Exponent-VRFs and Their Applications

Dan Boneh[1], Iftach Haitner[2,3], and Yehuda Lindell[3]

[1] Stanford University
[2] Tel-Aviv University
[3] Coinbase

Abstract. *Verifiable random functions (VRFs)* are pseudorandom functions with the addition that the function owner can prove that a generated output is correct (i.e., generated correctly relative to a committed key). In this paper we introduce the notion of an **exponent-VRF (eVRF)**: a VRF that does not provide its output y explicitly, but instead provides $Y = y \cdot G$, where G is a generator of some finite cyclic group (or $Y = g^y$ in multiplicative notation). We construct eVRFs from DDH and from the Paillier encryption scheme (both in the random-oracle model). We then show that an eVRF is a powerful tool that has many important applications in threshold cryptography. In particular, we construct (1) a one-round fully simulatable distributed key-generation protocol (after a single two-round initialization phase), (2) a two-round fully simulatable signing protocol for multiparty Schnorr with a deterministic variant, (3) a two-party ECDSA protocol that has a deterministic variant, (4) a threshold Schnorr signing protocol where the parties can later prove that they signed without being able to frame another group, and (5) an MPC-friendly and verifiable HD-derivation. All these applications are derived from this single new eVRF abstraction. The resulting protocols are concretely efficient.

Table of Contents

1 Introduction

A *pseudorandom function* (PRF) [34] $F(k, x)$ is a keyed function whose outputs are indistinguishable from random elements in the range of the PRF. In some applications, it is important to force the secret-key owner to always use the same key and to generate correct outputs. A *verifiable random function* (VRF) [50] associates a public verification key vk with the secret key k, and enables the owner to output a proof π, together with $y = F(k, x)$, that attests to the fact that y is correct.

In this paper, we introduce a VRF enhancement that we call an **exponent VRF**, or **eVRF**, which is a variant of a VRF that does not provide the VRF's output y explicitly, but rather provides $Y = y \cdot G$ where G is a group generator of some finite cyclic group $\mathbb{G}$, together with a proof π that Y was computed correctly by computing $y \leftarrow F(k, x)$ and $Y \leftarrow y \cdot G$. We use the term "exponent VRF", since the VRF output is provided in the exponent and not in the clear.[4]

eVRFs are useful in settings where the discrete log problem (or DDH) is hard over the group, and a party needs to generate a pseudorandom value r, and send $R \leftarrow r \cdot G$ to other parties. If the sender generates (r, R) using an eVRF, the receiving parties can verify that R is consistent with an initially committed key k. Concretely, consider a very basic setting where two parties wish to generate a random group element R in $\mathbb{G}$, where the parties hold shares r_1 and r_2 such that $(r_1 + r_2) \cdot G = R$, and no party knows $r = r_1 + r_2$. This basic building block is used in distributed key generation, and in ECDSA and Schnorr/EdDSA signing. The naive way of generating R is for each party P_i, for $i \in \{1, 2\}$, to choose a random r_i and send $R_i \leftarrow r_i \cdot G$ to the other party. In such a protocol, however, a cheating P_2 can wait to obtain R_1 from P_1, choose a random r, and send $R_2 \leftarrow r \cdot G - R_1$ back to P_1. This enables P_2 to single handedly determine the output $R := R_1 + R_2 = r \cdot G$, while knowing the discrete log r of R. This can be mitigated by forcing P_1 and P_2 to each provide a zero-knowledge proof of knowledge of the discrete log together with their values R_1 and R_2, respectively. A cheating P_2, however, can still receive R_1 and then locally try many values r_2 and set $R_2 \leftarrow r_2 \cdot G$ until $R = R_1 + R_2$ has some predetermined structure. For example, if P_2 wishes the ten least significant bits of R to equal zero, then it would need to try approximately 2^{10} values of r_2. In order to prevent such bias, protocols employ a commit-and-open approach: the parties first send *commitments* to their R_i values, and open them in the next round (for simulatability, proofs of knowledge are also included). This approach adds an extra round to the protocol.

An eVRF provides a much simpler construction for this basic building block. Suppose the parties have already generated and shared eVRF public verification keys vk_1 and vk_2. Then they can choose R_1 and R_2 as the eVRF outputs on some agreed-upon nonce, such as a simple counter. Neither party has any freedom in choosing its value R_i. This means that the first naive protocol described above becomes fully secure. In particular, each party sends R_i together with a proof that R_i is the output of its eVRF. The parties then set $R := R_1 + R_2$. This way they can generate R with a single message and using only one round, and no party can bias the output in any way, since they are already fully committed to their VRF value. Thus, an eVRF eliminates the need for commit-and-open and enables us to save a full round of communication in threshold signing and key generation protocols.

Applications. Using an eVRF, we provide several results in threshold signing using a unified framework. Specifically, we construct

[4] The use of the term *exponent* comes from multiplicative notation where g is the group generator and $Y = g^y$. However, throughout the paper we use additive notation for the group operation.

1. A concretely efficient, two-round, *fully simulatable* multiparty Schnorr signing protocol. Previous two-round protocols are either proven via a game-based definition (e.g., [40]) or are not concretely efficient (e.g., [31]). The resulting construction is a generalization of the MuSig-DN scheme of Nick, Ruffing, Seurin, and Wuille [54,55], discussed in related work below.
2. The first concretely efficient, two-round, two-party ECDSA signing protocol with full simulatability. Previous work achieving two-round two-party ECDSA signing did not use a standard signing functionality [26].
3. A concretely efficient, two-round, *deterministic* signing protocols for multiparty Schnorr and two-party ECDSA. Previous protocols use garbled circuits and so have more rounds and are less efficient [32]. Very recently Komlo and Goldberg [41] proposed a protocol for Schnorr signatures with similar properties, but using very different techniques.
4. A distributed key generation protocol that requires only a *single round* to generate a key, after a one-time two-round initialization phase. A recent work of Katz [38, §8] presents a protocol with similar properties using generic tools such a general NIZK proof system.
5. A hierarchical-deterministic (HD) key derivation method analogous to Bitcoin's BIP032 [63] that also enables parties to efficiently prove that a public key was derived correctly from the root secret. Unlike the standard BIP032, our derivation method is MPC friendly.

All these applications follow directly from the new eVRF abstraction.

Our multiparty (probabilistic) Schnorr and distributed key generation protocols also have *threshold* variants, with the above number of rounds as long as the set of participating parties is known at the onset. These protocols fulfill a new property that we call *proof of quorum identity*: the participating parties can later prove that they are the ones who participated, but are unable to frame any other subset. We achieve this while still generating a full standard Schnorr signature.

All of our protocols are UC secure [15] for static malicious adversaries and a dishonest majority, and are proven under standard assumptions in the random-oracle model. The use of simulation-based MPC definition, like UC security [15], has many advantages. In particular, security under composition with any protocol is guaranteed, as well as security even when related keys are used (like when BIP032 derivation is used) or when keys are generated with poor entropy. In such cases, the MPC protocol provides the same level of security as a locally computed signature, which is of course optimal. Our results provide an option to those who need two-round signing protocols, but still want to maintain full simulatability and composition.

We next construct two eVRF schemes, both in the random-oracle model. Here we give a brief overview of the constructions. Fix a "target group" $\mathbb{G}$ of order q and a generator G of $\mathbb{G}$. An eVRF is a triple of PPT algorithms $(\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$, where (i) $\mathsf{KGen} \to (k, \mathsf{vk})$ samples the (secret) key k and public verification key vk, (ii) $\mathsf{Eval}(k, x) \to (y, Y, \pi)$ outputs a pseudorandom $y \in \mathbb{Z}_q$, its value in the exponent $Y \leftarrow y \cdot G$, and a proof π, and (iii) $\mathsf{Verify}(\mathsf{vk}, x, Y, \pi) \to \{0, 1\}$ verifies that Y is consistent with vk and x.

A construction from Paillier. Assume for a moment that the eVRF key owner has a secret trapdoor that lets it efficiently compute discrete logarithms in the target group $\mathbb{G}$. In such a case, we could let H be a random oracle mapping arbitrary strings to uniform values in $\mathbb{G}$. Then, the eVRF evaluation would involve hashing the input x into a random group element, $Y \leftarrow \mathsf{H}(x)$, and then computing $y \leftarrow \log Y$. The verification procedure would simply verify that $\mathsf{H}(x) = Y$. This would be a perfect eVRF, where every party uses a different group $\mathbb{G}$.

Since we have no trapdoors to enable the efficient computation of the discrete log in groups of interest, we take a similar approach using an intermediate hard problem for which the secret-key owner has a trapdoor. Specifically, let H be a hash function $\mathsf{H} : \mathcal{X} \mapsto [N]$, for some $N \geq |\mathbb{G}|$. Now, we could take any trapdoor permutation f on the domain $[N]$, and have the secret key owner invert the permutation on $\mathsf{H}(x)$ to get $y \in [N]$ and set $Y := y \cdot G$. It would then prove that $f^{-1}(\mathsf{H}(x)) \cdot G = Y$. The challenge with this approach is finding an efficient zero-knowledge proof for this relation. To handle this challenge we use the additively homomorphic Paillier encryption scheme [57] instead of a trapdoor permutation. The key generation algorithm outputs a Paillier public and secret key pair $(\mathsf{sk}, \mathsf{pk})$. The evaluation algorithm $\mathsf{Eval}(\mathsf{sk}, x)$ acts as follows:

1. Hash the input x into a Paillier ciphertext ct, using a hash function $\mathsf{H} : \mathcal{X} \to \mathbb{Z}_{n^2}$, where $\mathbb{Z}_{n^2}$ is the set of Pallier ciphertexts for pk;
2. Decrypt ct using sk to get the plaintext $y \in [n]$, and set $Y := y \cdot G$ in $\mathbb{G}$;
3. Generate a zero-knowledge proof π that the Pallier decryption of $ct \in \mathbb{Z}_{n^2}$ modulo q is equal to the discrete log of $Y \in \mathbb{G}$ base G.

The value y is pseudorandom since ct is a random ciphertext, when H is modeled as a random oracle. Furthermore, since encryption is binding and the proof is sound, it is not possible to cheat and provide some different $Y' \neq Y$ and prove that it is consistent with pk and x. The challenge is to design an efficient zero-knowledge proof, which we do in Section 5.

A construction from DDH. Our second construction uses the classic DDH-based PRF [52] defined as $F(k, x) := k \cdot \mathsf{H}(x) \in \mathbb{G}_\mathrm{S}$, where $\mathbb{G}_\mathrm{S}$ is a finite cyclic group and H is a hash function $\mathsf{H} : \mathcal{X} \to \mathbb{G}_\mathrm{S}$. This PRF can be proved secure when the DDH assumption holds in $\mathbb{G}_\mathrm{S}$ and H is modeled as a random oracle. Concretely, $\mathbb{G}_\mathrm{S}$ can be the group of points of an elliptic curve E defined over some prime field $\mathbb{F}_q$, where q is the order of our target group $\mathbb{G}$. The eVRF will output $F(k, x)$ "in the exponent" of the target group $\mathbb{G}$. To do so we treat a point P in the group $\mathbb{G}_\mathrm{S} = E(\mathbb{F}_q)$ as a pair (x_P, y_P) in $\mathbb{F}_q^2$. Now, for a given key k and input $x \in \mathcal{X}$ the evaluation algorithm $\mathsf{Eval}(k, x)$ works as follows:

1. Compute $P := k \cdot \mathsf{H}(x)$ in $\mathbb{G}_\mathrm{S}$ and let x_P in $\mathbb{F}_q$ be the x-coordinate of P;
2. Set $Y := x_P \cdot G$ in the target group $\mathbb{G}$;
3. Generate a zero-knowledge proof π that $Y \in \mathbb{G}$ is computed correctly with respect to the input x and a commitment to k.

Then $\mathsf{Eval}(k, x)$ outputs (x_P, Y, π). As described, the PRF output x_P only ranges over about half of $\mathbb{F}_q$, namely the x-coordinates of points on the curve $E(\mathbb{F}_q)$. In Section 6 we show how to augment the construction using the left over hash lemma so that the range of the PRF is all of $\mathbb{F}_q$. The challenge is to design an efficient zero-knowledge proof that Y is computed correctly. We design such a proof in Section 6. Our proof is concretely practical; we estimate that it takes only a few tens of milliseconds to generate and verify the proof on a single thread on a modern processor.

Related work. The MuSig-DN scheme, proposed by Nick, Ruffing, Seurin, and Wuille [54,55], is an elegant two-round Schnorr multisignature scheme. Their scheme uses the DDH-based eVRF discussed above, but without abstracting it out as standalone primitive. We therefore attribute the DDH-based eVRF to Nick et al. [54,55]. The construction we present in Section 6 is a little simpler in that we avoid using quadratic extensions.

Katz [38, §8] presents a single round distributed key generation protocol that is structurely similar to the one obtained from an eVRF. That protocol uses a general NIZK proof system to prove that a general PRF is evlauated correctly in the exponent.

Recently, Kondi, Orlandi, and Roy [42,43] gave a two-round deterministic two-party Schnorr signing protocol. Their protocol is based on pseudorandom correlation functions (PCFs), and is quite different from our eVRF abstraction. Interestingly, their construction also uses the decryption of random Paillier ciphertexts.

Finally, we note that a classic one-round distributed key generation protocol, due to Fouque and Stern [29], requires a quadratic number of messages and relies on a synchronous network. Another one-round distributed key generation protocol, due to Groth [35], makes use of pairings and chunked encryption.

Overview of the rest of the paper. After some preliminary definitions in Section 2, we formally define exponent VRFs in Section 3. We give two definitions: a game based definition and an ideal-functionality based definition. The game based definition is useful for constructing an eVRF, while the simulation based definition is useful for describing and proving security of the applications. Theorem 2 proves that the two definitions are equivalent (assuming a zero-knowledge proof-of-knowledge of the private key). In Section 4, we present the many applications of eVRFs and prove their security. Then the Paillier-based eVRF is presented in Section 5, and the DDH-based eVRF is presented in Section 6. Our work leaves a number of important open questions for future work described in Section 7.

2 Preliminaries

Notation. We use $\lambda \in \mathbb{Z}^{(>0)}$ to denote the security parameter, and $\overset{c}{\approx}$ to denote computational indistinguishability. We write $x \leftarrow y$ to denote the assignment of the value of y to x, and $x \leftarrow\!\!\$\ S$ to denote sampling an element from the set S independently and uniformly at random. Similarly, for a randomized algorithm $\mathcal{A}$, we write $y \leftarrow\!\!\$\ \mathcal{A}(x)$ to denote that y is distributed according to the output of $\mathcal{A}(x)$ (over uniformly sampled random coins). For integers n, m we use $[n]$ for the set $\{1, \ldots, n\}$ and use $[n, m]$ for the set $\{n, n+1, \ldots, m\}$. We use additive notation for the group operation, and 0 for the group identity.

2.1 Pseudorandom Functions

We define pseudorandom functions and verifiable pseudorandom functions in a way that is convenient for the presentation in this paper. Since our eVRF constructions are given in the ROM, the following definitions are given for oracle-aided constructions. In such constructions, all entities (including the adversary) have oracle access to the same function, and they are secure in the ROM, if they are secure with respect to the all-function ensemble (the ensemble $\mathcal{O}$ according to the following definition). We start with formally defining ensemble of function families.

Definition 1 (Function families). *A **function family** with respect to domain/range $(\mathcal{X}, \mathcal{Y})$ is a family of functions $\mathcal{F} = \{f : \mathcal{X} \mapsto \mathcal{Y}\}$. We let $O_{\mathcal{X},\mathcal{Y}}$ denote the all-function family from $\mathcal{X}$ to $\mathcal{Y}$. A **function-family ensemble** with respect to domain/range ensemble $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ is an ensemble of function families $\{\mathcal{F}_\lambda\}$, where each $\mathcal{F}_\lambda$ is a subset of $O_{\mathcal{X}_\lambda,\mathcal{Y}_\lambda}$.*

Definition 2. *A **pseudorandom function (PRF)** with respect to the domain/range ensemble $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ and function-family ensemble $\mathcal{H}$, is a pair of oracle-aided* PPT *algorithms* (KGen, Eval) *such that for all $\lambda \in \mathbb{N}$ and $h \in \mathcal{H}_\lambda$:*

– $\mathsf{KGen}^h(1^\lambda) \to k$: *outputs a secret key* $k \in \mathcal{K}$.
– $\mathsf{Eval}^h(1^\lambda, k, x) \to y$: *on key* k *and input* $x \in \mathcal{X}_\lambda$, deterministically *outputs* $y \in \mathcal{Y}_\lambda$.

When clear from the context, we omit 1^λ *from the input list of* Eval. *The PRF is* ***secure*** *if for all oracle-aided* PPT A*:*

$$\Big|\Pr\big[\mathsf{A}^{h,\mathsf{Eval}^h(k,\cdot)}(1^\lambda) = 1\big] - \Pr\big[\mathsf{A}^{h,o(\cdot)}(1^\lambda) = 1\big]\Big| \le \mathsf{negl}(\lambda), \tag{1}$$

where $h \leftarrow\!\!\$\ \mathcal{H}_\lambda$, $k \leftarrow\!\!\$\ \mathsf{KGen}^h(1^\lambda)$, *and* $o \leftarrow\!\!\$\ \mathcal{O}_{\mathcal{X}_\lambda,\mathcal{Y}_\lambda}$.

We next define verifiable pseudorandom functions (VRFs). Our definition strengthens the standard definition of VRFs, by (naturally) demanding *simulatability*, meaning that the verifiability proof does not leak significant information beyond correctness.

Definition 3. *A* ***simulatable verifiable random function (VRF)*** *with respect to domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ *and function-family ensemble* $\mathcal{H}$*, is a triple of oracle-aided* PPT *algorithms* $(\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$ *such that for every* $\lambda \in \mathbb{N}$ *and* $h \in \mathcal{H}_\lambda$*:*

– $\mathsf{KGen}^h(1^\lambda) \to (k, \mathsf{vk})$.
– $\mathsf{Eval}^h(1^\lambda, k, x) \to (y, \pi)$. *We let* $\mathsf{Eval}_1^h(1^\lambda, k, x) \to y$ *be the same as* Eval*, but only outputs its first output (i.e.,* y*).*
– $\mathsf{Verify}^h(1^\lambda, \mathsf{vk}, x, y, \pi) \to \{0, 1\}$.

When clear from the context, we omit 1^λ *from the inputs to* Eval *and* Verify*. The VRF is* ***secure*** *if*

– ***Correctness.*** *For all* PPT A*:*

$$\Pr\big[\neg\mathsf{Verify}^h(\mathsf{vk}, x, y, \pi)\big] \le \mathsf{negl}(\lambda),$$

where $h \leftarrow\!\!\$\ \mathcal{H}_\lambda$, $(k, \mathsf{vk}) \leftarrow\!\!\$\ \mathsf{KGen}^h(1^\lambda)$, $x \leftarrow\!\!\$\ \mathsf{A}^{h,\mathsf{Eval}^h(k,\cdot)}(1^\lambda, \mathsf{vk})$, *and* $(y, \pi) \leftarrow\!\!\$\ \mathsf{Eval}^h(k, x)$*. Here the oracle* $\mathsf{Eval}^h(k, \cdot)$ *given to* A *takes as input* x' *and returns* $(y', \pi') \leftarrow\!\!\$\ \mathsf{Eval}^h(k, x')$.

– ***Pseudorandomness.*** $(\mathsf{KGen}, \mathsf{Eval}_1)$ *is a secure PRF with respect to the domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ *and* $\mathcal{H}$.
– ***Verifiability.*** *For all* PPT A*:* *if* $h \leftarrow\!\!\$\ \mathcal{H}_\lambda$ *and* $\big(\mathsf{vk}, x, (y_0, \pi_0), (y_1, \pi_1)\big) \leftarrow\!\!\$\ \mathsf{A}^h(1^\lambda)$ *then*

$$\Pr\big[\ y_0 \ne y_1 \quad\wedge\quad (\forall i \in \{0, 1\} \colon \mathsf{Verify}^h(\mathsf{vk}, x, y_i, \pi_i) = 1)\big] \le \mathsf{negl}(\lambda).$$

– ***Simulatability.*** *There exists a* PPT *simulator* Sim *such that for every* PPT *oracle-aided distinguisher* D*:*

$$\Big|\Pr\big[D^{h,\mathsf{Eval}^h(k,\cdot)}(\mathsf{vk}) = 1\big] - \Pr\big[D^{h,\mathcal{O}_{\mathsf{Sim}}(\cdot)}(\mathsf{vk}) = 1\big]\Big| \le \mathsf{negl}(\lambda) \tag{2}$$

where $h \leftarrow\!\!\$\ \mathcal{H}_\lambda$, $(k, \mathsf{vk}) \leftarrow\!\!\$\ \mathsf{KGen}^h(1^\lambda)$, *and* $\mathcal{O}_{\mathsf{Sim}}(x) := \mathsf{Sim}^h\big(\mathsf{vk}, x, \mathsf{Eval}_1^h(k, x)\big)$ *is* Sim*'s (simulated)* Eval *responses to* D.

For simulatability, the distinguisher D is given access to the random oracle h, and to a second oracle that either returns the output of the "real" $\mathsf{Eval}^h(k, \cdot)$ algorithm (that on input x returns the PRF output y and a real proof π), or returns the simulator's responses to evaluation queries. This therefore means that $\mathsf{Sim}^h\big(\mathsf{vk}, x, y\big)$ needs to output pairs (y, π) that are indistinguishable from real pairs, where $y = \mathsf{Eval}_1^h(k, x)$ is the correct VRF output with key k (where k is unknown to Sim).

2.2 Secure Computation

For our ideal-model definition of exponent VRFs and for our applications, we prove security for the stand-alone definition of secure multiparty computation [14,33] for security with abort (where some honest parties may have output and some may abort) and with no honest majority. In this model, all parties send their inputs to the ideal functionality (computed by a trusted party). The ideal functionality then sends the (ideal-model) adversary the corrupted parties' outputs, and the adversary then instructs the ideal functionality as to which honest parties should receive output. We denote the set of honest parties sent by the ideal adversary to the ideal functionality to receive output by $\mathcal{O}_{\mathsf{out}}$.

Although we prove security in the stand-alone model that guarantees security under sequential composition only, we are really interested in UC security [15]; i.e., security under concurrent general composition. This is achieved by all our protocols since they all have *straight-line simulation* (i.e., no rewinding). As shown in [44], this implies UC security if the protocol is *perfectly secure* or there is *start synchronization* (meaning that all parties have their input before the protocol begins).

Network model. In all of our protocols, we consider security with abort. As such, parties can just wait to receive a message, and "hang" if they do not (in practice, they can just abort if they wait too long, which is also fine). This means that we don't need to assume a synchronous network, as parties proceed to the next round only after receiving all messages from the previous round.

3 eVRFs

In this section, we formally define the concept of an eVRF. We begin by defining a game-based definition for the security of an eVRF. Next we define an eVRF ideal functionality, and prove that a simple protocol using the game-based definition, together with a zero-knowledge proof of knowledge of the private key, securely realizes the ideal functionality. The game-based definition will be used to argue that our eVRF constructions are secure, and the ideal functionality will be used for constructing our applications utilizing an eVRF.

3.1 Game-based Definition

Let $\mathbb{G}$ be a finite cyclic group with generator $G \in \mathbb{G}$. The evaluation algorithm $\mathsf{Eval}(k, x)$ of an eVRF outputs a triple (y, Y, π) such that $Y = y \cdot G$, with the property that $\mathsf{Eval}_1(k, x) := y$ is a pseudorandom function, and $\mathsf{Eval}_2(k, x) := (Y, \pi)$ is a (simulatable) VRF. Stated differently, the output y is pseudorandom, and there exists a (simulatable) proof that Y has been generated correctly from y as $Y \leftarrow y \cdot G$. The formal game-based definition, suited for constructions in the ROM, is given below.

Definition 4. *Let $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ be an ensemble of domains/ranges, where each $\mathcal{Y}_\lambda$ is a finite cyclic group with a specified generator G_λ. An **exponent verifiable random function (eVRF)** with respect to domain/range ensemble $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ and function-family ensemble $\mathcal{H}$, is a triple of oracle-aided* PPT *algorithms called* $(\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$ *such that for every $\lambda \in \mathbb{N}$ and $h \in \mathcal{H}_\lambda$:*

- $\mathsf{KGen}^h(1^\lambda) \to (k, \mathsf{vk})$.
- $\mathsf{Eval}^h(1^\lambda, k, x) \to (y, Y, \pi)$ *with $x \in \mathcal{X}_\lambda$, $y \in \mathbb{Z}_{|\mathcal{Y}_\lambda|}$, and $Y \in \mathcal{Y}_\lambda$. We define two auxiliary algorithms $\mathsf{Eval}_1(1^\lambda, k, x) \to y$ and $\mathsf{Eval}_2(1^\lambda, k, x) \to (Y, \pi)$ that are the same as* Eval*, but only output the first output (i.e., y) or the second and third outputs (i.e., (Y, π)) of* Eval*, respectively.*

– $\mathsf{Verify}^h(1^\lambda, \mathsf{vk}, x, Y, \pi) \to \{0,1\}$.

When clear from the context, we omit 1^λ *from the inputs to* Eval *and* Verify. *An eVRF is* ***secure*** *if*

– ***Consistency.*** *For every* PPT A*:*

$$\Pr\left[y \cdot G_\lambda \neq Y \;:\; \begin{array}{l} h \leftarrow\!\$\, \mathcal{H}_\lambda, \quad (k, \mathsf{vk}) \leftarrow \mathsf{KGen}^h(1^\lambda) \\ x \leftarrow\!\$\, \mathsf{A}^{h, \mathsf{Eval}^h(k,\cdot)}(1^\lambda, \mathsf{vk}), \quad (y, Y, \pi) \leftarrow\!\$\, \mathsf{Eval}^h(k, x) \end{array}\right] \leq \mathsf{negl}(\lambda).$$

– ***Pseudorandomness.*** $(\mathsf{KGen}, \mathsf{Eval}_1)$ *is a secure PRF with respect to the domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathbb{Z}_{|\mathcal{Y}_\lambda|})\}_{\lambda\in\mathbb{N}}$ *and function family ensemble* $\mathcal{H}$.

– ***Simulatable verifiability.*** $(\mathsf{KGen}, \mathsf{Eval}_2, \mathsf{Verify})$ *is a simulatable VRF with respect to the ensemble* $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$ *and function family ensemble* $\mathcal{H}$.

In some cases it will be convenient to define an eVRF where the output of Eval_1 is pseudorandom with respect to a subset $\mathcal{S}_\lambda$ of $\mathbb{Z}_{|\mathcal{Y}_\lambda|}$. The following definition captures the notion of a subset eVRF.

Definition 5. *Using the notation in Definition 4, let* $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$ *be an ensemble of subsets, where* $\mathcal{S}_\lambda \subseteq \mathbb{Z}_{|\mathcal{Y}_\lambda|}$ *for all* $\lambda \in \mathbb{N}$. *We say that an eVRF* $(\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$ *is* ***subset secure*** *with respect to domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)\}_{\lambda\in\mathbb{N}}$, *subset ensemble* $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$, *and function family ensemble* $\mathcal{H}$, *if the eVRF has consistency and simulatable verifiability as in Definition 4, and is pseudorandom in the following sense:* $(\mathsf{KGen}, \mathsf{Eval}_1)$ *is a secure PRF with respect to the domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathcal{S}_\lambda)\}_{\lambda\in\mathbb{N}}$ *and function family ensemble* $\mathcal{H}$.

3.2 Ideal Definition

We now define the ideal functionality for an eVRF and prove that it is implied by the game-based definition, together with a zero-knowledge proof of knowledge of the private key. To simplify the notation we refer to an explicit domain/range $(\mathcal{X}, \mathcal{Y})$ rather than an ensemble.

Definition 6 (eVRF functionality). *Let* $(\mathcal{X}, \mathcal{Y})$ *be as in Definition 4, where* $\mathcal{Y}$ *defines a group* $\mathbb{G}$ *of order* q *with generator* G. *The* ***eVRF ideal functionality*** *for* $(\mathcal{X}, \mathcal{Y})$, *denoted* $\mathcal{F}_{\mathsf{eVRF}}^{\mathcal{X},\mathcal{Y}}$ *or just* $\mathcal{F}_{\mathsf{eVRF}}$ *for short, is defined as follows:*

1. *Upon receiving* $(\mathsf{init}, i, *)$ *from some* P_i.
   (a) *If* P_i *is honest and the input is* (init, i), *then receive a value* sid *from the (ideal) adversary, verify that it's unique and store* (sid, i)
   (b) *If* P_i *is corrupted and the input is* $(\mathsf{init}, i, \mathsf{sid}, f)$ *where* f *is the description of a deterministic polynomial-time computable function* sid *has not been stored, then store* (sid, i, f)
   Send $(\mathsf{init}, i, \mathsf{sid})$ *to all parties*
2. *Upon receiving* $(\mathsf{Eval}, i, \mathsf{sid}, x)$ *from* P_i, *where* $x \in \mathcal{X}$*:*
   (a) *If* (sid, i) *or* (sid, i, f) *is not stored then ignore*
   (b) *If* P_i *is honest:*
      i. *If there does not exist a stored tuple* $(\mathsf{sid}, i, x, y, Y)$ *with* x, *then choose a random* $y \leftarrow\!\$\, \mathbb{Z}_q$, *compute* $Y \leftarrow y \cdot G$ *and store* $(\mathsf{sid}, i, x, y, Y)$
      ii. *Send* $(\mathsf{Eval}, i, \mathsf{sid}, x, y, Y)$ *to party* P_i *and* $(\mathsf{Eval}, i, \mathsf{sid}, x, Y)$ *to all parties*
   (c) *If* P_i *is corrupted, then compute* $Y \leftarrow f(x) \cdot G$ *and send* $(\mathsf{Eval}, i, \mathsf{sid}, x, Y)$ *to all parties*

In order to prove that the game-based definition implies the ideal functionality, we define a protocol that utilizes the game-based definition, and then prove that it securely realizes $\mathcal{F}_{\mathsf{eVRF}}$. Our protocol requires a ZK proof of knowledge for the relation $\mathcal{R}_{EF} := \{(\mathsf{vk}, (k, r)) : \mathsf{KGen}(1^\lambda; r) = (k, \mathsf{vk})\}$. We denote the ideal functionality for this proof by $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$; the functionality receives $(\mathsf{prove}, i, j, \mathsf{vk}, k, r)$ from P_i and sends $(\mathsf{prove}, i, j, \mathsf{vk}, 1)$ to P_j if $(\mathsf{vk}, (k, r)) \in \mathcal{R}_{EF}$. We note that by including this proof of knowledge, it is not possible for a party to copy an eVRF instance from another party (practically, the identity i of the prover is just included in the hash in the proof).

The ideal functionality $\mathcal{F}_{\mathsf{eVRF}}$ needs a sid in order to distinguish different eVRF instances. In the protocol, we use the verification key vk for this purpose and so do not require any additional sid.

The π_{EF} protocol. For any eVRF $EF = (\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$, we define the protocol π_{EF} with parameters $(\mathcal{X}, \mathcal{Y})$ as follows:

Protocol 1 (π_{EF})

- ***Initialize:***
  1. *Message 1 from P_i: Party P_i with input* (init, i),
     (a) $(k, \mathsf{vk}) \leftarrow \mathsf{KGen}(1^\lambda, \mathcal{X}, \mathcal{Y})$
     (b) *Send* $(\mathsf{init}, i, \mathsf{vk})$ *to all parties* $P_1, \ldots, P_n$
     (c) *Send* $(\mathsf{prove}, i, j, \mathsf{vk}, k, r)$ *to* $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$ *for all* $j \in [n]$
  2. *Message 2: Each party P_j upon receiving* $(\mathsf{init}, i, \mathsf{vk})$ *from party P_i*
     (a) *If* $(\mathsf{prove}, i, j, \mathsf{vk}, 1)$ *is received from* $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$ *then proceed; else ignore*
     (b) *Send* $(\mathsf{init}, i, \mathsf{vk})$ *to all parties* $P_1, \ldots, P_n$
  3. *Output: Upon receiving* $(\mathsf{init}, i, \mathsf{vk})$ *from all parties,*
     (a) *Party P_i: Output* $(\mathsf{init}, i, \mathsf{vk}, k)$
     (b) *All other parties P_j ($j \neq i$): if the same message* $(\mathsf{init}, i, \mathsf{vk})$ *is received from all parties then store* $(\mathsf{init}, i, \mathsf{vk})$*; else ignore*
- ***Evaluate:***
  1. *Message from P_i: Party P_i with input* $(\mathsf{Eval}, i, \mathsf{vk}, x)$,
     (a) $(y, Y, \pi) \leftarrow \mathsf{Eval}(k, x)$
     (b) *Send* $(\mathsf{Eval}, i, \mathsf{vk}, x, Y, \pi)$ *to all parties* $P_1, \ldots, P_n$
  2. *Output:*
     (a) *Party P_i: Output* $(\mathsf{Eval}, i, \mathsf{vk}, x, y, Y)$
     (b) *All other parties P_j ($j \neq i$): Upon receiving* $(\mathsf{Eval}, i, \mathsf{vk}, x, Y, \pi)$ *from P_i*
        i. *Verify that* $(\mathsf{init}, i, \mathsf{vk})$ *has been stored (i.e.,* vk *is associated with P_i)*
        ii. *If* $\mathsf{Verify}(\mathsf{vk}, x, Y, \pi) = 0$ *then ignore*
        iii. *Else, output* $(\mathsf{Eval}, i, \mathsf{vk}, x, Y)$

We stress that in the *two-party case*, the initialize phase consists only of P_i sending $(\mathsf{init}, i, \mathsf{vk})$ to P_j, who stores it (i.e., there is no need for a second message in order to obtain consensus). We now prove that π_{EF} securely realizes $\mathcal{F}_{\mathsf{eVRF}}$.

Theorem 2. *Let $EF = (\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$ be an exponent verifiable random function with respect to $(\mathcal{X}, \mathcal{Y})$. Then π_{EF} securely realizes with abort $\mathcal{F}_{\mathsf{eVRF}}$ with respect to $(\mathcal{X}, \mathcal{Y})$ in the $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$-hybrid model, in the presence of a static malicious adversary corrupting any number of parties.*

Proof. Let $P_1, \ldots, P_n$ be the parties, let $\mathcal{I} \subset [n]$ be the set of corrupted parties, and let $\mathcal{A}$ be a real-world adversary running protocol π_{EF}. We construct an ideal-model adversary/simulator $\mathcal{S}$ as follows:

1. **Initialize:**
   (a) Upon receiving (init, i) from $\mathcal{F}_{\mathsf{eVRF}}$ for an honest P_i, the simulator $\mathcal{S}$ runs $\mathsf{KGen}(1^\lambda, \mathcal{X}, \mathcal{Y})$ to obtain (k, vk), and sends $\mathsf{sid} = \mathsf{vk}$ to $\mathcal{F}_{\mathsf{eVRF}}$. Next, $\mathcal{S}$ simulates honest party P_i sending $(\mathsf{init}, i, \mathsf{vk})$ to all corrupted parties, and simulates $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$ sending $(\mathsf{prove}, i, j, \mathsf{vk}, 1)$ to all corrupted parties. $\mathcal{S}$ then simulates message 2 of the initialization protocol and defines $\mathcal{O}_{\mathsf{out}}$ to be the set of honest parties who would not abort (i.e., who all received the same correct messages from all corrupted parties). $\mathcal{S}$ sends $\mathcal{O}_{\mathsf{out}}$ to $\mathcal{F}_{\mathsf{eVRF}}$, instructing it to send output to the honest parties in $\mathcal{O}_{\mathsf{out}}$.
   (b) Upon a corrupted party receiving (init, i) for input, $\mathcal{S}$ obtains the message $(\mathsf{init}, i, \mathsf{vk})$ sent by $\mathcal{A}$ to all honest parties and the messages $(\mathsf{prove}, i, j, \mathsf{vk}, k, r)$ sent to $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$ for all $j \notin \mathcal{I}$. Then, $\mathcal{S}$ simulates the honest parties actions exactly according to the protocol, and defines $\mathcal{O}_{\mathsf{out}}$ to be the set of honest parties who would not abort (i.e., who all received the same correct messages from all corrupted parties, and who received correct proofs). If $\mathcal{O}_{\mathsf{out}}$ is not empty, then $\mathcal{S}$ defines $f(x) = \mathsf{Eval}_1(k, x)$ where k is from the message sent by $\mathcal{A}$ to $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$, and sends $(\mathsf{init}, \mathsf{vk}, i, f)$ to $\mathcal{F}_{\mathsf{eVRF}}$ together with $\mathcal{O}_{\mathsf{out}}$, instructing it to send output to the honest parties in $\mathcal{O}_{\mathsf{out}}$.
2. **Evaluate:**
   (a) Upon receiving $(\mathsf{Eval}, i, \mathsf{vk}, x, Y)$ from $\mathcal{F}_{\mathsf{eVRF}}$ for an honest P_i (vk is the sid as generated during initialize), the simulator $\mathcal{S}$ runs $\mathsf{Sim}(\mathsf{vk}, x, Y)$ to obtain π and simulates the honest P_i sending $(\mathsf{Eval}, i, \mathsf{vk}, x, Y, \pi)$ to all corrupted parties for the associated vk.
   (b) Upon a corrupted party receiving $(\mathsf{Eval}, i, \mathsf{vk}, x)$ for input (as above, vk is the sid), $\mathcal{S}$ sends $(\mathsf{Eval}, i, \mathsf{vk}, x)$ to $\mathcal{F}_{\mathsf{eVRF}}$ and obtains the messages $\left\{(\mathsf{Eval}, i, \mathsf{vk}, x, Y^j, \pi^j)\right\}_{j \notin \mathcal{I}}$ sent by $\mathcal{A}$ to the honest parties P_j for $j \notin \mathcal{I}$; observe that nothing forces $\mathcal{A}$ to send the same message to all parties and so we denote by (Y^j, π^j) denote the values received by honest P_j.[5] For each $j \notin \mathcal{I}$, $\mathcal{S}$ verifies that $\mathsf{Verify}(\mathsf{vk}, x, Y^j, \pi^j) = 1$. If no, it ignores the message. Else, it adds P_j to $\mathcal{O}_{\mathsf{out}}$, and sends $\mathcal{O}_{\mathsf{out}}$ to $\mathcal{F}_{\mathsf{eVRF}}$.

We separately consider the case that P_i is honest and that P_i is corrupted.

Let P_i be honest. Then, the simulation in the initalize phase is perfect (in the $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$-hybrid model) since $\mathcal{S}$ generates (k, vk) like an honest party and perfectly simulates the message from $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$ to the corrupted parties. Regarding the evaluation phase, there are two differences between the simulation and a real execution: **(a)** the value Y is truly random in the ideal execution (and in particular is independent of (k, vk) generated by $\mathcal{S}$ in the initialization phase) and equals $\mathsf{Eval}_1(k, x) \cdot G$ in the real execution, and **(b)** the proof π is simulated in the ideal execution and is output from $\mathsf{Eval}(k, x)$ in the real execution. We prove indistinguishability in three hybrid steps:

1. *Hybrid 1:* In this hybrid execution, we modify $\mathcal{F}_{\mathsf{eVRF}}$ so that upon receiving (init, i) from an honest P_i it computes $\mathsf{KGen}(1^\lambda, \mathcal{X}, \mathcal{Y})$ to obtain (k, vk), and sends $(\mathsf{init}, i, \mathsf{vk})$ to $\mathcal{S}$ setting $\mathsf{sid} = \mathsf{vk}$, instead of receiving sid from $\mathcal{S}$. Furthermore, $\mathcal{S}$ uses vk as it received from $\mathcal{F}_{\mathsf{eVRF}}$ instead of generating it by itself. Everything else remains the same, and in particular the evaluation is random.

   It is clear that the output distributions of this hybrid and the ideal execution are identical. The only difference is who chooses vk, which makes no difference.

[5] We consider different Y^j, π^j values but not different x values. This is because a different x is considered a different evaluation and is treated separately.

2. *Hybrid 2:* In this hybrid execution, we further modify $\mathcal{F}_{\mathsf{eVRF}}$ so that upon receiving $(\mathsf{Eval}, i, \mathsf{sid}, x)$ from an honest P_i, instead of choosing a random y, it computes $y \leftarrow \mathsf{Eval}_1(k, x)$ with the k generated after receiving (init, i). (The computation of $Y \leftarrow y \cdot G$ is unchanged.) Everything else remains the same as the first hybrid, including $\mathcal{S}$.
   The only difference between the first and second hybrid execution is how y is generated. In order to show that this is indistinguishable, we rely on the fact that Eval_1 is a pseudorandom function. Observe that it is possible to simulate both hybrid executions without ever receiving k (in particular, because the proof π is simulated). Thus, if it is possible to distinguish between hybrid 2 and hybrid 1, then it is possible to distinguish Eval_1 from random.
3. *Hybrid 3:* In this hybrid execution, we modify $\mathcal{F}_{\mathsf{eVRF}}$ so that upon receiving $(\mathsf{Eval}, i, \mathsf{sid}, x)$, instead of just sending $(\mathsf{Eval}, i, \mathsf{sid}, x, Y)$ to $\mathcal{S}$, it also sends π where $\mathsf{Eval}(k, x) = (y, Y, \pi)$. The simulator $\mathcal{S}$ is the same as for hybrid 2 except that instead of computing $\mathsf{Sim}(\mathsf{vk}, x, Y)$ to obtain π, it sends the proof π received from $\mathcal{F}_{\mathsf{eVRF}}$. Indistinguishability of hybrid 2 and hybrid 3 follows from Definition 4 and in particular that the output of the simulator is indistinguishable from $\mathsf{Eval}_2(k, x)$. We stress that in order to simulate these hybrids, the distinguisher needs to generate an ideal execution where all proofs are either real or simulated, where the inputs x are chosen dynamically by the adversary. This is achieved using the oracles given to the distinguisher in Definition 4.

Finally, we observe that hybrid 3 is identical to a real execution. The only difference is that $\mathcal{F}_{\mathsf{eVRF}}$ computes the honest parties' messages according to protocol π_{EF} instead of the parties themselves, but this does not affect the output distribution.

We now consider the case that P_i is corrupted. In this case, the simulator $\mathcal{S}$ perfectly detects who will receive output and who not in terms of receiving consistent messages during initialization and a valid proof via $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$. Furthermore, $\mathcal{S}$ perfectly detects who will receive output in the evaluation phase, based on the proof π being valid. However, in the ideal execution, the output received by honest parties during evaluation is always $f(x) = \mathsf{Eval}_1(k, x)$. Thus, there can only be a difference between the ideal and real executions if $\mathcal{A}$ sends a value $Y' \neq \mathsf{Eval}_1(k, x) \cdot G$ together with an *accepting proof* π. This would contradict the verifiability property of $(\mathsf{KGen}, \mathsf{Eval}_2, \mathsf{Verify})$ as a VRF. In particular, it is always possible to generate an accepting proof for $Y = \mathsf{Eval}_1(k, x) \cdot G$. If $\mathcal{A}$ generates an accepting proof also for some $Y' \neq \mathsf{Eval}_1(k, x) \cdot G$, then we could construct an adversary who outputs two distinct Y, Y' with respective accepting proofs π, π'. This completes the proof. □

UC security. The simulator in the proof of Theorem 2 is straight line. Therefore, by [44], the protocol is UC secure assuming start synchronization (all parties have their input before the protocol starts). However, in both the initialize and evaluate phases, the only party with input is P_i and therefore start synchronization holds always. We therefore have:

Corollary 1. *Let $EF = (\mathsf{KGen}, \mathsf{Eval}, \mathsf{Verify})$ be an exponent verifiable random function by Definition 4. Then π_{EF} UC realizes $\mathcal{F}_{\mathsf{eVRF}}$ with abort in the $\mathcal{F}_{\mathsf{zk}}^{\mathcal{R}_{EF}}$-hybrid model, in the presence of a static malicious adversary corrupting any number of parties.*

4 Applications

In this section, we present the many applications for eVRFs discussed in the introduction. All the protocols we present are "fully simulatable" meaning that they securely realize the *plain algorithm*

functionality (e.g., Schnorr signing), as opposed to some modified functionality. All of our protocols are concretely efficient, and are secure under standard assumptions in the random-oracle model.

4.1 One-Round Simulatable Distributed Key Generation

This protocol works by defining each party's key share to be the result of a pseudorandom function applied to a unique nonce. In order to ensure that each party uses its committed pseudorandom function, we use the $\mathcal{F}_{\mathsf{eVRF}}$ functionality to derive each party's key share. Intuitively, this enables us to generate a key in a single round since the only message the parties need to send is their single eVRF value. This suffices since each party can simply sum the public key-share values (we call $K_i = k_i \cdot G$ a party's $\mathsf{public\ share}$) to obtain the final public key. The crucial difference between this key generation and standard key generation protocols is that since each party is already committed to its value via the $\mathcal{F}_{\mathsf{eVRF}}$ (after running a single initialization step), the corrupted parties cannot bias the output key. In particular, if a key was generated by each party simply choosing a random k_i and sending $K_i = k_i \cdot G$ to all other parties, then a single corrupted party P_1 can completely determine the key by obtaining all the honest parties shares $K_2, \ldots, K_n$ and then setting $K_1 = K - \sum_{i=2}^{n} K_i$, where P_1 has chosen $K = k \cdot G$ where k is known to it. This trivial attack can be prevented by having each party add a zero-knowledge proof of knowledge of the discrete log of its K_i; this would prevent P_1 from carrying out this attack since it cannot know the discrete log of K_1. However, the protocol is still not simulatable since P_1 could bias the output. In particular, if P_1 wanted the public key K to have a certain property that holds in $1/1000$ keys then it can receive $K_2, \ldots, K_n$ and then repeatedly choose random k_1 and compute $K_1 = k_1 \cdot G$ and $K = \sum_{i=1}^{n} K_i$ until K has the required property. All of this isn't possible with our protocol since the corrupted parties are committed to the PRF output via $\mathcal{F}_{\mathsf{eVRF}}$.

We remark that generating a new key requires a unique nonce. In particular, if the same nonce is used twice then the same key will be output. This holds in both the ideal and real models, and therefore the security of the protocol does not rely on the nonce necessarily being unique. However, using the protocol to generate a new key does require ensuring a unique nonce (which could be a timestamp, a counter, etc.).

The additive DKG functionality $\mathcal{F}_{\mathsf{dkg}}$: Let $\mathbb{G}$ be a group of order q with generator G. The distributed key-generation functionality $\mathcal{F}_{\mathsf{dkg}}$ for $\mathbb{G}$ running with parties $P_1, \ldots, P_n$ and corrupted parties $\mathcal{I} \subseteq [n]$ is defined as follows:

1. Wait to receive $(\mathsf{gen}, \mathsf{nonce})$ from all honest parties and $(\mathsf{gen}, \mathsf{nonce}, k_i)$ from all corrupted parties P_i with $i \in \mathcal{I}$
2. If $(\mathsf{gen}, \mathsf{nonce}, k_1, \ldots, k_n)$ has already been stored
   - Retrieve $k_1, \ldots, k_n$

   Else
   - Choose random $k_j \leftarrow\!\!\$\, \mathbb{Z}_q$ for every $j \in [n] \setminus \mathcal{I}$
   - Store $(\mathsf{gen}, \mathsf{nonce}, k_1, \ldots, k_n)$
3. For $i = 1, \ldots, n$, compute $K_i = k_i \cdot G$
4. Compute $K = \sum_{i=1}^{n} K_i$
5. For $i = 1, \ldots, n$, send $(\mathsf{gen}, \mathsf{nonce}, k_i, K_1, \ldots, K_n, K)$ to P_i

We are now ready to present the protocol. Let $\mathcal{X} = \{0,1\}^\lambda$ and let $\mathcal{Y} = \mathbb{G}$ as desired for $\mathcal{F}_{\mathsf{dkg}}$.

Protocol 3 (Π_{dkg} for additive DKG)

– ***Initialize:***
  1. *In parallel, each P_i sends* (init, i) *to* $\mathcal{F}_{\mathsf{eVRF}}$ *(for* $\mathcal{X}, \mathcal{Y}$*)*
  2. *Wait to receive* $(\mathsf{init}, j, \mathsf{sid}_j)$ *from* $\mathcal{F}_{\mathsf{eVRF}}$ *for all* $j \in [n]$
– ***Generate (one round):*** *upon input* $(\mathsf{gen}, \mathsf{nonce})$*, each party* P_i
  - Message:
    1. *Send* $(\mathsf{Eval}, i, \mathsf{sid}_i, \mathsf{nonce})$ *to* $\mathcal{F}_{\mathsf{eVRF}}$ *and receive back* $(\mathsf{Eval}, i, \mathsf{sid}_i, \mathsf{nonce}, k_i, K_i)$
  - Output:
    1. *Wait to receive* $(\mathsf{Eval}, j, \mathsf{sid}_j, \mathsf{nonce}, K_j)$ *from* $\mathcal{F}_{\mathsf{eVRF}}$ *for all parties* P_j *(all with the correct* nonce*)*
    2. *Compute* $K = \sum_{i=1}^{n} K_i$
    3. *Output* $(\mathsf{gen}, \mathsf{nonce}, k_i, K_1, \ldots, K_n, K)$

The rationale behind the security of the protocol has already been described above and so we proceed directly to the proof of security.

Theorem 4. *Protocol 3 securely realizes* $\mathcal{F}_{\mathsf{dkg}}$ *with perfect security with abort, in the presence of a static malicious adversary corrupting up to* n *parties. Furthermore, after a single two-round initialization phase, each generation consists of a single round.*

Proof. If all n parties or no parties are corrupted, then the statement is trivial (if no parties are corrupted, then since the ideal functionality and honest parties communicate over ideal private channels, nothing is revealed). Let $\mathcal{A}$ be the adversary and let $\mathcal{I}$ be the set of corrupted parties with $0 < |\mathcal{I}| < n$. We construct a simulator $\mathcal{S}$ with "internal communication" to $\mathcal{A}$ and "external communication" to $\mathcal{F}_{\mathsf{dkg}}$, as follows:

– *Initialize:*
  1. Simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending $\{(\mathsf{init}, j)\}_{j \notin \mathcal{I}}$ to $\mathcal{A}$, and receive back $\{\mathsf{sid}_j\}_{j \notin \mathcal{I}}$ as $\mathcal{A}$ would send to $\mathcal{F}_{\mathsf{eVRF}}$
  2. Simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{init}, j, \mathsf{sid}_j)$ to all parties, for all $j \notin \mathcal{I}$
  3. For all $i \in \mathcal{I}$, internally receive $(\mathsf{init}, i, \mathsf{sid}_i, f_i)$ from $\mathcal{A}$ as sent to $\mathcal{F}_{\mathsf{eVRF}}$ from P_i, and simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending back $(\mathsf{init}, i, \mathsf{sid}_i)$ to all parties
– *Generate:* to simulate an execution of generate for nonce,
  1. For every $i \in \mathcal{I}$, compute $k_i = f_i(\mathsf{nonce})$, where f_i is as received for P_i in the initialize phase
  2. Send $(\mathsf{gen}, \mathsf{nonce}, k_i)$ externally to $\mathcal{F}_{\mathsf{dkg}}$ for every $i \in \mathcal{I}$, and receive back a set of tuples $\{(\mathsf{gen}, \mathsf{nonce}, k_i, K_1, \ldots, K_n, K)\}_{i \in \mathcal{I}}$
  3. Simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{Eval}, j, \mathsf{sid}_j, \mathsf{nonce}, K_j)$ to all corrupted parties, for every $j \in [n] \setminus \mathcal{I}$
  4. Define the set $\mathcal{O}_{\mathsf{out}}$ of honest parties to receive output by the set of parties for which $\mathcal{A}$ instructs $\mathcal{F}_{\mathsf{eVRF}}$ to provide output from all $i \in \mathcal{I}$
  5. Send $\mathcal{F}_{\mathsf{dkg}}$ the set $\mathcal{O}_{\mathsf{out}}$
  6. Output whatever $\mathcal{A}$ outputs

The initialization phase in the simulation is identical to the real execution in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model. Regarding the generate phase, in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model, the view of the adversary in the ideal execution is also identical to its view in a real execution. This is because in the $\mathcal{F}_{\mathsf{eVRF}}$ model all that it sees is the K_j values for the honest parties. Furthermore, these values are uniformly distributed in the real execution by the definition of $\mathcal{F}_{\mathsf{eVRF}}$, and uniformly distributed in the ideal

execution by the definition of $\mathcal{F}_{\mathsf{dkg}}$. Finally, the k_i values of the corrupted parties is the same, since $\mathcal{S}$ sends $\mathcal{F}_{\mathsf{dkg}}$ the same values that $\mathcal{A}$ is committed to by the definition of $\mathcal{F}_{\mathsf{eVRF}}$.

The round complexity statement in the theorem follows by observation that each $\mathcal{F}_{\mathsf{eVRF}}$ initialization operation is two rounds and each evaluation operation is just a single round (as described in Protocol 1, and these can all be sent in parallel. This completes the proof. □

4.2 One-Round Simulatable Threshold Distributed Key Generation

The protocol in Section 4.1 works for a set of n parties who all participate and wish to generate a key that is additively distributed amongst themselves. This can easily be extended to a threshold setting (and even to a more general access structure of a tree of AND, OR, and threshold gates) by simply having each party in a quorum generate a VSS sharing of its key share defined by the eVRF. This would work but would not be a single round only since a consensus round would be needed to ensure that all parties receive the same sharing (a simple echo-broadcast suffices, as shown in [21]). Fortunately, we can use the eVRF to achieve this as well, and have each party define *all the coefficients* in its polynomial for Feldman VSS [27] via the eVRF. Since each party is already committed to its values and therefore its polynomial via the eVRF, and since all parties can verify that the eVRF output values sent are correct, there is no need for an additional consensus round. Indeed, no party can send a value that isn't correct. We describe the protocol for a simple threshold only; the extension to a tree of AND, OR and threshold gates is immediate.

We describe the protocol with a quorum of $t+1$ online parties who generate the key for all n parties. We stress that there is no security benefit in having all n parties generate the key, since any quorum of $t+1$ parties will anyway have the entire key.

The quorum-specific threshold DKG functionality $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$: Let $\mathbb{G}$ be a group of order q with generator G, let $\alpha_1, \ldots, \alpha_n$ be fixed distinct elements in $\mathbb{Z}_q$, and let $t < n$. The distributed key-generation functionality $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$ for $\mathbb{G}$ running with parties $P_1, \ldots, P_n$, a quorum $\mathcal{Q} \subseteq [n]$ of $t+1$ online parties, and corrupted parties $\mathcal{I} \subseteq [n]$ with $|\mathcal{I}| \leq t$ is defined as follows:

1. Wait to receive $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q})$ from all $t+1$ parties
2. If $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q}, p(x))$ has already been stored
   - Retrieve $p(x)$
   - Store $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q}, p(x))$

   Else
   - Choose a random degree-t polynomial $p(x) \leftarrow\!\!\$\ \mathbb{F}_q[x]$
   - Store $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q}, p(x))$
3. For $j = 1, \ldots, n$, compute $k_j = p(\alpha_j)$
4. Let $a_0, \ldots, a_t$ be the coefficients of p; i.e., $p(x) = \sum_{i=0}^{t} a_i \cdot x^i$
5. For $i = 0, \ldots, t$, compute $A_i = a_i \cdot G$
6. Compute $k = p(0)$ and $K = k \cdot G$
7. For $i = 1, \ldots, n$, send $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q}, k_i, A_0, \ldots, A_t, K)$ to P_i

Observe that in the case of additive DKG, the functionality allowed each corrupted party to choose its own share (and the honest parties' shares were randomly chosen by the functionality). In contrast, here the functionality chooses all of the shares. The reason for this difference is that here each party's share is the sum of the shares received from all parties. Since each party is a priori committed to its values after the initialization phase, this means that no party can influence even its own share and so there is no need to give this extra power to the adversary.

Protocol 5 ($\Pi^{n,t}_{\mathsf{dkg}\mathcal{Q}}$ for threshold DKG)

- ***Initialize (all n parties):***
  1. *In parallel, each P_i sends* (init, i) *to* $\mathcal{F}_{\mathsf{eVRF}}$ *(for $\mathcal{X}, \mathcal{Y}$)*
  2. *Wait to receive* $(\mathsf{init}, j, \mathsf{sid}_j)$ *from* $\mathcal{F}_{\mathsf{eVRF}}$ *for all* $j \in [n]$
- ***Generate(a quorum $\mathcal{Q}$ of $t+1$ parties):*** *upon input* $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q})$ *with a* nonce *and with* $\mathcal{Q} \subseteq [n]$ *of size* $t+1$*, each party P_i with $i \in \mathcal{Q}$*
  1. Message (all parties in $\mathcal{Q}$): *Each party P_i with $i \in \mathcal{Q}$,*
     (a) *For* $\ell = 0, \ldots, t$*, send* $(\mathsf{Eval}, i, \mathsf{sid}_i, \mathsf{nonce} \| \mathcal{Q} \| \ell)$ *to* $\mathcal{F}_{\mathsf{eVRF}}$ *and receive back a six tuple* $(\mathsf{Eval}, i, \mathsf{sid}_i, \mathsf{nonce} \| \mathcal{Q} \| \ell, a_i^\ell, A_i^\ell)$*. Note that the nonce used for the evaluation includes the identities of the t participating parties.*
     (b) *Let* $p_i(x) = \sum_{\ell=0}^{t} a_i^\ell \cdot x^\ell$
     (c) *Compute* $k_{i \to j} = p_i(\alpha_j)$ *for* $j = 1, \ldots, n$
     (d) *Send* $k_{i \to j}$ *to* P_j *for* $j = 1, \ldots, n$
  2. Output (all n parties):
     (a) *Wait to receive* $\{(\mathsf{Eval}, j, \mathsf{sid}_j, \mathsf{nonce} \| \mathcal{Q} \| \ell, A_j^\ell)\}_{\ell=0}^{t}$ *from* $\mathcal{F}_{\mathsf{eVRF}}$ *for all parties P_j with $j \in \mathcal{Q}$ (all with the correct* nonce *and $\mathcal{Q}$)*
     (b) *Wait to receive* $k_{j \to i}$ *for all* $j \in \mathcal{Q}$
     (c) *Verify that* $k_{j \to i} \cdot G = \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_j^\ell$*, for all $j \in \mathcal{Q}$. Abort if any equality does not hold.*
     (d) *Compute* $k_i = \sum_{j \in \mathcal{Q}} k_{j \to i}$
     (e) *Compute* $A_\ell = \sum_{j \in \mathcal{Q}} A_j^\ell$ *for* $\ell = 0, \ldots, t$
     (f) *Compute* $K = A_0$
     (g) *Output* $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q}, k_i, A_0, \ldots, A_t, K)$

Theorem 6. *Protocol 5 securely realizes $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$ with perfect security with abort, in the presence of a static malicious adversary corrupting up to t parties. Furthermore, after a single two-round initialization phase, each generation consists of a single round only.*

Proof. The idea behind the proof is that the simulator can generate all of the corrupted parties' values itself (using f obtained in the initialization phase of $\mathcal{F}_{\mathsf{eVRF}}$). Then, for all but one honest party, the simulator chooses their values at random. Finally, for a single specified honest party, its polynomial (in the exponent; i.e., it's A_ℓ^j values) are computed by subtracting all the other parties' polynomials from the polynomial received from $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$. This ensures that all the polynomials of the parties add up to the random polynomial sent by $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$. This is identical to a real execution in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model since only "public values" are revealed.

We now proceed with the proof. If no parties are corrupted, then the statement is trivial. Let $\mathcal{A}$ be the adversary and let $\mathcal{I}$ be the set of corrupted parties with $0 < |\mathcal{I}| \leq t$. We construct a simulator $\mathcal{S}$ with "internal communication" to $\mathcal{A}$ and "external communication" to $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$, as follows:

- *Initialize (all n parties):*[6]
  1. Simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending $\{(\mathsf{init}, j)\}_{j \notin \mathcal{I}}$ to $\mathcal{A}$, and receive back $\{\mathsf{sid}_j\}_{j \notin \mathcal{I}}$ as $\mathcal{A}$ would send to $\mathcal{F}_{\mathsf{eVRF}}$
  2. Simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{init}, j, \mathsf{sid}_j)$ to all parties, for all $j \notin \mathcal{I}$
  3. For all $i \in \mathcal{I}$, internally receive $(\mathsf{init}, i, \mathsf{sid}_i, f_i)$ from $\mathcal{A}$ as sent to $\mathcal{F}_{\mathsf{eVRF}}$ from P_i, and simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending back $(\mathsf{init}, i, \mathsf{sid}_i)$ to all parties

[6] The simulation of this phase is exactly as in the proof of Theorem 4.

– *Generate (a quorum $\mathcal{Q}$ of $t+1$ parties):* to simulate an execution of generate for nonce with quorum $\mathcal{Q}$,
  1. Send $(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q})$ to $\mathcal{F}^{n,t}_{\mathsf{dkg}_\mathcal{Q}}$ for every $i \in \mathcal{I} \cap \mathcal{Q}$, and receive back $\{(\mathsf{gen}, \mathsf{nonce}, \mathcal{Q}, k_i, A_0, \ldots, A_t, K)\}_{i \in \mathcal{I}}$
  2. For every $i \in \mathcal{I} \cap \mathcal{Q}$, compute $a_i^\ell = f_i(\mathsf{nonce} \| \mathcal{Q} \| \ell)$ for $\ell = 0, \ldots, t$, where f_i is as received for P_i in the initialize phase, and set $p_i(x) = \sum_{\ell=0}^{t} a_i^\ell \cdot x^\ell$
  3. For every $i \in \mathcal{I} \cap \mathcal{Q}$, compute $A_i^\ell = a_i^\ell \cdot G$
  4. For every $i \in \mathcal{I} \cap \mathcal{Q}$ and every $j \in \mathcal{Q} \setminus \mathcal{I}$, compute $k_{i \to j} = p_i(\alpha_j)$
  5. Let $j' \in \mathcal{Q} \setminus \mathcal{I}$ be a specific honest party (since $|\mathcal{I}| \leq t$ there exists such a party)
  6. For all $j \in \mathcal{Q} \setminus \mathcal{I}$ with $j \neq j'$, choose a random $p_j(x) = \sum_{\ell=0}^{t} a_j^\ell \cdot x^\ell$ and compute $A_j^\ell = a_j^\ell \cdot G$ for $\ell = 0, \ldots, t$
  7. For all $j \in \mathcal{Q} \setminus \mathcal{I}$ with $j \neq j'$, and for all $i \in \mathcal{I}$, compute $k_{j \to i} = p_j(\alpha_i)$
  8. For every $i \in \mathcal{I}$, compute $k_{j' \to i} = k_i - \sum_{j \in \mathcal{Q} \setminus \{j'\}} k_{j \to i}$, where k_i is as received from $\mathcal{F}^{n,t}_{\mathsf{dkg}}$ (this ensures that for every corrupted party P_i it holds that $\sum_{j \in \mathcal{Q}} k_{j \to i} = k_i$)
  9. For $\ell = 0, \ldots, t$, compute $A_{j'}^\ell = A_\ell - \sum_{j \in \mathcal{Q} \setminus \{j'\}} A_\ell^j$, where $A_0, \ldots, A_t$ are as received from $\mathcal{F}^{n,t}_{\mathsf{dkg}}$
  10. Simulate $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{Eval}, j, \mathsf{sid}_j, \mathsf{nonce} \| \mathcal{Q} \| \ell, A_j^\ell)$ to all corrupted parties, for every $j \in \mathcal{Q} \setminus \mathcal{I}$ and $\ell = 0, \ldots, t$
  11. For each $j \in \mathcal{Q} \setminus \mathcal{I}$ and $i \in \mathcal{I}$, simulate honest P_j sending $k_{j \to i}$ as computed above to corrupted P_i
  12. Define the set $\mathcal{O}_{\mathsf{out}}$ of honest parties to receive output by the set of parties for which $\mathcal{A}$ instructs $\mathcal{F}_{\mathsf{eVRF}}$ to provide output from all $i \in \mathcal{I}$, and for which all $k_{i \to j}$ values are sent from corrupted parties to honest parties, and they are all valid ($\mathcal{S}$ can check validity as in the protocol)
  13. Send $\mathcal{F}^{n,t}_{\mathsf{dkg}_\mathcal{Q}}$ the set $\mathcal{O}_{\mathsf{out}}$
  14. Output whatever $\mathcal{A}$ outputs

First observe that **(a)** for every $i \in \mathcal{I}$ it holds that $\sum_{j \in \mathcal{Q}} k_{j \to i} = k_i$, and **(b)** for every $j \in \mathcal{Q} \setminus \mathcal{I}$ and $i \in \mathcal{I}$ it holds that $k_{j \to i} \cdot G = \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_j^\ell$. The former follows immediately from how $k_{j' \to i}$ is chosen. Regarding the latter, this also holds trivially for all $j \neq j'$. Regarding j', observe that $k_{j' \to i}$ is defined by $k_i - \sum_{j \in \mathcal{I} \setminus \{j'\}} k_{j \to i}$ and each $A_{j'}^\ell$ is defined by $A_{j'}^\ell = A_\ell - \sum_{j \in \mathcal{Q} \setminus \{j'\}} A_\ell^j$. Now, for all $k_{j \to i}$ with $j \neq j'$ we have that $k_{j \to i} \cdot G = \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_j^\ell$ and by the $\mathcal{F}^{n,t}_{\mathsf{dkg}}$ functionality definition we have that $k_i \cdot G = \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_\ell$. It therefore follows that

$$
\begin{aligned}
k_{j' \to i} \cdot G &= \left(k_i - \sum_{j \in \mathcal{Q} \setminus \{j'\}} k_{j \to i} \right) \cdot G = \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_\ell - \sum_{j \in \mathcal{Q} \setminus \{j'\}} \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_j^\ell \\
&= \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot \left(A_\ell - \sum_{j \in \mathcal{Q} \setminus \{j'\}} A_j^\ell \right) = \sum_{\ell=0}^{t} (\alpha_i)^\ell \cdot A_{j'}^\ell,
\end{aligned}
$$

as required. Furthermore, since $\mathcal{F}^{n,t}_{\mathsf{dkg}_\mathcal{Q}}$ chooses the polynomial to be random, and so do the honest parties, the distribution over these values is identical in the real and ideal execution. Finally, the simulation of the initialization phase is computationally indistinguishable from a real execution (as shown in Theorem 4) with the only difference being if there is a collision in the sid, and the simulation of the generate phase yields a distribution that is identical to the protocol (in the

$\mathcal{F}_{\mathsf{eVRF}}$ model) since the only thing that the corrupted parties sees are the A_ℓ^j values, and they are committed to their A_i^j and $k_{i\to j}$ values by $\mathcal{F}_{\mathsf{eVRF}}$.

The round complexity statement in the theorem follows by observation that each $\mathcal{F}_{\mathsf{eVRF}}$ initialization operation is two rounds and each evaluation operation is just a single round (as described in Protocol 1, and these can all be sent in parallel. This completes the proof. □

On knowing the set of generating parties $\mathcal{Q}$: Protocol $\Pi_{\mathsf{gen}}^{n,t}$ assumes that all parties know (and agree upon) the set of $t+1$ participants $\mathcal{Q}$ ahead of time. This is needed to ensure that an independent key is generated for each nonce and subset, which is needed since each party has a different eVRF, and so different subsets would generate different keys, even for the same nonce. The set $\mathcal{Q}$ is included to therefore ensure complete independence of keys generated with different subsets. (If we only include the nonce, then it is possible that two different keys will be generated for which the adversary knows the "difference" between them; e.g., consider a case of t honest parties running the execution twice on the same nonce, each time with a different corrupted party who is the $(t+1)$th party.) This limitation can be removed by simply having *all* parties participate in key generation, or by having a *fixed* set of parties who generate keys, or by adding an additional round to agree on the set of parties.

One-round threshold DKG without knowing $\mathcal{Q}$ ahead of time: As we have discussed, the DKG of Section 4.2 requires the parties to know the set $\mathcal{Q}$ ahead of time. In practice, this is not always a given, and the desired functionality is that after the first $t+1$ parties respond, the key is generated.[7] This can be achieved by using a threshold eVRF, which is an eVRF for which any authorized quorum of parties $\mathcal{Q}$ can compute the output Y and receive a sharing of y (where $Y = y \cdot G$). This solves the problem of knowing $\mathcal{Q}$ ahead of time since any $t+1$ parties can participate in computing the eVRF output, which can be used directly as the generated key. We leave the construction of a threshold eVRF as an open question (see Section 7). However, for a (very) small n, it is possible to achieve the desired result by simply calling $\mathcal{F}_{\mathsf{dkg}\mathcal{Q}}^{n,t}$ separately for every authorized subset $\mathcal{Q} \subset [n]$ and taking the result using the $\mathcal{Q}$ that defines the set of $t+1$ parties who actually participated (i.e., sent first round messages). This is not efficient for a large number of parties, but can certainly be used for thresholds of the type 2-out-of-3 (requiring 3 computations) or 3-out-of-5 (requiring 10 computations).

4.3 The Transformation Methodology for Signing Protocols

ECDSA and Schnorr signing both involve generating a random nonce k and revealing $R = k \cdot G$. In ECDSA, the signature is (r, s), where r is derived from the x-value of R and $s = k^{-1} \cdot (H(m) + r \cdot x)$ with x being the private key and m the message to be signed. In Schnorr (using comparable notation), the signature is (R, s) where $s = k - e \cdot x \bmod q$ (or some variant of this equation) and $e = H(Q, R, m)$, where $Q = x \cdot G$ is the public key.

Many protocols that achieve simulation – e.g., [45,26,25] for ECDSA and [46,48] for Schnorr – work by having each party choose a random k_i and commit to $R_i = k_i \cdot G$ (sometimes also including a zero-knowledge proof of knowledge of the discrete log) and then having the parties decommit and define $R = \sum_{i=1}^{n} R_i$ as the nonce. This methodology ensures that R is uniformly distributed (since

[7] Consider the case of human participants who receive an "invite" to participate in a DKG, and who connect to run the operation and then disconnect. It is much easier to not have to know who the parties that join are ahead of time, and whoever the first $t+1$ parties are, the DKG will go through.

no party can make their R_i depend on the others due to the commitments). When extractable and equivocal commitments are used this is also fully simulatable, since a simulator can choose the honest parties' R_i values after seeing the corrupted parties values, and can therefore make the sum equal the value R received externally in the ideal model. *Stated differently, many protocols generate R by running a simulatable distributed key generation protocol.* When looked at in this light, a natural transformation is to use our one-round DKG protocols (Protocols 3 and 5) to generate R (applying the eVRF to the message to be signed and/or a unique nonce), with the initialization phase being run together with the signing key generation protocol. This enables us to collapse two rounds (commit and open) into a single round, thereby reducing the number of rounds in the signing protocol from three to two, while still achieving simulatable DKG. This enables us to reduce the number of rounds from three to two, without sacrificing on full simulatability. Furthermore, by applying the eVRF to the message to be signed only, we have the all signatures on the same message will have the same nonce, and so *deterministic* signing is achieved at no additional cost.

On achieving a black-box transformation. Ideally, we would like to prove a theorem that says something like any protocol that has a "coin tossing phase" where parties exchange $R_1, \ldots, R_n$ (e.g., by committing and opening) and the nonce is $R = \sum_{i=1}^{n} R_i$ and transform it into a deterministic protocol where the R_i values are generated via the eVRF. However, this actually isn't possible when constructing deterministic signing. In order to see this, take *any* secure ECDSA or Schnorr protocol for probabilistic signing and modify it so that if any of the parties guesses in the first round what the nonce R will be, then all parties send them their private key share. Such a protocol will completely break when deterministically generating R from the message, like with our eVRF, when the same message is signed twice. This does not rule out the possibility of constructing probabilistic signing via a general transformation, but such a transformation is unlikely to be very useful since most existing protocols are not proven with the nonce generation as a separate modular operation. In the following, we therefore present specific protocols which are derived from taking existing protocols and replacing the commit-and-open phase with an eVRF computation.

4.4 Two-Round Simulatable Multiparty Schnorr Signing

In this section, we construct two-round multiparty Schnorr signing from the protocol by [46], by replacing the commit-and-open rounds by eVRF evaluations, as described above. We begin by constructing n-out-of-n deterministic signing, and then describe how to achieve probabilistic signing and threshold signing.

Deterministic signing with additive shares: We first consider the case where the parties hold additive shares of the signing key, and all n-of-n parties participate in signing. We begin by defining the signing functionality. This functionality computes the standard Schnorr signature for a set of parties with additively shared keys. The functionality does not mandate how the key is generated, and it works for any set of inputs held by the parties. This guarantees the same level of security as locally computed Schnorr no matter how keys are generated (using some HD scheme, poorly derived from passwords, or anything else).

Functionality 7 (Deterministic Schnorr Signing $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$)

Let $\mathbb{G}$ be a group of order q with generator G, and let H be the Schnorr hash function. Upon receiving $(\mathsf{Sign}, m, Q, Q_1, \ldots, Q_n, x_i)$ from all n different parties P_i, functionality $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$ works as follows:

1. *Verifies that all parties sent the same* $(m, Q, Q_1, \ldots, Q_n)$ *, that* $Q = \sum_{i=1}^{n} Q_i$ *and that* $Q_i = x_i \cdot G$ *for all* $i \in [n]$*. If no, then it does nothing. Else, it proceeds to the next step.*
2. *Computes* $x = \sum_{i=1}^{n} x_i \bmod q$*.*
3. *If some* (m, k) *is stored then retrieves* k*. Else, chooses a random* $k \leftarrow\!\!\$\ \mathbb{Z}_q$ *and stores* (m, k)*.*
4. *Computes* $R = k \cdot G$*,* $e = H(Q \| R \| m)$ *and* $s = k - e \cdot x \bmod q$*.*
5. *Sends* (m, e, s) *to all parties.*

Securely computing $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$. The idea behind the protocol for Schnorr is simple, due to the fact that the Schnorr signing equation is linear. Specifically, the parties use an eVRF to generate partial nonces (k_i, R_i) where $R_i = k_i \cdot G$ and to share all $R_1, \ldots, R_n$ with all parties. Then, each party can locally compute $R = \sum_{i=1}^{n} R_i$, $e = H(Q \| R \| m)$ and $s_i = k_i - e \cdot x_i \bmod q$. This implies that $\sum_{i=1}^{n} s_i = k - e \cdot x \bmod q$, and so (e, s) is a valid signature. The fact that the signing is deterministic is achieved by applying the eVRF to the (hash of the) message as input. Essentially, this is exactly the same as for the standard EdDSA signing scheme [4], except that a different pseudorandom function is used.

Protocol 8 (Multiparty Schnorr signing – n-out-of-n parties)

– **Input:**
  1. Group parameters: *Let* $\mathbb{G}$ *be a group of order* q *with generator* G*, and let* H *be the Schnorr hash function with output length* λ'
  2. Key shares: *Each party* P_i *holds* $(m, x_i, Q, Q_1, \ldots, Q_n)$ *where* $Q_i = x_i \cdot G$ *and* $\sum_{i=1}^{n} Q_i = Q$
  3. eVRF shares: *Each party* P_i *holds* $(\mathsf{sk}_i, \mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ *where* sk_i *is the eVRF private key associated with* vk_i*, for an eVRF with domain* $\{0,1\}^{\lambda'}$ *and range* $\mathbb{G}$*. These are generated in parallel using Protocol 1 (*π_{EF}*).*

– **The protocol:**
  1. Round 1: *Each party* P_i *computes* $(k_i, R_i, \pi_i) \leftarrow \mathsf{Eval}(\mathsf{sk}_i, H(m))$*, and sends* (R_i, π_i) *to all parties*
  2. Round 2: *Upon receiving* (R_j, π_j) *from all parties* $j \in [n]$*, each* P_i*:*
     (a) *Proceeds if* $\mathsf{Verify}(\mathsf{vk}_j, H(m), R_j, \pi_j) = 1$ *for all* $j \in [n]$ *and aborts otherwise*
     (b) *Computes* $R = \sum_{j=1}^{n} R_j$*,* $e = H(Q \| R \| m)$ *and* $s_i = k_i - x_i \cdot e \bmod q$
     (c) *Sends* s_i *to all parties*
  3. Output: *Upon receiving* $(s_1, \ldots, s_n)$*, each party computes* $s = \sum_{i \in S} s_i \bmod q$ *and checks that* $\mathsf{Verify}_Q(m, (s, e)) = 1$*. If yes, then it outputs* (s, e)*; otherwise it aborts.*

Our protocol assumes that the parties hold **valid and consistent inputs**, meaning that all parties hold the same vectors $(Q, Q_1, \ldots, Q_n)$ and $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$, and in addition $Q = \sum_{i=1}^{n} Q_i$, and $Q_i = x_i \cdot G$ and $\mathsf{vk}_i = \mathsf{sk}_i \cdot G$ for every $i \in [n]$. If this is not guaranteed from a previous protocol execution, then each P_i can simply verify that $x_i \cdot G = Q_i$ and $Q = \sum_{i=1}^{n} Q_i$ at the beginning of the protocol. In addition, all parties can send $h_q \leftarrow H(Q_1, \ldots, Q_n, \mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ to all other parties in the first round, and proceed in the second round only if the same hash value h_q is received from all.

The protocol for two parties. In the specific case of two parties, and where only one party needs to receive output, Protocol 8 can be converted into a protocol where P_1 sends a single message to P_2, and P_2 replies with a single message to P_1 (i.e., a single round trip), and P_1 can then generate output.

Security. We prove the protocol secure in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model (assuming that the initialize phase is carried out during key generation). We stress that the "perfect security" in the theorem statement only holds in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model, but when instantiating the protocol with a real eVRF, the protocol is computationally secure. In addition, we remark that security holds for *all* valid and consistent inputs $\{(x_i, Q, Q_1, \ldots, Q_n)\}_{i\in[n]}$ irrespective of how they are generated. In contrast, it is crucial that the eVRF inputs are securely generated; this is reflected in the proof by the fact that we consider the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model.

Theorem 9. *Assume that the parties hold valid and consistent inputs. Then, Protocol 8 securely computes functionality $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$ in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model with perfect security with abort, in the presence of a malicious static adversary controlling any subset of the parties.*

Proof. The idea behind the proof of security is simple. Using the simulatability of the eVRF (as demonstrated in the proof of Theorem 4) the simulator can force the sum of all R_i's to equal the R value it receives in the signature from the ideal functionality $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$. Then, the values s_i sent by the honest parties in the protocol can be derived perfectly by choosing the honest R_j values carefully for all but one honest party, and using the signature for the last honest party, exactly as in [46].

Let $\mathcal{A}$ be an adversary corrupting a (strict) subset of parties $I \subset [n]$ of size at most $n-1$ (if all n are corrupted, then the protocol is vacuously secure), and let J denote the set of honest parties (and so $I \cup J = [n]$). Without loss of generality, assume that $1 \in J$ (i.e., P_1 is an honest participant). We are now ready to construct the simulator $\mathcal{S}$, with input $\{(\mathsf{Sign}, m, Q, Q_1, \ldots, Q_n, x_i)\}_{i\in I}$, as follows:

1. $\mathcal{S}$ externally sends $(\mathsf{Sign}, m, Q, Q_1, \ldots, Q_n, x_i)$ to $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$ and receives back (m, e, s). $\mathcal{S}$ computes $R = s \cdot G + e \cdot Q$. Then, $\mathcal{S}$ invokes $\mathcal{A}$ in an execution of the protocol.
2. Let f_i be the stored function from the $\mathcal{F}_{\mathsf{eVRF}}$ initialization phase for party P_i, for every $i \in [n]$ (as in the proof of Theorem 4, the simulator $\mathcal{S}$ has these functions), and let sid be the identifier.
3. For all $i \in I$, simulator $\mathcal{S}$ computes $k_i = f_i(H(m))$ and $R_i = k_i \cdot G$.
4. For all $j \in J \setminus \{1\}$, simulator $\mathcal{S}$ chooses a random $s_j \leftarrow \mathbb{Z}_q$ and sets $R_j = s_j \cdot G + e \cdot Q_j$ (where e is from the signature received from $\mathcal{F}_{\mathsf{schnorr}}$). Then, $\mathcal{S}$ sets $R_1 = R - \sum_{i\in I} R_i - \sum_{j\in J\setminus\{1\}} R_j$, using R computed from the signature received from $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$.
5. $\mathcal{S}$ simulates $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{Eval}, \mathsf{sid}, j, H(m), R_j)$ to $\mathcal{A}$ for every $j \in J$, using the R_j values computed in the previous step.
6. $\mathcal{S}$ receives $(\mathsf{Eval}, \mathsf{sid}, i, H(m))$ from $\mathcal{A}$ as sent to $\mathcal{F}_{\mathsf{eVRF}}$ for every $i \in I$. $\mathcal{S}$ waits for all messages to be sent.
7. $\mathcal{S}$ computes $s_i = k_i - x_i \cdot e \bmod q$ for every $i \in I$ ($\mathcal{S}$ can do this since it knows the k_i values for each corrupted party from Step 3 above, and it is given the x_i values of the corrupted parties as input). Then, $\mathcal{S}$ computes
$$s_1 = s - \sum_{i\in I} s_i - \sum_{j\in J\setminus\{1\}} s_j \pmod q$$
using the s_j values chosen above.
8. $\mathcal{S}$ simulates P_j sending s_j to all parties, for every $j \in J$.
9. $\mathcal{S}$ receives $\{s_i\}_{i\in I}$ values sent by $\mathcal{A}$ to the honest parties. If the sum of all of the values sent to an honest P_j is correct (computed by $\sum_{i\in I} s_i$ where the s_i values are as above), then $\mathcal{S}$ adds P_j to $\mathcal{O}_{\mathsf{out}}$.

10. $\mathcal{S}$ sends $\mathcal{O}_{\mathsf{out}}$ to $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$ to instruct which honest parties should receive output.

This completes the simulation. We argue that the simulation is *perfect*. In order to see this, we show that the (R_j, s_j) values sent by the simulator to the adversary are identically distributed to the values sent by the honest parties to the corrupted parties in a real protocol execution. In order to see this, first note that for every $j \in J \setminus \{1\}$ the values are generated as follows:

- *Real:* $k_j \in_R \mathbb{Z}_q$ is random, $R_j = k_j \cdot G$, and $s_j = k_j - e \cdot x_j \bmod q$
- *Simulation:* $\tilde{s}_j \in_R \mathbb{Z}_q$ is random, $\tilde{R}_j = \tilde{s}_j \cdot G + e \cdot Q_j$ (we write $\tilde{s}_j$ and $\tilde{R}_j$ to differentiate from the real)

Let $\tilde{k}_j$ be such that $\tilde{R}_j = \tilde{k}_j \cdot G$. We remark that the simulator $\mathcal{S}$ does not know $\tilde{k}_j$, but the value is well defined. It follows that $\tilde{k}_j = \tilde{s}_j + e \cdot x_j \bmod q$ and so $\tilde{s}_j = \tilde{k}_j - e \cdot x_j \bmod q$, exactly like in a real execution. Furthermore, choosing $\tilde{k}_j$ at random and computing $\tilde{s}_j = \tilde{k}_j - e \cdot x_j \bmod q$ yields the exact same distribution as choosing $\tilde{s}_j$ at random and computing $\tilde{k}_j = \tilde{s}_j + e \cdot x_j \bmod q$.

Next, regarding (R_1, s_1), we have that

$$R_1 = R - \sum_{i \in I} R_i - \sum_{j \in J \setminus \{1\}} R_j = k \cdot G - \sum_{i \in I} k_i \cdot G - \sum_{j \in J \setminus \{1\}} k_j \cdot G$$

where k is the discrete log of R (as computed from the signature), $\{k_i\}_{i \in I}$ are the corrupted parties' values from the eVRF (enforced by the f_i functions), and $\{k_j\}_{j \in J \setminus \{1\}}$ are the implicit values defined above. This therefore defines $k_1 = k - \sum_{i \in I} k_i - \sum_{j \in J \setminus \{1\}} k_j \bmod q$. Similarly, we have

$$s_1 = s - \sum_{i \in I} s_i - \sum_{j \in J \setminus \{1\}} s_j = k - e \cdot x - \sum_{i \in I} (k_i - e \cdot x_i) - \sum_{j \in J \setminus \{1\}} (k_j - e \cdot x_j) \pmod q$$

which holds for $i \in I$ by how the simulator computes $\{s_i\}_{i \in I}$ and for $j \in J \setminus \{1\}$ by the above analysis. Writing $k = \sum_{\ell \in I \cup J} k_\ell$ and $d = \sum_{\ell \in I \cup J} x_\ell$ we have that

$$s_1 = \sum_{\ell \in I \cup J} k_\ell - e \cdot \sum_{i \in I \cup J} x_\ell - \sum_{i \in I} (k_i - e \cdot x_i) - \sum_{j \in J \setminus \{1\}} (k_j - e \cdot x_j) \pmod q$$

and so $s_1 = k_1 - e \cdot x_1 \bmod q$, as required. (We stress that $\mathcal{S}$ does not know these values, and in particular it does not know the x_j values of the honest parties including x_1, and yet is able to generate the correct distribution, as described above.)

Finally, since $\mathcal{S}$ is able to perfectly verify whether or not the corrupted parties send correct values, since it knows all of the corrupted (k_j, x_j) values and so can detect if the *sum* over all s_i values sent by $\mathcal{A}$ is correct. (Note that only the sum matters for $\mathcal{C}$ computing a correct signature.) Thus, the distribution over $\mathcal{C}$ receiving or not receiving output is exactly the same in the real and ideal executions. □

Security under concurrent composition. As shown by [44], perfect security without rewinding implies UC security. As such, assuming that $\mathcal{F}_{\mathsf{eVRF}}$ is implemented using a UC-secure protocol (as in Protocol 1), we have that the protocol is UC-secure and so secure under concurrent composition.

The final result. Using any two-round distributed key generation protocol (e.g., as described in [46]) in parallel with the two-round initialization in Protocol 1, we have the following corollary:

Corollary 2. *There exists a multiparty n-of-n protocol with two rounds for each of key generation and signing that UC computes the deterministic signing functionality $\mathcal{F}_{\mathsf{det\text{-}schnorr}}$ with abort, in the presence of a malicious static adversary controlling any subset of the parties.*

Probabilistic signing. We can achieve two-round probabilistic signing assuming that the parties hold the same *unique nonce*[8] before the protocol begins by having the parties apply the eVRF to $H(H(m), \mathsf{nonce})$. This will result in the eVRF output being (computationally) independent for every different nonce. Practically, the nonce can be a timestamp, with the observation that if two protocol executions use the same timestamp, then the result will just be the same and so no harm can be done. Formally, the functionality computed here $\mathcal{F}_{\mathsf{schnorr}}$ would either choose $k \leftarrow\!\!\$\, \mathbb{Z}_q$ randomly each time (when a unique nonce is guaranteed) or would receive a nonce from all parties in the input and would choose a new $k \leftarrow\!\!\$\, \mathbb{Z}_q$ for every unique nonce (in which case, the functionality works in the same way that the same nonce in different execution would yield the same result, while different nonces would yield a different random R in the signing).

Threshold probabilistic signing. In order to achieve probabilistic threshold signing, the parties can include the set of participating parties $\mathcal{Q}$ into the eVRF computation, together with $H(m)$ and the nonce, in the same way as in Protocol 5 for securely computing $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$. As long as the set of participating parties is known ahead of time (since $\mathcal{Q}$ needs to be input into the eVRF evaluation), this achieves a two-round protocol. In the same way as for $\mathcal{F}^{n,t}_{\mathsf{dkg}\mathcal{Q}}$, if n is very small then it is possible for the parties to provide eVRF values for all possible $\mathcal{Q}$ subsets, and so the set of parties need not be known ahead of time. However, this is only practical for very small n.

Proof of quorum identity. The threshold probabilistic signing protocol described above (that works by including $\mathcal{Q}$ into the eVRF) has a unique property that is of independent interest. The signature generated by the quorum of parties is a *standard signature* with no changes at all. However, the quorum of parties who signed can at a later time provide a proof that *they and only they* generated the signature, assuming a public record of the eVRF verifications keys. This proof for a signature (e, s) is simply the set $\{(R_i, \pi_i\}_{i\in\mathcal{Q}}$, and it is verified by checking that $\mathsf{Verify}(\mathsf{vk}_i, H(m), R_i, \pi_i) = 1$ for all $i \in \mathcal{Q}$ and that $\sum_{i\in\mathcal{Q}} R_i = R$, where $R = s \cdot G + e \cdot Q$. The fact that the proof is sound (i.e., it isn't possible to frame another subset of parties) follows from the fact that if a party P_j did not compute Eval on this input then their (R_j, π_j) value is unknown, and furthermore even if they did (at some later time) the probability that the sum of R_j's for any given subset $\mathcal{Q}'$ will equal a given R is $1/q$ (in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model).[9] Consider for example a setting, where like in a proof of stake, there is a reward for generating a signature and a penalty (slash) for generating a signature when you shouldn't. In such a setting, a proof of quorum identity provides a perfect solution. Furthermore, our protocol can enhance privacy by not revealing the identities of who signed except when needed, or except to entities who need to see it.

Precomputation for the nonce. Functionally speaking, the parties could precompute the nonce R before the message m is known (of course, this makes sense only for probabilistic signing). However, security wise, such a protocol would not securely realize $\mathcal{F}_{\mathsf{schnorr}}$ since the functionality provides

[8] Not to be confused with the "nonce" R in the Schnorr signing, here we mean a unique value nonce which can be a counter, timestamp, or anything.

[9] If t, n are very large, and so $\binom{t}{n}$ is not much smaller than q, then given all $\{(R_j, \pi_j)\}_{j\in[n]}$'s it may be possible to find an appropriate subset (this would require solving a type of subset sum problem, which may or may not be hard). However, if $\binom{t}{n} \ll q$, then the probability that there exists any such subset is negligible.

a full signature (e, s), which fully determines $R = s \cdot G + e \cdot Q$, only upon receiving m. Not only does the resulting protocol no realize $\mathcal{F}_{\mathsf{schnorr}}$, a recent attack due to Navot [53] suggests that precomputation of the nonce can compromise security.

Threshold deterministic signing. As we have mentioned, we do not achieve deterministic signing for the threshold setting. This is because different subsets of parties compute a different R value. In order to achieve this, we need to construct a threshold eVRF, which is left as an open question.

4.5 Two-Round Simulatable Two-Party ECDSA Signing

In a similar way to Section 4.4, in this section we construct a two-round two-party ECDSA signing protocol by replacing the commit-and-open phase in the protocol of [45] with an eVRF evaluation. The result is a protocol with a single round from P_1 to P_2, and a single response from P_2 to P_1.

Functionality 10 (Two-party deterministic ECDSA Signing $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$) *Let $\mathbb{G}$ be a group of order q with generator G, and let H be the ECDSA hash function. Upon receiving $(\mathsf{Sign}, m, Q, x_i)$ from both parties P_1 and P_2, functionality $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$ works as follows:*

1. *Verifies that both parties sent the same (m, Q) , and that $(x_1 + x_2) \cdot G = Q$. If no, then it does nothing. Else, it proceeds to the next step.*
2. *Computes $x = x_1 + x_2 \bmod q$.*
3. *If some (m, k) is stored then retrieves k. Else, chooses a random $k \leftarrow\!\!\$\ \mathbb{Z}_q$ and stores (m, k).*
4. *Computes $R = k \cdot G$ and $r = r_x \bmod q$ where $R = (r_x, r_y)$*
5. *Computes $s = k^{-1} \cdot (H(m) + r \cdot x) \bmod q$.*
6. *Sends (m, r, s) to party P_1 and to the ideal adversary $\mathcal{S}$.*[10]

We have defined the functionality so that only P_1 receives output. It is always possible to have P_1 send P_2 the output, if both are supposed to received the signature.

The protocol idea and differences from [45]. The protocol of [45] uses the Paillier additively homomorphic encryption scheme to enable the parties to generate a signature. In the key generation phase, the parties obtain x_1, x_2 such that $x = x_1 + x_2 \bmod q$ (although [45] refers to multiplicative sharing, the protocol works for additive sharing in the same way). In addition, P_1 generates a Paillier key, and sends an encryption c_{key} of x_1 to P_2, together with a proof that c_{key} is correctly formed.

Next in order to sign, the parties first generate $R = k_1 \cdot k_2 \cdot G$, where k_i is known to party P_i. Then, given the encryption c_{key} of x_1 and given R (and thus r), it is possible for P_2 to generate an encryption of an "almost" signature. In particular, it can compute an encryption of $k_2^{-1} \cdot (H(m) + r \cdot x)$ by adding x_2 to x_1 inside c_{key}, and then multiplying the result by r, adding $H(m)$, and finally multiplying again by k_2^{-1} (since the operations inside Pailier are over the integers, it also adds $\rho \cdot q$ for a random ρ of the appropriate size). Finally, given this ciphertext, party P_1 can decrypt and multiply the result by k_1^{-1}, yielding a "full" signature $k^{-1} \cdot (H(m) + r \cdot x)$.

The main difference between our protocol here and that of [45] is that we generate k_1 and k_2 (for $R = k_1 \cdot k_2 \cdot G$) using the eVRF. In this way, instead of doing commit-and-open, we are able to reduce the protocol to two messages (a single message in each direction), like Protocol 8 for Schnorr

[10] We need the functionality to provide the signature to $\mathcal{S}$ for the case that P_2 is corrupted since R is revealed to P_2 during the protocol execution, but only P_1 receives the signature for output. Thus, we give the signature to the adversary as well (in any case) in order to simulate.

signatures. In addition, the protocol of [45] achieves only a game-based definition of security under standard assumptions, and requires an ad-hoc assumption regarding Paillier to achieve simulation-based security. In addition, it requires P_1 to refuse to run additional executions with P_2 if it is caught cheating in the last message. This limits the ability to run concurrent independent executions with the same key, making some deployment scenarios difficult. (As pointed out in [49], this is necessary since it is possible to actually extract the key one bit at a time if executions are not halted upon cheating.) In order to avoid the requirement to halt if someone attempts to cheat, we add a zero-knowledge proof from P_2 to P_1 (as recommended in [49]) that the ciphertext c_{key} is correctly computed. We have implemented this proof, and it takes 17ms to compute and 11ms to verify (on a 2019 MacBook Pro with a 2.3 GHz 8-Core Intel Core i9 processor). This adds to the running time but is not a problem for most applications.

Protocol 11 (Two-party ECDSA signing)

– **Input:**
  1. Group parameters: *Let $\mathbb{G}$ be a group of order q with generator G, and let H be the ECDSA hash function with output length λ'*
  2. Key shares: *Each party P_i holds (x_i, Q). In addition, P_1 holds a Paillier key $(N, \phi(N))$, and P_2 holds $c_{\mathsf{key}} = \mathsf{Paillier\text{-}enc}_N(x_1)$; these are generated exactly as in [45]*
  3. eVRF shares: *Each party P_i holds $(\mathsf{sk}_i, \mathsf{vk}_1, \mathsf{vk}_2)$ where sk_i is the eVRF private key associated with vk_i, for an eVRF with domain $\{0,1\}^{\lambda'}$ and range $\mathbb{G}$. These are generated in parallel using Protocol 1 (π_{EF}).*

– **The protocol:**
  1. Round 1 – P_1 to P_2: *Party P_i computes $(k_1, R_1, \pi_1) \leftarrow \mathsf{Eval}(\mathsf{sk}_1, H(m))$, and sends (R_1, π_1) to party P_2*
  2. Round 2 – P_2 to P_1: *Upon receiving (R_1, π_1) from P_1, party P_2 works as follows,*
     (a) *Aborts if $\mathsf{Verify}(\mathsf{vk}_1, H(m), R_1, \pi_1) = 0$*
     (b) *Computes $(k_2, R_2, \pi_2) \leftarrow \mathsf{Eval}(\mathsf{sk}_2, H(m))$*
     (c) *Computes $R = k_2 \cdot R_1$ and $r = r_x \bmod q$ where $R = (r_x, r_y)$*
     (d) *Chooses a random $\rho \leftarrow \mathbb{Z}_{q^2}$ and random $\tilde{r} \in \mathbb{Z}_N^*$ (verifying explicitly that $\gcd(\tilde{r}, N) = 1$), and computes*

$$c = \mathsf{Paillier\text{-}enc}_N\left([k_2^{-1} \cdot H(m) \bmod q] + [k_2^{-1} \cdot r \bmod q] \cdot (x_2 + x_1) + \rho \cdot q\right)$$

     using the Paillier homomorphic operations, including ciphertext rerandomization.
     (e) *Sends a proof π_c that c is generated correctly, as in [49]*
     (f) *Sends (R_2, π_2, c, π_c) to P_1*
  3. Output: *Upon receiving (R_2, π_2, c, π_c), party P_2 works as follows,*
     (a) *Aborts if $\mathsf{Verify}(\mathsf{vk}_2, H(m), R_2, \pi_2) = 0$.*
     (b) *Aborts if π_c is not an accepting proof that c was generated correctly.*
     (c) *Computes $R = k_1 \cdot R_2$ and $r = r_x \bmod q$, where $R = (r_x, r_y)$.*
     (d) *Computes $s' = \mathsf{Paillier\text{-}dec}_{\phi(N)}(c)$ and $s'' = k_1^{-1} \cdot s' \bmod q$. (Also sets $s = \min\{s'', q - s''\}$ to ensures that the signature is always the smaller of the two possible values.)*
     (e) *Verifies that (r, s) is a valid signature on m with public key Q. If yes, outputs the signature (r, s); otherwise, outputs abort.*

Security. We now prove security of the protocol; the ideas are a combination of the proof of Theorem 9 (for the generation of R) together with the proof of the original protocol in [45].

Theorem 12. *Protocol 11 securely computes functionality $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$, in the presence of a malicious static adversary controlling any subset of the parties.*

Proof. We separately prove security for the case of a corrupted P_1 and a corrupted P_2. Let $\mathcal{A}$ be an adversary who has corrupted P_1; we construct a simulator $\mathcal{S}$. We prove only the signing protocol, as the key generation protocol is exactly as from [45] together with the eVRF key generation already proven in π_{EF} (Protocol 1). We prove security in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model, as in Theorem 9.

Simulating signing – corrupted P_1: The idea behind the security of the signing protocol is that a corrupted P_1 cannot do anything since all it does is participate in the generation of R and then decrypts the ciphertext c from P_2. Thus, the prove merely requires proving that a simulator can generate the corrupted P_1's view of the decryption of c, given only the signature (r, s) from $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$.

1. Upon input $(\mathsf{Sign}, m, Q, x_1)$, simulator $\mathcal{S}$ sends $(\mathsf{Sign}, m, Q, x_1)$ to $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$ and receives back a signature (r, s) on the message m.
2. $\mathcal{S}$ computes the point R from the signature (r, s), using the ECDSA verification procedure.
3. Let f_1 be the stored function from the $\mathcal{F}_{\mathsf{eVRF}}$ initialization phase with identifier sid ($\mathcal{S}$ has this function and sid).
4. $\mathcal{S}$ invokes $\mathcal{A}$ with input $(\mathsf{Sign}, m, Q, x_1)$ and receives $(\mathsf{Eval}, \mathsf{sid}, 1, H(m))$ from $\mathcal{A}$ as sent to $\mathcal{F}_{\mathsf{eVRF}}$.
5. $\mathcal{S}$ computes $k_1 = f_1(H(m))$, $R_1 = k_1 \cdot G$, and $R_2 = k_1^{-1} \cdot R$.
6. $\mathcal{S}$ simulates $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{Eval}, \mathsf{sid}, 2, H(m), R_2)$ to $\mathcal{A}$
7. $\mathcal{S}$ chooses a random $\rho \leftarrow \mathbb{Z}_{q^2}$, computes $c \leftarrow \mathsf{Paillier\text{-}enc}_N([k_1 \cdot s \bmod q] + \rho \cdot q)$, where s is the value from the signature received from $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$,
8. $\mathcal{S}$ generates a simulated proof π_c that the message c is generated correctly.
9. $\mathcal{S}$ internally hands (c, π_c) to $\mathcal{A}$.

The only difference between the view of $\mathcal{A}$ in a real execution and in the simulation in the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model is the way that c and π_c are generated. Specifically, R_2 is distributed identically in both cases due to the fact that R is randomly generated by $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$ in the signature generation (once for each m) and thus $k_1^{-1} \cdot R$ has the same distribution as $k_2 \cdot G$.

Regarding the ciphertext c, in the simulation it is an encryption of the value $[k_1 \cdot s \bmod q] + \rho \cdot q$, whereas in a real execution it is an encryption of the value $s' = k_2^{-1} \cdot (m' + r \cdot (x_1 + x_2)) + \rho \cdot q$, where $\rho \in \mathbb{Z}_{q^2}$ is random (we stress that all additions here are over the *integers* and not mod$q$, except for where it is explicitly stated in the protocol description). The fact that these two distributions of values are statistically close has been shown in the proof of security in [45].

Finally, regarding π_c, this is indistinguishable by the zero-knowledge property. (Formally, one replaces the generation of c wjth an honest generation, and then the only difference is the proof. This means that the ability to distinguish a real signing execution from a simulated one can be translated into the ability to distinguish a real proof from a simulated one. The other messages in the signing can be executed by the zero-knowledge distinguisher by providing it all secrets as auxiliary input.) This completes the proof of this simulation case.

Simulating signing – corrupted P_2: The simulator for the signing phase works as follows:

1. Upon input $(\mathsf{Sign}, m, Q, x_2)$, simulator $\mathcal{S}$ sends $(\mathsf{Sign}, m, Q, x_2)$ to $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$ and receives back a signature (r, s) on message m.
2. $\mathcal{S}$ computes the point R from the signature (r, s), using the ECDSA verification procedure.

3. Let f_2 be the stored function from the $\mathcal{F}_{\mathsf{eVRF}}$ initialization phase with identifier sid ($\mathcal{S}$ has this function and sid).
4. $\mathcal{S}$ computes $k_2 = f_2(H(m))$, $R_2 = k_2 \cdot G$, and $R_1 = k_2^{-1} \cdot R$.
5. $\mathcal{S}$ invokes $\mathcal{A}$ with input $(\mathsf{Sign}, m, Q, x_1)$ and simulates functionality $\mathcal{F}_{\mathsf{eVRF}}$ sending $(\mathsf{Eval}, \mathsf{sid}, 1, H(m), R_1)$ to $\mathcal{A}$, as the message from P_1 to P_2.
6. $\mathcal{S}$ receives $(\mathsf{Eval}, \mathsf{sid}, 1, H(m))$ from $\mathcal{A}$ as sent to $\mathcal{F}_{\mathsf{eVRF}}$.
7. $\mathcal{S}$ receives (c, π_c) from $\mathcal{A}$ as the message to be sent from P_2 to P_1.
8. $\mathcal{S}$ verifies π_c, and simulates P_1 aborting if it is not accepting (and instructs $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$ to not provide output to P_1).
9. If π_c is accepting, then $\mathcal{S}$ instructs $\mathcal{F}_{\mathsf{det\text{-}ecdsa}}$ to provide output to P_1 and outputs whatever $\mathcal{A}$ outputs.

In the $\mathcal{F}_{\mathsf{eVRF}}$-hybrid model, the message seen by P_2 in the simulation is *identical* to its view in the real execution. Furthermore, there can only be a difference in the result (whether P_1 outputs a valid signature (r, s) or aborts) if π_c is an accepting proof and yet c was not generated correctly. This contradicts the soundness of the zero-knowledge proof and thus occurs with negligible probability only. This concludes the proof. □

Security under concurrent composition. The simulator in the proof is straight-line (no rewinding) assuming straight-line simulation of the zero-knowledge proofs. As such, by [44], the protocol is UC secure assuming "start synchronization" (meaning that all parties have their input before the protocol begins).

Extensions. Probablistic signing can be achieved in the same way as for Schnorr, by providing an additional nonce as input. We remark that if the same nonce is used, then the same signature is obtained and there is no negative security ramification. As such, a timestamp or the like can be used, and this should be sufficient. Regarding precomputation of the first message as discussed for Schnorr, for ECDSA this is more problematic since ECDSA itself has no proof in any standard model. As such, it is not possible to justify (in a standard model) that ECDSA remains secure when m can be chosen after r is known.

Two-round multiparty ECDSA. We leave the question of achieving two-round *multiparty* ECDSA open. Our techniques can be used to remove a round of communication of commit-and-open in generating R. However, existing protocols require more than two rounds irrespective of this step (the protocol of [25] has only three rounds, but their first round requires additional steps beyond just committing to R_i values).

4.6 Verifiable and MPC-Friendly Hierarchical Key Derivation

BIP032/BIP044 [63,58] hierarchical-deterministic (HD) key derivation works by deriving multiple keys from a single root secret, utilizing a tree structure. The method includes hardened derivation and normal derivation. A hardened derivation takes a node's private key and path information and applies a pseudorandom function in order to derive a pseduorandom private key for the child node. In contrast, a normal derivation is applied to a node's public key and public path information only. Normal derivations enable anyone to generate new addresses that can be used, given a public key/address. In addition, it is possible to link different keys that have been normally derived from a single key, and delegation on normally derived keys is not possible (since given the private key

of one normally derived key, it is possible to find the private key of all of its siblings in the tree). In contrast, hardened derivations can only be computed by the private key owner, different keys derived via hard derivation from a single node cannot be linked, and given the private key of a hardened derived node it is not possible to find the private key of any of its siblings in the tree.

Although BIP032 prescribes a unified method for hardened and normal derivations, *any pseudorandom function* can be used in its place, and this does not affect the public method used for normal derivation. In this section, we propose a new paradigm for hardened derivations using an eVRF instead of a standard pseudorandom function. Concretely, hardened derivation in BIP032/BIP044 works by applying SHA512 to a node's private key and path information, and the output is used as the child's private key. (The exact details of how this is carried out is not important here, but this is the basic idea.) Instead of using this method, we propose adding an eVRF private key to the node of the tree, and deriving descendants in the tree by applying the eVRF to the path and taking the result as the private key. Concretely, for a given path (say, determined as in BIP044), we define the key associated with the node for that path by computing $\mathsf{Eval}(k, \mathsf{path}) = (x, Q, \pi)$, and take x to be the private key and Q to be the public key.[11] This guarantees that all hardened keys are pseudorandom, and cannot be linked. In addition, given any hardened derived key, it is not possible to find any other hardened derived key (by uniqueness of the eVRF output), and so hardened-derived keys can be delegated. This therefore makes it a suitable replacement to BIP032 derivation. We will now explain why this is advantageous, and what feature of BIP032 is lost. (We stress that normal derivations remain unchanged. As such, this method is indistinguishable from standard BIP032, since hardened-derived keys are indistinguishable from random in both methods.)

Before proceeding we remark that the use of HD wallets via BIP032/BIP044 is very popular since it enables parties to backup one seed, and to derive many keys from that seed for different purposes.

Derivation verifiability. Standard BIP032 derivation does not provide any validation that a public key in the tree has indeed been correctly derived from the initial seed. In contrast, by the properties of an eVRF, the public key of any derived key can be provably validated (since the eVRF outputs the public key and a proof, by definition). This can be useful in many settings. Consider an institutional wallet, for example, where keys are derived and public keys provided externally for deposit. The party transferring the funds to those addresses actually has no way of knowing that they are correct, barring being "told so". This opens the door to phishing and other attacks, where parties are fooled into transferring funds to malicious addresses. However, using our eVRF method, it suffices to generate a certificate on vk once and for all, and then any address derived can be linked cryptographically to the issuing institution.

MPC-friendly derivation. Consider a Blockchain wallet that uses BIP032/ BIP044, while backing up only the seed (almost all such wallets work this way). If one wishes to construct an MPC wallet so that the user's key is split between the user and an institution, then standard BIP032 derivation becomes very expensive. In particular, securely computing SHA512 operations using techniques like garbled circuits is possible, but expensive (especially, over a low bandwidth communication channel). In such cases, an eVRF based hardened derivation can be much more useful. Specifically, each of the user device and server can generate an eVRF instance, backing them both up, and then new keys can be derived by independently computing eVRF output; each party holds its own share of the private key, and the public key is obtained by adding the public output in both

[11] We stress that this is different to BIP032 where the input to SHA512 contains private data.

eVRF computations. This preserves the property that only the root eVRF keys need to be backed up (since given them it is possible to derive all keys), and each party can work independently and efficiently to derive a key that is additively shared between them. (Multiplicative sharing is possible in the same way.)

Of course, the above method could be achieved by just applying any local pseudorandom function to the path, and having the parties announce the public result to the other (P_1 generates $x_1 \leftarrow PRF_{k_1}(\mathsf{path})$ and announces $Q_1 \leftarrow x_1 \cdot G$, and likewise P_2). However, a malicious P_1 could just generate a random x_1 that is independent of k_1, and this cannot be detected. Such behavior would make the backup of k_1 useless, since the private key in this "derivation" cannot be obtained from it. In order to prevent such behavior, we want P_1 and P_2 to each provide a zero-knowledge proof that Q_1 and Q_2 are indeed generated correctly from the backed-up root keys. This ensures that the address $Q = Q_1 + Q_2$ can be used safely, since the private keys needed to derive them have been backed up. An eVRF provides exactly this property.

Delegation of a sub-tree. We conclude by noting that our method does *not* support delegating an entire subtree of hardened derivations. This is because the root key is needed to carry out any hardened derivation, and revealing this would reveal all keys. As such, a hardened-derived key can be delegated safely, but the party receiving that private key can only carry out normal derivations. This is not a limitation in the way current wallets work, but this difference from the standard BIP032 is worth noting.

5 An eVRF from Compatible Public-Key Encryption

In the following sections we turn to constructing efficient eVRFs. In the current section (Section 5), we show how to construct an efficient eVRF from *compatible* public-key encryption scheme. In Section 6, we present a different construction of eVRF from the (classic) PRF based on the Decision Diffie-Hellman (DDH) assumption.

Compatible encryption schemes are defined in Section 5.1. The construction itself is presented in Section 5.2. In Section 5.3, we show that some of the required properties hold (in particular, have efficient *equality proofs*) for any *linearly homomorphic encryption scheme*, and in Section 5.4 we exploit that to give an efficient construction using the Paillier encryption scheme.

5.1 Compatible Encryption Schemes

Recall that a **public-key encryption scheme** $\mathcal{E}$ is a tuple of algorithms $\mathcal{E} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$, where $\mathsf{KGen}(1^\lambda) \to (\mathsf{sk}, \mathsf{pk})$, $\mathsf{Enc}(\mathsf{pk}, m) \to ct$, and $\mathsf{Dec}(\mathsf{sk}, ct) \to m$ or $\perp$. By convention, we assume that the public key pk contains the description of the plaintext space, that the secret key sk contains the corresponding public key pk, and that both contain the security parameter 1^λ. For ease of notation we assume that the plaintext space is $\mathbb{Z}_n$ for some $n \in \mathbb{N}$. Let $\mathcal{R}_{\mathsf{pk}}$ be the randomness domain used by $\mathsf{Enc}(\mathsf{pk}, \cdot)$, and let $\mathcal{C}_{\mathsf{pk}}$ be the set of all valid ciphertexts associated with $\mathsf{pk} = (n, \cdot)$, namely

$$\mathcal{C}_{\mathsf{pk}} := \big\{ \mathsf{Enc}(\mathsf{pk}, m; r) \ : \ m \in \mathbb{Z}_n, \ r \in \mathcal{R}_{\mathsf{pk}} \big\}.$$

Our construction uses an encryption scheme that is compatible with the required eVRF domain in the following sense.

Definition 7 (Compatible encryption schemes). *Let $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda\in\mathbb{N}}$ be an ensemble of domains and ranges for an eVRF, where each $\mathbb{G}_\lambda$ is a finite cyclic group with generator $G_\lambda \in \mathbb{G}_\lambda$. We say that a public key encryption scheme $\mathcal{E} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ is **compatible** if it satisfies the following properties:*

- **Compatible domain***: for all $\lambda \in \mathbb{N}$, the plaintext space associated with every pk output by $\mathsf{KGen}(1^\lambda)$ is $\mathbb{Z}_{|\mathbb{G}_\lambda|}$. That is, the plaintext space is the same as the group of eVRF exponents.*

- **Perfectly binding***: There are no decryption errors, namely $\Pr\big[\mathsf{Dec}\big(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, m)\big) = m\big] = 1$ for all $(\mathsf{pk}, \mathsf{sk})$ in the support of $\mathsf{KGen}(1^\lambda)$ and all $m \in \mathbb{Z}_{|\mathbb{G}_\lambda|}$. This implies that $\mathcal{E}$ is a perfectly binding encryption scheme, that is, for every pk output by $\mathsf{KGen}(1^\lambda)$ and every distinct $m_1, m_2 \in \mathbb{Z}_{|\mathbb{G}_\lambda|}$, the support of $\mathsf{Enc}(\mathsf{pk}, m_1)$ is disjoint from the support of $\mathsf{Enc}(\mathsf{pk}, m_2)$.*

- **Uniform ciphertexts***: for every $(\mathsf{sk}, \mathsf{pk})$ output by KGen with corresponding plaintext space $\mathbb{Z}_{|\mathbb{G}_\lambda|}$, if m is uniform in $\mathbb{Z}_{|\mathbb{G}_\lambda|}$, then $\mathsf{Enc}(\mathsf{pk}, \mathsf{m})$ is statistically close to uniform in $\mathcal{C}_{\mathsf{pk}}$. Because $\mathcal{E}$ is perfectly binding, this property is equivalent to the dual property we call **uniform plaintexts**, which says that if c is uniform in $\mathcal{C}_{\mathsf{pk}}$, then $\mathsf{Dec}(\mathsf{sk}, \mathsf{c})$ is statistically close to uniform in $\mathbb{Z}_{|\mathbb{G}_\lambda|}$.*

- **Samplable ciphertexts***: There is a set ensemble $\{\mathcal{Z}_\lambda\}_{\lambda\in\mathbb{N}}$ of samplable sets such that for every $\lambda \in \mathbb{N}$ and every $(\mathsf{sk}, \mathsf{pk})$ output by $\mathsf{KGen}(1^\lambda)$, it holds that $\mathcal{C}_{\mathsf{pk}} = \mathcal{Z}_\lambda$. While this property simplifies the exposition, it does not always hold since in many public key encryption schemes the set $\mathcal{C}_{\mathsf{pk}}$ is different for every pk output by $\mathsf{KGen}(1^\lambda)$. Our eVRF construction works equally well with a relaxed notion of samplable ciphertexts given in Definition 8.*

- **Proof of valid public key***: there is a non-interactive zero-knowledge proof $(P_{\text{pub}}, V_{\text{pub}})$ for the instance-witness relation $\mathcal{R}_{\text{pub}}$ defined as*

$$\mathcal{R}_{\text{pub}} := \{(\mathsf{pk}; \mathsf{sk}) \ : \ (\mathsf{pk}, \mathsf{sk}) \in \mathcal{L}_{\text{pub}}\} \tag{3}$$

  where $\mathcal{L}_{\text{pub}}$ is the set of all pairs $(\mathsf{sk}, \mathsf{pk})$ that can be output by KGen.

- **Proof of equality***: there is a non-interactive zero knowledge proof $(P_{\text{eq}}, V_{\text{eq}})$ for the instance-witness relation $\mathcal{R}_{\text{eq}}$ defined as*

$$\mathcal{R}_{\text{eq}} := \left\{ \big((\mathsf{pk}, Y, ct) \ ; \ (\mathsf{sk}, y)\big) \ : \ \begin{array}{c} Y \in \mathbb{G}_\lambda, \quad y \in \mathbb{Z}_{|\mathbb{G}_\lambda|}, \quad (\mathsf{pk}, \mathsf{sk}) \in \mathcal{L}_{\text{pub}}, \\ Y = y \cdot G_\lambda, \quad \mathsf{Dec}(\mathsf{sk}, ct) = y \end{array} \right\} \tag{4}$$

We will construct a public key system with an efficient proof system for $(P_{\text{eq}}, V_{\text{eq}})$ in Section 5.3.

5.2 The Basic eVRF Construction

We now present the eVRF from a compatible public-key encryption scheme.

Construction 13 (an eVRF from compatible encryption) *Let $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda\in\mathbb{N}}$ be an ensemble of domains and ranges for an eVRF, where each $\mathbb{G}_\lambda$ is a finite cyclic group with generator $G_\lambda \in \mathbb{G}_\lambda$. Let $\mathcal{E} := (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ be a compatible public key encryption scheme with ciphertext space $\{\mathcal{Z}_\lambda\}_{\lambda\in\mathbb{N}}$. The **eVRF derived from** $\mathcal{E}$ using a hash function ensemble $\mathcal{H} = \{\mathcal{O}_{\mathcal{X}_\lambda, \mathcal{Z}_\lambda}\}_{\lambda\in\mathbb{N}}$, and a non-interactive proof system $(P_{\text{eq}}, V_{\text{eq}})$ for $\mathcal{R}_{\text{eq}}$, is defined for every $\lambda \in \mathbb{N}$ and hash function $\mathsf{H} : \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$ in $\mathcal{H}$ as follows:*

– $\mathsf{KGen}_{\mathrm{eVRF}}(1^\lambda)$: *output* $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\!\$\ \mathsf{KGen}(1^\lambda)$. // *then* $\mathcal{C}_{\mathsf{pk}} = \mathcal{Z}_\lambda$ *by the samplable ciphertexts property.*

– $\mathsf{Eval}^{\mathsf{H}}(\mathsf{sk}, x)$:

1. *Let* $ct \leftarrow \mathsf{H}(x) \in \mathcal{C}_{\mathsf{pk}}$, $y \leftarrow \mathsf{Dec}(\mathsf{sk}, ct) \in \mathbb{Z}_{|\mathbb{G}_\lambda|}$, $Y \leftarrow y \cdot G_\lambda \in \mathbb{G}_\lambda$.
2. *Let* $\pi \leftarrow\!\!\$\ P_{\mathrm{eq}}(\mathsf{pk}, Y, ct, \mathsf{sk}, y)$.
3. *Output* (y, Y, π).

– $\mathsf{Verify}^{\mathsf{H}}(\mathsf{pk}, x, Y, \pi)$: *Let* $ct \leftarrow \mathsf{H}(x)$ *and output* $V_{\mathrm{eq}}(\mathsf{pk}, Y, ct, \pi)$.

Theorem 14. *Let* $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda\in\mathbb{N}}$ *be an ensemble of domains and ranges for an eVRF, where each* $\mathbb{G}_\lambda$ *is a finite cyclic group with generator* $G_\lambda \in \mathbb{G}_\lambda$. *Suppose that* $\mathcal{E}$ *is a compatible and semantically secure public key encryption scheme, and has ciphertext space* $\{\mathcal{Z}_\lambda\}_{\lambda\in\mathbb{N}}$. *Then Construction 13 is a secure eVRF (as in Definition 4) with respect to the domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda\in\mathbb{N}}$ *and function-family ensemble* $\{\mathcal{O}_{\mathcal{X}_\lambda, \mathcal{Z}_\lambda}\}_{\lambda\in\mathbb{N}}$.

Proof. We argue that the four eVRF properties hold:

– *Consistency*: Holds by construction.

– *Verifiability* of Eval_2 as a VRF: This follows directly from the soundness of the proof system $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$ and the binding property of $\mathcal{E}$. Fix pk, H, and $x \in \mathcal{X}_\lambda$. By the binding property there is a unique $(y, r) \in \mathbb{Z}_{|\mathbb{G}_\lambda|} \times \mathcal{R}_{\mathsf{pk}}$ such that $\mathsf{Enc}(\mathsf{pk}, y; r) = \mathsf{H}(x)$. Then the soundness of $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$ implies that the probability of efficiently finding a pair (Y', π) with $Y' \neq y \cdot G_\lambda$ that makes the verifier accept is negligible.

– *Simulatability as a VRF*: The simulator $\mathsf{Sim}^{\mathsf{H}}(\mathsf{pk}, x, Y)$ works as follows: (1) compute $ct \leftarrow \mathsf{H}(x)$, (2) sample a proof π for the $\mathcal{R}_{\mathrm{eq}}$ statement (pk, Y, ct) using the zero knowledge simulator for the proof system $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$, and (3) output (Y, π).

It remains to prove the *Pseudorandomness* of Eval_1. We show that this follows from the semantic security of $\mathcal{E}$ and its uniform ciphertext property from Definition 7. Let $\mathcal{A}$ be a PRF adversary as in (1). We define the following sequence of hybrid distributions.

– *Game 0:* This game is the left hand side of (1). Recall that during the game $\mathcal{A}$ can query two oracles: a random oracle $\mathsf{H} : \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$ and an eVRF evaluation oracle $\mathsf{Eval}_1(\mathsf{sk}, \cdot) : \mathcal{X}_\lambda \to \mathbb{Z}_{|\mathbb{G}_\lambda|}$ defined as $\mathsf{Eval}_1(\mathsf{sk}, x) := \mathsf{Dec}(\mathsf{sk}, \mathsf{H}(x))$.

– *Game 1:* We replace the random oracle $\mathsf{H} : \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$ in (1) with an oracle that responds to a query for an $x \in \mathcal{X}_\lambda$ by sampling a random $y_x \leftarrow\!\!\$\ \mathbb{Z}_{|\mathbb{G}_\lambda|}$ and returning $\mathsf{H}(x) := \mathsf{Enc}(\mathsf{pk}, y_x)$. The oracle responds consistently to repeated queries for the same x. By the *uniform ciphertext* property of $\mathcal{E}$ from Definition 7, this *Game 1* is statistically indistinguishable from *Game 0*.

– *Game 2:* We replace the eVRF evaluation oracle in *Game 1* with an oracle that responds to a query for an $x \in \mathcal{X}_\lambda$ with $y_x \in \mathbb{Z}_{|\mathbb{G}_\lambda|}$, where y_x was sampled in response to a query for $\mathsf{H}(x)$. Since $\mathcal{E}$ is perfectly binding, $\mathcal{A}$'s view is identical in Games 1 and 2. Note that in *Game 2* the secret key sk is never used.

– *Game 3:* We replace the eVRF evaluation oracle in *Game 2* with an oracle that responds to a query for an $x \in \mathcal{X}_\lambda$ by sampling $y'_x \leftarrow\!\!\$\ \mathbb{Z}_{|\mathbb{G}_\lambda|}$ and responding with $\mathsf{Eval}(x) := y'_x$. The oracle responds consistently to repeated queries for the same x. Observe that the answer to $\mathsf{Eval}(x)$ is independent of the answer to $\mathsf{H}(x)$. We will argue that *Game 3* is indistinguishable from *Game 2* by the semantic security property of $\mathcal{E}$.

– *Game 4:* This game is the right hand side of (1). The only difference from *Game 3* is that $\mathsf{H}\colon \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$ is once again computed by a random function. This *Game 4* is statistically indistinguishable from *Game 3* by, once again, the uniform ciphertext property of $\mathcal{E}$.

It now follows that *Game 0* is indistinguishable from *Game 4* as required.

It remains to argue that Games 2 and 3 are indistinguishable. Suppose that adversary $\mathcal{A}$ can distinguish these games while making at most Q random oracle queries. We construct an adversary $\mathcal{B}$ that breaks semantic security of $\mathcal{E}$. It is convenient to use a semantic security game where $\mathcal{B}$ can request up to Q encryption challenges. This game is equivalent to the standard semantic security game [8, Thm 11.1]. The semantic security challenger is initialized with 1^λ and a bit $b \in \{0,1\}$ and then runs $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\$\ \mathsf{KGen}_{\mathrm{pal}}(1^\lambda)$. Now $\mathcal{B}(1^\lambda, \mathsf{pk})$ runs $\mathcal{A}(1^\lambda)$ and responds to its queries as follows:

whenever $\mathcal{A}$ issues a random oracle query for some $\mathsf{H}(x)$ our $\mathcal{B}$ does:

(1) $y_x^{(0)}, y_x^{(1)} \leftarrow\!\$\ \mathbb{Z}_{|\mathbb{G}_\lambda|}$

(2) issue an encryption query $(y_x^{(0)}, y_x^{(1)})$ to its semantic security challenger
and the challenger responds with $ct_x \leftarrow \mathsf{Enc}(\mathsf{pk}, y_x^{(b)})$

(3) $\mathcal{B}$ sends ct_x back to $\mathcal{A}$, meaning that $\mathsf{H}(x) := ct_x$

whenever $\mathcal{A}$ issues an eVRF evaluation query for some $x \in \mathcal{X}_\lambda$ our $\mathcal{B}$ responds with $y_x^{(0)}$.

eventually $\mathcal{A}$ outputs a bit $b' \in \{0,1\}$ and $\mathcal{B}$ outputs the same b'.

Observe that when $b = 0$ our $\mathcal{B}$ emulates a Game 2 challenger to $\mathcal{A}$. When $b = 1$ our $\mathcal{B}$ emulates a Game 3 challenger to $\mathcal{A}$. Therefore, $\mathcal{B}$ guesses its challenger's bit b with the same advantage that $\mathcal{A}$ distinguishes Game 2 from Game 3. Hence, since $\mathcal{E}$ is semantically secure, the two games are indistinguishable. This completes the proof of the theorem. □

A relaxed samplable ciphertext property. Recall that our eVRF in Construction 13 assumed the samplable ciphertext property in Definition 7, which says that the ciphertext space $\mathcal{C}_{\mathsf{pk}}$ is the same for all pk generated by $\mathsf{KGen}(1^\lambda)$. This property does not hold for cryptosystems where $\mathcal{C}_{\mathsf{pk}}$ is a different set for every pk. However, this is not a problem because Construction 13 can easily be adapted to work with the following relaxed samplable ciphertexts property.

Definition 8. *A public key encryption scheme* $\mathcal{E} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *satisfies the* ***relaxed samplable ciphertexts*** *property, if there a set ensemble* $\{\mathcal{Z}_\lambda\}_{\lambda\in\mathbb{N}}$ *and a pair of algorithms* $I(\mathsf{pk}, z) \to ct$ *and* $I^{-1}(\mathsf{pk}, ct) \to z$*, where* I *is deterministic poly-time and* I^{-1} *is a* PPT*. For every* λ *and every* $(\mathsf{sk}, \mathsf{pk})$ *output by* $\mathsf{KGen}(1^\lambda)$ *these algorithms must satisfy*

– *if* z *is uniform in* $\mathcal{Z}_\lambda$ *then* $I(\mathsf{pk}, \mathsf{z})$ *is statistically close to uniform on* $\mathcal{C}_{\mathsf{pk}}$*;*
– *if* c *is uniform in* $\mathcal{C}_{\mathsf{pk}}$ *then* $I^{-1}(\mathsf{pk}, \mathsf{c})$ *is statistically close to uniform on* $\mathcal{Z}_\lambda$*;*
– *for all* $c \in \mathcal{C}_{\mathsf{pk}}$*, the condition* $I(\mathsf{pk}, I^{-1}(\mathsf{pk}, c)) = c$ *holds with overwhelming probability.*

Now, in Construction 13 we can replace all occurrences of $\mathsf{H}(x)$ with $I\big(\mathsf{pk}, \mathsf{H}(x)\big)$, which outputs an element in $\mathcal{C}_{\mathsf{pk}}$, as required for the construction to work. The algorithm I^{-1} is only used in the security proof: we replace every response ct to a random oracle query in the proof of Theorem 14 with $I^{-1}(\mathsf{pk}, ct)$.

5.3 Public Key Encryption Scheme with Efficient Equality Proofs

It remains to construct a public key encryption scheme that has relaxed samplable ciphertexts and satisfies the other compatibility properties from Definition 7. The main challenge is to build an efficient NIZK for the relation $\mathcal{R}_{\text{eq}}$ from (4). Such a NIZK exists for a *linearly homomorphic* encryption scheme.

Recall that a **linearly homomorphic encryption scheme** is a public-key encryption scheme $\mathcal{E} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$, where the plaintext space is $\mathbb{Z}_n$ for some n, and there is a fourth algorithm called Eval that is invoked as $\mathsf{Eval}(ct_1, ct_2, a_1, a_2) \to ct$. The Eval algorithm takes as input two ciphertext $ct_1, ct_2 \in \mathcal{C}_{\mathsf{pk}}$, and two scalars $a_1, a_2 \in \mathbb{Z}_n$ and outputs a ciphertext ct, such that if $\mathsf{Dec}(\mathsf{sk}, ct_1) = m_1$ and $\mathsf{Dec}(\mathsf{sk}, ct_2) = m_2$, where $m_1, m_2 \in \mathbb{Z}_n$, then $\mathsf{Dec}(\mathsf{sk}, ct) = a_1 m_1 + a_2 m_2 \in \mathbb{Z}_n$. It will be convenient to use the notation $ct \leftarrow a_1 \cdot ct_1 + a_2 \cdot ct_2$ to mean $ct \leftarrow \mathsf{Eval}(ct_1, ct_2, a_1, a_2)$.

For a linearly homomorphic encryption scheme $\mathcal{E}$ there is an efficient Chaum-Pedersen [20] style honest verifier zero-knowledge interactive proof system for the relation $\mathcal{R}_{\text{eq}}$ from (4). Recall that the Chaum-Pedersen protocol proves in zero knowledge that a tuple $(P, \alpha G, Q, \beta Q) \in \mathbb{G}_\lambda^4$ satisfies $\alpha = \beta$. We use the fact that the protocol similarly applies when the group elements are replaced with linearly homomorphic ciphertexts. For completeness, we describe the proof system in Appendix A. The proof system is public coin and can be made non-interactive using the Fiat-Shamir transform.

Instantiating the linearly homomorphic encryption scheme. The next question is how to instantiate the linearly homomorphic encryption scheme with a compatible encryption scheme (as in Definition 7) that supports the relaxed samplable ciphertexts property. There are many linearly homomorphic encryption systems to choose from, such as [3,51,56,57,22,30,17,37,18] to name a few. The two that are most relevant to us are the Paillier encryption scheme [57] and the scheme of Castagnos and Laguillaumie [18]. However, neither one satisfies all the compatibility properties that we need.

- Paillier encryption satisfies all the compatibility requirements in Definition 7 except one: the plaintext domain is a large modulus that is a product of two large primes. For Construction 13 we need the plaintext domain to be the same as the domain of the eVRF (the compatible domain property in Definition 7).
- The Castagnos-Laguillaumie encryption scheme [18] has the required plaintext domain, but the ciphertext space is sparse in its ambient space. As a result we cannot map a random oracle output in $\mathcal{Z}_\lambda$ to a uniform ciphertext in the set $\mathcal{C}_{\mathsf{pk}}$ of valid ciphertexts, using a reversible map. More precisely, we do not know how to construct algorithms (I, I^{-1}) as needed for the relaxed samplable ciphertexts property (Definition 8).

Nevertheless, we show in Section 5.4 that by slightly tweaking Construction 13 we can make it work with Paillier encryption. First, let us show that Paillier satisfies the relaxed samplable ciphertexts property from Definition 8. Recall that the plaintext domain associated with a Paillier public key $(\mathsf{pk}, \mathsf{sk}) \leftarrow\!\$\ \mathsf{KGen}(1^\lambda)$ is $\mathbb{Z}_n$ for some integer n, and the ciphertext space is $\mathcal{C}_{\mathsf{pk}} := \mathbb{Z}_{n^2}$. Moreover n is in the set $[2^{\gamma\lambda-2},\ 2^{\gamma\lambda}]$ for some universal constant $\gamma \in \mathbb{N}$. To show that Paillier has the relaxed samplable ciphertexts property, let us set $\mathcal{Z}_\lambda := \{0, \ldots, B\}$, where $B := 2^{(2\gamma+1)\lambda}$. Then $B > 2^\lambda n^2$ and we can define

- $I_{\text{pal}}(\mathsf{pk}, z) = (z \bmod n^2)$ for $z \in \mathcal{Z}_\lambda$, and
- $I_{\text{pal}}^{-1}(\mathsf{pk}, c) = v \cdot n^2 + c$ for $c \in \mathbb{Z}_{n^2}$ and $v \leftarrow\!\$\ \big[0, \lfloor B/n^2 \rfloor\big]$.

It is not difficult to see that these functions satisfy the properties needed for the relaxed samplable ciphertexts property. We will argue this more precisely in the next section.

5.4 An Instantiation Using Paillier Encryption

In this section we adapt Construction 13 to use Paillier encryption, despite the fact that the plaintext domain of Paillier is different from the exponent group of the eVRF.

As usual, let $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda\in\mathbb{N}}$ be a domain and a range for an eVRF with security parameter λ. Here $\mathbb{G}_\lambda$ is a cyclic group of prime order q with generator $G_\lambda \in \mathbb{G}_\lambda$. Let $\mathcal{E}_{\text{pal}} = (\mathsf{KGen}_{\text{pal}}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be the Paillier linearly homomorphic encryption scheme, and let pk be a Paillier public key generated by $\mathsf{KGen}_{\text{pal}}(1^\lambda)$. We use $\mathbb{Z}_n$ to denote the plaintext domain associated with pk, and we will use the fact that n is much larger than q. The Paillier ciphertext space $\mathcal{C}_{\mathsf{pk}}$ associated with pk is the set $\mathcal{C}_{\mathsf{pk}} := \mathbb{Z}_{n^2}$. Strictly speaking, Paillier ciphertexts are in $\mathbb{Z}^*_{n^2}$, but we will adopt the convention that all ciphertexts in $\mathbb{Z}_{n^2}$ that are outside of $\mathbb{Z}^*_{n^2}$ decrypt to 0.

The eVRF derived from Paillier is the same as Construction 13, except that we reduce the decryption of $ct \leftarrow \mathsf{H}(x)$ modulo q and use the resulting value as the eVRF exponent. To do so we make a small, but important, modification to the relation $\mathcal{R}_{\text{eq}}$ from (4). Define the relation $\mathcal{R}'_{\text{eq}}$ as

$$\mathcal{R}'_{\text{eq}} := \left\{ \big((\mathsf{pk}, Y, ct)\ ;\ (\mathsf{sk}, y)\big)\ :\ \begin{array}{c} Y \in \mathbb{G}_\lambda, \quad y \in [0, q-1], \quad (\mathsf{pk}, \mathsf{sk}) \in \mathcal{L}_{\text{pub}}, \\ Y = y \cdot G_\lambda, \quad \mathsf{Dec}(\mathsf{sk}, ct) = y \end{array} \right\} \tag{5}$$

The difference from (4) is that now y is an integer and the equality $\mathsf{Dec}(\mathsf{sk}, ct) = y$ is interpreted as an equality of two integers in the set $[0, n-1]$, whereas in (4) this was an equality of elements in $\mathbb{Z}_{|\mathbb{G}_\lambda|}$. A proof system for $\mathcal{R}'_{\text{eq}}$ proves that the discrete log of Y base $\mathbb{G}_\lambda$, as an integer in $[0, q-1]$, is equal to the decryption of ct, as an integer in $[0, n-1]$. This problem is closely related to the problem of proving equality of discrete log across two finite cyclic groups of different sizes, one has order q and the other has order n. Protocols for this task were proposed by Camenisch and Lysyanskaya [13], by Agrawal, Ganesh, and Mohassel [1] using bit decomposition, and by Chase, Orrù, Perrin, and Zaverucha [19] using range proofs and rejection sampling. In Appendix B we adapt these techniques to give a proof system for $\mathcal{R}'_{\text{eq}}$.

Using a proof system for $\mathcal{R}'_{\text{eq}}$ we obtain the following concrete construction. We use the set ensemble $\{\mathcal{Z}_\lambda\}$ and functions $(I_{\text{pal}}, I^{-1}_{\text{pal}})$ defined at the end of Section 5.3.

Construction 15 (Paillier based eVRF) *The **eVRF derived from Paillier** using a non-interactive proof system $(P_{\text{eq}}, V_{\text{eq}})$ for $\mathcal{R}'_{\text{eq}}$ in (5) and a hash function ensemble $\mathcal{H} = \{\mathcal{O}_{\mathcal{X}_\lambda, \mathcal{Z}_\lambda}\}_{\lambda\in\mathbb{N}}$, is defined for every $\lambda \in \mathbb{N}$ and hash function $\mathsf{H} : \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$ in $\mathcal{H}$ as:*

- $\mathsf{KGen}_{\text{eVRF}}(1^\lambda)$*: Sample* $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\!\$\ \mathsf{KGen}_{\text{pal}}(1^\lambda)$ *and output* $(\mathsf{sk}, \mathsf{pk})$.
- $\mathsf{Eval}^{\mathsf{H}}(\mathsf{sk}, x)$*: Let* $q := |\mathbb{G}_\lambda|$ *and let* $\mathbb{Z}_n$ *be the plaintext space of* pk.
  1. *Let* $ct \leftarrow (\mathsf{H}(x) \bmod n^2)$ *and* $\alpha \leftarrow \mathsf{Dec}(\mathsf{sk}, ct) \in [0, n-1]$.
  2. *Let* $y \leftarrow \alpha \bmod q$ *and* $w \leftarrow \lfloor \alpha/q \rfloor \in \big[0, \lfloor n/q \rfloor\big]$ // *remainder and quotient mod q.*
  3. *Let* $ct_y \leftarrow ct - \mathsf{Enc}(\mathsf{pk}, w \cdot q; 0)$, $Y \leftarrow y \cdot G_\lambda \in \mathbb{G}_\lambda$ // *ct_y is an encryption of y.*
  4. *Let* $\pi \leftarrow\!\!\$\ P_{\text{eq}}(\mathsf{pk}, Y, ct_y, \mathsf{sk}, y)$ // *proof that* $\mathsf{Dec}(\mathsf{sk}, ct_y) = y$.
  5. *Output* $\big(y, Y, (w, \pi)\big)$.

– $\mathsf{Verify}^{\mathsf{H}}(\mathsf{pk}, x, Y, (w, \pi))$: *Let* $ct \leftarrow \mathsf{H}(x) \bmod n^2$, $ct_y \leftarrow ct - \mathsf{Enc}(\mathsf{pk}, w \cdot q; 0)$, *accept if* $0 \le w < \lfloor n/q \rfloor$ *and* $V_{\mathrm{eq}}(\mathsf{pk}, Y, ct_y, \pi)$.

Note that the evaluation algorithm outputs the quotient w in the clear as part of the evaluation proof. It can do that without compromising the secret y because w is almost independent of y whenever n is sufficiently larger than q. We show this more precisely in the proof of the following theorem. We assume without loss of generality that the n associated with every pk output by $\mathsf{KGen}(1^\lambda)$ is in the set $\left[|\mathbb{G}_\lambda| \cdot 2^\lambda,\ 2^{\gamma\lambda}\right]$, for some universal constant $\gamma > 1$.

Theorem 16. *Suppose that the Paillier scheme is semantically secure, that* $\mathsf{KGen}_{\mathrm{pal}}(1^\lambda)$ *outputs public keys for which the plaintext space has size in the range* $\left[|\mathbb{G}_\lambda| \cdot 2^\lambda,\ 2^{\gamma\lambda}\right]$ *for some universal constant* $\gamma > 1$*, and that* $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$ *is a non-interactive zero-knowledge proof system for* $\mathcal{R}'_{\mathrm{eq}}$*. Then Construction 15 is a secure eVRF with respect to the domain/range ensemble* $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda\in\mathbb{N}}$ *and function-family ensemble* $\{\mathcal{O}_{\mathcal{X}_\lambda, \mathcal{Z}_\lambda}\}_\lambda$*, where* $\mathcal{Z}_\lambda := [0, 2^{(2\gamma+1)\lambda}]$.

Proof. In the following, we fix the security parameter λ and set $q := |\mathbb{G}_\lambda|$. By assumption, if $(\mathsf{pk}, \mathsf{sk})$ is in the support of $\mathsf{KGen}(1^\lambda)$ and $\mathbb{Z}_n$ is the plaintext domain associated with pk, then $n \in [2^\lambda q,\ 2^{\gamma\lambda}]$ for some universal constant $\gamma \in \mathbb{N}$. Set $B := 2^{(2\gamma+1)\lambda}$ so that $B > 2^\lambda n^2$. Then H is a hash function $\mathsf{H} : \mathcal{X}_\lambda \to [0, B]$. Because $B > 2^\lambda n^2$ we know that if z is uniform in $[0, B]$ then $ct \leftarrow (\mathsf{z} \bmod n^2)$ is statistically close to uniform over $\mathbb{Z}_{n^2}$.

We argue that the four eVRF properties hold: consistency, verifiability, simulatability, and pseudo-randomness.

Consistency. Holds by construction, as long as the adversary cannot find an $x \in \mathcal{X}_\lambda$ such that $(\mathsf{H}(x) \bmod n^2)$ is outside of $\mathbb{Z}^*_{n^2}$. But this follows immediately from the sparseness of this set and the fact that $\mathsf{H} : \mathcal{X}_\lambda \to [0, B]$ is a random function.

Verifiability of Eval_2 as a VRF. Fix pk, H, and $x \in \mathcal{X}_\lambda$. Let $ct \leftarrow (\mathsf{H}(x) \bmod n^2)$. By the perfect binding property of $\mathcal{E}_{\mathrm{pal}}$ there is a unique $(\alpha, r) \in [0, n-1] \times \mathcal{R}_{\mathsf{pk}}$ such that $ct = \mathsf{Enc}(\mathsf{pk}, \alpha; r)$. Moreover, there is a unique quotient $0 \le w < \lfloor n/q \rfloor$ such that $y \leftarrow \alpha - w \cdot q$ is in the set $[0, q-1]$. Suppose that the adversary could find a triple (Y', w', π) such that $Y' \neq y \cdot G_\lambda$, but $\mathsf{Verify}^{\mathsf{H}}(\mathsf{pk}, x, Y', (w', \pi))$ accepts. We know that $0 \le w' < \lfloor n/q \rfloor$. Set $ct'_y \leftarrow ct - \mathsf{Enc}(\mathsf{pk}, w' \cdot q; 0)$. Then there are two cases.

- First, if $w = w'$ then ct'_y is an encryption of $y \in [0, q-1]$. But since $Y' \neq y \cdot G_\lambda$, the tuple $(\mathsf{pk}, Y', ct'_y, \pi)$ violates the soundness of $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$.
- Second, if $w \neq w'$ then ct'_y is an encryption of $y' \notin [0, q-1]$. But then again $(\mathsf{pk}, Y', ct'_y, \pi)$ violates the soundness of $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$.

Simulatability as a VRF. Consider the following simulator $\mathsf{Sim}^{\mathsf{H}}(\mathsf{pk}, x, Y)$ that outputs $\big(Y, (w, \pi)\big)$ and works as follows:

1 : $ct \leftarrow (\mathsf{H}(x) \bmod n^2)$
2 : $w \leftarrow\!\$ \left[0, \left\lfloor \frac{n-1}{q} \right\rfloor\right]$ // sample a uniform quotient w
3 : $ct_y \leftarrow ct - \mathsf{Enc}(\mathsf{pk}, w \cdot q; 0)$
4 : Sample a simulated proof π for the $\mathcal{R}'_{\mathrm{eq}}$ statement (pk, Y, ct_y)
// using the zero knowledge simulator for the proof system $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$
5 : Output $\big(Y, (w, \pi)\big)$.

To argue that this is a valid simulator, fix $(x, \mathsf{pk}, \mathsf{sk})$ and choose a uniform random oracle H. Next, define a hybrid (inefficient) simulator $\mathsf{Sim}^{\mathsf{H}}_0(\mathsf{pk}, x, Y, \mathsf{sk})$ that is the same as $\mathsf{Sim}^{\mathsf{H}}(\mathsf{pk}, x, Y)$ except

that we replace Line 2 of Sim with the following:

$2: \quad w \leftarrow \lfloor \mathsf{Dec}(\mathsf{sk}, ct)/q \rfloor$ // compute the correct quotient w

First, we show that there is no PPT distinguisher $D_0^{\mathsf{H},\mathcal{O}_{\mathsf{Sim}}(\cdot)}(\mathsf{pk})$ as in equation (2) of Definition 3 that can distinguish an oracle $\mathcal{O}_{\mathsf{Sim}}(x) := \mathsf{Sim}^{\mathsf{H}}\big(\mathsf{pk}, x, \mathsf{Eval}_3^{\mathsf{H}}(\mathsf{sk}, x)\big)$ from an oracle $\mathcal{O}_{\mathsf{Sim}}(x) := \mathsf{Sim}_0^{\mathsf{H}}\big(\mathsf{pk}, x, \mathsf{Eval}_3^{\mathsf{H}}(\mathsf{sk}, x), \mathsf{sk}\big)$. Here $\mathsf{Eval}_3^{\mathsf{H}}(\mathsf{sk}, x)$ runs $\mathsf{Eval}^{\mathsf{H}}(\mathsf{sk}, x)$ to get $\big(y, Y, (w, \pi)\big)$ and outputs only Y.
Let $D_0^{\mathsf{H},\mathcal{O}_{\mathsf{Sim}}(\cdot)}(\mathsf{pk})$ be such a distinguisher. We construct an adversary $\mathcal{B}_0$ that breaks semantic security of $\mathcal{E}_{\mathrm{pal}}$ with about the same advantage as D_0's distinguishing advantage. This $\mathcal{B}_0$ is given pk as input. It runs $D_0^{\mathsf{H},\mathcal{O}_{\mathsf{Sim}}(\cdot)}(\mathsf{pk})$ and responds to its oracle queries as follows:

whenever D_0 issues a query for $\mathsf{H}(x)$ do:

(1) sample $w_{x,0}, w_{x,1} \leftarrow\!\!\$\ \big[0,\ \lfloor (n-1)/q \rfloor\big]$

(2) sample $y \leftarrow\!\!\$\ [0, q-1]$ and set $Y_x \leftarrow y \cdot G_\lambda$

(3) set $\alpha_0 := w_{x,0} \cdot q + y$ and $\alpha_1 := w_{x,1} \cdot q + y$ // both are close to uniform in $[0, n-1]$

(4) $\mathcal{B}_0$ asks its semantic security challenger to encrypt α_0 or α_1,
gets back $ct_x \leftarrow\!\!\$\ \mathsf{Enc}(\mathsf{pk}, \alpha_b)$ for some $b \in \{0,1\}$

(4) return $z_x := I_{\mathrm{pal}}^{-1}(\mathsf{pk}, ct_x)$ to D_0 meaning that $\mathsf{H}(x) := z_x$

whenever D_0 issues a query for $\mathcal{O}_{\mathsf{Sim}}(x)$ do:

(1) $ct_y \leftarrow ct_x - \mathsf{Enc}(\mathsf{pk}, w_{x,0} \cdot q; 0)$ // $ct_x, w_{x,0}, Y_x$ were generated during a query for $\mathsf{H}(x)$

(2) sample a simulated proof π for the $\mathcal{R}'_{\mathrm{eq}}$ statement (pk, Y_x, ct_y)

(3) return $(Y_x, (w_{x,0}, \pi))$ to D_0

eventually D_0 outputs a bit $b' \in \{0,1\}$ and $\mathcal{B}_0$ outputs the same b'

Observe that when $b = 1$ then $\mathcal{B}_0$ is simulating an oracle $\mathcal{O}_{\mathsf{Sim}} = \mathsf{Sim}$. When $b = 0$ then $\mathcal{B}_0$ is simulating an oracle $\mathcal{O}_{\mathsf{Sim}} = \mathsf{Sim}_0$. Therefore, $\mathcal{B}_0$'s advantage in breaking semantic security of $\mathcal{E}_{\mathrm{pal}}$ is the same as D_0's advantage in distinguishing the two oracles. Hence, if $\mathcal{E}_{\mathrm{pal}}$ is semantically secure, then the two oracles are indistinguishable.
Second, suppose that an adversary $D_1^{\mathsf{H},\mathcal{O}(\cdot)}(\mathsf{pk})$ could distinguish an oracle $\mathcal{O}(x)$ for the function $\mathsf{Sim}_0^{\mathsf{H}}\big(\mathsf{pk}, x, \mathsf{Eval}_3^{\mathsf{H}}(\mathsf{sk}, x), \mathsf{sk}\big)$ from an oracle $\mathcal{O}(x)$ for $\mathsf{Eval}_2^{\mathsf{H}}(\mathsf{sk}, x)$. The only difference between these two oracles is that Sim_0 outputs a simulated proof, while Eval_2 outputs a real proof. Hence, if D_1 had non-negligible advantage, then it would break the ZK simulator of $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$. Since the ZK simulator of $(P_{\mathrm{eq}}, V_{\mathrm{eq}})$ is a valid simulator, it follows that an oracle for Sim_0 is indistinguishable from an oracle for the real Eval_2.
Putting these two steps together, we obtain that no PPT oracle-aided distinguisher can distinguish an oracle for $\mathcal{O}_{\mathsf{Sim}}(x) := \mathsf{Sim}^{\mathsf{H}}\big(\mathsf{pk}, x, \mathsf{Eval}_3^{\mathsf{H}}(\mathsf{sk}, x)\big)$ from an oracle for $\mathcal{O}(x) := \mathsf{Eval}_2^{\mathsf{H}}(\mathsf{sk}, x)$, as required by equation (2) of Definition 3. This completes the proof of simulatability.

Pseudorandomness of Eval_1. The proof from Theorem 14 carries over with only minor changes. The only difference is that we replace every response ct to a random oracle query with $I_{\mathrm{pal}}^{-1}(\mathsf{pk}, ct)$. Similarly, responses to evaluation queries are first reduced modulo q.

This completes the proof of the theorem. □

Knowledge of secret key. Uplifting the security of the scheme from a game-based one (Definition 4) to realizing the ideal functionality (Definition 6), requires a ZK-POK proof for the relation $\mathcal{R}_{\mathrm{pub}}$ from (3). Since the Paillier secret key is easy to deduce from the factors of the public key, one can use the ZK-POK proof from [16, Section 6.3.1].

6 A DDH-Based eVRF

In this section we show how to construct an eVRF from a classic secure PRF based on the Decision Diffie-Hellman (DDH) assumption. A variant of the eVRF in this section was previously used implicitly in MuSig-DN [54,55]. We first review the classic DDH PRF.

The DDH PRF. This PRF $F_{\text{DDH}} = (\mathsf{KGen}, \mathsf{Eval})$ is defined with respect to an ensemble of domains and ranges $\{(\mathcal{X}_\lambda, \mathbb{G}_\lambda)\}_{\lambda=1}^{\infty}$, where $\mathbb{G}_\lambda$ is a group of some prime order $s(\lambda)$ with generator $G_\lambda \in \mathbb{G}_\lambda$. In addition, the PRF uses a function-family ensemble $\{\mathsf{H}_\lambda : \mathcal{X}_\lambda \to \mathbb{G}_\lambda\}_{\lambda=1}^{\infty}$ and works as follows

- $\mathsf{KGen}(1^\lambda) \to k$: output $k \leftarrow\!\!\$\, \mathbb{Z}_{s(\lambda)}$.
- $\mathsf{Eval}^{\mathsf{H}_\lambda}(k, x) \to y$: for $x \in \mathcal{X}_\lambda$ output $y \leftarrow k \cdot \mathsf{H}_\lambda(x) \in \mathbb{G}_\lambda$.

Naor, Pinkas, and Reingold [52] showed that F_{DDH} is a secure PRF whenever DDH holds in the groups $\{\mathbb{G}_\lambda\}_{\lambda=1}^{\infty}$ and H_λ is sampled unifomly from $\mathcal{O}_{\mathcal{X}_\lambda, \mathbb{G}_\lambda}$, that is, H_λ is modeled as a random oracle. Papadopoulos et al. [59] observe that this PRF can be made into a VRF by publishing $\mathsf{vk} := k \cdot G_\lambda$, and attaching to every evaluation $y \leftarrow \mathsf{Eval}(k, x)$ a non-interactive zero-knowledge proof π that $\big(G_\lambda, \mathsf{vk}, \mathsf{H}_\lambda(x), y\big)$ is a DDH tuple.

From here on, when there is no confusion, we will drop the index λ and simply refer to the group $\mathbb{G}_\lambda$ as $\mathbb{G}$. We use s to denote its order and G its generator.

We construct an eVRF by embedding the output of the DDH PRF "in the exponent" of another group. To do so, we will need an explicit representation of the group $\mathbb{G}$ as an elliptic curve group. Such groups are parameterized by a prime field $\mathbb{F}_q$, where $q = q(\lambda)$, along with two scalars $a, b \in \mathbb{F}_q$. The group is the set of all pairs $(x, y) \in \mathbb{F}_q^2$ such that $y^2 = x^3 + ax + b$, along with the point at infinity. The group operation is defined using the cord-and-tangent rule discussed below [8].

Since we embed the group $\mathbb{G}$ in the exponent of another group, it is convenient to introduce the following notion of a group pair, or more precisely, a group pair ensemble.

Definition 9. *We say that $\big\{(\mathbb{G}_{\text{S}}^{(\lambda)}, \mathbb{G}_{\text{T}}^{(\lambda)})\big\}_{\lambda=1}^{\infty}$ is a **group pair ensemble** if for every $(\mathbb{G}_{\text{S}}, \mathbb{G}_{\text{T}})$ in the ensemble we have that*

- *The group $\mathbb{G}_{\text{T}}$, called the **target group**, has some prime order q. We use $\boldsymbol{G} = (G_{\text{T},1}, \ldots, G_{\text{T,n}})$ to denote a vector of n generators in $\mathbb{G}_{\text{T}}$. For a vector $\boldsymbol{x} \in \mathbb{F}_q^n$ we write*

$$\langle \boldsymbol{x}, \boldsymbol{G} \rangle = x_1 G_{\text{T},1} + \cdots + x_n G_{\text{T,n}} \in \mathbb{G}_{\text{T}}.$$

- *The group $\mathbb{G}_{\text{S}}$, called the **source group**, is a group of some prime order s with generator $G_{\text{S}} \in \mathbb{G}_{\text{S}}$. This group $\mathbb{G}_{\text{S}}$ is a group of points of an elliptic curve defined over the field $\mathbb{F}_q$, where q is the order of $\mathbb{G}_{\text{T}}$. Let $\mathbb{G}_{\text{S}}^* := \mathbb{G}_{\text{S}} \setminus \{0\}$. Elements in $\mathbb{G}_{\text{S}}^*$ are represented as pairs in $\mathbb{F}_q^2$ so that $\mathbb{G}_{\text{S}}^* \subseteq \mathbb{F}_q^2$.*

*We say that the group pair ensemble is **secure** if DLOG holds in $\{\mathbb{G}_{\text{T}}^{(\lambda)}\}_{\lambda=1}^{\infty}$ and DDH holds in $\{\mathbb{G}_{\text{S}}^{(\lambda)}\}_{\lambda=1}^{\infty}$. To simplify the notation, we will often drop the index λ and say that $(\mathbb{G}_{\text{S}}, \mathbb{G}_{\text{T}})$ is a **group pair** or a **secure group pair**.*

In our eVRF construction, we will use the source group, $\mathbb{G}_{\text{S}}$, for the output of the DDH PRF. We will use the target group, $\mathbb{G}_{\text{T}}$, to hide the PRF output "in the exponent" of an element in $\mathbb{G}_{\text{T}}$. To do so, we need an explicit representation of elements in $\mathbb{G}_{\text{S}}$. Requiring that $\mathbb{G}_{\text{S}}$ is an elliptic curve group is sufficient.

The challenge is to design a proof system that proves that the DDH PRF was evaluated correctly, despite being given the value of the PRF in the exponent. Towards this goal, we will make use of the following instance-witness relation $\mathcal{R}_{\mathsf{H}}$ defined with respect to a function $\mathsf{H} : \mathcal{X} \to \mathbb{G}_{\mathrm{S}}$ as

$$\begin{aligned}\mathcal{R}_{\mathsf{H}} := \Big\{ \big((Q, x, Y)\ ;\ k\big) \Big\} \subseteq (\mathbb{G}_{\mathrm{T}} \times \mathcal{X} \times \mathbb{G}_{\mathrm{T}}) \times [s-1] \quad \text{where} \\ (1)\ \ Q = k \cdot G_{\mathrm{T},1}, \\ (2)\ \ Y = x_P \cdot G_{\mathrm{T},2} \text{ for } P = (x_P, y_P) := k \cdot \mathsf{H}(x) \in \mathbb{G}_{\mathrm{S}}^* \subseteq \mathbb{F}_q^2.\end{aligned} \tag{6}$$

Here we are treating $G_{\mathrm{T},1}$ and $G_{\mathrm{T},2}$ as generators of $\mathbb{G}_{\mathrm{T}}$ that are part of the description of the group. As usual, we let $\mathcal{L}_{\mathcal{R}_{\mathsf{H}}}$ denote the language of all triples (Q, x, Y) for which there exists a witness $k \in [s-1]$ such that $\mathcal{R}_{\mathsf{H}}\big((Q, x, Y); k\big)$ is true.

In addition, to uplift the security of our eVRF from a game-based one (Definition 4) to realizing the ideal functionality (Definition 6), we require a non-interactive zero-knowledge proof-of-knowledge (ZK-POK) for the discrete log relation

$$\mathcal{R}_{\mathrm{dlog}} := \Big\{ \big((Q, G)\ ;\ k\big)\ :\ Q = k \cdot G \Big\} \subseteq \mathbb{G}_{\mathrm{T}}^2 \times [s-1] \tag{7}$$

One can use Schnorr's protocol [61] for $\mathcal{R}_{\mathrm{dlog}}$. However, to realize the ideal functionality we require a straight line (non-rewinding) knowledge extractor, for example, as presented by Fischlin [28].

The basic DDH eVRF. We can now present our eVRF. The construction uses a non-interactive zero-knowledge argument system (P, V) for the relation $\mathcal{R}_{\mathsf{H}}$ from (6). We will present the required argument system in Section 6.1. It also uses the Schnorr ZK-POK $(\mathsf{P}_{\mathrm{dlog}}, \mathsf{V}_{\mathrm{dlog}})$ for the relation $\mathcal{R}_{\mathrm{dlog}}$ from (7). As we explain below, the range of this eVRF is only about half of $\mathbb{G}_{\mathrm{T}}$. In Section 6.2 we enhance this eVRF so that its range is the full group $\mathbb{G}_{\mathrm{T}}$.

Construction 17 (The basic DDH-based eVRF) *Let $(\mathbb{G}_{\mathrm{S}}, \mathbb{G}_{\mathrm{T}})$ be a group pair where s is the size of $\mathbb{G}_{\mathrm{S}}$ and q is the size of $\mathbb{G}_{\mathrm{T}}$. Let $G_{\mathrm{T},1}, G_{\mathrm{T},2}$ be two generators of $\mathbb{G}_{\mathrm{T}}$. Let H be a function $\mathsf{H} : \mathcal{X} \to \mathbb{G}_{\mathrm{S}}^*$, where $\mathbb{G}_{\mathrm{S}}^* := \mathbb{G}_{\mathrm{S}} \setminus \{0\}$. The **DDH eVRF** with domain $\mathcal{X}$ and range $\mathbb{G}_{\mathrm{T}}$ is defined as follows:*

- $\mathsf{KGen}^{\mathsf{H}}(1^\lambda)$*: Sample $k \leftarrow\!\!\$\ [s-1]$ and set $Q \leftarrow k \cdot G_{\mathrm{T},1}$. Use the prover P_{dlog} for $\mathcal{R}_{\mathrm{dlog}}$ to generate a proof $\pi_Q \leftarrow\!\!\$\ \mathsf{P}_{dlog}(Q, G_{\mathrm{T},1}; k)$ and set $\mathsf{vk} := (Q, \pi_Q)$. Output (k, vk).*
- $\mathsf{Eval}^{\mathsf{H}}(k, x)$*: for $x \in \mathcal{X}$, let $P \leftarrow k \cdot \mathsf{H}(x) \in \mathbb{G}_{\mathrm{S}}^*$.*

  $$\text{with } P = (x_P, y_P) \subseteq \mathbb{F}_q^2 \ \text{ set } \ y \leftarrow x_P \in \mathbb{F}_q \quad \text{and} \quad Y \leftarrow y \cdot G_{\mathrm{T},2} \in \mathbb{G}_{\mathrm{T}}.$$

  Next, run the prover P for $\mathcal{R}_{\mathsf{H}}$ to construct a proof π that the triple $\big(\mathsf{vk}, x, Y\big)$ is in the language $\mathcal{L}_{\mathcal{R}_{\mathsf{H}}}$ from (6). The proof system for $\mathcal{R}_{\mathsf{H}}$ is described in Section 6.1. Output (y, Y, π). Recall that $\mathsf{Eval}_1^{\mathsf{H}}(k, x)$ is the same as algorithm $\mathsf{Eval}^{\mathsf{H}}(k, x)$ but only outputs $y \in \mathbb{F}_q$.
- $\mathsf{Verify}^{\mathsf{H}}(\mathsf{vk}, x, Y, \pi)$*: for $\mathsf{vk} = (Q, \pi_Q)$, the algorithm accepts if (1) π is a valid proof that (Q, x, Y) is in $\mathcal{L}_{\mathcal{R}_{\mathsf{H}}}$, and (2) π_Q is a valid proof for $(Q, G_{\mathrm{T},1})$ as an instance of $\mathcal{R}_{\mathrm{dlog}}$.*

At the heart of this eVRF construction is the non-interactive zero-knowledge proof for the language $\mathcal{L}_{\mathcal{R}_{\mathsf{H}}}$. Before we develop this proof system, let us first briefly argue that the eVRF is secure when $\mathsf{H} : \mathcal{X} \to \mathbb{G}_{\mathrm{S}}^*$ is modeled as a random oracle.

Algorithm $\mathsf{Eval}_1^{\mathsf{H}}(k, x)$ in Construction 17 outputs the x-coordinate of a point P in $\mathbb{G}_{\mathrm{S}}^*$. Let $\mathcal{S} \subseteq \mathbb{F}_q$ be the set of all x-coordinates of points in $\mathbb{G}_{\mathrm{S}}^*$. Then the size of $\mathcal{S}$ is about half the size of $\mathbb{F}_q$. Consequently, the range of the basic DDH eVRF is about half of $\mathbb{G}_{\mathrm{T}}$. The following theorem shows that it is a secure eVRF with respect to this subset. In Section 6.2 we enhance this basic eVRF so that its range is the full group $\mathbb{G}_{\mathrm{T}}$.

Theorem 18 (subset eVRF security). *Let $(\mathbb{G}_{\mathrm{S}}, \mathbb{G}_{\mathrm{T}})$ be a secure group pair, so that DDH holds in $\mathbb{G}_{\mathrm{S}}$. Let (P, V) be a non-interactive zero-knowledge argument system for the relation $\mathcal{R}_{\mathsf{H}}$ from (6). Let $\mathbb{G}_{\mathrm{S}}$ be defined over $\mathbb{F}_q$ and let $\mathcal{S} \subseteq \mathbb{F}_q$ be the set of all x-coordinates of points in $\mathbb{G}_{\mathrm{S}}^*$. Then the eVRF in Construction 17 is a subset secure eVRF (as in Definition 5) with respect to the domain/range $(\mathcal{X}, \mathbb{G}_{\mathrm{T}})$, subset $\mathcal{S}$, and function family $\mathcal{O}_{\mathcal{X},\mathbb{G}_{\mathrm{S}}^*}$.*

Proof. *Consistency* holds by construction. *Pseudorandomness* follows from the fact that the DDH PRF is a secure PRF when DDH holds in $\mathbb{G}_{\mathrm{S}}$ and H is sampled at random from $\mathcal{O}_{\mathcal{X},\mathbb{G}_{\mathrm{S}}^*}$. Therefore, the x-coordinate of the output of the DDH PRF is pseudorandom over the subset $\mathcal{S} \subseteq \mathbb{F}_q$. *Verifiability as a VRF* follows from the soundness of the proof system for $\mathcal{R}_{\mathsf{H}}$. *Simulatability as a VRF* follows from the zero knowledge property of the proof system for $\mathcal{R}_{\mathsf{H}}$. □

Looking ahead, soundness of our proof system for $\mathcal{R}_{\mathsf{H}}$ relies on the hardness of discrete log in $\mathbb{G}_{\mathrm{T}}$ and the knowledge soundness of the proof system for $\mathcal{R}_{\mathrm{dlog}}$. It also relies on $G_{\mathrm{T},1}, G_{\mathrm{T},2}$ being random generators of $\mathbb{G}_{\mathrm{T}}$ so that there is no known discrete log relation between them.

6.1 An Argument System for the Relation $\mathcal{R}_{\mathsf{H}}$

To complete the construction, we need an efficient non-interactive zero-knowledge argument for the relation $\mathcal{R}_{\mathsf{H}}$ from (6). There are several ways to proceed. One option is to use a generic zkSNARK [6] to produce a succinct proof. However, since $\mathcal{R}_{\mathsf{H}}$ uses arithmetic in both $\mathbb{G}_{\mathrm{S}}$ and $\mathbb{G}_{\mathrm{T}}$, this will require non-native arithmetic in the zkSNARK, which will result in an efficient prover.

Another option is to use the Bulletproofs argument system [9,12], which is especially well suited for proving statements about $\mathbb{F}_q$ field elements that are given "in the exponent." This is precisely what is needed for the relation $\mathcal{R}_{\mathsf{H}}$: the verifier is given k and x' in the exponent — they are provided as $Q = k \cdot G_{\mathrm{T},1}$ and $Y = x_P \cdot G_{\mathrm{T},2}$ — along with $x \in \mathcal{X}$, and we need a proof that x_P is the x-coordinate of $P := k \cdot \mathsf{H}(x) \in \mathbb{G}_{\mathrm{S}}$. We show how to use Bulletproofs to construct an efficient argument system for $\mathcal{R}_{\mathsf{H}}$. The resulting the proof size is only $O(\log \log s)$ group elements, and prover and verifier times are dominated by a $O(\log s)$ multi-scalar multiplication in $\mathbb{G}_{\mathrm{T}}$.

A brief overview of Bulletproofs. Bünz [11, §2.6] shows that Bulletproofs give an argument system for a rank-1 constraint system (R1CS), when the statement is given in the exponent. Specifically, Bulletproofs is well suited as an argument system for the following exponent R1CS relation:

$$\begin{aligned} \mathcal{R}_{\mathrm{eR1CS}} := \Big\{ (A, B, C, T) \ ; \ (\boldsymbol{x}, \boldsymbol{w}) \Big\} \quad \text{where} \\ (1) \quad A, B, C \in \mathbb{F}_q^{n \times m}, \quad T \in \mathbb{G}_{\mathrm{T}}, \quad \boldsymbol{x} \in \mathbb{F}_q^r, \quad \boldsymbol{w} \in \mathbb{F}_q^{m-r}, \\ (2) \quad (A\boldsymbol{z}) \circ (B\boldsymbol{z}) = (C\boldsymbol{z}) \quad \text{where} \quad \boldsymbol{z} := (\boldsymbol{x}, \boldsymbol{w}) \in \mathbb{F}_q^m, \\ (3) \quad T = \langle \boldsymbol{x}, \boldsymbol{G} \rangle. \end{aligned} \tag{8}$$

Here $\boldsymbol{G} \in \mathbb{G}_{\mathrm{T}}^r$ is a public tuple of r generators of $\mathbb{G}_{\mathrm{T}}$. The notation $\boldsymbol{u} \circ \boldsymbol{v}$ used on line (2) refers to the component-wise multiplication of the vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ (also called the Hadamard product of $\boldsymbol{u}$ and $\boldsymbol{v}$).

Note that the R1CS statement $\boldsymbol{x} \in \mathbb{F}_q^r$ is provided as a Pedersen commitment $T = \langle \boldsymbol{x}, \boldsymbol{G} \rangle$, exactly as in our settings. Indeed, for an $\mathcal{R}_{\mathsf{H}}$-instance (Q, x, Y) we will set $T := Q + Y \in \mathbb{G}_{\mathrm{T}}$. The vector $\boldsymbol{z}$ on line (2) is often called the **extended witness**. Each row in the matrices A, B, C is called a **constraint**, so that the R1CS above has n constraints. In the Bulletproofs argument, the prover sends to the verifier a Pedersen commitment $T' := \langle \hat{\boldsymbol{z}}, \boldsymbol{G}_z \rangle$ for a vector $\hat{\boldsymbol{z}}$ derived from $\boldsymbol{z}$. An inner-product argument is then used to prove that $\boldsymbol{z} := (\boldsymbol{x}, \boldsymbol{w})$ satisfies the condition on line (2).

The Bulletproofs argument system is complete, computationally sound, and zero knowledge. Here computational soundness means that either the system is sound, or there is a PPT algorithm that can find a non-trivial linear relation among the generators in $(\boldsymbol{G}, \boldsymbol{G}_z)$. The latter implies that discrete log is easy in $\mathbb{G}_{\mathrm{T}}$. The argument system can be made non-interactive using the Fiat-Shamir transform, and retains its soundness and zero knowledge properties in the random oracle model [2].

The length of the proof is $2\log_2(n+m) + 4$ group elements in $\mathbb{G}_{\mathrm{T}}$. The running times of the verifier is dominated by the time to compute a multi-scalar multiplication (MSM) for a vector dimension about $2(n+m)$. Using Pippenger's algorithm [60], computing such an MSM is faster than computing the exponentiations one by one. In addition, the Bulletproofs verifier can batch verify multiple proofs at once much faster than verifying the proofs one at a time [12]. We summarize these facts in the following theorem.

Theorem 19 ([11]). *Bulletproofs* $(\mathsf{P}_{\mathrm{BP}}, \mathsf{V}_{\mathrm{BP}})$ *is a zero knowledge non-interactive argument system for the relation* $\mathcal{R}_{eR1CS}$ *in the random oracle model, assuming discrete log in* $\mathbb{G}_{\mathrm{T}}$ *is hard. The length of the proof is* $2\lceil \log_2(n+m) \rceil + 3$ *group elements in* $\mathbb{G}_{\mathrm{T}}$.

The proof system for $\mathcal{R}_{\mathsf{H}}$. Given the above, it remains to design an efficient rank-1 constraint system (R1CS) — namely three matrices $A, B, C \in \mathbb{F}_q^{n \times m}$ — for the relation $\mathcal{R}_{\mathsf{H}}$ from (6). Consider an $\mathcal{R}_{\mathsf{H}}$ instance $\big((Q, x, Y)\ ;\ k\big)$ where k is in $[s-1]$. Let $(k_0, k_1, \ldots, k_\ell) \in \{0,1\}^{\ell+1}$ be the binary representation of k, so that $k = \sum_{i=0}^{\ell} 2^i \cdot k_i$. In addition, we use a fixed sequence of elements $c_0, \ldots, c_\ell \in \mathbb{Z}_s$ satisfying

$$\forall i \in [\ell - 1] : \quad \sum_{j=0}^{i-1} c_j \notin \{0, \pm c_i\} \qquad \text{and} \qquad \sum_{j=0}^{\ell} c_j = 0. \tag{9}$$

For example, when $s > \ell^2$ one can set $c_i := (i+2)$ for $i = 0, \ldots, \ell - 1$ and $c_\ell := -\binom{\ell+2}{2} + 1$.

Let $X := \mathsf{H}(x)$. Our plan for proving that (Q, x, Y) is in $\mathcal{L}_{\mathcal{R}_{\mathsf{H}}}$ is to have the prover compute a sequence of elements $P_0, \ldots, P_\ell \in \mathbb{G}_{\mathrm{S}}$ defined by

$$\begin{cases} P_0 := k_0 \cdot X + c_0 \cdot G_{\mathrm{S}} \\ P_i := P_{i-1} + \Delta_i \quad \text{where} \quad \Delta_i := (2^i k_i) \cdot X + c_i \cdot G_{\mathrm{S}} \quad (\text{for } i = 1, \ldots, \ell). \end{cases} \tag{10}$$

These points will become part of the extended witness $\boldsymbol{z}$ for our R1CS program. Observe that because $\sum_{i=0}^{\ell} 2^i k_i = k$ and $\sum_{j=0}^{\ell} c_j = 0$, the final point P_ℓ satisfies $P_\ell = k \cdot X \in \mathbb{G}_{\mathrm{S}}$. Now the R1CS program only needs to check that if $Y = x \cdot G_{\mathrm{T},2}$, then the x-coordinate of P_ℓ is equal to x, which is just one constraint in R1CS. The main challenge is to verify that the point P_ℓ was constructed correctly. We will do so by verifying inductively that P_i is correct given that all $P_0, \ldots, P_{i-1}$ are correct (for all $i = 1, \ldots, \ell$).

Looking ahead, the purpose of the term $c_i G_{\mathrm{S}}$ in the definition of Δ_i in (10) is to ensure that Δ_i is not equal to $\pm P_{i-1}$. This ensures that none of the P_i are the point at infinity in $\mathbb{G}_{\mathrm{S}}$, and that

the addition always adds distinct points in $\mathbb{G}_S$. This lets us always use the "cord" rule for addition, and never need the "tangent" rule.

Let us describe the R1CS program for $\mathcal{L}_{\mathcal{R}_H}$ in more detail. We begin by describing the extended witness $\boldsymbol{z}$. Recall that $(k_0, k_1, \ldots, k_\ell)$ is the binary representation of the secret key $k \in [s-1]$. For $i = 0, \ldots, \ell$ we let $P_i = (x_{P_i}, y_{P_i}) \in \mathbb{G}_S^* \subseteq \mathbb{F}_q^2$ be the points constructed in (10). Further, define

$$\boldsymbol{w}_i := \Big(k_i,\ x_{P_i}, x_{P_i}^2, x_{P_i}^3,\ y_{P_i}, y_{P_i}^2,\ t_1, t_2\Big)^\intercal \in \mathbb{F}_q^8 \quad \text{for } i = 1, \ldots, \ell. \tag{11}$$

We will explain what are $t_1, t_2 \in \mathbb{F}_q$ soon below. Then, with $P = (x_P, y_P) = kX \in \mathbb{G}_S$, the extended witness is defined as the column vector

$$\boldsymbol{z} := \Big(\underbrace{1, k, x_P}_{\substack{\text{the R1CS} \\ \text{statement}}},\ \ k_0, x_{P_0}, y_{P_0},\ \ \boldsymbol{w}_1, \ldots, \boldsymbol{w}_\ell\Big)^\intercal \in \mathbb{F}_q^{8\ell+6}. \tag{12}$$

Now, the R1CS program, which consists of three matrices $A, B, C \in \mathbb{F}_q^{n \times m}$, needs to verify three claims:

- *claim 1:* $k_0, \ldots, k_\ell$ is the binary representation of k, that is $k = \sum_{i=0}^{\ell} 2^i k_i$ and $k_i \in \{0, 1\}$ for $i = 0, \ldots, \ell$.
- *claim 2:* The points $P_0, \ldots, P_\ell$ are constructed according to (10).
- *claim 3:* The point $P_\ell = (x_{P_\ell}, y_{P_\ell})$ satisfies $x_{P_\ell} = x_P$, where $P := kX$.

If all three claims hold then the verifier is convinced that (Q, x, Y) is in $\mathcal{L}_{\mathcal{R}_H}$.

Claims 1 and 3 are easy to check in an R1CS: checking that $k = \sum_{i=0}^{\ell} 2^i k_i$ takes one constraint (i.e., one row in A, B, C); checking that $k_i \in \{0, 1\}$ for $i = 0, \ldots, \ell$ is done by checking that $k_i(1 - k_i) = 0$, and this takes $\ell + 1$ constraints, one for each k_i; checking that $x_{P_\ell} = x_P$ takes one constraint. Hence, Claims 1 and 3 require a total of $\ell + 3$ constraints in A, B, C.

Checking Claim 2 in R1CS is more interesting. First, note that while the verifier has $X, G_S \in \mathbb{G}_S$, it does not know the bits k_i of k and therefore cannot construct the points Δ_i in (10) by itself. However, we observe that given X and G_S, all the $\Delta_i \in \mathbb{G}_S$ can be expressed as a public linear function of k_i. Indeed, since k_i is binary, we know that Δ_i takes one of two values:

$$\Delta_i = \Delta := 2^i X + c_i G_S \qquad \text{or} \qquad \Delta_i = \Delta' := c_i G_S.$$

Let $(x, y) \in \mathbb{F}_q^2$ be the elliptic curve point representing Δ and let $(x', y') \in \mathbb{F}_q^2$ be the point representing Δ'. Then we can express Δ_i as

$$\Delta_i = k_i(x - x', y - y') + (x', y') = (k_i \delta_x + x',\ \ k_i \delta_y + y') \in \mathbb{F}_q^2 \tag{13}$$

where $\delta_x := x - x'$ and $\delta_y := y - y'$. The verifier can construct $x', y', \delta_x, \delta_y \in \mathbb{F}_q$ on its own. Hence, Δ_i can be expressed as a public linear function of k_i. Since $P_o = \Delta_0$ this method also lets us express P_0 as a public linear function of k_0.

Now, to verify claim 2, the R1CS program will verify that $P_i = P_{i-1} + \Delta_i$ for all $i = 1, \ldots, \ell$. The program does so by checking two things:

– First, for $i = 1, \ldots, \ell$ verify that the point $P_i = (x_{P_i}, y_{P_i})$ in $\boldsymbol{z}$ is a point in $\mathbb{G}_{\mathrm{S}}^*$. That is, $(x_{P_i}, y_{P_i}) \in \mathbb{F}_q^2$ satisfies $y_{P_i}^2 = x_{P_i}^3 + a x_{P_i} + b$ where a and b are the constants that define the curve $\mathbb{G}_{\mathrm{S}}$. This takes four constraints in the matrices A, B, C: two constraints to verify that $x_{P_i}^3$ in $\boldsymbol{z}$ is computed correctly (i.e., it is the cube of x_{P_i}); one constraint to verify that $y_{P_i}^2$ in $\boldsymbol{z}$ is computed correctly (i.e., it is the square of y_{P_i}); and one linear constraint to verify the elliptic curve relation $y_{P_i}^2 = x_{P_i}^3 + a x_{P_i} + b$.
– Second, for $i = 1, \ldots, \ell$ verify that the points P_{i-1}, Δ_i and $-P_i = (x_{P_i}, -y_{P_i})$ are co-linear. When $P_{i-1} \neq \pm\Delta_i$ this is sufficient to prove that $P_i = P_{i-1} + \Delta_i$ in $\mathbb{G}_{\mathrm{S}}$. Moreover, we show in Theorem 21 that indeed $P_{i-1} \neq \pm\Delta_i$ must always hold. To check that P_{i-1}, Δ_i and $-P_i$ are co-linear, the R1CS program verifies that the slope of the line through P_{i-1} and $-P_i$ is equal to the slope of the line through Δ_i and $-P_i$. That is, the program verifies that

$$(y_{P_{i-1}} + y_{P_i}) \cdot (x_{\Delta_i} - x_{P_i}) = (y_{\Delta_i} + y_{P_i}) \cdot (x_{P_{i-1}} - x_{P_i}) \tag{14}$$

  This is where the values t_1 and t_2 in (11) are used. The R1CS verifies that t_1 is equal to the left hand side of (14), that t_2 is equal to the right hand side, and that $t_1 = t_2$. This takes a total of three constraints in A, B, C. Recall that x_{Δ_i} and y_{Δ_i} used in (14) are expressed as a linear function of k_i using (13).

The explicit matrices A, B, C are shown in Figure 3 in the appendix. Together, the two checks above prove that $P_1, \ldots, P_\ell$ in $\boldsymbol{z}$ are computed as in (10). It remains to verify that P_0 is constructed correctly, and this is done using (13), which takes two constraints.

Putting it all together. We can now describe the complete proof system for $\mathcal{R}_{\mathsf{H}}$ from (6). The proof system uses a proof system $(\mathsf{P}_{\mathrm{BP}}, \mathsf{V}_{\mathrm{BP}})$ for the relation $\mathcal{R}_{\mathrm{eR1CS}}$ from (8), and a proof system $(\mathsf{P}_{\mathrm{dlog}}, \mathsf{V}_{\mathrm{dlog}})$ for the relation $\mathcal{R}_{\mathrm{dlog}}$ from (7).

Construction 20 *Let $\big((Q, x, Y); k\big)$ be an $\mathcal{R}_{\mathsf{H}}$ instance-witness pair with respect to a hash function $\mathsf{H} : \mathcal{X} \to \mathbb{G}_{\mathrm{S}}$, so that $Q = k \cdot G_{\mathrm{T},1}$ and $Y = x_P \cdot G_{\mathrm{T},2}$, as in (6). The proof system (P, V) for $\mathcal{R}_{\mathsf{H}}$ works as follows:*

– *Prover $\mathsf{P}^{\mathsf{H}}(Q, x, Y; k)$: Calculate the R1CS matrices (A, B, C) constructed from $\mathsf{H}(x)$ and G_{S} as described above. Set $T := Q + Y$ and $\boldsymbol{x} := (k, x_P) \in [s-1] \times \mathbb{F}_q$. Calculate the witness $\boldsymbol{w}$ for the $\mathcal{R}_{eR1CS}$ instance (A, B, C, T), and generate proofs*

$$\pi_0 \leftarrow\!\$\ \mathsf{P}_{\mathrm{BP}}\big(A, B, C, T; \boldsymbol{x}, \boldsymbol{w}\big), \quad \pi_Q \leftarrow\!\$\ \mathsf{P}_{\mathrm{dlog}}(Q, G_{\mathrm{T},1}; k), \quad \pi_Y \leftarrow\!\$\ \mathsf{P}_{\mathrm{dlog}}(Y, G_{\mathrm{T},2}; x_P).$$

  Output $\pi := (\pi_0, \pi_Q, \pi_Y)$.

– *Verifier $\mathsf{V}^{\mathsf{H}}(Q, x, Y; \pi_0, \pi_Q, \pi_Y)$: Calculate the R1CS matrices (A, B, C) constructed from $\mathsf{H}(x)$ and G_{S} as described above. set $T := Q + Y$ and accept if $\mathsf{V}_{\mathrm{BP}}\big(A, B, C, T; \pi_0\big)$ and $\mathsf{V}_{\mathrm{dlog}}\big(Q, G_{\mathrm{T},1}; \pi_Q\big)$ and $\mathsf{V}_{\mathrm{dlog}}\big(Y, G_{\mathrm{T},2}; \pi_Y\big)$.*

The running time of the proof system is determined by the dimensions of the matrices A, B, C. All the linear constraints in the R1CS (A, B, C) can be collapsed into a single constraint by taking a random linear combination of the constraints using verifier randomness. Consequently, the total number of the constraints in our R1CS is: $(\ell+1)$ constraints for claims 1 and 3; another 3ℓ constraints to verify that all P_i are in $\mathbb{G}_{\mathrm{S}}^*$; and 2ℓ constraints to verify co-linearity. This is a total for $6\ell + 1$ constraints, plus one more constraint to verify all the linear constraints at once. An observation in Figure 3 in the appendix reduces the number of constraints to $5\ell + 2$. Therefore, the matrices $A, B, C \in \mathbb{F}_q^{n \times m}$ have dimension $n = 5\ell + 2$ and $m = 8\ell + 6$.

Remark 1 (an optimization). It is not difficult to generalize (13) and process two bits of the key k at every iteration, instead of one bit as in (13). This will halve the number of iterations, at the cost of two additional constraints per iteration.

Theorem 21. *Suppose that discrete log is hard in both $\mathbb{G}_S$ and $\mathbb{G}_T$, and that $G_{T,1}$ and $G_{T,2}$ are independent generators of $\mathbb{G}_T$. Moreover, suppose that the proof system $(\mathsf{P}_{BP}, \mathsf{V}_{BP})$ is a non-interactive zero-knowledge argument for $\mathcal{R}_{eR1CS}$, and $(\mathsf{P}_{dlog}, \mathsf{V}_{dlog})$ is a non-interactive zero-knowledge proof of knowledge for the relation $\mathcal{R}_{dlog}$. Then the proof system (P, V) from Construction 20 is a non-interactive zero knowledge argument system for $\mathcal{R}_{\mathsf{H}}$, when the hash function $\mathsf{H} : \mathcal{X} \to \mathbb{G}_S$ is sampled uniformly from $O_{\mathcal{X},\mathbb{G}_S}$. The length of the proof is $2\lceil \log_2(\ell) \rceil + 12$ group elements in $\mathbb{G}_T$.*

Concretely, for $\ell = 256$ we obtain a non-interactive proof that contains about 28 group elements in $\mathbb{G}_T$. For a 256-bit elliptic curve this comes out to about 900 bytes.

Proof of Theorem 21. The proof system is complete by construction and inherits its zero knowledge property from the proofs for $\mathcal{R}_{\text{eR1CS}}$ and $\mathcal{R}_{\text{dlog}}$. It remains to prove computational soundness.

Since discrete log in $\mathbb{G}_T$ is hard and the generators $G_{T,1}, G_{T,2} \in \mathbb{G}_T$ are sampled independently, the proof system for $\mathcal{R}_{\text{eR1CS}}$ is sound. Therefore if the verifier accepts, then (A, B, C, T) is in $\mathcal{L}_{\mathcal{R}_{\text{eR1CS}}}$. We already showed that our R1CS (A, B, C) verifies that the three claims about the points $P_0, \ldots, P_\ell$ hold, and this implies that (Q, x, Y) is in $\mathcal{L}_{\mathcal{R}_{\mathsf{H}}}$, as required.

The only part that remains to argue is that when verifying that $P_i = P_{i-1} + \Delta_i$ it suffices to use the cord rule, as we did in (14). That is, we need to argue that $P_{i-1} \neq \pm\Delta_i$ for $i = 1, \ldots, \ell$. To do so, consider an adversary $\mathcal{A}$ that outputs $(Q, x, Y, \pi_0, \pi_Q, \pi_Y)$ such that (1) the $\mathcal{R}_{\mathsf{H}}$ verifier accepts, and (2) if $Q = k \cdot G_{T,1}$, then the derived points $P_0, \ldots, P_\ell$ in $\mathbb{G}_S$ satisfy that $P_{i-1} = \pm\Delta_i$ for some i. We use this adversary $\mathcal{A}$ to break discrete log in $\mathbb{G}_S$.

Let us construct an algorithm $\mathcal{B}$ that computes discrete log in $\mathbb{G}_S$. This $\mathcal{B}$ takes as input $R \in \mathbb{G}_S$ and needs to output an $\alpha \in \mathbb{Z}_s$ such that $R = \alpha G_S$. Algorithm $\mathcal{B}$ runs $\mathcal{A}$ and responds to its random oracle queries for $\mathsf{H}(x_i)$ by choosing a random $r_i \leftarrow\$ \mathbb{Z}_s$ and responding with $\mathsf{H}(x_i) := r_i R$. Eventually, algorithm $\mathcal{A}$ outputs $(Q, x, Y, \pi_0, \pi_Q, \pi_Y)$. Since the $\mathcal{R}_{\mathsf{H}}$ verifier accepts this tuple, we obtain three things:

- First, thanks to the proof π_Q, our $\mathcal{B}$ can run the extractor for the proof system for $\mathcal{R}_{\text{dlog}}$ to extract from $\mathcal{A}$ an integer $k \in [s-1]$ such that $Q = kG_{T,1}$. As usual, we let $(k_0, \ldots, k_\ell)$ be the binary representation of k.
- Second, if we set $T := Q + Y$, then we know that (A, B, C, T) is in $\mathcal{L}_{\mathcal{R}_{\text{eR1CS}}}$. This is because the Bulletproofs proof π_0 is accepted by the verifier.
- Third, since $\mathcal{A}$ must have queried for $\mathsf{H}(x)$, it follows that $\mathcal{B}$ knows an $r \in \mathbb{Z}_s$ such that $\mathsf{H}(x) = rR$.

Next, by assumption on $\mathcal{A}$ we know that the points $P_0, \ldots, P_\ell$, derived from $\mathsf{H}(x)$ and G_S via (10) satisfy that $P_0 = 0$ or $P_{i-1} = \pm\Delta_i$ for some $i \in [\ell]$. If $P_0 = 0$ then $k_0(rR) + c_0 G_S = 0$. Then since $c_0 \neq 0$ it must be that $k_0 r \neq 0$, and this immediately reveals the discrete log α such that $R = \alpha G_S$.

Next, if $P_0 \neq 0$, let i^* be the smallest index i in $[\ell]$ such that $P_{i^*-1} = \pm\Delta_{i^*}$. For now, let us assume that $P_{i^*-1} = -\Delta_{i^*}$. The case $P_{i^*-1} = \Delta_{i^*}$ is handled the same way. Since (A, B, C, T) is in

$\mathcal{L}_{\mathcal{R}_{\text{eR1CS}}}$ we know that $P_0, \ldots, P_{i^*-1}$ must be computed as in (10). Therefore,

$$2^{i^*} k_{i^*}(rR) + c_{i^*} G_{\text{S}} = \Delta_{i^*} = -P_{i^*-1} = -\left[\sum_{j=0}^{i^*-1} (2^j k_j)(rR) + \sum_{j=0}^{i^*-1} c_j G_{\text{S}}\right]$$

which leads to

$$r \sum_{j=0}^{i^*} (2^j k_j) \cdot R = \sum_{j=0}^{i^*} c_j \cdot G_{\text{S}}. \tag{15}$$

When $i^* < \ell$ we know by construction of c_i in (9) that $\sum_{j=0}^{i^*} c_j \neq 0$. It follows that the left hand side is non-zero, and then (15) reveals the discrete log α. The case $i^* = \ell$ is not possible because then the left hand side of (15) is non-zero since $k \neq 0$, but the right hand side is zero because $c_0 + \cdots + c_\ell = 0$. The case $P_{i^*-1} = \Delta_{i^*}$ is the same and relies on the fact that $c_{i^*} \neq \sum_{j=0}^{i^*-1} c_j$.

Hence, we conclude that if discrete log in $\mathbb{G}_{\text{S}}$ is hard, then $P_0 \neq 0$ and $P_{i-1} \neq \pm\Delta_i$ for all $i \in [\ell]$, as required for (14) to properly verify addition in $\mathbb{G}_{\text{S}}$. This completes the proof of the theorem. □

Practical considerations. For the applications to ECDSA discuss in Section 4, the target group $\mathbb{G}_{\text{T}}$ needs to be the standard group of points on the elliptic curve Secp256k1. This curve is defined by the elliptic curve equation $y^2 = x^3 + 7$ in $\mathbb{F}_p$ for a specific prime p. This curve has q point in $\mathbb{F}_p$ for some prime q. Remarkably, for such curves, a theorem due to Silverman and Stange [62, Cor. 22], shows that the same curve $y^2 = x^3 + 7$, but this time defined over $\mathbb{F}_q$, has prime order p. Hence, we can take as our source group the curve $y^2 = x^3 + 7$ defined over $\mathbb{F}_q$.

Suppose that instead we use the curve Ed25519 [5] as the target group $\mathbb{G}_{\text{T}}$, as needed for EdDSA. The curve is defined over $\mathbb{F}_p$ where $p := 2^{255} - 19$ and has order $8q$ for some prime q. In this case we would choose some prime order elliptic curve defined over $\mathbb{F}_q$ and use it as our source group.

The running time of the verifier is dominated by the time to do a multiscalar multiplication (MSM) of a vector of dimension $2 \times (13\ell)$. For $\ell = 256$ this gives an MSM of length about 6,600. This drops to about 4,000 using Remark 1. The running time of the prover is comparable.

The cost of multiscalar multiplications (MSM) for these dimensions is about a sixth of a sequence multiplications done one at a time. In addition, the bases are all fixed ahead of time, and so the multiplications can be sped up pre-computing the required tables. We therefore estimate the cost of the MSM by considering the cost of multiplications and dividing by six. A crypto library for the curve `secp256k1` computes 140,000 generator multiplications per second, and for `Ed25519` it computes 55,000 generator multiplications per second (running a single thread on a 2.3 GHz 8-Core Intel Core i9). This yields an estimated running time for proving and verifying of approximately just 5ms for `secp256k1` and 14ms for `Ed25519`. Of course, this can be further accelerated by employing multiple threads in parallel.

6.2 The full DDH eVRF

Recall that the output of the eVRF in Construction 17 is restricted to about half the group $\mathbb{G}_{\text{T}}$. In particular, the output of $\mathsf{Eval}_1^{\mathsf{H}}(k, x)$ is only pseudorandom over half of $\mathbb{F}_q$, namely the set $\mathcal{S}$ of x-coordinates of points $\mathbb{G}_{\text{S}}^*$. While this is fine for our distributed key generation application in Section 4.1, it is insufficient for the threshold Schnorr protocol in Section 4.4 where the generated nonce must be uniform over all of $\mathbb{F}_q$.

In this section we show how to enhance the basic DDH eVRF so that the output of $\mathsf{Eval}_1^{\mathsf{H}}(k,x)$ is pseudorandom over all of $\mathbb{F}_q$. There are two ways to do it. The MuSig-DN [54,55] scheme does so by working with a product of the elliptic curve $\mathbb{G}_{\mathrm{S}}^*$ and its twist. This ensures that the final x-coordinate is distributed over all of $\mathbb{F}_q$, but at the cost of making the intermediate elliptic curve points constructed in the proof for $\mathcal{R}_{\mathsf{H}}$ live in a quadratic extension of $\mathbb{F}_q$.

Instead, to avoid working in a quadratic extension, we simply evaluate the basic DDH PRF twice, and extract entropy from the two x-coordinates of the resulting points. That is, we use two independent hash functions $\mathsf{H}_1, \mathsf{H}_2 : \mathcal{X} \to \mathbb{G}_{\mathrm{S}}$ and define the PRF value as

$$Eval_1^{\mathsf{H}}(k,x) := ext(x_1,x_2) \quad \text{where } k \cdot \mathsf{H}_1(x) = (x_1,y_1) \text{ and } k \cdot \mathsf{H}_2(x) = (x_2,y_2).$$

Here $ext : \mathbb{F}_q^2 \to \mathbb{F}_q$ is an entropy extraction map that takes as input two elements, each uniformly distributed in $\mathcal{S}$, and outputs a statistically close to uniform element in $\mathbb{F}_q$. One can treat x_1 and x_2 as two independent samples from a biased source, and construct *ext* as a deterministic 2-source extractor [10]. However, the resulting extractors are not strong enough to ensure that the output is statistically close to uniform in $\mathbb{F}_q$ without additional assumptions on the structure of $\mathcal{S}$.

Instead, we construct *ext* as a simple randomized extractor using the leftover hash lemma [36]. First, let us review the lemma. For a random variable r distributed in a finite set $\mathcal{R}$, we define the **guessing probability** of r as $\max_{z \in \mathcal{R}} \Pr[\mathsf{r} = z]$. In addition, a function $ext : \mathcal{K} \times \mathcal{R} \to \mathbb{F}_q$ is said to be a **universal hash** if $\Pr_{k \leftarrow\$ \mathcal{K}}\big[ext(k,z) = ext(k,z')\big] \le 1/q$ for all distinct $z, z' \in \mathcal{R}$. Finally, the **statistical distance** between two distributions $\mathcal{P}_1$ and $\mathcal{P}_2$ defined over $\mathcal{R}$ is defined as $\Delta := \frac{1}{2}\sum_{z \in \mathcal{R}} |\mathcal{P}_1(z) - \mathcal{P}_2(z)|$. Then $\Delta \in [0,1]$.

Lemma 1 (Leftover Hash Lemma [36]). *Let $ext : \mathcal{K} \times \mathcal{R} \to \mathbb{F}_q$ be a universal hash. Let $\mathsf{k}, \mathsf{r}_1, \ldots, \mathsf{r}_m$ be mutually independent random variables, where k is uniformly distributed over $\mathcal{K}$, and each r_i is distributed over $\mathcal{R}$ with guessing probability at most γ. Let Δ be the statistical distance between $\big((\mathsf{k},\ ext(\mathsf{k},\mathsf{r}_1), \ldots, ext(\mathsf{k},\mathsf{r}_m)\big)$ and the uniform distribution on $\mathcal{K} \times \mathcal{R}^m$. Then*

$$\Delta \le \frac{m}{2}\sqrt{\gamma \cdot q} \qquad \square$$

In our setting, we have $\mathcal{R} := \mathbb{F}_q \times \mathbb{F}_q$ and each random variable r_i is a pair (x_1,x_2) uniformly distributed over $\mathcal{S} \times \mathcal{S} \subseteq \mathcal{R}$. Moreover, since the number of points in $\mathbb{G}_{\mathrm{S}}^*$ is at least $q - 2\sqrt{q}$, we know that $|\mathcal{S}| \ge (q - 2\sqrt{q})/2$. Therefore, the guessing probability of r_i is

$$\gamma = 1/|\mathcal{S}|^2 \le 4/(q - 2\sqrt{q})^2 = 4/q(\sqrt{q} - 2)^2$$

We will use a hash function $ext : \mathbb{F}_q \times \mathbb{F}_q^2 \to \mathbb{F}_q$ defined as

$$ext\big(k', (x_1,x_2)\big) := k' \cdot x_1 + x_2.$$

It is not difficult to show that this hash function is a universal hash. Then by the leftover hash lemma, if (x_1,x_2) is uniform in $\mathcal{S} \times \mathcal{S}$ and k' is uniform in $\mathbb{F}_q$ then $(k',\ k'x_1 + x_2)$ is statistically close to uniform over $\mathbb{F}_q^2$ with statistical distance

$$\Delta \le 0.5\sqrt{\gamma q} \le 1/(\sqrt{q} - 2). \tag{16}$$

which is negligible in the security parameter. If we extract from m samples, then the statistical distance increases by at most a factor of m. Therefore, if m is polynomial in security parameter, then the statistical distance to uniform of all the samples remains negligible.

In our construction, the universal hash key k' will be sampled by the KGen algorithm and become a part of the eVRF verification key vk. In addition, the relation $\mathcal{R}_{\mathsf{H}}$ from (6) is replaced by the following two-time relation

$$\begin{aligned}\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2} := \Big\{\big((Q,k',x,Y)\ ;\ k\big)\Big\} \subseteq (\mathbb{G}_{\mathrm{T}} \times \mathbb{F}_q \times \mathcal{X} \times \mathbb{G}_{\mathrm{T}}) \times [s-1] \quad \text{where}\\ (1)\ \ Q = k \cdot G_{\mathrm{T},1},\\ (2)\ \ Y = y \cdot G_{\mathrm{T},2} \quad \text{for } y := k' \cdot x_{P_1} + x_{P_2} \in \mathbb{F}_q \text{ where}\\ P_1 = (x_{P_1}, y_{P_1}) := k \cdot \mathsf{H}_1(x) \text{ and } P_2 = (x_{P_2}, y_{P_2}) := k \cdot \mathsf{H}_2(x).\end{aligned} \tag{17}$$

Here $P_1, P_2 \in \mathbb{G}_{\mathrm{S}}^*$ are the two evaluations of the DDH PRF at input x, and $y := k' \cdot x_{P_1} + x_{P_2}$ is leftover hash extractor applied to the x-coordinates of P_1 and P_2. The proof system for $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$ contains twice as many constraints as the proof system for $\mathcal{R}_{\mathsf{H}}$ in Section 6.1: one set of constraints to prove that x_{P_1} is computed correctly and another to prove x_{P_2}. There is one more constraint to prove that $y = k' \cdot x_{P_1} + x_{P_2}$. The resulting construction for a DDH eVRF is as follows.

Construction 22 (The full DDH-based eVRF) *Let $(\mathbb{G}_{\mathrm{S}}, \mathbb{G}_{\mathrm{T}})$ be a group pair where s is the size of $\mathbb{G}_{\mathrm{S}}$ and q is the size of $\mathbb{G}_{\mathrm{T}}$. Let $G_{\mathrm{T},1}, G_{\mathrm{T},2}$ be two generators of $\mathbb{G}_{\mathrm{T}}$. Let $\mathsf{H}_1, \mathsf{H}_2$ be two functions $\mathsf{H}_1, \mathsf{H}_2 : \mathcal{X} \to \mathbb{G}_{\mathrm{S}}^*$, where $\mathbb{G}_{\mathrm{S}}^* := \mathbb{G}_{\mathrm{S}} \setminus \{0\}$. The **full DDH eVRF** with domain $\mathcal{X}$ and range $\mathbb{G}_{\mathrm{T}}$ is defined as follows:*

- $\mathsf{KGen}(1^\lambda)$*: Sample $k \leftarrow\!\!\$\ [s-1]$ and set $Q \leftarrow k \cdot G_{\mathrm{T},1}$. Sample $k' \leftarrow\!\!\$\ \mathbb{F}_q$. Use the prover P_{dlog} for $\mathcal{R}_{\mathrm{dlog}}$ to generate a proof $\pi_Q \leftarrow\!\!\$\ \mathsf{P}_{dlog}(Q, G_{\mathrm{T},1}; k)$. Set $\mathsf{vk} := (Q, k', \pi_Q)$ and $\mathsf{sk} := (k, k')$. Output the pair $(\mathsf{sk}, \mathsf{vk})$.*
- $\mathsf{Eval}^{\mathsf{H}_1,\mathsf{H}_2}\big((k,k'),x\big)$*: for $x \in \mathcal{X}$, let $P_1 \leftarrow k \cdot \mathsf{H}_1(x)$ and $P_2 \leftarrow k \cdot \mathsf{H}_2(x)$ so that $P_1, P_2 \in \mathbb{G}_{\mathrm{S}}^*$.*

  $$\begin{aligned}&\text{with } P_1 = (x_{P_1}, y_{P_2}) \text{ and } P_2 = (x_{P_1}, y_{P_2}) \text{ both in } \mathbb{F}_q^2\\ &\qquad \text{set } \ y \leftarrow k' \cdot x_{P_1} + x_{P_2} \in \mathbb{F}_q \quad \text{and} \quad Y \leftarrow y \cdot G_{\mathrm{T},2} \in \mathbb{G}_{\mathrm{T}}.\end{aligned}$$

  Next, run the prover P for $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$ to construct a proof π that the triple $\big(Q, k', x, Y\big)$ is in the language of the relation $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$ from (17). Output (y, Y, π).
- $\mathsf{Verify}^{\mathsf{H}_1,\mathsf{H}_2}(\mathsf{vk}, x, Y, \pi)$*: for $\mathsf{vk} = (Q, k', \pi_Q)$, the algorithm accepts if (1) π is a valid proof that (Q, k', x, Y) is in the language of the relation $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$, and (2) π_Q is a valid proof for $(Q, G_{\mathrm{T},1})$ as an instance of $\mathcal{R}_{\mathrm{dlog}}$.*

Recall that soundness of our proof system for $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$ relies on the hardness of discrete log in $\mathbb{G}_{\mathrm{T}}$ and the knowledge soundness of the proof system for $\mathcal{R}_{\mathrm{dlog}}$. It also relies on $G_{\mathrm{T},1}, G_{\mathrm{T},2}$ being random generators of $\mathbb{G}_{\mathrm{T}}$ so that there is no known discrete log relation between them. If we take the soundness and zero knowledge of $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$ as a given, then the following theorem shows the security of Construction 22 as an eVRF.

Theorem 23 (eVRF security). *Let $(\mathbb{G}_{\mathrm{S}}, \mathbb{G}_{\mathrm{T}})$ be a secure group pair, so that DDH holds in $\mathbb{G}_{\mathrm{S}}$. Let (P, V) be a non-interactive zero-knowledge argument system for the relation $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$ from (17). Then the eVRF in Construction 22 is a secure eVRF (as in Definition 4) with respect to the domain/range $(\mathcal{X}, \mathbb{G}_{\mathrm{T}})$ and function family $\mathcal{O}_{\mathcal{X},(\mathbb{G}_{\mathrm{S}}^*)^2}$.*

Proof. *Consistency* holds by construction. *Verifiability* and *Simulatability* as a VRF follow from the soundness and zero knowledge properties of the proof system for $\mathcal{R}^2_{\mathsf{H}_1,\mathsf{H}_2}$. It remains to argue *Pseudorandomness*, which we do with a sequence of hybrid games with an adversary $\mathcal{A}$.

Let Game 0 be the usual pseudorandomness game from Definition 4 with respect to the PRF $(\mathsf{KGen}, \mathsf{Eval}_1^{\mathsf{H}_1,\mathsf{H}_2})$. Recall that $\mathsf{Eval}_1^{\mathsf{H}_1,\mathsf{H}_2}\big((k,k'),x\big)$ makes use of two DDH PRFs with domain/range $(\mathcal{X},\mathcal{S})$, where $\mathcal{S} \subseteq \mathbb{F}_q$ is the set of x-coordinates of points in $\mathbb{G}_{\mathrm{S}}^*$. In Game 1 we replace both these PRFs by truly random functions. That is, in Game 1 we define $\mathsf{Eval}_1^{\mathsf{H}_1,\mathsf{H}_2}\big((k,k'),x\big)$ as

$$\mathsf{Eval}_1^{\mathsf{H}_1,\mathsf{H}_2}\big((k,k'),x\big) := k' \cdot f_1(x) + f_2(x) \in \mathbb{F}_q$$

where f_1 and f_2 are sampled at random from $\mathcal{O}_{\mathcal{X},\mathcal{S}}$. Since DDH holds in $\mathbb{G}_{\mathrm{S}}$ and the functions $\mathsf{H}_1, \mathsf{H}_2$ are sampled at random from $\mathcal{O}_{\mathcal{X},\mathbb{G}_{\mathrm{S}}^*}$, we know that the both DDH PRFs are secure over the range $\mathcal{S} \subseteq \mathbb{F}_q$. Therefore, replacing these two PRFs in $\mathsf{Eval}_1^{\mathsf{H}_1,\mathsf{H}_2}$ by truly random functions f_1, f_2, changes the adversary's advantage by at most a negligible amount.

In Game 2 we replace $\mathsf{Eval}_1^{\mathsf{H}_1,\mathsf{H}_2}\big((k,k'),x\big)$ by a truly random function f with domain/range $(\mathcal{X},\mathbb{F}_q)$. If the adversary makes at most m queries to Eval_1, then by (16), the statistical distance between the adversary's view in Game 1 and Game 2 is at most $m/(\sqrt{q}-2)$, which is negligible. Hence, the adversary's advantage in Game 2 is at most negligibly different from its advantage in Game 1. Now, since adversary $\mathcal{A}$ has advantage zero in Game 2, it must also have at most negligible advantage in Game 0, as required. □

7 Conclusions and Open Problems

In this paper, we have introduced a new primitive called an exponent VRF (eVRF), and shown that it has many applications in the field of threshold cryptography and signing. In particular, it enables us to achieve one-round simulatable key generation, two-round signing for Schnorr (multiparty) and ECDSA (two-party), and it provides a hierarchical key derivation method like BIP032 with additional properties like MPC friendliness and public verifiability. We also provided constructions under both the DDH and Paillier (DCRA) assumptions.

Our work leaves open a number of interesting questions. An important open question raised by this work is the construction of an efficient key homomorphic eVRF, namely an eVRF that satisfies

$$\mathsf{Eval}_1(k_1 + k_2, x) \cdot G = \mathsf{Eval}_1(k_1, x) \cdot G + \mathsf{Eval}_1(k_2, x) \cdot G$$

where G is a generator of a standard cryptographic group used for Schnorr signatures. This will enable *deterministic* two round threshold Schnorr signing, by constructing a threshold eVRF itself so that any subset $\mathcal{Q}$ of the parties will compute the same Eval result on the same message. In addition, it will enable threshold Schnorr signing without knowing the set of parties ahead of time. Finally, it will enable the parties to refresh their secrets and achieve (static) proactive security.

One possible approach to constructing a key homomorphic eVRF is to explore building an efficient eVRF from a lattice-based random oracle PRF, described in [7], whose security is based on the learning with rounding problem (LWR). This PRF is almost key homomorphic, which is sufficient for the applications in Section 4. One would build an eVRF by encoding the output of this PRF in the exponent of another group, as we did in the constructions in this paper. The challenge then is to devise an efficient non-interactive zero-knowledge proof that the PRF was evaluated correctly.

Another important open question is to construct a simulatable two-round *multiparty* protocol for ECDSA (our ECDSA protocol is only for two parties).

Acknowledgments. We thank the authors of MuSig-DN for pointing out a mistake in the initial description of our DDH-based eVRF. The first author was supported by NSF, DARPA, the Simons Foundation, and NTT Research. Opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA.

References

1. Agrawal, S., Ganesh, C., Mohassel, P.: Non-interactive zero-knowledge proofs for composite statements. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 643–673. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2018). https://doi.org/10.1007/978-3-319-96878-0_22
2. Attema, T., Fehr, S., Klooß, M.: Fiat-shamir transformation of multi-round interactive proofs. In: Kiltz, E., Vaikuntanathan, V. (eds.) TCC 2022, Part I. LNCS, vol. 13747, pp. 113–142. Springer, Heidelberg, Germany, Chicago, IL, USA (Nov 7–10, 2022). https://doi.org/10.1007/978-3-031-22318-1_5
3. Benaloh, J.: Verifiable secret-ballot elections. Ph.D thesis (1988)
4. Bernstein, D.J., Duif, N., Lange, T., Schwabe, P., Yang, B.Y.: High-speed high-security signatures. Cryptology ePrint Archive, Report 2011/368 (2011), https://eprint.iacr.org/2011/368
5. Bernstein, D.J., Duif, N., Lange, T., Schwabe, P., Yang, B.Y.: High-speed high-security signatures. Journal of Cryptographic Engineering **2**(2), 77–89 (Sep 2012). https://doi.org/10.1007/s13389-012-0027-1
6. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: Goldwasser, S. (ed.) ITCS 2012. pp. 326–349. ACM, Cambridge, MA, USA (Jan 8–10, 2012). https://doi.org/10.1145/2090236.2090263
7. Boneh, D., Lewi, K., Montgomery, H.W., Raghunathan, A.: Key homomorphic PRFs and their applications. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 410–428. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2013). https://doi.org/10.1007/978-3-642-40041-4_23
8. Boneh, D., Shoup, V.: A graduate course in applied cryptography (version 0.6). Cambridge University Press (2023), cryptobook.us
9. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 327–357. Springer, Heidelberg, Germany, Vienna, Austria (May 8–12, 2016). https://doi.org/10.1007/978-3-662-49896-5_12
10. Bourgain, J.: More on the sum-product phenomenon in prime fields and its applications. International Journal of Number Theory **1**(01), 1–32 (2005)
11. Bünz, B.: Improving the privacy, scalability, and ecological impact of blockchains. Ph.D. thesis, Stanford University (2023)
12. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy. pp. 315–334. IEEE Computer Society Press, San Francisco, CA, USA (May 21–23, 2018). https://doi.org/10.1109/SP.2018.00020
13. Camenisch, J., Lysyanskaya, A.: Dynamic accumulators and application to efficient revocation of anonymous credentials. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 61–76. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2002). https://doi.org/10.1007/3-540-45708-9_5
14. Canetti, R.: Security and composition of multiparty cryptographic protocols. Journal of Cryptology **13**(1), 143–202 (Jan 2000). https://doi.org/10.1007/s001459910006
15. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: 42nd FOCS. pp. 136–145. IEEE Computer Society Press, Las Vegas, NV, USA (Oct 14–17, 2001). https://doi.org/10.1109/SFCS.2001.959888
16. Canetti, R., Gennaro, R., Goldfeder, S., Makriyannis, N., Peled, U.: UC non-interactive, proactive, threshold ECDSA with identifiable aborts. Cryptology ePrint Archive, Report 2021/060 (2021), https://eprint.iacr.org/2021/060
17. Castagnos, G., Chevallier-Mames, B.: Towards a DL-based additively homomorphic encryption scheme. In: Garay, J.A., Lenstra, A.K., Mambo, M., Peralta, R. (eds.) ISC 2007. LNCS, vol. 4779, pp. 362–375. Springer, Heidelberg, Germany, Valparaíso, Chile (Oct 9–12, 2007)

18. Castagnos, G., Laguillaumie, F.: Linearly homomorphic encryption from DDH. In: Nyberg, K. (ed.) CT-RSA 2015. LNCS, vol. 9048, pp. 487–505. Springer, Heidelberg, Germany, San Francisco, CA, USA (Apr 20–24, 2015). https://doi.org/10.1007/978-3-319-16715-2_26
19. Chase, M., Orrù, M., Perrin, T., Zaverucha, G.: Proofs of discrete logarithm equality across groups. Cryptology ePrint Archive, Report 2022/1593 (2022), https://eprint.iacr.org/2022/1593
20. Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: Brickell, E.F. (ed.) CRYPTO'92. LNCS, vol. 740, pp. 89–105. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 1993). https://doi.org/10.1007/3-540-48071-4_7
21. Chen, Y.H., Lindell, Y.: Feldman's verifiable secret sharing for a dishonest majority. Cryptology ePrint Archive, Paper 2024/031 (2024), https://eprint.iacr.org/2024/031, https://eprint.iacr.org/2024/031
22. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 119–136. Springer, Heidelberg, Germany, Cheju Island, South Korea (Feb 13–15, 2001). https://doi.org/10.1007/3-540-44586-2_9
23. Devevey, J., Fallahpour, P., Passelègue, A., Stehlé, D.: A detailed analysis of fiat-shamir with aborts. Cryptology ePrint Archive, Report 2023/245 (2023), https://eprint.iacr.org/2023/245
24. Devevey, J., Libert, B., Peters, T.: Rational modular encoding in the DCR setting: Non-interactive range proofs and paillier-based naor-yung in the standard model. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) PKC 2022, Part I. LNCS, vol. 13177, pp. 615–646. Springer, Heidelberg, Germany, Virtual Event (Mar 8–11, 2022). https://doi.org/10.1007/978-3-030-97121-2_22
25. Doerner, J., Kondi, Y., Lee, E., abhi shelat: Threshold ECDSA in three rounds. Cryptology ePrint Archive, Paper 2023/765 (2023), https://eprint.iacr.org/2023/765, https://eprint.iacr.org/2023/765
26. Doerner, J., Kondi, Y., Lee, E., shelat, a.: Secure two-party threshold ECDSA from ECDSA assumptions. Cryptology ePrint Archive, Report 2018/499 (2018), https://eprint.iacr.org/2018/499
27. Feldman, P.: A practical scheme for non-interactive verifiable secret sharing. In: 28th FOCS. pp. 427–437. IEEE Computer Society Press, Los Angeles, CA, USA (Oct 12–14, 1987). https://doi.org/10.1109/SFCS.1987.4
28. Fischlin, M.: Communication-efficient non-interactive proofs of knowledge with online extractors. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 152–168. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2005). https://doi.org/10.1007/11535218_10
29. Fouque, P.A., Stern, J.: One round threshold discrete-log key generation without private channels. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 300–316. Springer, Heidelberg, Germany, Cheju Island, South Korea (Feb 13–15, 2001). https://doi.org/10.1007/3-540-44586-2_22
30. Galbraith, S.D.: Elliptic curve paillier schemes. Cryptology ePrint Archive, Report 2001/050 (2001), https://eprint.iacr.org/2001/050
31. Garg, S., Srinivasan, A.: Two-round multiparty secure computation from minimal assumptions. J. ACM **69**(5) (2022). https://doi.org/10.1145/3566048, https://doi.org/10.1145/3566048
32. Garillot, F., Kondi, Y., Mohassel, P., Nikolaenko, V.: Threshold Schnorr with stateless deterministic signing from standard assumptions. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part I. LNCS, vol. 12825, pp. 127–156. Springer, Heidelberg, Germany, Virtual Event (Aug 16–20, 2021). https://doi.org/10.1007/978-3-030-84242-0_6
33. Goldreich, O.: Foundations of Cryptography: Basic Applications, vol. 2. Cambridge University Press, Cambridge, UK (2004)
34. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. Journal of the ACM **33**(4), 792–807 (Oct 1986). https://doi.org/10.1145/6490.6503
35. Groth, J.: Non-interactive distributed key generation and key resharing. Cryptology ePrint Archive, Report 2021/339 (2021), https://eprint.iacr.org/2021/339
36. Impagliazzo, R., Levin, L.A., Luby, M.: Pseudo-random generation from one-way functions (extended abstracts). In: 21st ACM STOC. pp. 12–24. ACM Press, Seattle, WA, USA (May 15–17, 1989). https://doi.org/10.1145/73007.73009
37. Joye, M., Libert, B.: Efficient cryptosystems from 2^k-th power residue symbols. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 76–92. Springer, Heidelberg, Germany, Athens, Greece (May 26–30, 2013). https://doi.org/10.1007/978-3-642-38348-9_5
38. Katz, J.: Round optimal fully secure distributed key generation. Cryptology ePrint Archive, Paper 2023/1094 (2023), https://eprint.iacr.org/2023/1094, https://eprint.iacr.org/2023/1094
39. Kiltz, E., Lyubashevsky, V., Schaffner, C.: A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 552–586. Springer, Heidelberg, Germany, Tel Aviv, Israel (Apr 29 – May 3, 2018). https://doi.org/10.1007/978-3-319-78372-7_18

40. Komlo, C., Goldberg, I.: FROST: Flexible round-optimized Schnorr threshold signatures. In: Dunkelman, O., Jr., M.J.J., O'Flynn, C. (eds.) SAC 2020. LNCS, vol. 12804, pp. 34–65. Springer, Heidelberg, Germany, Halifax, NS, Canada (Virtual Event) (Oct 21-23, 2020). https://doi.org/10.1007/978-3-030-81652-0_2
41. Komlo, C., Goldberg, I.: Arctic: Lightweight and stateless threshold schnorr signatures. Cryptology ePrint Archive, Paper 2024/466 (2024), https://eprint.iacr.org/2024/466, https://eprint.iacr.org/2024/466
42. Kondi, Y., Orlandi, C., Roy, L.: Two-round stateless deterministic two-party Schnorr signatures from pseudorandom correlation functions. In: CRYPTO 2023, Part I. pp. 646–677. LNCS, Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 2023). https://doi.org/10.1007/978-3-031-38557-5_21
43. Kondi, Y., Orlandi, C., Roy, L.: Two-round stateless deterministic two-party schnorr signatures from pseudorandom correlation functions. Cryptology ePrint Archive, Report 2023/216 (2023), https://eprint.iacr.org/2023/216
44. Kushilevitz, E., Lindell, Y., Rabin, T.: Information-theoretically secure protocols and security under composition. In: Kleinberg, J.M. (ed.) 38th ACM STOC. pp. 109–118. ACM Press, Seattle, WA, USA (May 21–23, 2006). https://doi.org/10.1145/1132516.1132532
45. Lindell, Y.: Fast secure two-party ECDSA signing. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 613–644. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 20–24, 2017). https://doi.org/10.1007/978-3-319-63715-0_21
46. Lindell, Y.: Simple three-round multiparty schnorr signing with full simulatability. Cryptology ePrint Archive, Report 2022/374 (2022), https://eprint.iacr.org/2022/374
47. Lyubashevsky, V.: Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer, Heidelberg, Germany, Tokyo, Japan (Dec 6–10, 2009). https://doi.org/10.1007/978-3-642-10366-7_35
48. Makriyannis, N.: On the classic protocol for MPC schnorr signatures. Cryptology ePrint Archive, Report 2022/1332 (2022), https://eprint.iacr.org/2022/1332
49. Makriyannis, N., Yomtov, O., Galansky, A.: Practical key-extraction attacks in leading mpc wallets. Cryptology ePrint Archive, Paper 2023/1234 (2023), https://eprint.iacr.org/2023/1234, https://eprint.iacr.org/2023/1234
50. Micali, S., Rabin, M.O., Vadhan, S.P.: Verifiable random functions. In: 40th FOCS. pp. 120–130. IEEE Computer Society Press, New York, NY, USA (Oct 17–19, 1999). https://doi.org/10.1109/SFFCS.1999.814584
51. Naccache, D., Stern, J.: A new public key cryptosystem based on higher residues. In: Gong, L., Reiter, M.K. (eds.) ACM CCS 98. pp. 59–66. ACM Press, San Francisco, CA, USA (Nov 2–5, 1998). https://doi.org/10.1145/288090.288106
52. Naor, M., Pinkas, B., Reingold, O.: Distributed pseudo-random functions and KDCs. In: Stern, J. (ed.) EUROCRYPT'99. LNCS, vol. 1592, pp. 327–346. Springer, Heidelberg, Germany, Prague, Czech Republic (May 2–6, 1999). https://doi.org/10.1007/3-540-48910-X_23
53. Navot, S.: Insecurity of musig and bn multi-signatures with delayed message selection. Cryptology ePrint Archive, Paper 2024/437 (2024), https://eprint.iacr.org/2024/437, https://eprint.iacr.org/2024/437
54. Nick, J., Ruffing, T., Seurin, Y., Wuille, P.: MuSig-DN: Schnorr multi-signatures with verifiably deterministic nonces. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 1717–1731. ACM Press, Virtual Event, USA (Nov 9–13, 2020). https://doi.org/10.1145/3372297.3417236
55. Nick, J., Ruffing, T., Seurin, Y., Wuille, P.: MuSig-DN: Schnorr multi-signatures with verifiably deterministic nonces. Cryptology ePrint Archive, Report 2020/1057 (2020), https://eprint.iacr.org/2020/1057
56. Okamoto, T., Uchiyama, S.: A new public-key cryptosystem as secure as factoring. In: Nyberg, K. (ed.) EUROCRYPT'98. LNCS, vol. 1403, pp. 308–318. Springer, Heidelberg, Germany, Espoo, Finland (May 31 – Jun 4, 1998). https://doi.org/10.1007/BFb0054135
57. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT'99. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg, Germany, Prague, Czech Republic (May 2–6, 1999). https://doi.org/10.1007/3-540-48910-X_16
58. Palatinus, M., Rusnak, P.: Multi-account hierarchy for deterministic wallets (2014), https://github.com/bitcoin/bips/blob/master/bip-0044.mediawiki
59. Papadopoulos, D., Wessels, D., Huque, S., Naor, M., Včelák, J., Reyzin, L., Goldberg, S.: Making NSEC5 practical for DNSSEC. Cryptology ePrint Archive, Report 2017/099 (2017), https://eprint.iacr.org/2017/099
60. Pippenger, N.: On the evaluation of powers and monomials. SIAM Journal on Computing **9**(2), 230–250 (1980)
61. Schnorr, C.P.: Efficient signature generation by smart cards. Journal of Cryptology **4**(3), 161–174 (Jan 1991). https://doi.org/10.1007/BF00196725
62. Silverman, J.H., Stange, K.E.: Amicable pairs and aliquot cycles for elliptic curves. Experimental Mathematics **20**(3) (2011). https://doi.org/10.1080/10586458.2011, link

63. Wuille, P.: Hierarchical deterministic wallets (2012), https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki

A A Chaum-Pedersen style ZK proof system for the relation $\mathcal{R}_{\text{eq}}$

Fig. 1 gives a standard Chaum-Pedersen [20] style honest verifier zero knowledge (HVZK) interactive proof system for the relation $\mathcal{R}_{\text{eq}}$ from (4). The proof system uses a non-interactive zero knowledge (NIZK) proof system $(P_{\text{zero}}, V_{\text{zero}})$ for the relation $\mathcal{R}_{\text{zero}}$ defined as

$$\mathcal{R}_{\text{zero}} := \left\{ \big((\mathsf{pk}, ct)\ ;\ \mathsf{sk}\big)\ :\ \mathsf{Dec}(\mathsf{sk}, ct) = 0 \text{ AND } (\mathsf{pk}, \mathsf{sk}) \in \mathcal{L}_{\text{pub}} \right\} \tag{18}$$

It also uses an NIZK proof system $(P_{\text{pub}}, V_{\text{pub}})$ for the relation $\mathcal{R}_{\text{pub}}$ from (3).

Theorem 24 ([20]). *The proof system $(P_{\text{eq}}, V_{\text{eq}})$ in Fig. 1 is an HVZK for the relation $\mathcal{R}_{\text{eq}}$, assuming $(P_{\text{zero}}, V_{\text{zero}})$ is a NIZK for for the relation $\mathcal{R}_{\text{zero}}$, and $(P_{\text{pub}}, V_{\text{pub}})$ is a NIZK for the relation $\mathcal{R}_{\text{pub}}$.*

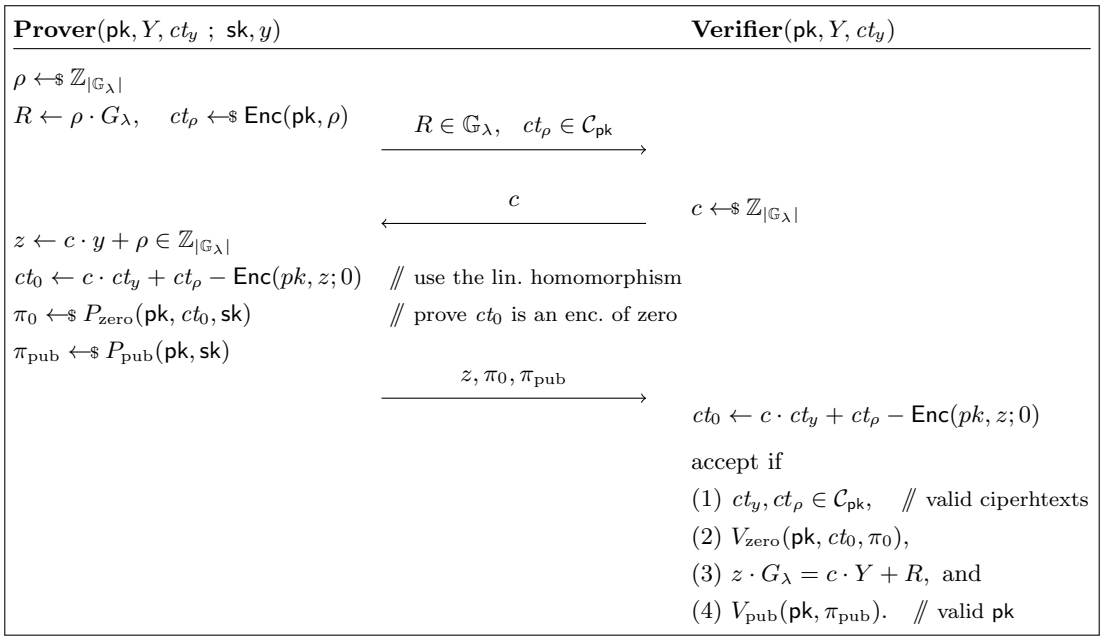

Fig. 1. A Chaum-Pedersen style ZK proof system for the relation $\mathcal{R}_{\text{eq}}$ from (4).

B A proof system for the relation $\mathcal{R}'_{\text{eq}}$

We construct a proof system for the relation $\mathcal{R}'_{\text{eq}}$ from (5) by adapting the protocol of Chase, Orrù, Perrin, and Zaverucha [19] to our settings. The resulting proof system is shown in Fig. 2. It can be made non-interactive using the Fiat-Shamir transform.

The proof system in Fig. 2 uses a range proof for Paillier ciphertexts, namely a non-interactive zero-knowledge proof system for the instance-witness relation

$$\mathcal{R}_{\text{range}} := \left\{ \big((\mathsf{pk}, ct, B)\ ;\ \mathsf{sk}\big)\ :\ 0 \le \mathsf{Dec}(\mathsf{sk}, ct) < B \right\} \tag{19}$$

There are several ways to build such a range proof for Paillier ciphertexts. One approach uses bit decomposition. Another approach, by Devevey, Libert, and Peters [24], generates a range proof that contains only a constant number of Paillier ciphertexts, but requires a trusted setup.

In addition, The proof system in Fig. 2 uses a non-interactive zero knowledge proof system $(P_{\text{zero}}, V_{\text{zero}})$ for the relation $\mathcal{R}_{\text{zero}}$ from (18), and a non-interactive zero knowledge proof system $(P_{\text{pub}}, V_{\text{pub}})$ for the relation $\mathcal{R}_{\text{pub}}$ from (3).

Rejection sampling. The proof system uses rejection sampling to reduce the size of the transcript, which may cause the prover to abort. By [19, Lemma 2], the probability that the prover aborts is exactly $1/A$, where A is a parameter used in the proof system. Setting $A := 256$ is a reasonable choice so that aborting is infrequent. After applying Fiat-Shamir, an abort simply causes the prover to try again with different randomness ρ. Moreover, when the protocol does not abort, the same lemma from [19] shows that the quantity z is uniform in the set $[q^2, q^2A - 1]$.

A proof system where the prover can abort [47,23] often fails to be honest-verifier zero-knowledge (HVZK) because it may not be possible to simulate aborted transcripts. Nevertheless, such protocols can satisfy a weaker notion called *no-abort honest-verifier zero-knowledge*, or **naHVZK**, where the simulator returns a valid transcript or $\perp$, and indistinguishability need only hold for non aborted transcripts [39]. This is a useful notion because naHVZK is sufficient to simulate non-interactive proofs after the Fiat-Shamir transform is applied.

Security. The following theorem states the security property of the proof system in Fig. 2. Recall that q is the order of the group $\mathbb{G}_\lambda$ and n is the size of the Paillier plaintext space.

Theorem 25. *For every $A > 0$, the proof system $(P_{\text{eq}}, V_{\text{eq}})$ in Fig. 2 is an naHVZK for the relation $\mathcal{R}'_{\text{eq}}$ from (5), provided that $2q^2A < n$, $(P_{\text{zero}}, V_{\text{zero}})$ is a NIZK for the relation $\mathcal{R}_{\text{zero}}$, $(P_{\text{pub}}, V_{\text{pub}})$ is a NIZK for the relation $\mathcal{R}_{\text{pub}}$, and $(P_{\text{range}}, V_{\text{range}})$ is a NIZK for the relation $\mathcal{R}_{\text{range}}$.*

Proof. We need to prove completeness, zero knowledge, and soundness.

Completeness. As explained above, an honest prover aborts with probability $1/A$, and if it doesn't abort then the verifier accepts the proof. Therefore, the protocol has $(1/A)$-completeness.

Zero knowledge. We construct a simulator $\mathsf{Sim}(\mathsf{pk}, Y, ct_y) \to (R, ct_\rho, c, z, \pi_0, \pi_y, \pi_{\text{pub}})$ that matches the distribution of an accepting transcript between the honest prover and honest verifier, when the prover does not abort. As explained above, when the honest prover does not abort, the quantity z is uniform in the set $[q^2, q^2A - 1]$. Then $\mathsf{Sim}(\mathsf{pk}, Y, ct_y)$ works as follows:

1 : $c \leftarrow\!\$\ [0, q-1], \quad z \leftarrow\!\$\ [q^2, q^2A - 1]$
2 : $R \leftarrow zG_\lambda - c \cdot Y$
3 : $ct_\rho \leftarrow\!\$\ \mathsf{ReRand}\big(\mathsf{pk},\ \mathsf{Enc}(\mathsf{pk}, z; 0) - c \cdot ct_y\big)$ // re-randomize the ciphertext
4 : $ct_0 \leftarrow (c \cdot ct_y + ct_\rho) - \mathsf{Enc}(pk, z; 0)$
5 : $\pi_0 \leftarrow\!\$\ \mathsf{Sim}_{\text{zero}}(\mathsf{pk}, ct_0), \quad \pi_y \leftarrow\!\$\ \mathsf{Sim}_{\text{range}}(\mathsf{pk}, ct_y, q), \quad \pi_{\text{pub}} \leftarrow\!\$\ \mathsf{Sim}_{\text{pub}}(\mathsf{pk})$
6 : **return** $(R, ct_\rho, c, z, \pi_0, \pi_y, \pi_{\text{pub}})$

This simulator generates the required distribution.

Soundness. Let (pk, Y, ct_y) be an $\mathcal{R}'_{\text{eq}}$ instance where $Y \in \mathbb{G}_\lambda$ and $ct_y \in \mathcal{C}_{\mathsf{pk}}$ such that $Y = y_q \cdot G_\lambda$ and $y_n = \mathsf{Dec}(\mathsf{sk}, ct_y)$ for some $y_q \in [0, q-1]$ and $y_n \in [0, n-1]$. Let $(R, ct_\rho, c, z, \pi_0, \pi_y, \pi_{\text{pub}})$ be an

accepting transcript, where $R = \rho_q \cdot G_\lambda$ and $\mathsf{Dec}(\mathsf{sk}, ct_\rho) = \rho_n$ for some integers $\rho_q \in [0, q-1]$ and $\rho_n \in [0, n-1]$. Then by soundness of $(P_{\text{range}}, V_{\text{range}})$ and $(P_{\text{zero}}, V_{\text{zero}})$ we know that

$$z = c \cdot y_q + \rho_q + w_q \cdot q, \quad z = c \cdot y_n + \rho_n + w_n \cdot n, \quad y_n \in [0, q-1]$$

for some $w_q, w_n \in \mathbb{Z}$. Since $cy_n < z < n$ and $\rho_n \in [0, n-1]$ it follows that $z - cy_n - \rho_n$ is in $[-(n-1), n-1]$ and therefore $w_n = 0$. Then equating the right hand sides of the first two equalities leads to

$$c \cdot y_q + \rho_q + w_q \cdot q = c \cdot y_n + \rho_n$$

Reducing this equality modulo q gives

$$c(y_q - y_n) \equiv \rho_n - \rho_q \pmod{q}. \tag{20}$$

Now, if $y_q \neq y_n \pmod{q}$ then there is a unique $c \in [0, q-1]$ for which (20) holds. The probability that the verifier chooses that c is $1/q$ which is negligible. Therefore, if the verifier accepts, then with high probability we have $y_q \equiv y_n \pmod{q}$. But since both y_q and y_n are in $[0, q-1]$, they must be equal as integers, as required. □

Prover$(\mathsf{pk}, Y, ct_y \ ; \ \mathsf{sk}, y)$		**Verifier**(pk, Y, ct_y)
$\pi_y \leftarrow\!\$\ P_{\text{range}}(\mathsf{pk}, ct_y, q, \mathsf{sk})$	// prove that $y \in [0, q-1]$	
$\rho \leftarrow\!\$\ [0, q^2 A - 1]$		
$R \leftarrow \rho \cdot G_\lambda, \quad ct_\rho \leftarrow\!\$\ \mathsf{Enc}(\mathsf{pk}, \rho)$		
	$R \in \mathbb{G}_\lambda, \quad ct_\rho \in \mathcal{C}_{\mathsf{pk}}$ →	
	← c	$c \leftarrow\!\$\ [0, q-1]$
$z \leftarrow c \cdot y + \rho \in \mathbb{Z}$		
abort if $z \notin [q^2, q^2 A - 1]$		
$ct_0 \leftarrow (c \cdot ct_y + ct_\rho) - \mathsf{Enc}(pk, z; 0)$	// enc. of $cy + \rho - z$	
$\pi_0 \leftarrow\!\$\ P_{\text{zero}}(\mathsf{pk}, ct_0, \mathsf{sk})$	// prove ct_0 is enc. of zero	
$\pi_{\text{pub}} \leftarrow\!\$\ P_{\text{pub}}(\mathsf{pk}, \mathsf{sk})$		
	$z, \pi_0, \pi_y, \pi_{\text{pub}}$ →	
		$ct_0 \leftarrow (c \cdot ct_y + ct_\rho) - \mathsf{Enc}(pk, z; 0)$
		accept if
		(1) $Y, R \in \mathbb{G}_\lambda, \quad ct_y, ct_\rho \in \mathcal{C}_{\mathsf{pk}}$,
		(2) $V_{\text{zero}}(\mathsf{pk}, ct_0, \pi_0), \quad V_{\text{range}}(\mathsf{pk}, ct_y, q, \pi_y)$,
		(3) $z \cdot G_\lambda = c \cdot Y + R$,
		(4) $z \in [q^2, q^2 A - 1]$,
		(5) $V_{\text{pub}}(\mathsf{pk}, \pi_{\text{pub}})$. // valid pk

Fig. 2. A ZK proof system for the relation $\mathcal{R}'_{\text{eq}}$ from (5). The system is parameterized by $A \in [n]$ that determines the abort probability. Recall that q is the order of the group $\mathbb{G}_\lambda$ and G_λ is its generator.

C The R1CS matrices A, B, C used in Section 6.1

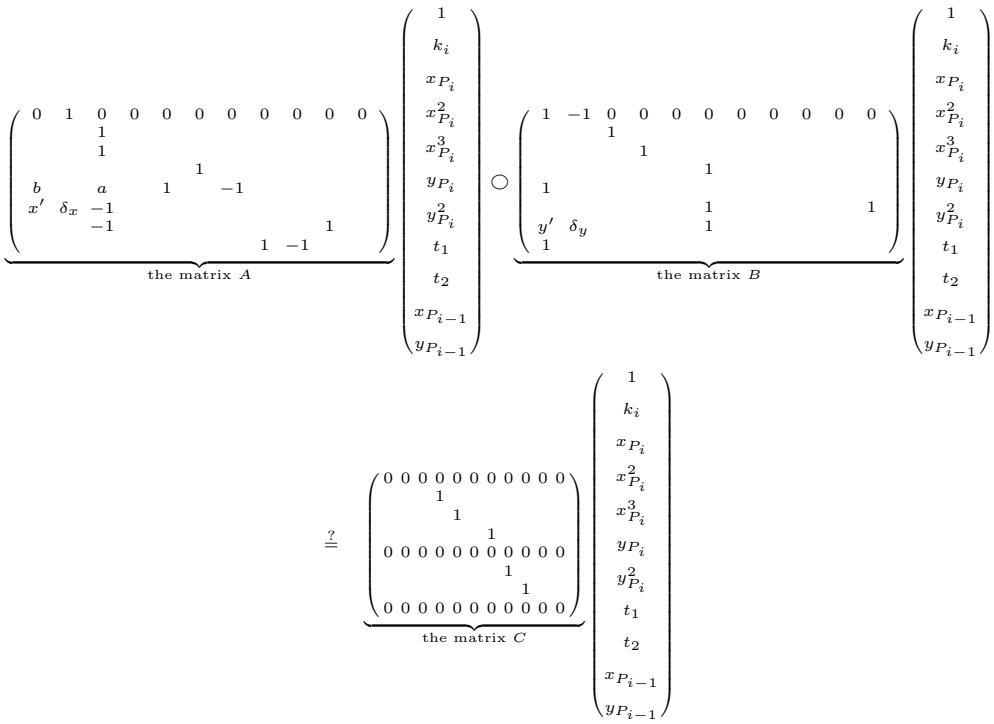

Fig. 3. The R1CS matrices A, B, C used in Section 6.1 for the i'th block of checks for some $i \in [\ell]$. Empty cells are set to 0. The first row confirms that $k_i \times (1 - k_i) = 0$. The second row confirms that $x_{P_i} \times x_{P_i} = x^2_{P_i}$. The third row confirms that $x_{P_i} \times x^2_{P_i} = x^3_{P_i}$. The fourth row confirms that $y_{P_i} \times y_{P_i} = y^2_{P_i}$. The fifth row confirms that $x^3_{P_i} + a x_{P_i} + b - y^2_{P_i} = 0$. The sixth row confirms that $(\delta_x k_i + x' - x_{P_i}) \times (y_{P_i} + y_{P_{i-1}}) = t_1$. The seventh row confirms that $(x_{P_{i-1}} - x_{P_i}) \times (\delta_y k_i + y' + y_{P_i}) = t_2$. The final row confirms that $t_1 - t_2 = 0$. The prover and verifier compute the values $\delta_x, \delta_y, x', y'$ using (13). Note that the 5th and 8th rows check a linear relation and can be combined into a single row by taking a linear combination using verifier randomness. The same holds for the second and third rows.