{B}$ behaves as follows

1. Choose random $\bar{a}_j \xleftarrow{\$} \mathbb{F}$ for each $j \in \mathbb{I}_B$ and interpolate in the exponent to find $(X_{k,0}, \ldots, X_{k,t})$ such that

$$X_{k,0} = g_1^{c_0}, \ldots, X_{k,t} = g_1^{c_t}$$

where $\sum_{i=0}^{t} c_i X^i$ evaluates to $\bar{x}_{k,j}$ at ω_j for $j \in \mathbb{I}_B$ and $\log_{g_1}(A_1)$ at ω_1. Note that these c_i values are unknown to the simulator.

This simulation is perfect. If $\mathcal{A}$ broadcasts a complaint then reveal the relevant $\bar{x}_{k,j}$.

When the adversary returns their share

$$X_{i,0}, \ldots, X_{i,t}$$

If they do not send verifying $\bar{x}_{i,j}$ to any participant in $\mathbb{I}_G$ then $\mathcal{B}$ broadcasts a complaint. We have that $\mathcal{A}$ will be disqualified if they do not publicly send this value. In that case we set $a_i = 0$. Otherwise, $\mathcal{B}$ interpolates the $t+1$ values $\bar{x}_{i,j}$ to get $f_i(X)$ and set a_i to equal the 0 coefficient. Note that $|\mathbb{I}_G| \geq t+1$ because $t < n/2$.

When the DKG phase terminates one is left with the public key

$$\mathsf{pk} = y_1 \cdots y_n$$

Then $\mathcal{B}$ extracts

$$\begin{aligned} \mathsf{pk}_2 &= \textstyle\prod_{j=1, j\neq k}^{n} y_j \\ \mathsf{sk}_2 &= \textstyle\sum_{j=1, j\neq k}^{n} a_j \end{aligned}$$

such that $\mathsf{pk}_2 = g_1^{\mathsf{sk}_2}$ and such that

$$\begin{aligned} \mathsf{pk} &= g_2 \cdot \mathsf{pk}_2 = \mathsf{pk}_1 \cdot \mathsf{pk}_2 \\ \mathsf{sk} &= \mathsf{sk}_1 + \mathsf{sk}_2 \end{aligned}$$

Thus $\mathcal{B}$ returns $(1, \mathsf{sk}_2)$.

E.3 Security proof for Theorem 4

We design an adversary $\mathcal{B}$ that takes as input pk_1 such that whenever the DKG outputs pk, $\mathcal{B}$ outputs a, sk_2 such that the associated public key pk_2 satisfies $\mathsf{pk} = \alpha \mathsf{pk}_1 + \mathsf{pk}_2$. Suppose $\mathcal{B}$ receives input $\mathsf{pk}_1 = g_2$.

First $\mathcal{B}$ runs the DKG with $\mathcal{A}$. Let $\mathbb{I}_B \subset [1, n]$ be the set of corrupted (i.e. "bad") parties and $\mathbb{I}_G \subset [1, n]$ be the set of uncorrupted ("good") parties. Assume without loss of generality that party P_k is uncorrupted i.e. $k \in \mathbb{I}_G$. The

adversary $\mathcal{B}$ follows the DKG protocol on behalf of the parties $P_i \in \mathbb{I}_G/\{k\}$ as prescribed, but for party P_k $\mathcal{B}$ *simulates* the adversarial view of this party's output, so that public view y_k broadcasted by P_k is equal to g_2.

When queried on P_i for $i \in \mathbb{I}_G/\{k\}$ the simulator samples $a_i \xleftarrow{\$} \mathbb{F}$ and computes $y_i = g_1^{a_i}$. They continue to return the public values $X_{i,j}$ and the private values $\bar{x}_{i,j}$ honestly.

When queried on P_k the simulator is required to output

$$g_2, X_{k,j}, A_{k,j}, Y_{k,j}, \pi_{k,j}$$

that are indistinguishable from a valid output. Assume without loss of generality that $1 \notin \mathbb{I}_B$ and that $|\mathbb{I}_B| = t$. The simulator then behaves as follows

1. Choose random $\bar{a}_j \xleftarrow{\$} \mathbb{F}$ for each $j \in \mathbb{I}_B$ and interpolate in the exponent to find $(X_{k,0}, \ldots, X_{k,t})$ such that
$$X_{k,0} = g^{c_0}, \ldots, X_{k,t} = g^{c_t}$$
where $\sum_{i=0}^{t} c_i X^i$ evaluates to $\bar{x}_{k,j}$ at ω_j for $j \in \mathbb{I}_B$ and $\log_{g_1}(A_1)$ at ω_1. Note that these c_i values are unknown to the simulator.
2. Set
$$A_{k,i} = \prod_{j=0}^{t} X_{k,j}^{\omega_i^j}$$
for $1 \leq i \leq n$.
3. Encrypt $\bar{x}_{k,j}$ for $j \in \mathbb{I}_B$ honestly i.e. $Y_{k,j} = \mathsf{Encrypt}(\mathsf{pk}_j, \bar{x}_{k,j})$. For $j \in \mathbb{I}_G$ encrypt a random value i.e. choose $m_j \xleftarrow{\$} \mathbb{F}$ and set $Y_{k,j} = \mathsf{Encrypt}(\mathsf{pk}_j, m_j)$.
4. For $j \in \mathbb{I}_B$ compute $\pi_{k,j} = \mathsf{Prove}(\mathbb{G}_1, (\mathsf{pk}_j, A_{k,j}, Y_{k,j}), (\bar{x}_{k,j}, u_{k,j}))$ honestly with respect to the encryptions $Y_{k,j}$. For each $j \in \mathbb{I}_G$ we must simulate the proof. Choose $e_j, z_j, w_j \xleftarrow{\$} \mathbb{F}^2 \times \mathbb{Z}_{N^2}$. Program $\mathsf{Hash}(g, G_j, A_{k,j}, Y_{k,j}, g^{z_j} A_{k,j}^{-e_j}, G_j^{z_j} w_j^{N_j} Y_{k,j}^{-e_j})$ to return e_j and set $\pi_{k,j} = (e_j, z_j, w_j)$.

By the decisional composite residuosity assumption, the adversary cannot distinguish the simulated encryptions from real encryptions. The simulated proof is a perfect simulation provided that $\mathcal{A}$ has not already queried $\mathsf{Hash}()$ on the relevant values. Because these values are randomised, this happens with probability no greater than q_H/N_j^2.

When the adversary returns their share

$$y_i, X_i, \ldots, X_{i,t}, A_{i,1}, \ldots, A_{i,n}, Y_{i,1}, \ldots, Y_{i,n}, \pi_{i,1}, \ldots, \pi_{i,n}$$

for each $j \in \mathbb{I}_G$ decrypt $\bar{x}_{i,j} = \mathsf{Decrypt}(\mathsf{sk}_j, Y_{i,j})$. The probability that $A_{i,j} \neq g_1^{m_{i,j}}$ is equal to the probability that $\mathcal{A}$ forges a proof, which is negligible by the decisional composite residuosity assumption. Set $a_i = \bar{x}_{i,0}$.

By the synchrony assumption the DKG phase will eventually terminate, and one is left with the public key

$$\mathsf{pk} = y_1 \cdots y_n$$

Then $\mathcal{B}$ extracts

$$\begin{aligned}\mathsf{pk}_2 &= \textstyle\prod_{j=1,j\neq k}^{n} y_j \\ \mathsf{sk}_2 &= \textstyle\sum_{j=1,j\neq k}^{n} a_j\end{aligned}$$

such that $\mathsf{pk}_2 = g_1^{\mathsf{sk}_2}$ and such that

$$\begin{aligned}\mathsf{pk} &= \mathsf{pk}_1 \cdot \mathsf{pk}_2 \\ \mathsf{sk} &= \mathsf{sk}_1 + \mathsf{sk}_2\end{aligned}$$

for $\mathsf{pk} = g_1^{\mathsf{sk}}$. Thus $\mathcal{B}$ returns $(1, \mathsf{sk}_2)$.

F Proofs of VUF Security

In this section, we provide proofs of the security of our VUF from Section 7. We begin with a proof of Theorem 5.

Proof. We must show that no adversary $\mathcal{A}$, even an adversary who may choose the public key, can convince a verifier unless their two outputs derive the same values. To do this, we construct an adversary $\mathcal{B}$ such that whenever $\mathcal{A}$ breaks uniqueness, $\mathcal{B}$ breaks the SXDH assumption. Thus

$$\mathsf{Adv}^{unique}_{\mathcal{A}}(1^\lambda) \leq \mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{B}}(1^\lambda).$$

The adversary $\mathcal{B}$ receives an SXDH challenge $(\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$ that is either equal to $(\hat{h}_1, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi})$ or is randomly generated. They sample $\hat{u}_1 \xleftarrow{\$} \mathbb{G}_2$ and run $\mathcal{A}$ on the input $\mathsf{crs}_{\mathsf{vuf}} = (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{u}_1, \hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$.

When $\mathcal{A}$ queries the random oracle $\mathsf{Hash}_{\mathbb{G}_1}()$ on m, $\mathcal{B}_2$ will first check whether the response at m is already defined, and if yes return that response. Else $\mathcal{B}$ chooses z at random and returns g_1^z.

When $\mathcal{A}$ returns the public key pk, the VUF input m, and the signatures σ_1 and σ_2, $\mathcal{B}$ checks whether $\mathcal{A}$'s responses derive different values and both verify. If yes they return 1 indicating that they believe $(\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$ is an SXDH challenge. If not they return 0 indicating they believe $(\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$ is not an SXDH challenge.

We argue that if $\mathcal{B}$'s inputs are SXDH instances then it is statistically impossible for $\mathcal{A}$ to output two verifying responses that derive different values. This implies that $\mathsf{Adv}^{unique}_{\mathcal{A}}(1^\lambda) \leq \mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{B}}(1^\lambda)$.

Formally, $\mathcal{B}$ behaves as follows.

$\underline{\mathcal{B}(\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{h}_2, \hat{h}_3, \hat{h}_4)}$
$\hat{u}_1 \xleftarrow{\$} \mathbb{G}_2$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{u}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$
$(\mathsf{pk}, m, \sigma_1, \sigma_2) \xleftarrow{\$} \mathcal{A}(\mathsf{crs}_{\mathsf{vuf}})$
if $\mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma_1) = \mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma_2)$ return 0
if $\mathsf{VUF.Ver}(\mathsf{pk}, m, \sigma_1) = 0$ or $\mathsf{VUF.Ver}(\mathsf{pk}, m, \sigma_2) = 0$ return 0
else return 1

If $(\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$ are random then this is a perfect simulation of the uniqueness game. If instead

$$(\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4) = (\hat{h}_1, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi})$$

then whenever $\mathcal{A}$ returns a verifying signature $\sigma = (\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)$, we have that $\mathsf{Hash}_{\mathbb{G}_1}(m) = g_1^z$ and $e(A, \hat{h}_1)^z = e(g_1^z, \hat{\pi}_2)e(\pi_2, \hat{h}_3)e(\pi_4, \hat{h}_4)$. Hence if $\mathcal{A}$ outputs two verifying signatures then they must derive the same value.

To see that $\mathcal{A}$ cannot cheat if the input is an SXDH challenge, observe that by dividing the third verification equation by ξ times the first verification equation, *i.e.*,

$$\mathsf{VE}.3 - \xi\mathsf{VE}.1$$

we have that

$$e(\mathsf{pk}, \hat{h}_1) = e(g_1, \hat{\pi}_2\hat{\pi}_1^{-\xi})e(\pi_1, \hat{h}_3\hat{h}_1^{-\xi})e(\pi_3, \hat{h}_4\hat{h}_2^{-\xi}).$$

Also $\hat{h}_3\hat{h}_1^{-\xi} = 1$ and $\hat{h}_4\hat{h}_2^{-\xi} = 1$ implying that

$$e(\mathsf{pk}, \hat{h}_1) = e(g_1, \hat{\pi}_2\hat{\pi}_1^{-\xi})$$

and that $\hat{\pi}_2\hat{\pi}_1^{-\xi} = \hat{h}_1^a$ where a is such that $\mathsf{pk} = g_1^a$. Thus

$$e(Z, \hat{\pi}_2\hat{\pi}_1^{-\xi}) = e(Z, \mathsf{sk}).$$

Now observe that by multiplying the second verification equation by ξ we have that

$$1 = e(Z, \hat{\pi}_1^{\xi})e(\pi_2, \hat{h}_1^{\xi})e(\pi_4, \hat{h}_2^{\xi})$$

and

$$e(Z, \hat{\pi}_1^{-\xi}) = e(\pi_2, \hat{h}_3)e(\pi_4, \hat{h}_4).$$

Putting the two together we have that

$$e(Z, \mathsf{sk}) = e(Z, \hat{\pi}_2\hat{\pi}_1^{-\xi}) = e(Z, \hat{\pi}_2)e(\pi_2, \hat{h}_3)e(\pi_4, \hat{h}_4)$$

and $\mathcal{A}$ has output the unique outcome.

Next, we also prove Theorem 6.

Proof. Let $\mathcal{A}$ be an adversary playing in the unpredictability game. In particular, $\mathcal{A}$ is given a public key $\mathsf{pk} = g_1^a$ as input and aims to output some m and $e(\mathsf{Hash}_{\mathbb{G}_1}(m), \hat{h}_1)^a$ that they have not queried the signer on previously. We first transition into a game where $\mathcal{A}$ receives a simulated reference string using an adversary $\mathcal{B}_0$ against SXDH. We then demonstrate that there exists an adversary $\mathcal{B}_1$ that can program the oracle $\mathsf{Hash}_{\mathbb{G}_1}$ and that succeeds in breaking the BDH assumption with advantage $\frac{1}{q_H}\mathsf{Adv}^{sig}_{\mathcal{A}}(1^\lambda)$ where q_H is the total number of queries $\mathcal{A}$ makes to the oracle. Thus

$$\mathsf{Adv}^{unpredictable}_{\mathcal{A}}(1^\lambda) \leq \mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{B}_0}(1^\lambda) + q_H\mathsf{Adv}^{\mathsf{BDH}}_{\mathcal{B}_1}(1^\lambda).$$

Game unpredictable to Game 1: Let $\mathsf{Game}^1_{\mathcal{A}}(1^\lambda)$ denote $\mathcal{A}$ playing a modified unforgeability game where, instead of $\mathcal{A}$ receiving a randomly distributed reference string, $\mathcal{A}$ receives a simulated reference string of the form

$$(\hat{h}_1, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi+1}).$$

We construct an adversary $\mathcal{B}_0$ against the SXDH assumption in the first source group such that

$$\Pr[\mathsf{Game}^{unpredictable}_{\mathcal{A}}(1^\lambda)] - \Pr[\mathsf{Game}^1_{\mathcal{A}}] \leq \mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{B}_0}.$$

The adversary $\mathcal{B}_0$ uses an SXDH challenge to determine the form of the common reference string that it passes to $\mathcal{A}$. They receive an SXDH challenge $(\hat{h}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$ that is either equal to $(\hat{h}_1, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi})$ or is randomly generated. If $\mathcal{A}$ succeeds it returns 0, else it returns 1. Formally $\mathcal{B}_0$ behaves as follows.

$\underline{\mathcal{B}_0(\mathsf{bp}, \hat{h}_2, \hat{h}_3, \hat{h}_4)}$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{h}_2, \hat{h}_3, \hat{h}_4\hat{h}_1)$
$(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{VUF.Gen}(\mathsf{crs}_{\mathsf{vuf}})$
$(m, y) =\xleftarrow{\$} \mathcal{A}^{\mathsf{VUF.Sign}}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk})$
if $y = \mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)$ and $m \notin Q$ return 1
else return 0

$\underline{\mathsf{VUF.Sign}(m)}$
$\sigma \leftarrow \mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m)$
add m to query set Q
return σ

If $\hat{h}_4$ is random then $\mathcal{B}_0$ perfectly simulates $\mathsf{Game}^{unpredictable}_{\mathcal{A}}(1^\lambda)$ and succeeds whenever $\mathcal{A}$ fails. If $(\hat{h}_2, \hat{h}_3, \hat{h}_4) = (\hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi})$ then $\mathcal{B}_0$ perfectly simulates $\mathsf{Game}^1_{\mathcal{A}}(1^\lambda)$ and succeeds whenever $\mathcal{A}$ succeeds.

Game 1:
We design an adversary $\mathcal{B}_1$ against the BDH assumption that works as follows. As input $\mathcal{B}_1$ receives some

$$(g_1, g_2, g_3, \hat{h}_1, \hat{u}_1, \hat{u}_2) = (g_1, g_1^{\alpha}, g_1^{\beta}, \hat{h}_1, \hat{h}_1^{\gamma}, \hat{h}_1^{\alpha\gamma})$$

and aims to compute $e(g_1, \hat{h}_1)^{\alpha\beta}$. Let q_H be the maximum number of unique queries that $\mathcal{A}$ makes to $\mathsf{Hash}_{\mathbb{G}_1}$. Then $\mathcal{B}_1$ chooses i at random from $[1, q_H]$ and samples a simulated reference string $\mathsf{crs}_{\mathsf{vuf}}$ and retains a trapdoor (ρ, ξ) for simulating signatures. Then $\mathcal{B}_1$ runs $\mathcal{A}$ on input $\mathsf{crs}_{\mathsf{vuf}}$ and $\mathsf{pk} = (g_2, \hat{u}_2)$.

When $\mathcal{A}$ queries the random oracle $\mathsf{Hash}_{\mathbb{G}_1}(\cdot)$ on m, $\mathcal{B}_2$ will first check whether the response at m is already defined, and if yes return that response. If m is the ith unique query, $\mathcal{B}$ returns g_3. Else $\mathcal{B}$ chooses z at random and returns g_1^z.

When $\mathcal{A}$ queries the signing oracle on m, $\mathcal{B}_1$ will first query $\mathsf{Hash}_{\mathbb{G}_1}(m)$ and if the response is g_3 then $\mathcal{B}_1$ aborts. Else $\mathcal{B}_1$ extracts z from the oracle such that $\mathsf{Hash}_{\mathbb{G}_1}(m) = g_1^z$ and uses (ρ, ξ, z) to compute a signature.

When $\mathcal{A}$ returns an evaluation (m, y), $\mathcal{B}_1$ aborts if m is not the ith query to $\mathsf{Hash}_{\mathbb{G}_1}(\cdot)$. Else $\mathcal{B}_1$ returns y.

Formally, $\mathcal{B}_1$ behaves as follows.

$\underline{\mathcal{B}_1(\mathsf{bp}, g_2, g_3, \hat{u}_1, \hat{u}_2)}$
$i \xleftarrow{\$} [1, q_H]$, $H \leftarrow \emptyset$
$\rho, \xi \xleftarrow{\$} \mathbb{F}$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{u}_1, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi+1})$
$(m, y) \xleftarrow{\$} \mathcal{A}^{\mathsf{VUF.Sign}}(\mathsf{crs}_{\mathsf{vuf}}, (g_2, \hat{u}_2))$
if m is not the ith query to $\mathsf{Hash}_{\mathbb{G}_1}(\cdot)$ return $\perp$
return y

$\underline{\mathsf{Hash}_{\mathbb{G}_1}(m)}$
if $(m, Z, z) \in H$ return Z
if ith unique query $z \leftarrow \perp$, $Z \leftarrow g_3$
else $z \xleftarrow{\$} \mathbb{F}$, $Z \leftarrow g_1^z$
add (m, Z, z) to query set H
return Z

$\underline{\mathsf{VUF.Sign}(m)}$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
select (m, Z, z) from H
if $z = \perp$ return $\perp$
$\alpha, \beta \xleftarrow{\$} \mathbb{F}$
$\hat{\pi}_1, \hat{\pi}_2 \leftarrow \hat{h}_1^{\alpha}, \hat{h}_1^{\beta}$
$\pi_1 \leftarrow g_1^{\beta\rho - \alpha - \alpha\rho\xi} g_2^{-\rho}$
$\pi_2 \leftarrow \pi_1^z$
$\pi_3 \leftarrow g_1^{\alpha\xi - \beta} g_2$
$\pi_4 \leftarrow \pi_3^z$
return $(\pi_1, \pi_2, \pi_3, \pi_4, \hat{\pi}_1, \hat{\pi}_2)$

We see that $\mathsf{VUF.Sign}$ outputs signatures that are distributed identically to those in $\mathsf{Game}^1_{\mathcal{A}}(1^\lambda)$. This is because $\hat{\pi}_1, \hat{\pi}_2$ are distributed uniformly at random. Given $\hat{\pi}_1, \hat{\pi}_2$ there exists unique π_1 and π_3 that satisfy both the first and third verification equations. Similarly, given $\hat{\pi}_1, \hat{\pi}_2$ there exists unique π_2 and π_4 that satisfy both the second and fourth verification equations.

Since $\mathsf{VUF.Sign}$ outputs signatures that are distributed identically to those in $\mathsf{Game}^1_{\mathcal{A}}(1^\lambda)$ we have that

$$\mathsf{Adv}^{\mathsf{Game}^1}_{\mathcal{A}}(1^\lambda) \leq q_H \mathsf{Adv}^{\mathsf{BDH}}_{\mathcal{B}_1}(1^\lambda).$$

Indeed $\mathcal{B}_1$ succeeds whenever $\mathcal{A}$ returns an evaluation on the message programmed to equal g_3 i.e. when $\mathcal{A}$ returns

$$\mathsf{VUF.Eval}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{sk}, m_i) = e(\mathsf{Hash}_{\mathbb{G}_1}(m), \mathsf{sk}) = e(g_3, \hat{h}_1^{\alpha}) = e(g_1, \hat{h}_1)^{\alpha\beta}$$

This happens with probability $\frac{1}{q_H}$.

G Optimized VUF

The construction described in Section 7 has 4 $\mathbb{G}_1$ and 2 $\mathbb{G}_2$ elements for each VUF, and requires the verifier to perform 10 pairings to check correctness. An additional 3 pairings are performed by the deriver in obtaining the unpredictable component. Here we discuss an optimisation that reduces VUF to just 2 $\mathbb{G}_1$ elements, which is only twice as expensive as the BLS signature scheme, and unlike BLS we can thresholdise our scheme using our more efficient DKG that outputs group elements as secret key shares.

Our optimisation works by transferring some elements of the signature into the public key, and allowing the private key to contain some field elements. The observant reader might be thinking that this removes a lot of the motivation for our VUF because by having a private key contain field elements we are no longer fully structure preserving. Nonetheless, we argue that our optimised construction is useful, because these field elements are entirely independent of the original private key. Thus in our threshold scheme, a user can split their public key into components derived in the DKG and components they chose themselves. The components derived in the DKG consist entirely of group elements.

G.1 Construction

Our optimised VUF construction is given in Figure 8. We prove in Theorem 7 and Theorem 8 that the optimised construction retains its uniqueness and unpredictability. The intuition behind our optimisations stems from the fact that the following checks by the verifier do not involve message-specific elements:

$$1 = e(g_1, \hat{\pi}_1)e(\pi_1, \hat{h}_1)e(\pi_3, \hat{h}_2) \tag{5}$$

$$e(\mathsf{pk}, \hat{h}_1) = e(g_1, \hat{\pi}_2)e(\pi_1, \hat{h}_3)e(\pi_3, \hat{h}_4) \tag{6}$$

and thus the signature elements $\hat{\pi}_1, \hat{\pi}_2, \pi_1$ and π_3 can be a part of the public key instead of the signature.

Our public key consists of the additional elements

$$(p_1, p_2, \hat{p}_1, \hat{p}_2) = (g_1^{\alpha}, g_1^{\beta}, \hat{h}_1^{-\alpha}\hat{h}_2^{-\beta}, \hat{h}_1^{a}\hat{h}_3^{-\alpha}\hat{h}_4^{-\beta})$$

where $\hat{h}_1^a, \alpha, \beta$ are kept secret. Our VUF consists only of the π_1 and π_3 components from the unoptimised VUF, which we relabel to be π_1 and π_2.

Our verification algorithm is now split into two separate algorithms: one to check that the public key is well-formed and the other to check that the VUF is well-formed. Specifically we introduce a new $\mathsf{VUF.KeyVer}$ function. When a verifier receives a public key pk for the first time, they run $\mathsf{VUF.KeyVer}(\mathsf{crs_{vuf}}, \mathsf{pk})$ and cache the result, so subsequent executions of $\mathsf{VUF.Ver}(\mathsf{crs_{vuf}}, \mathsf{pk}, m, \sigma)$ receive a smaller VUF of 2 $\mathbb{G}_1$ elements and perform only 3 pairings.

The benefits in how this can speed up our threshold VUF are mostly seen by the aggregator. We will still require that the final aggregated VUF is unoptimised because the verifier does not know which public keys contributed to the outcome.

However the shares of the VUF that are communicated to the aggregator by the signing parties can be optimised. In particular, the aggregator aggregates the public key components (which it only sees once) together with the per-message VUF shares (which are approximately 1/4 the size). They can also check the well-formedness of the contributions using 3 pairings rather than 10.

$\underline{\mathsf{Setup}(\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1})}$
$\hat{u}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4 \xleftarrow{\$} \mathbb{G}_2$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{u}_1, \hat{h}_2, \hat{h}_3, \hat{h}_4)$
return $\mathsf{crs}_{\mathsf{vuf}}$

$\underline{\mathsf{VUF.Gen}(\mathsf{crs}_{\mathsf{vuf}})}$
$a, \alpha, \beta \xleftarrow{\$} \mathbb{F}$
$\mathsf{pk} = \begin{pmatrix} A \in \mathbb{G}_1, \hat{u}_2 \in \mathbb{G}_2, \\ p_1, p_2 \in \mathbb{G}_1, \hat{p}_1, \hat{p}_2 \in \mathbb{G}_2 \end{pmatrix}$
$\leftarrow \begin{pmatrix} g_1^a, \hat{u}_1^a, g_1^\alpha, g_1^\beta, \\ \hat{h}_1^{-\alpha}\hat{h}_2^{-\beta}, \hat{h}_1^a\hat{h}_3^{-\alpha}\hat{h}_4^{-\beta} \end{pmatrix}$
$\mathsf{sk} \leftarrow \left(\hat{h}_1^a, \alpha, \beta\right) \in \mathbb{G}_2 \times \mathbb{F}^2$
return $(\mathsf{pk}, \mathsf{sk})$

$\underline{\mathsf{VUF.Sign}(\mathsf{crs}_{\mathsf{vuf}}, \hat{\mathsf{sk}}, m)}$
$(\hat{h}_1^a, \alpha, \beta \in \mathbb{G}_2 \times \mathbb{F}^2) \leftarrow \mathsf{parse}(\mathsf{sk})$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
$\pi_1, \pi_2 \leftarrow Z^\alpha, Z^\beta$
return (π_1, π_2)

$\underline{\mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)}$
$\begin{pmatrix} A \in \mathbb{G}_1, \hat{u}_2 \in \mathbb{G}_2, \\ p_1, p_2 \in \mathbb{G}_1, \hat{p}_1, \hat{p}_2 \in \mathbb{G}_2 \end{pmatrix} \leftarrow \mathsf{parse}(\mathsf{pk})$
$(\pi_1, \pi_2 \in \mathbb{G}_1^2) \leftarrow \mathsf{parse}(\sigma)$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
return $e(Z, \hat{p}_2)e(\pi_1, \hat{h}_3)e(\pi_2, \hat{h}_4)$

$\underline{\mathsf{VUF.KeyVer}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk})}$
$\begin{pmatrix} A \in \mathbb{G}_1, \hat{u}_2 \in \mathbb{G}_2, \\ p_1, p_2 \in \mathbb{G}_1, \hat{p}_1, \hat{p}_2 \in \mathbb{G}_2 \end{pmatrix} \leftarrow \mathsf{parse}(\mathsf{pk})$
check:
$1 = e(g_1, \hat{p}_1)e(p_1, \hat{h}_1)e(p_2, \hat{h}_2)$
$e(A, \hat{h}_1) = e(g_1, \hat{\pi}_2)e(\pi_1, \hat{h}_3)e(\pi_3, \hat{h}_4)$
return 1 if all checks pass, else return 0

$\underline{\mathsf{VUF.Ver}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk}, m, \sigma)}$
$\begin{pmatrix} A \in \mathbb{G}_1, \hat{u}_2 \in \mathbb{G}_2, \\ p_1, p_2 \in \mathbb{G}_1, \hat{p}_1, \hat{p}_2 \in \mathbb{G}_2 \end{pmatrix} \leftarrow \mathsf{parse}(\mathsf{pk})$
$(\pi_1, \pi_2 \in \mathbb{G}_1^2) \leftarrow \mathsf{parse}(\sigma)$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
check $1 = e(Z, \hat{p}_1)e(\pi_1, \hat{h}_1)e(\pi_2, \hat{h}_2)$
return 1 if all checks pass, else return 0

Fig. 8. Optimized verifiable unpredictable function

G.2 Security Proofs

Theorem 7. *The algorithm in Figure 8 is a unique VUF under the* SXDH *assumption in the random oracle model.*

Proof. The proof goes through exactly the same as for Theorem 5 except for the following changed. When distinguishing an SXDH challenge, the adversary $\mathcal{B}$ first calls both VUF.KeyVer and VUF.Ver (instead of just VUF.Ver), and returns 0 if either checks fail.

Theorem 8. *The algorithm in Figure 8 is an unpredictable VUF under the* SXDH *and the* BDH *assumption in the random oracle model.*

Proof. The first part of the proof remains the same. We transition from $\mathsf{Game}^{\mathcal{A}}_{unpredictable}$ to $\mathsf{Game}^{\mathcal{A}}_{1}$, where $\mathcal{A}$ receives a simulated reference string of the form

$$(\hat{h}_1, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi+1}).$$

$\mathcal{B}_1$ is also defined differently, since the pk it runs $\mathcal{A}$ with contains more elements. Formally $\mathcal{B}_1$ behaves as follows.

$\underline{\mathcal{B}_1(\mathsf{bp}, g_2, g_3, \hat{u}_1, \hat{u}_2)}$
$i \xleftarrow{\$} [1, q_H]$, $H \leftarrow \emptyset$
$\rho, \xi, \alpha, \beta \xleftarrow{\$} \mathbb{F}$
$\mathsf{crs}_{\mathsf{vuf}} \leftarrow (\mathsf{bp}, \mathsf{Hash}_{\mathbb{G}_1}, \hat{h}_1^{\rho}, \hat{h}_1^{\xi}, \hat{h}_1^{\rho\xi+1})$
$\mathsf{pk} = (g_2, \hat{u}_1, \hat{u}_2, p_1, p_2, \hat{p}_1, \hat{p}_2)$
$\quad \leftarrow (g_2, \hat{u}_1\hat{u}_2, g_1^{\beta\rho-\alpha-\alpha\rho\xi} g_2^{-\rho}, g_1^{\alpha\xi-\beta}, \hat{h}_1^{\alpha}, \hat{h}_2^{\beta})$
$(m, \pi) \xleftarrow{\$} \mathcal{A}^{\mathsf{VUF.Sign}}(\mathsf{crs}_{\mathsf{vuf}}, \mathsf{pk})$
if m is not the ith query to $\mathsf{Hash}_{\mathbb{G}_1}()$ return $\bot$
return $\mathsf{VUF.Derive}(\mathsf{crs}_{\mathsf{vuf}}, m, \pi)$

$\underline{\mathsf{Hash}_{\mathbb{G}_1}(m)}$
if $(m, Z, z) \in H$ return Z
if ith unique query $z \leftarrow \bot$, $Z \leftarrow g_3$
else $z \xleftarrow{\$} \mathbb{F}$, $Z \leftarrow g_1^z$
$H \leftarrow H \cup \{(m, Z, z)\}$
return Z

$\underline{\mathsf{VUF.Sign}(m)}$
$Z \leftarrow \mathsf{Hash}_{\mathbb{G}_1}(m)$
select (m, Z, z) from H
if $z = \bot$ return $\bot$
$\pi_1 \leftarrow p_1^z$
$\pi_2 \leftarrow p_2^z$
return (π_1, π_2)

As in Theorem 6, we see that $\mathsf{VUF.Sign}$ outputs signatures that are distributed identically to those in $\mathsf{Game}^{\mathcal{A}}_{1}(1^{\lambda})$. This is because given $\hat{p}_1, \hat{p}_2$, there are unique π_1 and π_2 that satisfy the verification equation. The rest of the proof is the same, arriving at the following conclusion.

$$\mathsf{Adv}^{unpredictable}_{\mathcal{A}}(1^{\lambda}) \leq \mathsf{Adv}^{\mathsf{SXDH}}_{\mathcal{B}_0}(1^{\lambda}) + q_H \mathsf{Adv}^{\mathsf{BDH}}_{\mathcal{B}_1}(1^{\lambda})$$