

Simple Three-Round Multiparty Schnorr Signing with Full Simulatability

Yehuda Lindell  

Coinbase, USA

Abstract. In a multiparty signing protocol, also known as a threshold signature scheme, the private signing key is shared amongst a set of parties and only a quorum of those parties can generate a signature. Research on multiparty signing has been growing in popularity recently due to its application to cryptocurrencies. Most work has focused on reducing the number of rounds to two, and as a result: (a) are not fully simulatable in the sense of MPC real/ideal security definitions, and/or (b) are not secure under concurrent composition, and/or (c) utilize non-standard assumptions of different types in their proofs of security. In this paper, we describe a simple three-round multiparty protocol for Schnorr signatures that is secure for any number of corrupted parties; i.e., in the setting of a dishonest majority. The protocol is fully simulatable, secure under concurrent composition, and proven secure in the standard model or random-oracle model (depending on the instantiations of the commitment and zero-knowledge primitives). The protocol realizes an ideal Schnorr signing functionality with perfect security in the ideal commitment and zero-knowledge hybrid model (and thus the only assumptions needed are for realizing these functionalities).

In our presentation, we do not assume that all parties begin with the message to be signed, the identities of the participating parties and a unique common session identifier, since this is often not the case in practice. Rather, the parties achieve consensus on these parameters as the protocol progresses.

1 Introduction

Multiparty signing, or threshold signature schemes, enable a quorum of parties to sign on a message, while preventing any subset of parties that does not form a quorum from doing so. This can be used to digitally emulate processes like multiple signers on a cheque or to protect a signing key from being stolen or misused by ensuring that multiple machines (and possibly humans) approve a transaction before it is signed. Threshold cryptography (for signing and encryption) were studied in the late 1980s and 1990s (cf. [7, 13, 16, 17, 23, 39, 40, 35, 33, 36, 3]), but has recently gained a lot of interest in both academia and industry – primarily for ECDSA and Schnorr signatures – due to its application to the protection of cryptocurrencies [32, 22, 28, 37, 1]. In this setting, signing keys protect large sums of money, and the motivation to misuse a key is very high.

Threshold Schnorr signatures. Schnorr signatures [38] have been around for over three decades, and are in wide use in the form of EdDSA and in the cryptocurrency space. There have been many threshold signature schemes designed for Schnorr signatures. However, to the best of our knowledge, these schemes are *all* either for the honest majority setting (e.g. [23, 42]), are proven secure via a reduction to the discrete log problem in the random-oracle model (e.g., [35, 1, 3]), or use non-standard assumptions like the one-more discrete

E-mail: yehuda.lindell@gmail.com (Yehuda Lindell)

log assumption and/or idealized generic group model assumptions (e.g., [28, 1, 37, 14]). When reducing to the discrete log problem directly, these proofs of security rely on variants of the forking lemma. This impacts the tightness of the reduction, as well as the ability to achieve concurrent security. Although it is possible to prove concurrent security, this is very tricky and error prone, as was shown in [18, 4].

Full simulation. In the standard MPC paradigm, protocols are proven secure by showing that they securely realize a certain “ideal functionality” [8, 9, 25]. In the context of threshold signatures, this can be formalized by defining an ideal functionality that simply generates a standard Schnorr signature on a message m after receiving an approval to sign on the message from a quorum of parties. Such an approach is actually agnostic to the question of whether or not Schnorr signatures themselves are secure. This has the advantage of not requiring the forking lemma, rewinding and so on in order to prove the protocol secure. In addition, the assumptions and model needed to prove the protocol secure need not be the same as for the signature scheme itself (e.g., the protocol may or may not rely on a random oracle). This may sound strange: if Schnorr signatures are proven secure in the random-oracle model, then what value is there in proving a threshold-signing protocol secure in the standard model? The answer is that after over three decades in existence one can reasonably just assume that Schnorr signatures are secure as an assumption in itself (similarly to what we do for ECDSA since we have no reasonable proof). In this light, it is of value to construct a secure threshold signature scheme for Schnorr that reduces to the *assumption that Schnorr signatures are secure* and nothing else.

We argue that there are significant advantages to this approach. First, it significantly reduces the assumptions required.¹ Second, simulation can be achieved without rewinding, providing tight security and concurrent composition. Third, as we see here, the proof of security is far more straightforward and thus less error prone.

Our results. In this paper, we construct a multiparty signing (and key generation) protocol that securely realizes the ideal functionality that computes Schnorr signatures, in the presence of a malicious static adversary corrupting any number of parties. Since we consider the dishonest-majority setting, we do not achieve guaranteed output delivery or fairness. Our protocol has three rounds of communication, supports quorum thresholds (including AND/OR combinations of threshold sets), and is extremely simple and efficient. We prove the security of our protocol under the standard real/ideal paradigm for MPC [8, 25]. In a hybrid model with an ideal zero-knowledge and commitment functionality, our simulator does not rewind the adversary, and the simulation is perfect. Thus, if the zero-knowledge and commitment functionalities used are UC-secure [9], then the entire protocol is UC-secure as proven in [29]. This means that the protocol is secure under composition, when run concurrently with arbitrary other secure and insecure protocols.

In many protocols in the literature, the parties are assumed to all know the message being signed upon, a unique session identifier (where needed) and possibly also the identity of the participating parties, *before* the protocol begins. However, in practice, this information needs to be agreed upon by the parties, and even generated by them (as in the case of a session identifier). Furthermore, parties are assumed to communicate via private point-to-point channels, whereas in practice communication is often carried out by a “central (untrusted) coordinator”. In the case of threshold cryptography, this is usually the case since different subsets of parties carry out different operations, and it’s unlikely that they have anything more than the public-keys of the other parties. Of course, one can use these keys to encrypt and sign, and can even set up secure channels if needed. However, this can add overhead to the protocol and is often not required. We therefore describe

¹The most efficient instantiations of the protocol (specifically, for the commitments and zero-knowledge proofs) will require a random oracle. However, its usage is limited to these primitives (making its use not integral to the protocol) and non-random oracle instantiations can be used if desired.

our protocol in a very basic setup where the parties only have a public-key infrastructure, and all communication is via an untrusted coordinator. This means that our three-round protocol is truly three rounds without assuming any prior consensus on the message to be signed, the participating parties, or a unique session identifier.

Protocol idea. Our protocol is very simple, and is indeed what one would expect for Schnorr signatures (in fact, we were surprised not to find this protocol in the literature). Recall that a Schnorr signature (over the group \mathbb{G} of order q with generator G) is of the form (e, s) where $e = H(m\|Q\|R)$ for $R = k \cdot G$, and $s = k - e \cdot d \pmod{q}$, where the private key is $d \in \mathbb{Z}_q$ and the public key is $Q = d \cdot G$. Our protocol works by the parties first running a simulatable coin tossing by each choosing a random k_i and committing to $R_i = k_i \cdot G$. Then, the parties decommit to reveal R_i and provide a zero-knowledge proof of knowledge of the discrete log k_i . Finally, each party computes $R = \sum R_i$, $e = H(m\|Q\|R)$, and $s_i = k_i - e \cdot d_i \pmod{q}$, where the d_i values are an additive sharing of the private key d . The resulting Schnorr signature on m is (e, s) where $s = \sum s_i \pmod{q}$. The reason that this is fully simulatable is that the coin tossing enables the simulator to “force” the R value to be that defined by the signature (e, s) received from the ideal functionality (by computing $R = s \cdot G + e \cdot Q$), and the zero-knowledge proofs enable the simulator to extract all of the corrupted k_i values and exactly compute the s_i value that each honest party would send in a real protocol. In more detail, the simulator can know all of the corrupted k_i and d_i values, as well as all of the honest k_i and d_i values except for one. This means that it can exactly compute the s_i value that all but one honest party would compute in the protocol. Then, given the real signature s from the ideal functionality, the simulator can compute the final honest party’s s_j value as $s - \sum_{i \neq j} s_i$. This strategy is slightly harder to implement when there are different parties in each execution (since the “one” honest party can be different each time, and all honest parties’ values have to be consistent across all executions), but the values can be obtained by applying Lagrange interpolation appropriately.

On efficiency, optimizations and conservative design. As we have discussed, the literature is full of Schnorr threshold signature protocols, and many of those are more efficient than our protocol. These protocols sometimes have only two rounds of communication, and do not use zero-knowledge proofs. We argue that although such optimizations are important in some cases,² they are not needed in many other cases. For example, for consumer cryptocurrency wallets that run MPC between a user’s mobile and a server, it makes no difference if a signature takes 10ms or 100ms. This is also true of most custody use cases for cryptocurrency where the number of transactions a day is not counted in the millions (and even if they are, extremely low latency isn’t required, and throughput can easily be increased by adding machines). Likewise, if such protocols are used for code signing (albeit unlikely today since RSA or ECDSA are typically used for that), then the number of signatures generated is small. In such cases, two rounds versus three rounds does not make any difference,³ and we have no need to overly optimize the number of exponentiations (e.g., by not having a zero-knowledge proof). In these cases, especially where the keys are extremely valuable (e.g., protecting large amounts of cryptocurrencies or used to sign highly valuable code), we argue that a conservative approach makes the most sense. It is interesting that for standard cryptographic operations like signing and encryption, industry is extremely conservative, adopting new schemes only after many years of careful evaluation. However, when it comes to threshold signing, many protocols

²For example, consider the application of using MPC for the signing operation in a TLS handshake on a popular web server. In this case, reducing latency and increasing throughput without requiring additional computing resources can be important.

³There are some use cases where disconnected machines run some of the MPC parties, and a person has to physically deliver the MPC messages to those machines. In such cases, two versus three rounds may be significant. Whether or not that warrants less conservative design and assumptions depends on the risk appetite of the organization deploying them.

rely on non-standard assumptions, some of them new and unstudied, and sometimes even heuristic arguments of security. This is not prudent for deployment in practice when so much is at stake. The protocol in this paper is aimed at use cases where shaving a few milliseconds off is not of significance, and being highly confident in the security of the protocol takes priority.

The simulation versus game-based approach. As we have mentioned, the fact that our protocol is UC-secure means that it is guaranteed to remain secure under concurrent composition with arbitrary other protocols. Another very significant advantage of this approach is that our protocol remains secure even if related keys are used (something which cannot be claimed for game-based proofs, unless specifically taken into account). This holds since the signing protocol is proven secure independently of the key generation protocol. Thus, in the UC model, the environment can provide related keys to the parties and the proof of security remains unchanged. Of course, the use of related keys may result in Schnorr itself no longer being secure, but this is not due to the protocol. Although seemingly strange, this is actually an extremely relevant issue since many (or most) cryptocurrency wallets use key derivation techniques that result in related keys (e.g., the normal derivations in [43]). With our approach, it suffices to prove that the Schnorr signature scheme remains secure with these derivation methods, and this automatically implies that the threshold protocol remains secure as well.

A note on EdDSA. The EdDSA signing scheme [5] is a Schnorr variant where the randomness used to generate the signature is derived deterministically from the key, using a pseudorandom function. This design is intended to prevent biased and reused nonces which are catastrophic in Schnorr (and ECDSA) signatures. Our protocol can be used for EdDSA, but the result is not actually compliant with the standard since it is possible to generate two different signatures on the same message since it is not deterministic. As long as only one signature is generated, this is indistinguishable. However, if a higher-level protocol relies on the signature being deterministic, then this is a problem. (Having said that, if a higher-level protocol relies on it being deterministic, then a party can easily break it by generating two different signatures, so higher-level protocols should not rely on this.)

Related work. As described, there has been considerable prior work on constructing threshold signatures. Some of this dates back to the early 2000s (e.g., [35, 36]), while other work is very recent (e.g. [28, 1, 37, 14]). However, to the best of our knowledge, no protocol in the literature has been proven secure under the standard ideal/real model paradigm for MPC, with a functionality just computing Schnorr signatures. Furthermore, most recent work is based on non-standard assumptions, some of them interactive, non-falsifiable, or utilizing generic and algebraic group models. This is the primary novelty of this paper.

Subsequent work. In subsequent work to this paper, [34] show that three-round Schnorr is secure with simulation-based security, even without zero-knowledge proofs of knowledge (which we require). However, their ideal functionality does not compute a Schnorr signature, but rather an ideal signing functionality that hides the signature contents (and is secure assuming the discrete log problem is hard). This provides a different security guarantee, and it isn't clear that security is maintained when keys are not generated randomly, and in particular when BIP032 key derivation [43] or the like is used. To an extreme, assume that a poor key generation method is used so that signatures can be forged with some non-negligible probability. Then, our protocol will maintain this probability, and security will be the same as for locally generated signatures. In contrast, it isn't clear that any security is achieved by the protocol in [34] in such a case. Another subsequent work [12] also considers three-round Schnorr without zero-knowledge proofs. They also study adaptive security, but work with game-based definitions, which are fundamentally different as discussed above. In particular, security is not implied for related keys, poorly generated keys, and so on.

Organization. In Section 2, we present preliminary definitions and notation. Then, in Section 3 we show how an ideal commitment functionality can be UC-realized without a unique session identifier. This is needed in order to achieve a three-round protocol without assuming that the parties already hold a unique session identifier. Next, in Section 4, we present and prove the protocol for the case that the parties are given ahead of time a session identifier, the message to be signed and the set of participating parties. In Section 5 we extend the protocol to the case that the parties have only their shares of the private key (and a PKI of public signing keys) at the onset of the protocol, in Section 6 we show how identifiable abort is added, and in Section 7 we show how to achieve efficient UC zero-knowledge. Finally, in Appendix A we show how distributed key generation can be achieved using a similar methodology.

2 Preliminaries

We denote the computational security parameter by κ , and so all parties run in time that is polynomial in κ .

The Schnorr signing algorithm. The Schnorr signing algorithm is defined as follows. Let \mathbb{G} be an Elliptic curve group of order q with base point (generator) G . The private key is a random value $d \leftarrow \mathbb{Z}_q$ and the public key is $Q = d \cdot G$. The Schnorr signing operation on a message $m \in \{0, 1\}^*$ is defined as follows:

1. Choose a random $k \leftarrow \mathbb{Z}_q$
2. Compute $R = k \cdot G$.
3. Compute $e = H(m \| Q \| R)$; different Schnorr variants hash different values and in different orders, but this is inconsequential.
4. Compute $s = k - e \cdot d \pmod{q}$; some Schnorr variants add rather than subtract here, but this is also inconsequential.
5. Output (e, s)

In order to verify a signature given the public-key Q , compute $R = s \cdot G + e \cdot Q$ and accept if and only if $e = H(m \| Q \| R)$.

The ideal multiparty zero knowledge functionality \mathcal{F}_{zk} . We use the standard ideal zero-knowledge functionality defined by $((x, w), \lambda) \rightarrow (\lambda, (x, R(x, w)))$, where λ denotes the empty string, with the only difference being that the proof is sent to all parties. For a relation R , the functionality is denoted by $\mathcal{F}_{\text{zk}}^R$.

FIGURE 1 (The Zero-Knowledge Functionality $\mathcal{F}_{\text{zk}}^R$ for Relation R).

Functionality \mathcal{F}_{zk} works with parties P_1, \dots, P_n , as follows:

- Upon receiving $(\text{prove}, \text{sid}, S, x, w)$ with $S \subset [n]$ from a party P_i (for $i \in [n]$): if $(x, w) \notin R$ or sid has been previously used then ignore the message. Otherwise, send $(\text{zkproof}, \text{pid}_i, \text{sid}, x)$ to party P_j for every $j \in S$, where pid_i is P_i 's unique party identifier.

Note that any zero-knowledge proof of knowledge securely realizes the \mathcal{F}_{zk} functionality [27, Section 6.5.3]; non-interactive versions can be achieved in the random-oracle model via the Fiat-Shamir paradigm with security under sequential composition. In order to achieve this under concurrent composition, the rewinding of the random oracle causes a problem. This can be solved (in general, and specifically for proving knowledge of the discrete log) using the methods described in Section 7. We therefore assume Functionality 1 and UC-security in our presentation. We use this specifically for the discrete log relation

defined by $R_{\text{dl}} = \{((\mathbb{G}, G, P), w) \mid P = w \cdot G\}$, where \mathbb{G} is the description of a group with generator G . In this case, we denote the functionality by $\mathcal{F}_{\text{zk}}^{\text{dl}}$. In Appendix A, we also refer to a batch discrete log proof, which is simply the proof of knowledge of many discrete logs in parallel. This can be achieved at almost the same cost as a single discrete log proof using the protocol of [24], but not for UC security (for which the method of Section 7 must be used, and for which the batch proof is more costly).

Security with abort, the hybrid model and composition. We prove the security of our protocol under simulation-based ideal/real model definitions that are standard for MPC [8, 25, 9]. We consider a dishonest majority here, and therefore security with abort. This is modeled in the ideal model by the adversary first receiving output, and then instructing the trusted party computing the functionality who should receive output. Our protocol is proven secure in a hybrid model with zero-knowledge and commitment ideal functionalities. The soundness of working in this model is justified in [8, 25] (for stand-alone security) and in [9] (for UC security under concurrent composition). Specifically, as long as subprotocols that securely compute the functionalities are used (under the definition of [8] or [9], respectively), it is guaranteed that the output of the honest and corrupted parties when using real subprotocols is computationally indistinguishable to when calling a trusted party that computes the ideal functionalities. We note that any protocol that is perfectly secure (or computationally secure with input synchronization) and is proven via a simulator that does not rewind is UC secure [29].

The communication model and PKI. The standard real model for communication in MPC is that of point-to-point authenticated (or even private) channels between all parties. This can be achieved in any model of communication (as long as we don't assume guaranteed delivery of messages) by signing (and encrypting if privacy is needed) each message sent from each party. In the case of broadcast (with abort) over point-to-point authenticated communication channels, it is possible to use a simple echo-broadcast, as shown in [26]. Having said the above, in some cases, it is possible to achieve simpler and more efficient communication patterns than these generic transformations. In many real-world settings where there is a threshold of participants, all messages are sent via a central coordinator. This is due to the fact that the effort required in agreeing on the participating parties and then setting up direct secure channels is larger than just sending all messages via such a coordinator machine. Since communication via a coordinator machine is more complex than assuming broadcast and pairwise point-to-point channels, and since we view this as a likely real-world deployment model, we adopt this as the model of communication in our protocol description. We therefore also include an explicit PKI for signing all messages (our protocol does not require private channels, and thus signing is enough). As we will see, we also achieve implicit broadcast in the protocol by running a signed echo-broadcast on the messages sent, where needed.

We remark that since we consider security with abort (which is inherent to the dishonest majority setting), the communication can be completely asynchronous as parties can simply wait to receive all of the current-round messages before sending their next-round message. (This isn't always the case in protocols where not all parties communicate in each round; however, in our protocols, all parties communicate in all rounds.)

3 Multiparty Commitments Without SIDs

We use commitments in the ‘‘coin tossing’’ phase of signing in order to choose the nonce for the signature, and in distributed key generation to generate a random key. As discussed, we prove our protocol secure assuming access to an ideal commitment functionality. For this purpose, we can use any UC-secure commitment scheme (e.g., [30, 6, 21]). In the random-oracle model, the basic UC-secure commitment can be trivially realized with static

security by simply defining $\text{Com}(\text{sid}, x) = H_{\text{ro}}(\text{sid} \| x \| r)$ where $r \leftarrow \{0, 1\}^\kappa$ is random (and κ is a computational security parameter of a fixed and known length, and sid is a unique session identifier of known length).

In our protocol for Schnorr signing, the parties need to commit in the first round. However, in a setting with no prior setup, the parties need to generate the sid . This takes at least one round, and so the protocol will be four and not three rounds.⁴ In order to overcome this, we show that in the random-oracle model, it is possible to securely compute the ideal commitment functionality with UC security *without* a session identifier. In general, session identifiers are used to prevent copying from one session to another and from other parties. However, copying from other parties can easily be prevented by including *party identifiers*, and copying from previous executions is actually not necessarily a problem. In order to see why an sid isn't needed, note that when the committing party's identifier is included in the commitment, a corrupted party can only "copy" a commitment by **(a)** taking the content from a previous already opened commitment from a different honest or corrupted party, or **(b)** using a commitment that it itself sent in parallel or concurrently. However, an adversary can do these anyway in the perfect $\mathcal{F}_{\text{mcom}}$ -hybrid model (i.e., with an ideal multiparty commitment functionality). Regarding **(a)**: once a commitment has been opened, any party can commit to that value with $\mathcal{F}_{\text{mcom}}$. Regarding **(b)**: a party can commit to the same value concurrently using $\mathcal{F}_{\text{mcom}}$ with different sid 's and this will have the same effect as achieved here. We therefore conclude that a multiparty commitment functionality $\mathcal{F}_{\text{mcom}}$ can be securely realized in the standard way, with the exception that the commitment does *not* need to include an sid (as long as the committing party's pid is included, to prevent unopened commitments from being copied). In particular, if a corrupted P_i sends a commitment value that has not been seen before, then it is dealt with in the usual way by the simulator for $\mathcal{F}_{\text{mcom}}$ (either extracting successfully or concluding that the adversary will not be able to provide a valid decommitment).

Having said the above, there is a technical challenge in the ideal-model simulation to *match* commitments to decommitments with no identifier (in the real protocol this isn't a problem since the commitment value can be recomputed from the decommitment and matched). If party P_i sends two commitments to P_j , then how does the simulator distinguish which decommitment is associated with which commitment? If the committing party is corrupted, then the simulator can actually do that without any problem since it received the commitment value from the adversary and so can compare. However, if the committing party is honest, then the simulator is unable to differentiate since all it received from the ideal functionality is a blank receipt message. In order to solve this, we have each party P_i input a random/unique sid_i that they chose themselves when committing. When the committing party is honest, this identifier will be unique and so will enable the simulator to connect commitments to decommitments. Of course, when the committing party is corrupted, nothing stops them reusing an identifier, but as we have discussed, this doesn't affect security. Furthermore, if a corrupted party uses the same identifier for two *different* commitment values, in the real protocol we differentiate these by the commitment values themselves.

In our protocol, we also wish to ensure that all honest parties receive the *same* commitment from the committer. As such, it is actually a *broadcast commitment* functionality. This is achieved by running an echo-broadcast on the commitment. That is, the committer sends its commitment to all parties, and all parties then send what they received to all others. If a party sees the same commitment value from all, then it knows that all honest parties have the same commitment, and it accepts. Otherwise, it aborts. The fact that this achieves a UC-secure broadcast with non-unanimous abort (i.e., there exists a single value so that each honest party either outputs that value or aborts) was shown in [26]. As

⁴It is possible to agree on an sid in one round and not more since it needs to be unique only, and does not need to be random. Thus, it suffices for each party to send a random identifier and to set the sid to be the collision-resistant hash of the concatenation of all of the sent identifiers.

we have discussed, in some settings, the subset of participating parties is known only in the second round. Therefore, the consensus on who is participating and who received the broadcast is only obtained at the end. The ideal multiparty broadcast commitment $\mathcal{F}_{\text{mcom}}$ is defined in Functionality 2.

FIGURE 2 (The Broadcast Commitment Functionality $\mathcal{F}_{\text{mcom}}$).

Functionality $\mathcal{F}_{\text{mcom}}$ works with parties P_1, \dots, P_n and adversary \mathcal{S} , as follows:

- Upon receiving $(\text{commit}, \text{sid}_i, \text{pid}_i, S, x)$ with $S \subset [n]$ from party P_i : store $(\text{sid}_i, \text{pid}_i, S, x)$ and send $(\text{receipt}, \text{sid}_i, \text{pid}_i, S)$ to \mathcal{S} and every P_j with $j \in S$.
- Upon receiving $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x)$ from party P_i , if $(\text{sid}_i, \text{pid}_i, S, x)$ is recorded then send $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x)$ to \mathcal{S} and to P_j for every $j \in S$.

Note that the decommitment also needs to include the commitment value x since there is no other way to identify exactly what is being decommitted when the committing party is corrupted and so may use the same sid_i for two different commitments (unlike regular UC-commitments, the functionality is willing to accept multiple commitments with the same sid_i). We now prove that the standard UC-commitment in the random-oracle model together with an echo-broadcast securely realizes $\mathcal{F}_{\text{mcom}}$. The protocol assumes that the lengths of the bit representation of sid_i , pid_i and S are fixed and known, but any unambiguous encoding of the strings would be equivalent. For simplicity, we present the commitment protocol assuming point-to-point channels, and later explain how it is modified for the case of communication via a central coordinator, using a PKI of public signing keys.

Our protocol is proven secure in the random-oracle model. We model this via an ideal random-oracle functionality \mathcal{F}_{rom} that works as follows. Upon receiving (ROM, x) , functionality \mathcal{F}_{rom} checks if a pair $(x; y)$ is stored for some $y \in \{0, 1\}^\kappa$. If yes, it returns y . Else, it chooses a random $y \leftarrow \{0, 1\}^\kappa$, stores $(x; y)$, and returns y . We prove that Protocol 3 securely realizes $\mathcal{F}_{\text{mcom}}$ for the case that honest parties always input unique sid_i values. This is achieved in practice by simply choosing $\text{sid}_i \in \{0, 1\}^\kappa$ randomly.⁵

PROTOCOL 3 (Broadcast Commitment Protocol).

- **Commit:** Upon input $(\text{commit}, \text{sid}_i, \text{pid}_i, S, x)$, party P_i works as follows:
 1. *Message 1 (broadcast)* – P_i to all: P_i chooses a random $r \leftarrow \{0, 1\}^\kappa$, computes $c = \text{Com}(\text{sid}_i, \text{pid}_i, S, x; r) = H_{\text{ro}}(\text{sid}_i \parallel \text{pid}_i \parallel S \parallel x \parallel r)$ and sends $(\text{sid}_i, \text{pid}_i, S, c)$ to P_j for every $j \in S$.
 2. *Message 2 (echo)* – all to all: Upon receiving $(\text{sid}_i, \text{pid}_i, S, c)$, if $j \in S$ then party P_j sends $(\text{sid}_i, \text{pid}_i, S, c)$ to P_ℓ for every $\ell \in S$. Else, it ignores the message.
 3. *Output decision:* If P_j received the same $(\text{sid}_i, \text{pid}_i, S, c)$ from all parties (including itself), then it outputs and stores $(\text{sid}_i, \text{pid}_i, S, c)$. Otherwise, it ignores the commitment.
- **Decommit:** Upon input $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x)$, party P_i works as follows:
 1. P_i sends $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x, r)$ to P_j for every $j \in S$.
 2. Upon receiving $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x, r)$, party P_j computes $c = H_{\text{ro}}(\text{sid}_i \parallel \text{pid}_i \parallel S \parallel x \parallel r)$. If it had previously stored $(\text{sid}_i, \text{pid}_i, S, c)$ then it outputs $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x)$. Otherwise, it ignores the message.

⁵Note that this can be enforced by the functionality rather than assuming it, by having the ideal functionality reject a repeated sid_i from an honest party, while allowing a corrupted party to use the same sid_i multiple times. However, this involves the functionality to be aware of which parties are honest and which are corrupted. This can be modeled, but unnecessarily complicates the treatment here.

Proposition 1. *Let H_{ro} be modeled by the ideal random oracle functionality \mathcal{F}_{rom} . Then, Protocol 3 UC-securely realizes \mathcal{F}_{mcom} in the \mathcal{F}_{rom} -hybrid model, in the presence of a malicious static adversary, in the case that honest parties always use unique local sid_i identifiers.*

Proof. We prove security according to [10] where \mathcal{S} delivers all messages between \mathcal{F}_{mcom} and the parties. (Note that in this model, the public header for \mathcal{F}_{mcom} contains the type of message (commit or decommit) along with sid_i, pid_i, S , and the private content is the message being committed to.) Let \mathcal{A} be a real (dummy) adversary who simply forwards messages to and from the environment \mathcal{Z} . Let $I \subset [n]$ be the set of corrupted parties. We now describe the ideal-model adversary/simulator \mathcal{S} :

1. *Random-oracle queries:* When \mathcal{A} sends any message $(ROM, sid_i, pid_i, S, x, r)$ to \mathcal{F}_{rom} , simulator \mathcal{S} works exactly as the real \mathcal{F}_{rom} except that it doesn't allow any collision. That is, \mathcal{S} checks if it has stored $(ROM, sid_i, pid_i, S, x, r; c)$ for some $c \in \{0, 1\}^\kappa$. If yes, it returns c . Else, it chooses a random $c \leftarrow \{0, 1\}^\kappa$. If the chosen c has already been chosen previously, then \mathcal{S} outputs $fail_1$. Else, it stores $(ROM, sid_i, pid_i, S, x, r; c)$, and returns c .
2. *Commitments from corrupted parties:* Upon receiving (sid_i, pid_i, S, c) from \mathcal{A} for some $i \in I$, as a message sent from a corrupted P_i to some honest party P_j , simulator \mathcal{S} checks if a tuple $(ROM, sid_i, pid_i, S, x, r; c)$ has been stored for some $x \in \{0, 1\}^*$ ($x \neq \perp$). If no, then \mathcal{S} stores (ROM, \perp, c) and ignores the message. If yes, but the stored tuple doesn't represent a valid commitment string, e.g., not having a correct pid_i , then it is also ignored. If yes and the stored tuple is of valid format, then \mathcal{S} sends $(commit, sid_i, pid_i, S, x)$ to \mathcal{F}_{mcom} and proceeds as follows:

\mathcal{S} plays the honest parties in the echo-broadcast to \mathcal{A} for this commitment.
For every honest party P_j who concludes successfully and would store (sid_i, pid_i, S, c) , simulator \mathcal{S} delivers the message $(receipt, sid_i, pid_i, S)$ from \mathcal{F}_{mcom} to P_j .

3. *Decommitments from corrupted parties:* Upon receiving a message $(decommit, sid_i, pid_i, S, x, r)$ from \mathcal{A} for some $i \in I$ to be sent to some honest P_j , simulator \mathcal{S} checks if some $(ROM, sid_i, pid_i, S, x, r; c)$ has been stored. If yes, then it sends $(decommit, sid_i, pid_i, S, x)$ to \mathcal{F}_{mcom} and delivers the decommit message to P_j . If no, then it chooses a random $c \leftarrow \{0, 1\}^\kappa$. If (ROM, \perp, c) has been stored, then \mathcal{S} outputs $fail_2$. Else, \mathcal{S} ignores this decommit message.
4. *Commitments from honest parties:* Upon receiving $(commit, sid_j, pid_j, S)$ from \mathcal{F}_{mcom} for some honest $j \notin I$, simulator \mathcal{S} chooses a random c , stores $(ROM, sid_i, pid_j, S, *, *, c)$ and simulates the honest parties running an echo broadcast with P_j sending (sid_j, pid_j, S, c) to all parties. (Note that if such a c has already been stored, then \mathcal{S} outputs $fail$.)
5. *Decommitments from honest parties:* Upon receiving $(decommit, sid_j, pid_j, S, x)$ from \mathcal{F}_{mcom} for some honest $j \notin I$, simulator \mathcal{S} chooses a random $r \leftarrow \{0, 1\}^\kappa$ and updates the stored tuple $(ROM, sid_j, pid_j, S, *, *, c)$ to the full tuple $(ROM, sid_j, pid_j, S, x, r; c)$. Then, \mathcal{S} simulates P_j sending the decommitment message $(decommit, sid_j, pid_i, S, x, r)$ to all parties.

We claim that as long as \mathcal{S} does not output $fail$, the view of the environment in the ideal and real executions are identical. In order to see this, observe that if \mathcal{S} does not output $fail_1$ then there are never any collisions in c . Furthermore, if \mathcal{S} does not output $fail_2$ then \mathcal{A} must have queried the random oracle with a commitment value before it decommits. Together, this implies that the commitment and decommitment messages sent by \mathcal{S} to \mathcal{F}_{mcom} map exactly to the commitment and decommitment messages of \mathcal{A} in a real execution.

Furthermore, commitments and decommitments made by honest parties are perfectly simulated. In particular, the commitment values have the same distribution since \mathcal{S} runs the random oracle. Thus, it can first choose c and later on “fill in” the commitment value $(\text{sid}_j, \text{pid}_j, S, x, r)$ by programming the oracle, and this has exactly the same distribution of a real execution. (Note that the correct mapping of commitments to decommitments is achieved under the assumption that an honest party P_j never uses the same sid_j in two different commitments.)

It is easy to see that \mathcal{S} outputs fail with negligible probability (since a collision with a random oracle happens with negligible probability, and the probability that \mathcal{S} chooses a random c that equals a commitment value previously sent by \mathcal{A} is also negligible). We therefore conclude that the view of the environment in a real execution is statistically close to its view in an ideal execution, as required. ■

Realizing $\mathcal{F}_{\text{mcom}}$ via a coordinator. In our main communication model, all messages are sent via a coordinator. Since we also need to ensure consistency of commitments sent to all honest parties, this is achieved by having the committer send its commitment message signed to the coordinator, who forwards it to all parties. Then, each party signs on what it received and sends it back to the coordinator, who forwards all signatures to all parties. The commitment is accepted only if all parties have signed on the same commitment string. Since all messages are signed, any modification made by the coordinator or any party sending different messages to different parties will result in an abort. This is valid since the second round of the commitment protocol is only used to validate that the committer sent the same commitment to all parties.

Parallel commit and decommit. The commitment phase of the protocol takes two rounds of communication. However, if we need to run a “coin tossing” flow, where all parties commit and then all parties decommit, this can all be achieved in two rounds as well. This works by verifying that all signatures are to the same commitment strings of all parties in the decommitment round, and accepting the commitment and decommitment only in that case.

Instantiating the random-oracle commitment. A good instantiation is to compute $\text{Com}(\text{sid}_i, \text{pid}_i, S, x; r) = H_{\text{ro}}(\text{sid}_i \parallel \text{pid}_i \parallel S \parallel x \parallel r) = \text{HMAC}_r(\text{sid}_i \parallel \text{pid}_i \parallel S \parallel x)$, where r is of a fixed (predetermined) length. That is, use r as the HMAC key and the rest as input. This has the advantage that it is a secure commitment scheme in the standard model. In particular, hiding holds under the assumption that HMAC is a pseudorandom function, and binding follows from the collision resistance of the hash function. Of course, in order for it to be a UC-secure commitment, we need the random-oracle properties so that the simulator can extract committed values from corrupted parties and equivocate simulated commitments from honest parties.

4 Multiparty Schnorr Signing – Assuming Initialization

4.1 The Protocol

In this section, we present the basic protocol for a set of parties who hold Shamir shares of the private key. In this presentation, we assume that the set of $t + 1$ participating parties $S \subset [n]$ is known, and that all parties hold the same message m . Furthermore, we assume that a global sid is known to all parties in the second round. In Section 5 we will present the general protocol where parties begin from nothing except their share of the private key and a PKI. The protocol can be presented more succinctly assuming pairwise communication between all parties, and assuming the ideal $\mathcal{F}_{\text{mcom}}$ and \mathcal{F}_{zk} functionalities (see Protocol 7 in the proof of Theorem 6). Nevertheless, as mentioned in the introduction, we present this more detailed and specific description here where all communication is via

a coordinator and signatures are used to achieve consistency of views of all honest parties, since this is how it will most likely be implemented in practice.

PROTOCOL 4 (Multiparty Schnorr Signing).

Input: Each party in the set S of parties has the Schnorr public key Q , the set of participating parties S , a session identifier sid (from the second round), the message to be signed m , a PKI of signing keys $\{pk_i\}_{i \in S}$, its private signing key sk_i , and its private input which is a Shamir share of the private key d where $d \cdot G = Q$.

The protocol:

Before beginning, each party P_i updates its Shamir share to be an additive share d_i of the private key, using the Lagrange coefficients for the set S .

1. **Message 1 – all to \mathcal{C} :** Each party P_i works as follows:
 - (a) P_i chooses a random $k_i \leftarrow \mathbb{Z}_q$ and sets $R_i = k_i \cdot G$.
 - (b) P_i chooses a random $\text{sid}_i \leftarrow \{0, 1\}^\kappa$ and a random $r_i \leftarrow \{0, 1\}^\kappa$ and sets $c_i = H(\text{sid}_i \| \text{pid}_i \| S \| R_i \| r_i)$.
 - (c) P_i sends $(\sigma_i^1, \text{sid}_i, c_i)$ to the coordinator \mathcal{C} , where $\sigma_i^1 = \text{sign}_{sk_i}(1, \text{sid}_i, c_i)$ using its signing-key sk_i from the PKI.
2. **Transmission 1 – \mathcal{C} to all:** \mathcal{C} receives all $(\sigma_i^1, \text{sid}_i, c_i)$ messages, and sends $\{(\sigma_i^1, c_i)\}_{i \in S}$ to all parties.
3. **Message 2 – all to \mathcal{C} :** Each party P_i works as follows:
 - (a) P_i verifies that it received $(\sigma_j^1, \text{sid}_j, c_j)$ for every $j \in S$, that c_i as it sent in the first message appears in the set, and that all signatures are valid. If not, it aborts.
 - (b) P_i computes $\pi_i \leftarrow \text{ZKDL}_P(\text{sid}, \text{pid}_i, R_i; k_i)$ (where ZKDL denotes a Fiat-Shamir proof of knowledge of the discrete log, pid_i is the known identity or public-key of P_i , and sid is the session identifier).
 - (c) P_i sends $(\sigma_i^2, R_i, r_i, \pi_i)$ to the coordinator \mathcal{C} , where $\sigma_i^2 = \text{sign}_{sk_i}(\text{sid}, 2, \{c_i\}_{i \in S}, R_i, r_i, \pi_i)$ using its signing-key sk_i from the PKI.
4. **Transmission 2 – \mathcal{C} to all:** \mathcal{C} receives all $(\sigma_i^2, R_i, r_i, \pi_i)$ messages, and sends $\{(\sigma_i^2, R_i, r_i, \pi_i)\}_{i \in S}$ to all parties.
5. **Message 3 – all to \mathcal{C} :** Each party P_i works as follows:
 - (a) After receiving all $\{(\sigma_j^2, R_j, r_j, \pi_j)\}_{j \in S}$, party P_i verifies all signatures (using its sid and the series of commitments it received from \mathcal{C} in the first round). If not valid, it aborts. If valid, it proceeds.
 - (b) For every $j \in S$ ($j \neq i$):
 - i. P_i verifies that $H(\text{sid}_j \| \text{pid}_j \| S \| R_j \| r_j) = c_j$, and that all values are valid (i.e., it has the correct $\text{sid}_j, \text{pid}_j$ and S , and overall structure).
 - ii. P_i verifies $\text{ZKDL}_V(\text{sid}, \text{pid}_j, R_j; \pi_j) = 1$.
 - iii. If the commitment is not valid, or R_j is not a valid point in the curve subgroup, or R_j is equal to the identity point, or if $\text{ZKDL}_V(\text{sid}, \text{pid}_j, R_j; \pi_j) = 0$ (i.e., the ZK verification fails) then P_i aborts. Else, it proceeds.
 - (c) P_i sets $R = \sum_{j \in S} R_j$, $e = H(m \| Q \| R)$ and $s_i = k_i - d_i \cdot e \pmod{q}$, with the exact hash and s_i formula as needed for the Schnorr variant.
 - (d) P_i sends s_i to \mathcal{C}
6. **Output:** \mathcal{C} sets $s = \sum_{i \in S} s_i \pmod{q}$ and checks that $\text{verify}_Q(m, (s, e)) = 1$. If yes, then it outputs (e, s) ; otherwise it aborts.

Communication pattern. Observe that the communication pattern in both the consensus phase and in Protocol 4 is *all parties send a message to \mathcal{C} , and \mathcal{C} sends all messages together to all parties*. This simple communication pattern is easy to implement.

4.2 Proof of Security

We begin by defining the ideal functionality \mathcal{F} for signing. We wish our protocol to be secure for *any* key generation (trusted key generation, distributed key generation, hierarchical-deterministic key generation via BIP-032 [43] and so on). As such, we define the ideal functionality for the signing process only, and where parties input their key shares, however they may have been generated. This guarantees that the signing protocol provides the same level of security as a local signing operation, irrespective of how the keys are generated.

The natural way to model this is that each party inputs their private share d_i and the public key Q , with the guaranteed that all $t + 1$ input d_i values lie on a single polynomial $p(x)$ such that $p(0) \cdot G = Q$. (That is, the input shares define the correct private key.) However, in many cases, especially when distributed key generation is used (e.g., with Feldman VSS), the parties also all have the vector (Q_1, \dots, Q_n) where P_i 's private share is d_i and $Q_i = d_i \cdot G$. We therefore prove security even when this additional information is held by the parties (and the adversary).

We prove that our protocol is secure for *any* set of inputs, as long as all parties hold consistent and correct inputs. This holds as long as **(a)** all parties hold the same vector (Q, Q_1, \dots, Q_n) , **(b)** there exists a polynomial $p(x)$ such that $p(0) \cdot G = Q$ and $p(i) \cdot G = Q_i$ for every $i = 1, \dots, n$, and **(c)** each party's private input share d_i is correct, meaning that $Q_i = d_i \cdot G$. If this holds, we say that the inputs are **valid and consistent**. We stress that if the parties are not guaranteed to have valid and consistent inputs, then each party should verify that their input share is correct, and the parties should run an echo-broadcast on the vector (Q, Q_1, \dots, Q_n) . We do not include this in our protocol since in practical settings, this input is the result of a secure distributed key generation or the like, and so the parties' inputs are guaranteed to indeed be valid and consistent.

We define the functionality such that only \mathcal{C} receives output. This is sufficient since \mathcal{C} can always forward the signature to any other party, and since the signature can be verified the output cannot be modified. Note that if \mathcal{C} is not corrupted, then the adversary can cause it to not receive output since we consider security with abort.

FUNCTIONALITY 5 (Schnorr Signing Functionality \mathcal{F}).

- **Signing:** Upon receiving $(\text{sign}, \text{sid}, m, Q, Q_1, \dots, Q_n, d_i)$ from $t+1$ different parties P_i , functionality \mathcal{F} verifies that all parties sent the same $(\text{sid}, m, Q, Q_1, \dots, Q_n)$, that $Q_i = d_i \cdot G$ for all inputs received, and that there exists a degree- t polynomial $p(x)$ such that $p(i) \cdot G = Q_i$ for $i = 1, \dots, n$, and $p(0) \cdot G = Q$. If no, then it does nothing. Else, it chooses a random $k \leftarrow \mathbb{Z}_q$, and computes $R = k \cdot G$, $e = H(m \| Q \| R)$ and $s = k - e \cdot d \pmod{q}$. Then, \mathcal{F} sends (sid, e, s) to \mathcal{C} .

We are now ready to state the theorem and prove it secure. We stress that the “perfect security” in the theorem statement only holds in the hybrid model with ideal commitment and zero-knowledge functionalities (and assuming honest relay of messages by \mathcal{C}). In reality, it is computationally secure (e.g., the signatures generated by the parties to authenticate their messages can be broken, and the instantiations of $\mathcal{F}_{\text{mcom}}$ and \mathcal{F}_{zk} may have computational security). We note that the fact that the security of Schnorr signatures is itself a computational assumption is of no consequence here since our protocol securely computes the signature as is. Thus, our MPC protocol can be perfectly secure, even while computing a computationally secure object.

Theorem 6. *Assume that the parties have valid and consistent inputs, and are invoked with the same (S, m, sid) , where $S \subset [n]$ (with $|S| = t + 1$) is the set of participating parties, m is the message to be signed, and sid is a unique session identifier. Then, Protocol 4 securely computes Functionality 5 in the $(\mathcal{F}_{\text{mcom}}, \mathcal{F}_{\text{zk}})$ -hybrid model with perfect security with abort, in the presence of a malicious static adversary \mathcal{A} controlling any subset of the parties and the coordinator \mathcal{C} .*

Proof. We prove the theorem in the $(\mathcal{F}_{\text{mcom}}, \mathcal{F}_{\text{zk}})$ -hybrid model. This means that instead of a party sending $(\sigma_i^1, \text{sid}_i, c_i)$ in the first message where c_i is a commitment to R_i (where coordinator \mathcal{C} sends this to all parties), it sends $(\text{commit}, \text{sid}_i, \text{pid}_i, S, R_i)$ to $\mathcal{F}_{\text{mcom}}$ and the functionality forwards $(\text{receipt}, \text{sid}_i, \text{pid}_i, S)$ to all parties. Likewise, instead of sending the opening and zero-knowledge proof $(\sigma_i^2, R_i, r_i, \pi_i)$, party P_i sends $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, x)$ to $\mathcal{F}_{\text{mcom}}$ and sends $(\text{prove}, \text{sid}, S, R_i, k_i)$ to $\mathcal{F}_{\text{zk}}^{\text{dl}}$. Note that the signatures on the messages are only needed in order to ensure that \mathcal{C} forwards all values without change, and that all parties receive the same set of committed values. This is guaranteed by communicating directly via $\mathcal{F}_{\text{mcom}}$ and \mathcal{F}_{zk} ; recall that these functionalities ensure consistency. The formal protocol description in this hybrid model appears in Protocol 7.

PROTOCOL 7 (Multiparty Schnorr Signing).

Input: Each party in the set S of parties has the Schnorr public key Q , the set of participating parties S , a session identifier sid , the message to be signed m , and its private input which is a Shamir share of the private key d where $d \cdot G = Q$.

The protocol: Before beginning, each party P_i derives an additive share d_i of the private key, using the Lagrange coefficients for the set S .

1. **Message 1:** Each party P_i chooses a random $k_i \leftarrow \mathbb{Z}_q$, computes $R_i = k_i \cdot G$, and sends $(\text{commit}, \text{sid}_i, \text{pid}_i, S, R_i)$ to $\mathcal{F}_{\text{mcom}}$.
2. **Message 2:** After receiving $(\text{receipt}, \text{sid}_j, \text{pid}_j, S)$ for every $j \in S$, each party P_i sends $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, R_i)$ to $\mathcal{F}_{\text{mcom}}$ and sends $(\text{prove}, \text{sid}, S, (G, R_i), k_i)$ to $\mathcal{F}_{\text{zk}}^{\text{dl}}$.
3. **Message 3:** After receiving $(\text{decommit}, \text{sid}_j, \text{pid}_j, S, R_j)$ from $\mathcal{F}_{\text{mcom}}$ and $(\text{zkproof}, \text{pid}_j, \text{sid}, R_j)$ from $\mathcal{F}_{\text{zk}}^{\text{dl}}$ for every $j \in S$, party P_i computes $R = \sum_{j \in S} R_j$, $e = H(R \| Q \| m)$ and $s_i = k_i - d_i \cdot e \pmod{q}$, and sends s_i to \mathcal{C} .
4. **Output:** \mathcal{C} sets $s = \sum_{i \in S} s_i \pmod{q}$ and checks that $\text{verify}_Q(m, (s, e)) = 1$. If yes, then it outputs (s, e) ; otherwise it aborts and reports a corruption.

Let \mathcal{F} denote Functionality 5. Let \mathcal{A} be an adversary corrupting a (strict) subset of parties $I \subset S$ of size at most t (if $t + 1$ are corrupted, then the protocol is vacuously secure), and let J denote the set of honest parties (note that $I \cup J = S$). 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 $\{(\text{sign}, \text{sid}, m, Q, Q_1, \dots, Q_n, d_i)\}_{i \in I}$, as follows:

1. \mathcal{S} externally sends $(\text{sign}, \text{sid}, m, Q, Q_1, \dots, Q_n)$ to \mathcal{F} and receives back (sid, 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. \mathcal{S} sends $(\text{receipt}, \text{sid}_j, \text{pid}_j, S)$ to \mathcal{A} for every $j \in J$ (for each honest party), as if coming from $\mathcal{F}_{\text{mcom}}$.
3. \mathcal{S} receives $(\text{commit}, \text{sid}_i, \text{pid}_i, S, R_i)$ as sent by each corrupted party to $\mathcal{F}_{\text{mcom}}$. (Note that if any illegal commitments are sent, with incorrect values or format, then this is just ignored by \mathcal{S}).
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}). 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} .
5. \mathcal{S} simulates $\mathcal{F}_{\text{mcom}}$ sending $(\text{decommit}, \text{sid}_j, \text{pid}_j, S, R_j)$ to \mathcal{A} for every $j \in S$ (using the R_j values computed in the previous step). In addition, \mathcal{S} simulates \mathcal{F}_{zk} sending $(\text{zkproof}, j, \text{sid}, (G, R_j))$ to \mathcal{A} for every $j \in S$.
6. \mathcal{S} receives $(\text{decommit}, \text{sid}_i, \text{pid}_i, S, R_i)$ and $(\text{prove}, \text{sid}, S, (G, R_i), k_i)$ messages sent by \mathcal{A} to $\mathcal{F}_{\text{mcom}}$ and \mathcal{F}_{zk} for every corrupted party. If any message is not sent or if $R_i \neq k_i \cdot G$, then \mathcal{S} waits.

7. \mathcal{S} computes $s_i = k_i - d_i \cdot e \pmod{q}$ for every $i \in I$ (\mathcal{S} can do this since it knows the k_i values for each corrupted party from the ZK proof and the d_i values from the key generation). Then, \mathcal{S} computes

$$s_1 = s - \sum_{i \in I} s_i - \sum_{j \in J \setminus \{1\}} s_j \pmod{q}.$$

8. If \mathcal{C} is corrupted, then \mathcal{S} simulates each P_j sending s_j to \mathcal{C} , for every $j \in J$. If \mathcal{C} is honest, then \mathcal{S} receives the s_i values sent by \mathcal{A} to \mathcal{C} for the corrupted parties. \mathcal{S} verifies that $\sum_{i \in I} s_i = \sum_{i \in I} (k_i - d_i \cdot e) \pmod{q}$ where the s_i values are received from \mathcal{A} and the (k_i, d_i) values are as above (k_i from the ZK proofs and d_i from the input). If yes, then it instructs \mathcal{F} to provide output to \mathcal{C} in the ideal model; if no, then it instructs \mathcal{F} to send **abort** to \mathcal{C} in the ideal model.

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 $s_j = k_j - e \cdot d_j \pmod{q}$
- *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 $\tilde{k}_j = \tilde{s}_j + e \cdot d_j \pmod{q}$ and so $\tilde{s}_j = \tilde{k}_j - e \cdot d_j \pmod{q}$, exactly like in a real execution. Furthermore, choosing \tilde{k}_j at random and computing $\tilde{s}_j = \tilde{k}_j - e \cdot d_j \pmod{q}$ yields the exact same distribution as choosing \tilde{s}_j at random and computing $\tilde{k}_j = \tilde{s}_j + e \cdot d_j \pmod{q}$.

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 ZK-proofs, and $\{k_j\}_{j \in J \setminus \{1\}}$ are the implicit values defined above. This therefore defines $k_1 = k - \sum_{i \in I} k_i - \sum_{j \in J \setminus \{1\}} k_j \pmod{q}$. Similarly, we have

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

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 $d = \sum_{\ell \in I \cup J} d_\ell$ we have that

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

and so $s_1 = k_1 - e \cdot d_1 \pmod{q}$, as required. (We stress that \mathcal{S} does not know these values, and in particular it does not know the d_j values of the honest parties including d_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, d_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 [29], perfect security without rewinding implies UC security. As such, assuming that the commitments and zero-knowledge implementations are UC secure, we have that the protocol is UC-secure and so secure under concurrent composition.

On the necessity of the ZK proofs. It may appear that the zero-knowledge proofs are not required since \mathcal{S} can extract the R_i values of the corrupted parties from the commitments, and can set $R_1 = R - \sum_{i \in I} R_i - \sum_{j \in J \setminus \{1\}} R_j$ without knowing the k_i values. Indeed, this is correct, and there is no need for zero-knowledge proofs in order to achieve simulation of the R_i values so that they sum up to R . However, when the simulator needs to generate the signature part $s_1 = s - \sum_{i \in I} s_i - \sum_{j \in J \setminus \{1\}} s_j \pmod{q}$, it does use the k_i (discrete log of R_i) and d_i (share of private key) values. This is so that it can compute $\sum_{i \in I} s_i$ itself. Note that \mathcal{S} cannot rely on the s_i values sent by the corrupted parties for this computation since they may be incorrect. In such a case where \mathcal{A} does send incorrect s_i values, \mathcal{S} is unable to generate the correct s_1 value with the correct distribution. This would therefore result in the real and ideal distributions being distinguishable.

5 Threshold Signing Without Trusted Initialization

5.1 Achieving Consensus on the Parties, Message and SID

We begin by showing how parties can agree on the list of participating parties, a unique session identifier, and the message to be signed. The idea of the protocol is based on the fact that an echo-broadcast where every party sends a message to all others, and then they echo them and accept if they only see the same message from all, is a secure broadcast protocol with abort [26]. As above, the description here assumes a coordinator \mathcal{C} and parties P_1, \dots, P_n of which $t + 1$ need to participate. In practice, the message to be signed is typically sent by some coordinator to all parties, and thus this is the way that we start. We consider a setting where the coordinator chooses any subset of $t + 1$ parties amongst those who responded to participate. Our method assumes a PKI for signing, and so each party P_i has a signing/verification key-pair (sk_i, pk_i) , where the public keys (pk_1, \dots, pk_n) are known to all other parties. See Protocol 8.

Security. Observe that the above method is essentially an echo-broadcast: after the second transmission from \mathcal{C} , all parties have a message that consists of the concatenation of the identities of the parties and their session identifiers. In the next message, they all echo this message by signing on it, and all accept if they receive valid signatures on the string. As we have mentioned, it has been shown in [26] (Protocol 1 in Section 3) that echo-broadcast of this type constitutes a secure *broadcast with abort*, guaranteeing that there exists a string such that every honest parties either outputs that string or aborts. Here, one can view \mathcal{C} as the general or leader, and the message being broadcast as that sent in \mathcal{C} 's second transmission. The formalization in [26] is also one where all parties directly communicate with each other. Here this is replaced with signatures which is the same up to the possibility of replay. However, replay is prevented by each party verifying that their own sid_i is present. As a result, we can conclude that all honest parties agree on the message, the set of the participants and on the sid (since all of these are included in the hash to generate the sid). In addition, since each honest party verifies that its fresh sid_i is part of the set of identifiers, it is also guaranteed that the resulting sid is unique (except with negligible probability).

PROTOCOL 8 (Input Consensus Protocol).

Input: Each party P_i has its private signing key sk_i , and the set of public keys (pk_1, \dots, pk_n) for all parties.

The protocol:

1. **Transmission 1 – \mathcal{C} to all:** \mathcal{C} sends a request to sign on a message m to all parties.
2. **Message 1 – all to \mathcal{C} :** Upon receiving a request to participate, party P_i chooses a random $\text{sid}_i \leftarrow \{0, 1\}^\kappa$ and sends (i, sid_i) to \mathcal{C} . Party P_i stores (m, sid_i) for use below.
3. **Transmission 2 – \mathcal{C} to all:** After receiving a quorum of $t + 1$ responses from parties, \mathcal{C} sends the concatenated list of identities and sid 's to all chosen parties. The list is sorted in a canonical ordering (e.g., lexicographically increasing based on ID if such an identity is known, or based on the parties' public keys). Let S denote the set of $t + 1$ participating parties.
4. **Message 2 – all to \mathcal{C} :** After receiving the list from \mathcal{C} , each P_i verifies that its identity is in the list, that its sid_i chosen initially is present in the list, and that the list is a valid authorized quorum (i.e., there are exactly $t + 1$ parties overall). If yes:
 - (a) P_i sets the session identifier sid for this execution to be a collision-resistant hash applied to the message m , the list of identities S , and the identifiers $\{\text{sid}_i\}_{i \in S}$ that were signed.
 - (b) P_i prepares a signature σ_i on the session identifier sid using its signing key sk_i . Party P_i stores the list of identities and identifiers for use below.
 - (c) P_i sends σ_i to \mathcal{C} .
5. **Transmission 3 – \mathcal{C} to all:** \mathcal{C} sends each party in S the signatures $\{\sigma_j\}_{j \in S}$.
6. **Output:** After receiving the series of signatures $\{\sigma_j\}_{j \in S}$ from \mathcal{C} , party P_i verifies that each of the participating parties (in its list received previously) signed on the same sid (using their public keys available via the PKI). If no, it aborts. If yes, it output (m, sid, S) .

5.2 The Full Three-Round Protocol

In this section, we present the protocol that is derived from running Protocols 4 and 8 in parallel. The only deviation from above here is that in Protocol 4, the message and set of participating parties is known at the onset, and the sid is known in the second round. In contrast, in this parallel execution, the message m is known at the onset, but the set of participating parties and the sid are only known in the second round (i.e., the difference is with respect to the set of participating parties which is only known in the second round). Furthermore, consensus on all of these values is only achieved at the beginning of the third round (i.e., before computing message 3). In order to see why this suffices, note that the consensus on m and sid do not matter before the third message is sent. This is due to the fact that only the third message is dependent on the private key shares. Thus, everything else is easily simulatable up to that point even when consensus will not be reached due to cheating (note also that the simulator easily detects this cheating). Regarding the set of participating parties S , we use this information in Protocol 4 in order to ensure that the commitment is actually a commitment broadcast to the set S , with the guarantee that all honest parties in S who do not abort have the same commitment value. In Protocol 9, we achieve this same effect by having all parties sign on all of the commitment values $\{c_i\}_{i \in S}$ that they receive. This therefore trivially achieves the same effect. See Protocol 9 for the full description of the 3-round protocol including consensus on the input message, participating parties and sid .

PROTOCOL 9 (Multiparty Schnorr Signing Without Initialization).**Input:**

1. Each of P_1, \dots, P_n has the Schnorr public key Q , a PKI of signing keys $\{pk_1, \dots, pk_n\}$, its signing key sk_i , and its private input which is a Shamir share of the private key d where $d \cdot G = Q$.
2. The coordinator \mathcal{C} has a message m to be signed.

The protocol:

1. **Transmission 1 – \mathcal{C} to all:** \mathcal{C} sends a request to sign on a message m to all parties.
2. **Message 1 – all to \mathcal{C} :** Each party P_i works as follows:
 - (a) P_i chooses a random $\text{sid}_i \leftarrow \{0, 1\}^\kappa$.
 - (b) P_i chooses a random $k_i \leftarrow \mathbb{Z}_q$ and sets $R_i = k_i \cdot G$.
 - (c) P_i chooses a random $r_i \leftarrow \{0, 1\}^\kappa$ and sets $c_i = H(\text{sid}_i \parallel \text{pid}_i \parallel R_i \parallel r_i)$.
 - (d) P_i sends $(\sigma_i^1, \text{sid}_i, c_i)$ to the coordinator \mathcal{C} , where $\sigma_i^1 = \text{sign}_{sk_i}(1, \text{sid}_i, c_i)$.
3. **Transmission 2 – \mathcal{C} to parties in S :** After receiving a quorum of $t + 1$ responses $(\sigma_i^1, \text{sid}_i, c_i)$, \mathcal{C} sets S to be the set of responding parties (assume a canonical ordering of parties). \mathcal{C} sends $\{(\sigma_i^1, \text{sid}_i, c_i)\}_{i \in S}$ to P_i for all $i \in S$.
4. **Message 2 – all to \mathcal{C} :** Each party P_i works as follows:
 - (a) P_i verifies that it received $(\sigma_j^1, \text{sid}_j, c_j)$ for $t + 1$ parties, that it is included in the list of participants, that the sid_i that it chose is in the list, that c_i as it sent in the first message appears in the set, and that all signatures are valid. If not, it aborts. If yes, it sets sid to be a collision-resistant hash of m , S and all $\{\text{sid}_j\}_{j \in S}$.
 - (b) P_i computes $\pi_i \leftarrow \text{ZKDL}_P(\text{sid}, \text{pid}_i, R_i; k_i)$ (where ZKDL denotes a Fiat-Shamir proof of knowledge of the discrete log and i is the known identity or public-key of P_i).
 - (c) P_i sends $(\sigma_i^2, \text{sid}, R_i, r_i, \pi_i)$ to the coordinator \mathcal{C} , where $\sigma_i^2 = \text{sign}_{sk_i}(\text{sid}, 2, \{c_i\}_{i \in S}, R_i, r_i, \pi_i)$.
5. **Transmission 3 – \mathcal{C} to parties in S :** \mathcal{C} receives all $(\sigma_i^2, \text{sid}, R_i, r_i, \pi_i)$ messages, and sends $\{(\sigma_i^2, \text{sid}, R_i, r_i, \pi_i)\}_{i \in S}$ to all parties P_i with $i \in S$.
6. **Message 3 – all to \mathcal{C} :** Each party P_i works as follows:
 - (a) After receiving all $\{(\sigma_j^2, \text{sid}, R_j, r_j, \pi_j)\}_{j \in S}$, party P_i verifies that all signatures are valid and are computed on the same sid that it computed.
 - (b) For every $j \in S$ ($j \neq i$):
 - i. P_i verifies that $H(\text{sid}_j \parallel \text{pid}_j \parallel R_j \parallel r_j) = c_j$ and that all values are valid (i.e., has the correct sid_j and pid_j and the overall structure).
 - ii. P_i verifies $\text{ZKDL}_V(\text{sid}, \text{pid}_j, R_j; \pi_j) = 1$.
 - iii. If the commitment is not valid, or R_j is not a valid point in the curve subgroup, or R_j is equal to the identity point, or if $\text{ZKDL}_V(\text{sid}, \text{pid}_j, R_j; \pi_j) = 0$ (i.e., the ZK verification fails) then P_i aborts. Else, it proceeds.
 - (c) P_i updates its Shamir share to be an additive share d_i of the private key, using the Lagrange coefficients for the set S .
 - (d) P_i sets $R = \sum_{j \in S} R_j$, $e = H(m \parallel Q \parallel R)$, and $s_i = k_i - d_i \cdot e \pmod{q}$.
 - (e) P_i sends s_i to \mathcal{C} .
7. **Output:** \mathcal{C} sets $s = \sum_{i \in S} s_i \pmod{q}$ and checks that $\text{verify}_Q(m, (s, e)) = 1$. If yes, then it outputs (s, e) ; otherwise it aborts.

The following theorem is proven in the same way as Theorem 6, with the addition of the guarantees provided by the echo-broadcast:

Theorem 10. *Assume that the parties have valid and consistent inputs. Then, Protocol 9 securely computes Functionality 5 in the $(\mathcal{F}_{\text{mcom}}, \mathcal{F}_{\text{zk}})$ -hybrid model with perfect security with abort, in the presence of a malicious static adversary \mathcal{A} controlling any subset of the parties and the coordinator \mathcal{C} .*

6 Identifiable Abort

In the standard model of security with abort for MPC, the honest parties may not necessarily know what caused the abort and who is “to blame”. Protocols with identifiable abort have the advantage of adding deterrent to a party to misbehave and cause the protocol to fail, since they will be caught.

In Protocol 9, since all messages are signed, it is possible to detect who cheated in the first two rounds. The only exception to this is regarding the consensus in that a party can claim that the coordinator sent it a different set of parties. This can be solved by also having the coordinator sign on all messages sent (using the sid that can already be computed after the first-round messages have been received).

In contrast, the way it is described, it is not possible to know if some party P_i sent an incorrect s_i at the end of the protocol. Rather, the coordinator \mathcal{C} will fail to obtain a valid signature, but that is all. However, in many cases (e.g., when distributed key generation is used), all parties are given (Q_1, \dots, Q_n) where $Q_i = d_i \cdot G$. This additional information can be used to achieve identifiable abort in the following way. Upon receiving s_i from party P_i , it is possible to check that $s_i \cdot G = R_i - e \cdot \lambda_i^S \cdot Q_i$, where λ_i^S is the Lagrange coefficient for converting P_i 's Shamir share to an additive share amongst the set S . Since s_i is supposed to equal $k_i - e \cdot \lambda_i^S \cdot d_i$, this validates with certainty whether or not P_i sent a correct value. In order to determine whether \mathcal{C} or P_i is cheating, we can simply have P_i sign on the last message s_i sent to \mathcal{C} (together with the sid). Then, \mathcal{C} can also prove to other parties that P_i cheated.

7 UC NIZK for Discrete Log in the Random Oracle Model

In order to instantiate the \mathcal{F}_{zk} functionality with UC security, it is necessary to extract the witness without rewinding. One may just “assume” that the Fiat-Shamir transform provides a zero-knowledge proof of knowledge that fulfills the UC requirements. However, there is no justification for this. Fortunately, an efficient transformation from Sigma protocols to zero-knowledge with non-rewinding extraction was given in the random-oracle model in [20]. The transformation requires some parallel repetition of the proof, but the additional expense is relatively mild (the example parameters provided in [20] are such that the proof needs to be repeated approximately 10 times to get a good balance of completeness and extraction). For the discrete log proof we need here, the computational cost is approximately 2.1ms for the prover and 0.7ms for the verifier (averaged over 1000 repetitions), when run on a Macbook Pro with a 2.3 GHz 8-Core Intel Core i9 CPU. This therefore does not add significant overhead.

References

- [1] H.K. Alper and J. Burdges. Two-Round Trip Schnorr Multi-Signatures via Delin-earized Witnesses. In *CRYPTO 2021*, Springer (LNCS 12825), pages 157–188

-
- [2] M. Bellare and A. Palacio. The Knowledge-of-Exponent Assumptions and 3-Round Zero-Knowledge Protocols. In *CRYPTO 2004*, Springer (LNCS 3152), pages 273–289, 2004.
- [3] M. Bellare and G. Neven. Multi-Signatures in the Plain Public-Key Model and a General Forking Lemma. In *ACM CCS 2006*, pages 390–399, 2006.
- [4] F. Benhamouda, T. Lepoint, J. Loss, M. Orru and M. Raykova. On the (in)security of ROS. In *EUROCRYPT 2021*, Springer (LNCS 12696), pages 33–53, 2021.
- [5] D.J. Bernstein, N. Duif, T. Lange, P. Schwabe, and Y. Bo-Yin. High-speed high-security signatures. In *Journal of Cryptographic Engineering*, 2(2):77–89, 2012.
- [6] O. Blazy, C. Chevalier, D. Pointcheval and D. Vergnaud. Analysis and Improvement of Lindell’s UC-Secure Commitment Schemes. In *ACNS 2013*, Springer (LNCS 7954), pages 534–551, 2013.
- [7] C. Boyd. Digital Multisignatures. In *Cryptography and Coding*, pages 241–246, 1986.
- [8] R. Canetti. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology*, 13(1):143–202, 2000.
- [9] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd FOCS*, pages 136–145, 2001. Full version available at <http://eprint.iacr.org/2000/067>.
- [10] R. Canetti, A. Cohen and Y. Lindell. A Simpler Variant of Universally Composable Security for Standard Multiparty Computation. In *CRYPTO 2015*, Springer (LNCS 9216), pages 3–22, 2015.
- [11] R. Canetti and A. Herzberg. Maintaining Security In The Presences Of Transient Faults. In *CRYPTO’94*, Springer-Verlag (LNCS 839), pages 425–438, 1994.
- [12] E.C. Crites, C. Komlo and M. Maller. Fully Adaptive Schnorr Threshold Signatures. In *CRYPTO 2023*, Springer (LNCS 14081), pages 678–709, 2023.
- [13] R.A. Croft and S.P. Harris. Public-Key Cryptography and Reusable Shared Secrets. In *Cryptography and Coding*, pages 189–201, 1989.
- [14] E. Crites, C. Komlo and Mary Maller. How to Prove Schnorr Assuming Schnorr: Security of Multi- and Threshold Signatures. *Cryptology ePrint Archive*, 2021/1375, 2021.
- [15] I. Damgård, M. Jurik and J.B. Nielsen. A Generalization of Paillier’s Public-Key System with Applications to Electronic Voting. In *International Journal of Information Security*, 9(6):371–385, 2010.
- [16] Y. Desmedt. Society and Group Oriented Cryptography: A New Concept. In *CRYPTO’87*, Springer (LNCS 293), pages 120–127, 1988.
- [17] Y. Desmedt and Y. Frankel. Threshold Cryptosystems. In *CRYPTO’89*, Springer (LNCS 435), pages 307–315, 1990.
- [18] M. Drijvers, K. Edalatnejad, B. Ford, E. Kiltz, J. Loss, G. Neven and I. Stepanovs. On the Security of Two-Round Multi-Signatures. In *IEEE Symposium on Security and Privacy*, pages 1084–1101, 2019.

- [19] P. Feldman. A Practical Scheme for Non-interactive Verifiable Secret Sharing. In *FOCS 1987*, pages 427–437, 1987.
- [20] M. Fischlin. Communication-Efficient Non-interactive Proofs of Knowledge with Online Extractors. In *CRYPTO 2005*, Springer (LNCS 3621), pages 152–168, 2005.
- [21] E. Fujisaki. Improving Practical UC-Secure Commitments Based on the DDH Assumption. In *SCN 2016*, Springer (LNCS 9841), pages 257–272, 2016.
- [22] R. Gennaro, S. Goldfeder and A. Narayanan: Threshold-Optimal DSA/ECDSA Signatures and an Application to Bitcoin Wallet Security. In *ACNS 2016*, Springer (LNCS 9696), pages 156–174, 2016.
- [23] R. Gennaro, S. Jarecki, H. Krawczyk and T. Rabin. Robust Threshold DSS Signatures. In *EUROCRYPT’96*, Springer (LNCS 1070), pages 354–371, 1996.
- [24] R. Gennaro, D. Leigh, R. Sundaram, and W.S. Yezauris. Batching Schnorr Identification Scheme with Applications to Privacy-Preserving Authorization and Low-Bandwidth Communication Devices. In *ASIACRYPT 2004*, Springer (LNCS 3329), pages 276–292, 2004.
- [25] O. Goldreich. *Foundations of Cryptography: Volume 2 – Basic Applications*. Cambridge University Press, 2004.
- [26] S. Goldwasser and Y. Lindell. Secure Computation Without Agreement. In the *Journal of Cryptology*, 18(3):247–287, 2005. (Extended abstract appeared at *DISC 2002*.)
- [27] C. Hazay and Y. Lindell. *Efficient Secure Two-Party Protocols: Techniques and Constructions*. Springer, November 2010.
- [28] C. Komlo and I. Goldberg. FROST: Flexible Round-Optimized Schnorr Threshold Signatures. In *SAC 2020*, Springer (LNCS 12804), pages 34–65, 2020.
- [29] E. Kushilevitz, Y. Lindell and T. Rabin. Information-Theoretically Secure Protocols and Security Under Composition. In the *SIAM Journal on Computing* (SICOMP), 39(5):2090—2112, 2010. (Preliminary version in the *38th STOC*, pages 109–118, 2006.)
- [30] Y. Lindell: Highly-Efficient Universally-Composable Commitments Based on the DDH Assumption. In *EUROCRYPT 2011*, Springer (LNCS 6632), pages 446–466, 2011.
- [31] Y. Lindell. An Efficient Transform from Sigma Protocols to NIZK with a CRS and Non-Programmable Random Oracle. In *TCC 2015*, Springer (LNCS 9014), pages 93–109, 2015.
- [32] Y. Lindell. Fast Secure Two-Party ECDSA Signing. In *Journal of Cryptology*, 34:44, 2021. (An extended abstract appeared at *CRYPTO 2017*.)
- [33] P.D. MacKenzie and M.K. Reiter. Two-party generation of DSA signatures. *International Journal of Information Security*, 2(3-4):218–239, 2004. (An extended abstract appeared at *CRYPTO 2001*.)
- [34] N. Makriyannis. On the Classic Protocol for MPC Schnorr Signatures. *Cryptology ePrint Archive*, Paper 2022/1332, 2022.
- [35] S. Micali, K. Ohta and L. Reyzin. Accountable-Subgroup Multisignatures. In *ACM CCS 2001*, pages 245–254, 2001.

- [36] A. Nicolosi, M.N. Krohn, Y. Dodis and D. Mazieres: Proactive Two-Party Signatures for User Authentication. *NDSS 2003*.
- [37] J. Nick, T. Ruffing and Y. Seurin. MuSig2: Simple Two-Round Schnorr Multi-signatures. In *CRYPTO 2021*, Springer (LNCS 12825), pages 189–221, 2021.
- [38] C.P. Schnorr. Efficient Identification and Signatures for Smart Cards. In *CRYPTO 1989*, Springer (LNCS 435), pages 239–252, 1989.
- [39] V. Shoup and R. Gennaro. Securing Threshold Cryptosystems against Chosen Ciphertext Attack. In *EUROCRYPT 1998*, Springer (LNCS 1403), pages 1–16, 1998.
- [40] V. Shoup. Practical Threshold Signatures. In *EUROCRYPT 2000*, Springer (LNCS 1807), pages 207–220, 2000.
- [41] V. Shoup. Private communication, 2019.
- [42] D.R. Stinson and R. Strobl. Provably Secure Distributed Schnorr Signatures and a (t, n) Threshold Scheme for Implicit Certificates. In *ACISP 2001*, Springer (LNCS 2119), pages 417–434, 2001.
- [43] BIP 32: Hierarchical Deterministic Wallets. <https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki>.

A Distributed Key Generation

Observe that the first stage of Schnorr signing is the generation of a random element $R \in \mathbb{G}$ where its discrete log is additively shared amongst the participants. As such, the exact same protocol can be used for generating a random public-key Q whose private key (the discrete log of Q) is additively shared amongst the participants. Since we wish to use threshold sharing, the same strategy works for choosing a random Feldman VSS sharing [19] where each coefficient of the polynomial is generated in this way. See Protocol 11 for details. In this protocol, since each party needs to receive a private share on the generated polynomials, the parties also need to have a PKI of encryption keys.

Security: The proof of security for Protocol 11 follows the same lines as for signing. In particular, the A_j^0 values are simulated in exactly the same way as the R_i values so that the result is the public key Q given to the simulator \mathcal{S} by the ideal functionality. In more detail, for all but one honest party, the simulator \mathcal{S} works honestly. Regarding the last honest party, denote it P_ℓ , simulator \mathcal{S} sets $A_\ell^0 = Q - \sum_{i \in \mathcal{S}} A_i^0$. Then, \mathcal{S} chooses random points $a_\ell(\text{pid}_i)$ for each corrupted party P_i , and additional random points for honest parties up to t (if less than t parties are corrupted). Finally, \mathcal{S} interpolates “in the exponent” to find the polynomial $A_\ell(x)$ with coefficients $(A_\ell^0, A_\ell^1, \dots, A_\ell^t)$ so that $a_\ell(\text{pid}_i) \cdot G = A(\text{pid}_i)$ for every corrupted P_i . The rest of the protocol – echo-broadcasting the commitments and the ciphertexts encrypting the shares – all ensure consensus on the shared polynomial. The important point to note is that if at most t parties are corrupted, then at least one participating party is honest. In this case, the honest party only signs on one set of polynomials and commitments when generating its σ_i^2 . Furthermore, by the echo-broadcast, all honest parties see the same set of commitments from round 1 and thus the same set of polynomials \bar{A}_i . This guarantees that there is only one possible polynomial \bar{Q} , and that all honest parties agree on that polynomial. Finally, each party locally verifies that its local share that it decrypts matches that polynomial.

PROTOCOL 11 (Multiparty Schnorr Distributed Key Generation).

Input: Each of P_1, \dots, P_n has a PKI of signing keys $\{pk_1, \dots, pk_n\}$ and encryption keys $\{ek_1, \dots, ek_n\}$, and its own signing key sk_i and decryption key dk_i .

The protocol:

1. **Transmission 1 – C to all:** C sends a request to generate a key to all parties.
2. **Message 1 – all to C:** Each party P_i works as follows:
 - (a) P_i chooses a random $\text{sid}_i \leftarrow \{0, 1\}^\kappa$.
 - (b) For $k \in \{0, \dots, t\}$, P_i chooses a random $a_i^k \leftarrow \mathbb{Z}_q$ and sets $A_i^k = a_i^k \cdot G$. Let $\bar{A}_i = (A_i^0, \dots, A_i^t)$, $\bar{a}_i = (a_i^0, \dots, a_i^t)$, $a_i(x) = \sum_{k=0}^t a_i^k \cdot x^k$, and $A_i(x) = a_i(x) \cdot G$.
 - (c) For every $j \in \{1, \dots, n\}$, party P_i sets P_j 's share in $A_i(x)$ to be $d_{i \rightarrow j} = a_i(\text{pid}_j)$ and encrypts it under P_j 's public key ek_j ; denote the ciphertext by $c_{i \rightarrow j}$. Denote the set of all these ciphertexts by \bar{c}_i .
 - (d) P_i chooses a random $r_i \leftarrow \{0, 1\}^\kappa$ and sets $\text{com}_i = H(\text{sid}_i \parallel \text{pid}_i \parallel \bar{A}_i \parallel \bar{c}_i \parallel r_i)$.
 - (e) P_i computes $\sigma_i^1 = \text{sign}_{sk_i}(1, \text{sid}_i, \text{com}_i)$
 - (f) P_i sends $(\sigma_i^1, \text{sid}_i, \text{com}_i)$ to the coordinator C .
3. **Transmission 2 – C to parties in S :** After receiving a quorum of $t + 1$ responses $(\sigma_i^1, \text{sid}_i, \text{com}_i)$, C sets S to be the set of responding parties (assume a canonical ordering of parties). C sends $\{(\sigma_i^1, \text{sid}_i, \text{com}_i)\}_{i \in S}$ to P_i for all $i \in S$.
4. **Message 2 – all to C:** Each party P_i works as follows:
 - (a) P_i verifies that it received $(\sigma_j^1, \text{sid}_j, \text{com}_j)$ for $t + 1$ parties, that it is included in the list of participants, that the sid_i that it chose is in the list, that com_i as it sent in the first message appears in the set, and that all signatures are valid. If not, it aborts. If yes, it sets sid to be a collision-resistant hash of S and all $\{\text{sid}_j\}_{j \in S}$.
 - (b) P_i computes $\pi_i \leftarrow \text{ZKDL}_P^{t+1}(\text{sid}, \text{pid}_i, \bar{A}_i; \bar{a}_i)$ (where ZKDL^{t+1} denotes a batch Fiat-Shamir proof of knowledge of the discrete log of $t + 1$ values, and i is the known identity or public-key of P_i).
 - (c) P_i computes $\sigma_i^2 = \text{sign}_{sk_i}(\text{sid}, 2, \{\text{com}_i\}_{i \in S}, \bar{A}_i, \bar{c}_i, r_i, \pi_i)$
 - (d) P_i sends $(\sigma_i^2, \text{sid}, \bar{A}_i, \bar{c}_i, r_i, \pi_i)$ to the coordinator C .
5. **Transmission 3 – C to parties in S :** C receives all $(\sigma_i^2, \text{sid}, \bar{A}_i, \bar{c}_i, r_i, \pi_i)$ messages, and sends $\{(\sigma_i^2, \text{sid}, \bar{A}_i, \bar{c}_i, r_i, \pi_i)\}_{i \in S}$ and $\{(\text{sid}_i, \text{com}_i)\}_{i \in S}$ to all parties P_i for $i = 1, \dots, n$.
6. **Output:** Each party P_i for $i = 1, \dots, n$ works as follows:
 - (a) P_i verifies that it received $\{(\sigma_j^2, \text{sid}, \bar{A}_j, \bar{c}_j, r_j, \pi_j)\}_{j \in S}$ and $\{(\text{sid}_j, \text{com}_j)\}_{j \in S}$ for a set S of size $t + 1$, and that all signatures are valid and are computed on the same sid and set $\{\text{com}_j\}_{j \in S}$, where sid equals the hash of S and $\{\text{sid}_j\}_{j \in S}$.
 - (b) For every $j \in S$:
 - i. P_i verifies that $\text{com}_j = H(\text{sid}_j \parallel \text{pid}_j \parallel \bar{A}_j \parallel \bar{c}_j \parallel r_j)$ and that all values are valid (i.e., has the correct sid_j and pid_j and the overall structure).
 - ii. P_i verifies $\text{ZKDL}_V^{t+1}(\text{sid}, \text{pid}_j, \bar{A}_j; \pi_j) = 1$.
 - iii. If the commitment is not valid or any A_j^k is not a valid point in the curve subgroup of is equal to the identity point, or if the zero-knowledge verification fails, then P_i aborts. Else, it proceeds.
 - (c) P_i decrypts all $\{c_{j \rightarrow i}\}_{j \in S}$ and sets $d_i = \sum_{j \in S} d_{j \rightarrow i} \pmod{q}$.
 - (d) P_i sets the VSS sharing polynomial to be $Q(x) = \sum_{j \in S} A_j(x)$. That is, the k th coefficient Q_k of $Q(x)$ is set to $\sum_{j \in S} A_j^k$. Denote $\bar{Q} = (Q_0, \dots, Q_t)$.
 - (e) P_i verifies that $(d_{j \rightarrow i}) \cdot G = A_j(\text{pid}_i)$ for all $j \in S$ and that $d_i \cdot G = Q(\text{pid}_i)$. If not, it aborts.
 - (f) P_i sets the output public key to be $Q = Q_0$.
 - (g) P_i outputs the public-key Q , the polynomial $Q(x)$, and its share d_i .