Threshold ECDSA Pseudocode for Coinbase

1 Overview

This document contains pseudocode for ECDSA threshold signatures based on the recent paper of Gennaro and Goldfeder [4]. The exact details of the protocol presented emerged from conversations with Coinbase as to what best suits their needs (e.g. not necessary to attribute misbehavior to a single server). Throughout the document, we will reference specific design decisions and briefly mention the other options available, should it be of interest at some point for us to elaborate further.

Our goal is for this to be a living document that we update with requested clarifications. During the second phase of this engagement, we will also add pseudocode for a dealerless protocol as well as a protocol to dynamically re-share the key.

1.1 Sequential vs. Concurrent Security

The protocol described here is proven secure under sequential composition (i.e. with signatures issued one at a time sequentially) against a static adversary. We conjecture that the protocol is secure sequentially as well as all known attacks of this form are non-practical, although this is not provable in the current protocol. We could provide a variant of our protocol with Universally Composable (UC) concurrent provable security, although this would take an efficiency hit, and based discussions with Coinbase, we are providing the simple protocol which has been proven secure in the standalone sequential model.

Should this decision change once implemented, please let us know and we can provide a document for the UC protocol.

1.2 On our use of the Fiat-Shamir Heuristic

In our academic paper, we assume that the zero knowledge proofs in the protocol are interactive, but in practice these can be made non-interactive using the Fiat-Shamir heuristic in the random oracle model.

However, there is one irregularity in our use of Fiat-Shamir in that we use it for proofs of knowledge, which means that in the security proof, the simulator must extract the adversary's private values. For interactive proofs this is done via "rewinding". Our proofs are three move sigma protocols, which involve the prover committing, the verifier issuing a challenge, and the prover responding. The sigma protocols have a "special soundness" property which means that if the prover commits to its first message and then responds to two different challenges, the verifier can extract the witness. In a "rewinding" argument, the simulator rewinds the adversary back to its first commitment in the sigma protocol and provides a different challenge. The adversary, unaware that it was rewound, responds allowing the prover to extract the witness.

In Fiat-Shamir proofs, the challenge is not supplied by the adversary but by the Random Oracle, and thus all three steps of the protocol are done by the prover: It commits, queries the random oracle, and responds, and then sends the transcript to the verifier. In order to extract by a typical rewinding argument, it is crucial to be able to rewind back to the point in between when the prover commits and gets a challenge. But if the protocol is fully non-interactive the simulator cannot do that since it can only rewind the adversary to the point before it committed. And at this point, the adversary can change its commitment, preventing extraction.

The simplest way to deal with this is by adding an additional round for the FS proofs – in which the adversary sends and commits to its first message. However, in an effort not to add rounds, we

conjecture that there exists a non black-box extractor that allows us to extract the secret values from the adversary.

We are happy to discuss this more, and if you are not comfortable with this assumption, we can add commitment rounds to the proofs. We also mention that related to the point above: the UC version of the protocol does not extract from the FS proofs and thus this issue doesn't arise.

1.3 Accessibility

Throughout this document, we have often appealed to color in the pseudocode to convey important data and simplify the presentation. Upon request, we would be happy to provide an accessible version of this document that does not use color to convey information.

2 Preliminaries

In this section, we will discuss the primitives on which we rely, and provide pseudocode when appropriate.

2.1 ECDSA signatures

The Digital Signature Algorithm (DSA) was proposed by Kravitz in 1991, and adopted by NIST in 1994 as the Digital Signature Standard (DSS) [1, 5]. ECDSA, the elliptic curve variant of DSA, has become quite popular in recent years, especially in cryptocurrencies. Our focus in this document will be on ECDSA.

We will for completeness cover the details of ECDSA here, and indeed the standard ECDSAVerify function will be compatible with the signatures produced by the threshold signing protocol. We do not include pseudocode as we strongly recommend that you choose an existing centralized ECDSA implementation and ensure that you are compatible with the verifier, and we do not recommend implementing this from scratch as there are numerous open-source high quality implementations.

The Public Parameters consist of a cyclic group $\mathcal{G}$ of prime order q, a generator g for $\mathcal{G}$, a hash function $H : \{0,1\}^* \to Z_q$, and another hash function $H' : \mathcal{G} \to Z_q$.

Key-Gen On input the security parameter λ, outputs a private key x chosen uniformly at random in Z_q, and a public key $y = g^x$ computed in $\mathcal{G}$.

Sig On input an arbitrary message M,

- Compute $m = H(M) \in Z_q$
- choose $k \in_R Z_q$
- compute $R = g^{k^{-1}}$ in $\mathcal{G}$ and $r = H'(R) \in Z_q$
- compute $s = k(m + xr) \bmod q$
- output $\sigma = (r, s)$

Ver On input M, σ and y,

- check that $r, s \in Z_q$
- compute $R' = g^{ms^{-1} \bmod q} y^{rs^{-1} \bmod q}$ in $\mathcal{G}$
- Accept (output 1) iff $H'(R') = r$.

The traditional DSA algorithm is obtained by choosing large primes p, q such that $q|(p-1)$ and setting $\mathcal{G}$ to be the order q subgroup of Z_p^*. In this case the multiplication operation in $\mathcal{G}$ is multiplication modulo p. The function H' is defined as $H'(R) = R \bmod q$.

The ECDSA scheme is obtained by choosing $\mathcal{G}$ as a group of points on an elliptic curve of cardinality q. In this case the multiplication operation in $\mathcal{G}$ is the group operation over the curve. The function H' is defined as $H'(R) = R_x \bmod q$ where R_x is the x-coordinate of the point R.

The specific choice of the curve and hash function H will depend on the implementation, and the threshold signing protocol is agnostic to these choices.

2.2 Cryptographic Hash functions

Aside from the hash function used to hash the ECDSA message, our protocol employs a hash function to be used in the commitment scheme as well as for the Random Oracle for making our zero-knowledge proofs non-interactive using Fiat-Shamir. In the pseudocode, we refer to SHA256, but other cryptographically secure hash functions could be used as well (e.g. SHA3).

Notation. When the pseudocode contains a call to SHA256(a, b, c) we mean that the values a, b, and c should be concatenated with an appropriate delimiter in between them to ensure that the separation between protocol values can never be ambiguous. The delimiter should be a character that can never appear in the hashed values (e.g. \$).

2.3 Fiat Shamir Hash function

In our zero knowledge proofs, we require a hash function to non-interactively compute the Fiat-Shamir challenge. We could use a hash function with an appropriate delimiter as discussed in the previous paragraph, but in order to make encoding simple without need to worry about message lengths, we propose the following simple hash function in which we first hash the internal messages and then combine the intermediate hashes into a single digest. Since the intermediate hashes are fixed length, no delimiters are needed for the final digest, and we indicate this using the simple concatenation operator, $||$.

$h \leftarrow$ FS-HASH$(m_1, m_2, \ldots, m_n)$

1. `For` $i = [1, \ldots, n]$
2. Compute $h_i =$ SHA256(m_i)
3. Compute $h =$ SHA256$(h_1 || h_2 || \cdots || h_n)$
4. `Return` h

2.4 Network assumptions

We assume that each pair of players is connected by a point-to-point authenticated channel. In the pseudocode, sending a message over this channel is denoted by the P2PSend function, and this will be used when players need to send unique messages to other players.

We also make use of a Broadcast function which is called when players need to send the same message to all other players. We do not require a full broadcast channel, but a simple echo broadcast suffices. During an echo broadcast, each players utilizes the point-to-point channel to send each player a hash of its view of the the messages sent from all players. If any party receives an inconsistent hash from some other party, it aborts.

2.5 Commitment scheme

Throughout the protocol, we use a commitment scheme so that players can commit to values in a manner that will bind them to their choice but also hide their choice until they choose to reveal it. We require that the commitment schemes is non-malleable and concurrently secure, roughly meaning that given a set of commitments one cannot compute a commitment to a different but related value. In the Random Oracle model, we can achieve this using the simple "canonical" commitment function.

For the commitment function, we could use a cryptographic hash function and separate the two inputs by an appropriate delimeter. Here, we use HMAC, and in our notation, the first input is the HMAC key and the second input is the message.

$[C, D] \leftarrow \mathsf{Commit}(m)$

1. Choose $r \xleftarrow{\$} \{0,1\}^{256}$ // r is a 256-bit random nonce
2. Compute $C = \mathsf{HMAC}(r, m)$
3. Set $D = (m, r)$
4. `Return` $[C, D]$

$m \leftarrow \mathsf{Open}(C, D = (m, r))$

1. `If` $(D \neq \texttt{nil})$,
2. Compute $C_{\texttt{check}} = \mathsf{HMAC}(r, m)$
3. `If` $(C_{\texttt{check}} = C)$
4. `Return` m
5. `Return` $\perp$

2.6 Paillier Cryptosystem

The threshold ECDSA protocol makes heavy use of an encryption scheme that is additively homomorphic modulo a large integer N, and we instantiate it with Paillier's cryptosystem [6].

Let $E_{pk}(\cdot)$ denote the encryption algorithm for $\mathcal{E}$ using public key pk. Given ciphertexts $c_1 = E_{\mathsf{pk}}(a)$ and $c_2 = E_{\mathsf{pk}}(b)$, there is an efficiently computable function $+_E$ such that

$$c_1 +_E c_2 = E_{\mathsf{pk}}(a + b \bmod N)$$

The existence of a ciphertext addition operation also implies a scalar multiplication operation, which we denote by $\times_E$. Given an integer $a \in N$ and a ciphertext $c = E_{\mathsf{pk}}(m)$, then we have

$$a \times_E c = E_{\mathsf{pk}}(am \bmod N)$$

We include the pseudocode for Paillier's cryptosystem, but if a suitable library is identified, it may be preferable to use it rather than re-implement it.

On the security of Paillier. We note that while the Paillier cryptosystem is widely used in practice, it has not been standardized, relies on non-standard assumptions and is less studied than standardized primitives. We recommend instantiating our protocol with Paillier, but following the approach of [2], we can instantiate our protocol using a sub-protocol that makes use of Oblivious Transfer instead of Paillier, at the cost of bandwidth efficiency.

Our recommendation is to use Paillier as we are comfortable with its security and it yields a significantly better overall protocol with respect to the amount of data communicated between signers, but we can provide the alternative if it is desired.

$(pk, sk) \leftarrow$ PaillierKeyGen(1^κ)
1. Choose a 1024-bit prime P
2. Choose a 1024-bit prime Q
3. Compute $N = P \cdot Q$
4. Compute $\lambda(N) = lcm(P-1, Q-1)$
5. Compute $u = \mathsf{L}((N+1)^{\lambda(N)} \bmod N^2, N)^{-1} \bmod N$
6. Compute $\phi(N) = (P-1) \cdot (Q-1)$
7. Set $pk = N$
8. Set $sk = [N, \lambda(N), \phi(N), u]$
9. `Return` (pk, sk)

$c \leftarrow$ PaillierEncrypt(pk, m)
1. Set $N = pk.N$
2. `If` $m \notin \mathbb{Z}_N$, `Return` $\perp$
3. Choose $r \xleftarrow{\$} \mathbb{Z}_N^*$
4. Compute $c = (N+1)^m r^N \bmod N^2$
5. `Return` c

$(c, r) \leftarrow$ PaillierEncryptAndReturnRandomness(pk, m)
1. Set $N = pk.N$
2. `If` $m \notin \mathbb{Z}_N$, `Return` $\perp$
3. Choose $r \xleftarrow{\$} \mathbb{Z}_N^*$
4. Compute $c = (N+1)^m r^N \bmod N^2$
5. `Return` (c, r)

$m \leftarrow$ PaillierDecrypt(sk, c)
1. Set $N = sk.N$
2. `If` $c \notin \mathbb{Z}_{N^2}$, `Return` $\perp$
3. Compute $m = \mathsf{L}(c^{\lambda(N)} \bmod N^2, N) \cdot u \bmod N$
4. `Return` m

$c_3 \leftarrow$ PaillierAdd(pk, c_1, c_2)
1. Set $N = pk.N$
2. `If` $c_1, c_2 \notin \mathbb{Z}_{N^2}$, `Return` $\perp$
3. `Return` $c_1 \cdot c_2 \bmod N^2$

$c_2 \leftarrow$ PaillierMultiply(pk, a, c_1)
1. Set $N = pk.N$
2. `If` $a \notin \mathbb{Z}_N$, `Return` $\perp$
3. `If` $c_1 \notin \mathbb{Z}_{N^2}$, `Return` $\perp$
4. `Return` $c_1^a \bmod N^2$

$c_3 \leftarrow \mathsf{L}(u, N)$
1. `If` $u \notin \mathbb{Z}_{N^2}$, `Return` $\perp$
2. `If` $u \neq 1 \bmod N$, `Return` $\perp$
3. `Return` $(u-1)/N$

Fig. 1: **The Paillier Cryptosystem**. Paillier's cryptosystem is an additively homomorphic encryption scheme, which supports the addition of two ciphertexts and the multiplication of a ciphertext and a scalar. In a standalone instantiation of Paillier's cryptosystem, we would not need the PaillierEncryptAndReturnRandomness function, but it is often useful (and required in our threshold signing protocol) to return the randomness to facilitate proving statements about the ciphertext. In all cases, it is important that the randomness r is kept private and not passed around as part of the ciphertext.

2.7 Shamir's Secret Sharing (SSS) Scheme

The underlying secret sharing scheme used in the threshold signing protocol is Shamir Secret Sharing (SSS) [7]. For a prime p, Shamir's secret sharing allows one to share a secret $x \in Z_p$ by distributed points on a random degree t polynomial $p(\cdot)$ with x as the constant term:

$$p(x) = x + a_1 x + a_2 x^2 + \cdots + a_t x^t \bmod q$$

.

Each player $\mathcal{P}_i$ is associated with a unique non-zero index p_i and the player's share is $p(p_i)$, the evaluation of the polynomial at p_i. Given $t+1$ shares, the polynomial can be reconstructed using *Lagrange interpolation*, and the secret is obtained by evaluating the polynomial at 0.

We now provide pseudocode for both sharing and revealing a secret using Shamir's secret sharing. We note that we will not use Reveal anywhere in our protocol and have only included it for informational completeness, but it does not need to be implemented.

$[x_1, \ldots, x_n] \leftarrow$ ShamirShare$(x, t, [p_1, \ldots, p_n])$

1. `For` $i = [1, \ldots, t]$ // Generate random coefficients
2. $\quad a_i \xleftarrow{\$} \mathbb{Z}_p$
3. `For` $i = [1, \ldots, n]$
4. $\quad$ `If` $p_i = 0$, `Abort`
5. $\quad x_i = x + a_1 p_i + a_2 p_i^2 + \cdots + a_t p_i^t \bmod p$
6. `Return` $[x_1, \ldots, x_n]$

$x \leftarrow$ Reveal$([p_1, x_1], \ldots, [p_{t+1}, x_{t+1}])$

1. Set $x = 0$
2. `For` $i = [1, \ldots, t+1]$
3. $\quad \ell = x_i$
4. $\quad$ `For` $j = [1, \ldots, t+1]$
5. $\quad\quad$ `If` $i = j$, `Continue`
6. $\quad\quad \ell = \ell \times \frac{p_j}{p_j - p_i} \bmod p$
7. $\quad x = x + \ell \bmod p$
8. `Return` x

Fig. 2: **Shamir's secret sharing scheme.** In the ShamirShare algorithm, the dealer generates a random polynomial of degree t and evaluates it at p_i, the index associated with each player $\mathcal{P}_i$. The dealer then distributed the appropriate share to each player, which they are to keep secret. The Reveal algorithm takes any $t+1$-sized subset of shares, reconstructs the polynomial using Lagrange interpolation, and outputs the secret.

2.8 Security with Abort

Since we are in the dishonest majority model, it is not possible to guarantee that the protocol will always successfully generate a signature as some parties may refuse to participate or send incorrect values. The protocol as described here is secure in the presence of aborts. This means that the protocol may abort without a signature, but security is maintained and the adversary will remain unable to forge messages.

In the protocol described here, it is not always possible to identify which player caused the protocol to fail. In [4], there is a variant of the protocol that supports identification of misbehavior. However, this requires a cryptographic broadcast channel or a bulletin board (/blockchain). Based on our discussions, we did not include the identifiability feature in this document.

When a protocol aborts, the pseudocode contains the `Abort` call. When this is called, the signer should broadcast a message that it has aborted and refuse to participate further in the aborted instantiation of the protocol.

2.9 Notation

In the pseudocode, each function invocation marked in red is a call to a sub-function. Values marked in green should be stored persistently for the duration of the signing protocol as they will be needed in subsequent rounds. We also note such values as return parameters of the function in which they are initially computed.

We use the notation $[x_j]_{j \neq i}$ to denote an array of values $[x_1, x_2, \cdots, x_{i-1}, x_{i+1}, \cdots, x_n]$ with the value x_j corresponding to player $\mathcal{P}_j$ for every other player $\mathcal{P}_{j \neq i}$.

3 Key Generation with a trusted dealer

We now describe the key generation protocol in the presence of a trusted dealer. The protocol assumes that each player has a unique, public, non-zero index p_i that is assigned to that player. As a practical matter, we recommend assigning each player a sequential index beginning at 1.

The high-level purpose of the key generation procedure is to output (1) a public ECDSA signing key, (2) private key shares for each signer corresponding to the public signing key, (3) Paillier key pairs for each player and (4) trusted parameters for instantiating the zero-knowledge proofs that will be employed during the signing protocol.

As input, the Dealer receives the index of each player, the signing threshold t such that $t+1$ players[1] must participate to sign, and the curve parameters. The dealer also needs to choose the size of the proof modulus, but we have hard-coded suitable parameters.

We note that in the event that multiple key generations are performed with different signing keys, the proof parameters can be reused as long as the dealer is trusted by all of the signers. For this reason we have separated the dealer into two functions, the first of which is a global setup for generating proof parameters and the second is a per-signing group/key setup.

We note that having a trusted dealer generate the proof parameters (i.e. DealerGenerateProofParams()) will lead to a significant performance increase in the signing protocol. By contrast, having the dealer generate the other parameters (DealerKeyGen()) simplifies the key generation but will not simplify the signing protocol. It thus might make sense to restrict the dealer to running DealerGenerateProofParams() which as noted above can be run once for all key generation instances, and use the distributed key generation procedure for all subsequent instantiations.

$(\tilde{N}, h_1, h_2) \leftarrow$ DealerGenerateProofParams()

1. Choose a 1024-bit safe prime P // $P = 2p + 1$ where both P and p are prime
2. Choose a 1024-bit safe prime Q // $Q = 2q + 1$ where both Q and q are prime
3. Compute $\tilde{N} = P \cdot Q$
4. Choose $f \xleftarrow{\$} \mathbb{Z}_{\tilde{N}^*}$
5. Choose $\alpha \xleftarrow{\$} \mathbb{Z}_{\tilde{N}^*}$
6. Compute $h_1 = f^2 \bmod \tilde{N}$ // square f
7. Compute $h_2 = h_1^{\alpha} \bmod \tilde{N}$
8. `Return` $\tilde{N}, h_1, h_2$

$y, [x_1, x_2, \ldots, x_n], [X_1, X_2, \ldots, X_n] \leftarrow$DealerKeyGen$(g, q, [p_1, p_2, \ldots, p_n], t)$

1. Choose $x \xleftarrow{\$} \mathbb{Z}_q$
2. Compute $y = g^x$
3. $[x_1 \ldots, x_n] \leftarrow$ ShamirShare$(x, t, \{p_1, \ldots, p_n\})$
4. `For` $i = [1, \ldots, n]$
5. Compute $X_i = g^{x_i}$
6. Compute $sk_i, pk_i =$ PaillierKeyGen(1^{κ}) // generate a Paillier key pair for each player
7. `Return` $y, [x_1, x_2, \ldots, x_n], [X_1, X_2, \ldots, X_n]$

Fig. 3: **Threshold key generation with a trusted dealer.** We split the function into two parts: one for creating the proof parameters which can be run once to establish global parameters for all key generations that trust the dealer running it (DealerGenerateProofParams), and one for generating the ECDSA key shares, ECDSA public key, and the Paillier key for each player (DealerKeyGen).

At the end of the key generation, the dealer distributes the following public parameters to each player:

[1] Note that in this document a threshold value of t implies that $t+1$ signer are needed to generate a signature. This notation is consistent with the academic literature, but inconsistent with e.g. Bitcoin multi-signatures in which (k, n) signatures require the participation of k parties (and not $k+1$).

1. The ECDSA public key: y
2. The index of every player: $p_i, \ldots, p_n$
3. The ECDSA public key share for all players: $[X_i, \ldots, X_n]$
4. The Paillier public keys of all players: $[pk_1, \ldots, pk_n]$
5. The global proof parameters: $\tilde{N}, h_1, h_2$

Additionally, for every player $\mathcal{P}_i$, the dealer distributes the following which is to be kept privately by that player:

1. The ECDSA private key share: x_i
2. The Paillier private key: sk_i

We stress that the ECDSA verification algorithm will only need the public key y. The other parameters are required to securely generate the signature using the distributed signing protocol.

4 Going from key generation to signing

After the key generation, each player has a key share that will allow them to participate in the signing protocol. But in a (t, n) configuration, the key generation is done with all n players whereas only a subset of $t+1$ players participate in the signing protocol. Once the group of $t+1$ signers has been identified, the players must convert their Shamir private key shares that were output during the signing protocol to additive shares in which the secret is shared additively among the $t+1$ active participants. Doing this is straightforward and simply requires the players to multiply their shares by the Lagrange coefficients of the active signers. We describe the conversion function here. The function is non-interactive and run locally by each player.

$(w_i, [W_1, \ldots, W_{t+1}]) \leftarrow$ ConvertToAdditive($x_i, [p_1, \ldots, p_{t+1}], [X_1, \ldots, X_{t+1}]$)

```
For  j = [1, ..., t + 1] {
     Set ℓ = 1
     For  k = [1, ..., t + 1] {
          If j = k, Continue
          Compute ℓ = ℓ × p_k / (p_k − p_j) mod q
     }
     Compute W_j = X_j^ℓ in G
     If j = i {
          Compute w_i = x_i × ℓ mod q
     }
}
Return w_i, [W_1, ..., W_{t+1}]
```

Fig. 4: **Converting (t, n) key shares to $t+1$ additive shares.** Each player $\mathcal{P}_i$ runs this function locally. Recall that during the key generation protocol, each player obtains a private key share x_i as well as a public key share X_j for every player $\mathcal{P}_j$. The inputs to the function are x_i, the private key share of $\mathcal{P}_i$, the set of indexes for $(t+1)$ players that will actively participate in the signing protocol, as well as their public key shares X_j. The protocol outputs w_i the additive private key share for player $\mathcal{P}_i$ as well as $W_j = g^{w_j}$ for each other player $\mathcal{P}_j$.

5 Signing protocol a with trusted dealer

Below, we give pseudocode for the first signing round of our protocol. The code is symmetric for every player and represents the view point of player $\mathcal{P}_i$.

Although the signing protocol doesn't directly make use of the trusted dealer, the fact that the dealer generated trusted parameters during the key generation protocol simplifies the protocol. In particular, without a trusted dealer, each player would generate their own proof parameters for the zero knowledge proofs (i.e. $\tilde{N}, h_1, h_2$). When proving, player $\mathcal{P}_i$ would need to generate a different proof for each player. If there is a trusted dealer, however, all players can use the same proving parameters reducing the number of proofs.

We note that in the pseudocode we mix together the networking code (i.e. calls to Broadcast and Send) together with the cryptographic code. When organizing the actual code though, it is probably simpler to separate these functionalities by having the cryptographic functions return both the values to store locally and the values to send, and have the calling code take care of the networking/storage.

$(k_i, \gamma_i, D_i, c_i, r_i) \leftarrow$SignRound1$(w_i, g, q, pk_i, \tilde{N}, h_1, h_2)$ // the function parameters are output during KeyGen

1. Choose $k_i \xleftarrow{\$} \mathbb{Z}_q$ // $\mathbb{Z}_q$ are the integers from 0 to $q-1$
2. Choose $\gamma_i \xleftarrow{\$} \mathbb{Z}_q$
3. Compute $\Gamma_i = g^{\gamma_i}$ in $\mathcal{G}$
4. Compute $[C_i, D_i] =$ Commit(Γ_i)
5. Compute $(c_i, r_i) =$ PaillierEncryptAndReturnRandomness(pk_i, k_i)
6. Compute $\pi_i^{\text{Range1}} =$ MtAProveRange1$(g, q, pk_i, \tilde{N}, h_1, h_2, k_i, c_i, r_i)$
7. Broadcast $(C_i, c_i, \pi_i^{\text{Range1}})$ to all other players
8. Return $k_i, \gamma_i, D_i, c_i, r_i$ // The returned values are stored locally for use in later rounds of the signing protocol

$[C_j, \beta_{ji}, \nu_{ji}]_{j\neq i} \leftarrow$SignRound2$([C_j, (c_j, \pi_j^{\text{Range1}})]_{j\neq i})$ // The C_j values are received now, but used in subsequent rounds

1. For $j = [1, \ldots, t+1]$
2. If $i = j$, Continue
3. If MtAVerifyRange1$(\pi_j^{\text{Range1}}, g, q, \tilde{N}, h_1, h_2, c_j) =$ False, Abort
4. Compute $(c_{ji}^{\gamma}, \beta_{ji}, \pi_{ji}^{\text{Range2}}) =$ MtAResponse$(\gamma_i, g, q, pk_j, \tilde{N}, h_1, h_2, c_j)$
5. Compute $(c_{ji}^{w}, \nu_{ji}, \pi_{ji}^{\text{Range3}}) =$ MtAResponse_wc$(w_i, W_i, g, q, pk_j, \tilde{N}, h_1, h_2, c_j)$
6. P2PSend$(c_{ji}^{\gamma}, c_{ji}^{w}, \pi_{ji}^{\text{Range2}}, \pi_{ji}^{\text{Range3}})$ to player $\mathcal{P}_j$
7. Return $[C_j, \beta_{ji}, \nu_{ji}]_{j\neq i}$

$(\delta_i, \sigma_i) \leftarrow$SignRound3$([c_{ij}^{\gamma}, c_{ij}^{w}, \pi_{ij}^{\text{Range2}}, \pi_{ij}^{\text{Range3}}]_{j\neq i})$

1. Compute $\delta_i = k_i \gamma_i \bmod q$
2. Compute $\sigma_i = k_i w_i \bmod q$
3. For $j = [1, \ldots, t+1]$
4. If $i = j$, Continue
5. Compute $\alpha_{ij} =$MtAFinalize$(g, q, sk_i, pk_i, \tilde{N}, h_1, h_2, c_i, c_{ij}^{\gamma}, \pi_{ij}^{\text{Range2}})$
6. If $\alpha_{ij} = \perp$, Abort // $\alpha_{ij} = \perp$ if proof verification failed
7. Compute $\mu_{ij} =$MtAFinalize_wc$(g, q, sk_i, pk_i, \tilde{N}, h_1, h_2, c_i, c_{ij}^{w}, \pi_{ij}^{\text{Range3}}, W_j)$
8. If $\mu_{ij} = \perp$, Abort // $\mu_{ij} = \perp$ if proof verification failed
9. Compute $\delta_i = \delta_i + \alpha_{ij} + \beta_{ji} \bmod q$
10. Compute $\sigma_i = \sigma_i + \mu_{ij} + \nu_{ji} \bmod q$
11. Broadcast δ_i to all other players
12. Return δ_i, σ_i

$\delta \leftarrow$SignRound4$([\delta_j]_{j\neq i})$

1. Set $\delta = \delta_i$
2. For $j = [1, \ldots, t+1]$
3. If $i = j$, Continue
4. Compute $\delta = \delta + \delta_j \bmod q$
5. Broadcast D_i to all other players
6. Return δ

$(r, \bar{R}_i) \leftarrow$SignRound5$([D_j]_{j\neq i})$

1. Compute $R = g^{\gamma_i}$ in $\mathcal{G}$
2. For $j = [1, \ldots, t+1]$
3. If $i = j$, Continue
4. Compute $\Gamma_j =$ Open(C_j, D_j)
5. If $\Gamma_j = \perp$, Abort
6. Else Compute $R = R \cdot \Gamma_j$ in $\mathcal{G}$
7. Compute $R = R^{\delta^{-1}}$ in $\mathcal{G}$
8. Set $r = R_x$ // R_x denotes the x coordinate of the elliptic curve point R
9. Compute $\bar{R}_i = R^{k_i}$
10. Compute $\pi_i^{k\text{CONSIST}} =$ProvePDL$(g, q, R, pk_i, \tilde{N}, h_1, h_2, k_i, \bar{R}_i, c_i, r_i)$
11. Broadcast $(\bar{R}_i, \pi_i^{k\text{CONSIST}})$ to all other players
12. Return $r, \bar{R}_i$

$s_i \leftarrow$SignRound6$(M, [R_j]_{j\neq i})$

1. Set $V = \bar{R}_i$
2. For $j = [1, \ldots, t+1]$
3. If $i = j$, Continue
4. If VerifyPDL$(\pi_j^{k\text{CONSIST}}, g, q, R, pk_j, \tilde{N}, h_1, h_2, c_j, \bar{R}_j) =$ False, Abort
5. Compute $V = V \cdot \bar{R}_j$ in $\mathcal{G}$
6. If $V \neq g$, Abort
7. Compute $m = H(M) \in \mathbb{Z}_q$ // Hash the message with the hash function used by the centralized ECDSA signer/verifier
8. Compute $s_i = mk_i + r\sigma_i \bmod q$
9. Broadcast s_i to all other players
10. Return s_i

$\sigma \leftarrow$ SignOutput$([s_j]_{j\neq i})$

1. Set $s = s_i$
2. For $j = [1, \ldots, t+1]$
3. If $i = j$, Continue
4. Compute $s = s + s_j \bmod q$
5. Set $\sigma = (r, s)$
6. If ECDSAVerify$(y, \sigma, M) =$ False, Abort
7. Return σ

Fig. 5: **The ECDSA threshold signature generation protocol.**

6 One round signing

Notice that in the pseudocode above, the message M is first input to the SignRound6 function. To achieve non-interactive signing, the players simply run the first five rounds of the protocol offline, and indeed they can run many such instantiations in parallel (see introduction regarding concurrent security in our protocol). Then later upon receiving the message to sign, they can complete the protocol by beginning from SignRound6, leading to only a single round in which every player broadcasts a single message.

Indeed, when executing the protocol in the manner, even part of SignRound6 can be performed during the pre-processing phase, and doing so will require the players to store less information for the online phase. We now break up SignRound6 into two functions: one which can be run offline during the pre-processing, and one which is message dependent.

$r, k_i, \sigma_i \leftarrow$ SignRound6Offline$([R_j]_{j \neq i})$

1. Set $V = R_i$
2. `For` $j = [1, \ldots, t+1]$
3. `If` $i = j$, `Continue`
4. `If` VerifyPDL$(\pi_j^{k^{\text{CONSIST}}}, g, q, R, pk_j, \tilde{N}, h_1, h_2, c_j, \bar{R}_j) =$ `False`, `Abort`
5. Compute $V = V \cdot \bar{R}_j$ in $\mathcal{G}$
6. `If` $V \neq g$, `Abort`
7. `Return` r, k_i, σ_i

$s_i \leftarrow$ SignRound6Online(M, r, k_i, σ_i)

1. Compute $m = H(M) \in \mathbb{Z}_q$ // Hash the message with the hash function used by the centralized ECDSA signer/verifier
2. Compute $s_i = mk_i + r\sigma_i \bmod q$
3. Broadcast s_i to all other players
4. `Return` s_i

Fig. 6: **Modification for one round signing.** To achieve one round signing, the first five rounds and part of the sixth round can be run during an offline message-independent pre-processing phase. Here we split SignRound6 into two functions showing which parts can be run during pre-processing, and which must be run online once the message is known. Notice that the values returned by SignRound6Offline are the values that must be stored persistently for the online phase, and indeed these are listed as function parameters to SignRound6Online.

To execute the protocol with one online round, the first five rounds as well as SignRound6Offline are run during a pre-processing phase. Then, during the online phase, only SignRound6Online and SignOutput are run. Notice that the only values that need to be kept for the online round are r, k_i, σ_i. All other values can be discarded.

Keeping state. We stress that just like during the fully online protocol, we require an independent execution of the protocol for each generated signature. The values r, k_i, σ_i must only be ever used in one call of the signing protocol (and even if the protocol aborts, these values must not be re-used).

For safety, we've presented the protocol such that the pre-processing phase is done with a specific set of $t + 1$ signers. This minimizes the chance that a pre-processed value tuple will be accidentally re-used since every signer needs to participate and thus they cannot be partitioned such that different subgroups re-use a tuple. Nevertheless, much care should be taken so that every signers discards the tuple after it is used (no matter whether a signature was successfully generated or the protocol aborted).

7 Multiplicative-to-Additive Share Conversion Protocol (MtA)

Perhaps the most involved part of the signing protocol is the sub-protocol for converting multiplicative secret shares to additive shares of their product. We present the details of the MtA protocol here as well as complete pseudocode.

The setting consists of two players, $\mathcal{P}_1$ and $\mathcal{P}_2$, who hold multiplicative shares of a secret x. In particular, $\mathcal{P}_1$ holds a share $a \in Z_q$, and $\mathcal{P}_2$ holds a secret share $b \in Z_q$ such that $x = ab \bmod q$. The goal of the MtA protocol is to convert these multiplicative shares into additive shares. $\mathcal{P}_1$ receives private output $\alpha \in Z_q$ and $\mathcal{P}_2$ receives private output $\beta \in Z_q$ such that $\alpha + \beta = x = ab \bmod q$.

In the basic MtA protocol, the player's inputs are not verified, and indeed the players can cause the protocol to produce an incorrect output by inputting the wrong values $\hat{a}, \hat{b}$. In the case that $B = g^b$ is public, the protocol can be enhanced to include an extra check that ensures that $\mathcal{P}_2$ inputs the correct value $b = \log_g(B)$. This enhanced protocol is denoted as MtAwc (for MtA "with check").

In the protocol, $\mathcal{P}_1$ first encrypts the value a under its own key and sends the ciphertext c_1 to $\mathcal{P}_2$. $\mathcal{P}_2$ then uses the homomorphic property of the encryption scheme to multiply its value b into the ciphertext c_1 and add a mask β' resulting in a new ciphertext c_2 that it sends back to $\mathcal{P}_1$. $\mathcal{P}_1$ decrypts the ciphertext to reveal its share α, and $\mathcal{P}_2$ sets its share to $\beta = -\beta'$.

Note that $\alpha = \mathsf{PaillierDecrypt}(pk_1, c_2) = a \cdot b + \beta'$ and thus $\alpha + (\beta) = a \cdot b$ as desired.

7.1 MtA and MtAwc initiation

In the first phase of the MtA protocol the initiator encrypts its message a and sends the ciphertext c together with a range proof π^{Range1} to the other player.

In this document, we present the range proof here as a sub-function, but we implement the Paillier encryption of the first message directly into the first round of the signing protocol (i.e. SignRound1). We do this primarily for clarity and the ability to write generic functions that will work both for the dealer and dealerless protocols, since in the dealerless version of our protocol we require a single ciphertext for all players, but a unique proof for every other player.

We now present pseudocode for the prover and verifier functions for the initiator's range proof.

As one of the inputs to both MtAProveRange1/MtAVerifyRange1, the prover's Paillier's public key is supplied. Similarly, in MtAResponse/MtAProveRange2/MtAVerifyRange2 the verifier's public key is supplied. In practice this is N, but for clarity, we refer to it as pk in the function header and use the notation $pk.N$ when accessing N to demonstrate that N is stored as part of the respective player's public key.

$\pi \leftarrow$ MtAProveRange1$(g, q, pk, \tilde{N}, h_1, h_2, a, c, r)$

1. Set $N = pk.N$
2. Choose $\alpha \xleftarrow{\$} \mathbb{Z}_{q^3}$
3. Choose $\beta \xleftarrow{\$} \mathbb{Z}_N^*$
4. Choose $\gamma \xleftarrow{\$} \mathbb{Z}_{q^3\tilde{N}}$
5. Choose $\rho \xleftarrow{\$} \mathbb{Z}_{q\tilde{N}}$
6. Compute $z = h_1^a h_2^\rho \bmod \tilde{N}$
7. Compute $u = (N+1)^\alpha \beta^N \bmod N^2$
8. Compute $w = h_1^\alpha h_2^\gamma \bmod \tilde{N}$
9. Compute $e =$ FS-HASH$(g, q, pk, \tilde{N}, h_1, h_2, c, z, u, w)$
10. Compute $s = r^e \beta \bmod N$
11. Compute $s_1 = ea + \alpha$ // computed over the integers
12. Compute $s_2 = e\rho + \gamma$ // computed over the integers
13. Set $\pi = [z, e, s, s_1, s_2]$
14. Return π

`True/False` $\leftarrow$MtAVerifyRange1$(\pi = [z, e, s, s_1, s_2]), g, q, pk, \tilde{N}, h_1, h_2, c)$

1. Set $N = pk.N$
2. `If` $s_1 > q^3$, `Return False` // Check range
3. Compute $\hat{u} = (N+1)^{s_1} s^N c^{-e} \bmod N^2$
4. Compute $\hat{w} = h_1^{s_1} h_2^{s_2} z^{-e} \bmod \tilde{N}$
5. Compute $\hat{e} =$ FS-HASH$(g, q, pk, \tilde{N}, h_1, h_2, c, z, \hat{u}, \hat{w})$
6. `If` $\hat{e} \neq e$, `Return False`
7. `Return True`

Fig. 7: **Initiator proof in the** MtA **and** MtAwc **protocols.** This proof is run by the initiator and is identical in both the MtA and MtAwc protocols.

7.2 MtA Response

We now describe the MtA/MtAwc procedures for the respondent that is executed upon receiving the proof and ciphertext from the initiator. The respondent first verifies the received proof (using the procedure from Section 7.1) and then crafts its own ciphertext as well as its own proof using the MtAResponse and MtAResponse_wc functions. The respondent's functions for MtA and MtAwc are quite similar but subtly different. To avoid duplicating text, we mark in blue the parts of the protocol that are only run during MtAwc.

$(c_2, \beta, \pi^{\text{Range2}}) \leftarrow$ MtAResponse_wc$(b, B, g, q, pk, \tilde{N}, h_1, h_2, c_1)$

1. Compute $c_b =$ PaillierMultiply(pk, b, c_1)
2. Set $N = pk.N$
3. Choose $\beta' \xleftarrow{\$} \mathbb{Z}_N$
4. Compute $(c_\beta, r_\beta) =$ PaillierEncryptAndReturnRandomness(pk, β')
5. Compute $c_2 =$ PaillierAdd(pk, c_b, c_β)
6. Compute $\beta = -\beta' \bmod q$
7. Compute $\pi^{\text{Range2}} =$ MtAProveRange2_wc $(g, q, pk, \tilde{N}, h_1, h_2, b, \beta', r_\beta, c_1, c_2, B)$
8. `Return` $c_2, \beta, \pi^{\text{Range2}}$

Fig. 8: **The** MtA **and** MtAwc **response function.** This function is run by the respondent in the MtA protocol after successfully verifying the range proof sent by the initiator. By running this function, the respondent obtains the ciphertext c_2 and the range proof π^{Range2} that he sends back to the initiator as well as the private additive share β which is kept secret and not sent. The protocols for MtA and MtAwc are nearly identical with the only different being that in MtA, the caller invokes the MtAProveRange2 sub-function, whereas in MtAwc, the caller invokes the MtAProveRange2_wc sub-function, which takes an extra parameter: $B = g^b$. To avoid duplicating text, we mark in blue the parts of the protocol that are only run during MtAwc.

7.3 The Respondent's zero knowledge proof

We now present the proving and verification pseudocode for the zero knowledge proofs run by the respondent. As before, we mark in blue the parts of the protocol that are only run during MtAwc. We use the notation from [3] for this proof, but note that x in this proof corresponds to b in the protocol, X corresponds to B, and y corresponds to β.

$\pi \leftarrow$ MtAProveRange2_wc$(g, q, pk, \tilde{N}, h_1, h_2, x, y, r, c_1, c_2, X)$

1. Set $N = pk.N$
2. Choose $\alpha \xleftarrow{\$} \mathbb{Z}_{q^3}$
3. Choose $\rho \xleftarrow{\$} \mathbb{Z}_{q\tilde{N}}$
4. Choose $\rho' \xleftarrow{\$} \mathbb{Z}_{q^3\tilde{N}}$
5. Choose $\sigma \xleftarrow{\$} \mathbb{Z}_{q\tilde{N}}$
6. Choose $\beta \xleftarrow{\$} \mathbb{Z}_N^*$
7. Choose $\gamma \xleftarrow{\$} \mathbb{Z}_N^*$
8. Choose $\tau \xleftarrow{\$} \mathbb{Z}_{q\tilde{N}}$
9. Compute $u = g^{\alpha}$
10. Compute $z = h_1^x h_2^{\rho} \bmod \tilde{N}$
11. Compute $z' = h_1^{\alpha} h_2^{\rho'} \bmod \tilde{N}$
12. Compute $t = h_1^y h_2^{\sigma} \bmod \tilde{N}$
13. Compute $v = c_1^{\alpha}(N+1)^{\gamma}\beta^N \bmod N^2$
14. Compute $w = h_1^{\gamma} h_2^{\tau} \bmod \tilde{N}$
15. Compute $e =$ FS-HASH$(g, q, pk, \tilde{N}, h_1, h_2, X, c_1, c_2, u, z, z', t, v, w)$
16. Computes $s = r^e\beta \bmod N$
17. Compute $s_1 = ex + \alpha$ // computed over the integers
18. Compute $s_2 = e\rho + \rho'$ // computed over the integers
19. Compute $t_1 = ey + \gamma$ // computed over the integers
20. Compute $t_2 = e\sigma + \tau$ // computed over the integers
21. Set $\pi = [z, z', t, e, s, s_1, s_2, t_1, t_2]$
22. `Return` π

`True/False` $\leftarrow$ MtAVerifyRange2_wc$(\pi = [z, z', t, e, s, s_1, s_2, t_1, t_2], g, q, pk, \tilde{N}, h_1, h_2, c_1, c_2, X)$

1. Set $N = pk.N$
2. `If` $s_1 > q^3$, `Return False` // check range
3. Compute $s_1' = s_1 \bmod q$
4. Compute $\hat{u} = g^{s_1'} \cdot X^{-e}$ in $\mathcal{G}$
5. Compute $\hat{z'} = (h_1)^{s_1} \cdot (h_2)^{s_2} z^{-e} \bmod \tilde{N}$
6. Compute $\hat{v} = (c_1)^{s_1} \cdot s^N \cdot (N+1)^{t_1} \cdot c_2^{-e} \bmod N^2$
7. Compute $\hat{w} = (h_1)^{t_1} \cdot (h_2)^{t_2} \cdot t^{-e} \bmod \tilde{N}$
8. Compute $\hat{e} =$ FS-HASH$(g, q, pk, \tilde{N}, h_1, h_2, X, c_1, c_2, \hat{u}, z, \hat{z'}, t, \hat{v}, \hat{w})$
9. `If` $\hat{e} \neq e$, `Return False`
10. `Return True`

Fig. 9: **Respondent proof in the MtA and MtAwc protocols.** This figure shows the proofs for both the MtA and MtAwc protocol. In both, the proof includes a range proof, but in the MtAwc there is an additional check for consistency with $X = g^x$. The proofs are mostly the same, with some extra checks added in the MtAwc checks. For simplicity, we have marked the items in blue that are only run during MtAwc and omitted during MtA.

7.4 The initiator's finalization function

In the last phase of the MtA/MtAwc protocols, upon receiving the respondent's message, the initiator checks the zero knowledge proof and calculates its own additive secret share. We present the pseudocode for this here.

$\alpha \leftarrow$ MtAFinalize_wc$(g, q, sk, pk, \tilde{N}, h_1, h_2, c_1, c_2, \pi^{\text{Range2}}, B)$

1. If MtAVerifyRange2_wc$(\pi^{\text{Range2}}, g, q, pk, \tilde{N}, h_1, h_2, c_1, c_2, B) =$ False, Return $\perp$
2. Compute $\alpha = \text{Decrypt}(sk, c_2) \bmod q$
3. Return α

Fig. 10: **The MtA finalization function.** Upon receiving $(c_2, \pi^{\text{Range2}})$, $\mathcal{P}_1$ calls MtAFinalize to check the proof's validity compute its additive share α.

8 The zero knowledge proof in SignRound5

In SignRound5 the protocol requires a zero knowledge proof of consistency between a discrete logarithm and a value encrypted in a Paillier ciphertext. In the protocol this is used to prove that the value k_i that a player input to the MtA protocol is the same k_i that they used to compute $\bar{R} = R^{k_i}$.We present the pseudocode for this proof here. In particular, we need to show consistency between a ciphertext c such that $(c, r) =$ PaillierEncryptAndReturnRandomness(pk, x) and a curve point X such that $X = R^x$.

$\pi \leftarrow$ProvePDL$(g, q, R, pk, \tilde{N}, h_1, h_2, x, X, c, r)$

1. Set $N = pk.N$
2. Choose $\alpha \xleftarrow{\$} \mathbb{Z}_{q^3}$
3. Choose $\beta \xleftarrow{\$} \mathbb{Z}_N^*$
4. Choose $\gamma \xleftarrow{\$} \mathbb{Z}_{q^3\tilde{N}}$
5. Choose $\rho \xleftarrow{\$} \mathbb{Z}_{q\tilde{N}}$
6. Compute $u = R^\alpha$ in G
7. Compute $z = h_1^x h_2^\rho \bmod \tilde{N}$
8. Compute $v = (N+1)^\alpha \beta^N \bmod N^2$
9. Compute $w = h_1^\alpha h_2^\gamma \bmod \tilde{N}$
10. Compute $e =$ FS-HASH$(pk, \tilde{N}, h_1, h_2, g, q, R, X, c, u, z, v, w)$
11. Computes $s = r^e \beta \bmod N$
12. Compute $s_1 = ex + \alpha$ // computed over the integers
13. Compute $s_2 = e\rho + \gamma$ // computed over the integers
14. Set $\pi = [z, e, s, s_1, s_2]$
15. Return π

True/False $\leftarrow$ VerifyPDL$(\pi = [z, e, s, s_1, s_2], g, q, R, pk, \tilde{N}, h_1, h_2, c, X)$

1. Set $N = pk.N$
2. If $s_1 > q^3$, Return False // check range
3. Compute $s_1' = s_1 \bmod q$
4. Compute $\hat{u} = R^{s_1'} \cdot X^{-e}$ in $\mathcal{G}$
5. Compute $\hat{v} = s^N \cdot (N+1)^{s_1} \cdot c^{-e} \bmod N^2$
6. Compute $\hat{w} = (h_1)^{s_1} \cdot (h_2)^{s_2} \cdot z^{-e} \bmod \tilde{N}$
7. Compute $\hat{e} =$ FS-HASH$(pk, \tilde{N}, h_1, h_2, g, q, R, X, c, \hat{u}, z, \hat{v}, \hat{w})$
8. If $\hat{e} \neq e$, Return False
9. Return True

Fig. 11: **Consistency proof between a Paillier encrypted value and a discrete logarithm.** This figure shows the proof of consistency between a ciphertext c such that $(c, r) =$ PaillierEncryptAndReturnRandomness(pk, x) and a curve point X such that $X = R^x$.

References

1. Boneh, D.: Digital signature standard. In: Encyclopedia of cryptography and security, pp. 347–347. Springer (2011)

2. Doerner, J., Kondi, Y., Lee, E., et al.: Secure two-party threshold ecdsa from ecdsa assumptions. In: IEEE Symposium on Security and Privacy. p. 0. IEEE (2018)
3. Gennaro, R., Goldfeder, S.: Fast multiparty threshold ecdsa with fast trustless setup. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. pp. 1179–1194. ACM (2018)
4. Gennaro, R., Goldfeder, S.: One round threshold ecdsa with identifiable abort. In: IACR eprint Report 2020/540 (2020)
5. Kravitz, D.W.: Digital signature algorithm (Jul 27 1993), uS Patent 5,231,668
6. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: International Conference on the Theory and Applications of Cryptographic Techniques. pp. 223–238. Springer (1999)
7. Shamir, A.: How to share a secret. Communications of the ACM **22**(11), 612–613 (1979)