Threshold Receipt-Free Voting with Server-Side Vote Validation

Thi Van Thao Doan, Olivier Pereira, and Thomas Peters

UCLouvain – ICTEAM – Crypto Group
B-1348 Louvain-la-Neuve – Belgium
{thi.doan, olivier.pereira, thomas.peters}@uclouvain.be

Abstract. Proving the validity of ballots is a central element of verifiable elections. Such proofs can however create challenges when one desires to make a protocol receipt-free.
We explore the challenges raised by validity proofs in the context of protocols where threshold receipt-freeness is obtained by secret sharing an encryption of a vote between multiple authorities. In such contexts, previous solutions verified the validity of votes by decrypting them after passing them through a mix-net. This approach however creates subtle privacy risks, especially when invalid votes leak structural patterns that threaten receipt-freeness.
We propose a different approach of threshold receipt-free voting in which authorities re-randomize ballot shares then jointly compute a ZK proof of ballot validity before letting the ballots enter a (possibly homomorphic) tallying phase. Our approach keeps the voter computational costs limited while offering verifiability and improving the ballot privacy of previous solutions.
We present two protocols that enable a group of servers to verify and publicly prove that encrypted votes satisfy some validity properties: MiniMix, which preserves prior voter-side behavior with minimal overhead, and HomoRand, which requires voters to submit auxiliary data to facilitate validation over large vote domains. We show how to use our two protocols within a threshold receipt-free voting framework. We provide formal security proofs and efficiency analyses to illustrate trade-offs in our designs.

1 Introduction

Receipt-freeness (RF), is a central property of voting systems that ensures that voters cannot prove how they voted to a third party. Since its formal introduction by Benaloh and Tuinstra [2], RF has been the subject of extensive research, especially because of the tension that it creates with the transparency and verifiability of the election process.

A standard approach first explored by Hirt [13] consists in asking voters to encrypt their vote and submit the ciphertext to a ballot processing server that privately re-randomizes the ciphertext before posting it on a public bulletin-board. Since the re-randomization removes the voter's knowledge of the encryption randomness, voters cannot offer any evidence of the ciphertext content to

any third party. In Hirt's solution, the ballot processing server must also send a designated verifier proof that it re-randomized the ballot without modifying its content. Subsequent works have aimed to eliminate the need for sending such a proof while preserving RF using various types of randomizable signatures [3,4,6–8]. However, in all these works, receipt-freeness depends on a single trusted ballot processing server: if the ballot processing server leaks the randomness used to re-randomize the ballot, RF is lost (though privacy and verifiability would still be preserved).

In order to remove this single point of failure, Doan et al. [9] introduced a threshold receipt-free voting framework, leveraging multiple ballot processing servers (randomizers) to decentralize trust. Their system ensures RF and correctness as long as the number of corrupted randomizers remains below a threshold. Moreover, it maintains universal verifiability via mixnets.

However, Doan's protocol does not include any ballot validity proof in the ballot submission process: the validity of the votes is only verified when the mixed ciphertexts are decrypted. As a result, a coercer or a vote buyer could ask to see a very unusual invalid vote among the decrypted ballots (it could simply ask to encrypt a large randomly chosen value). This strategy, akin to the Italian attack [2], provides a limited form of receipt, in the sense that the voter can only demonstrate that he submitted an invalid vote, but it would still be desirable to avoid it whenever possible.

Worst situations may happen when the set of valid votes is very large. For instance, if a vote is by approval among 100 candidates, a voter could simply demonstrate how he voted by submitting a highly unusual approval pattern, which will come uniquely in the tally. This limitation highlights the need for alternative solutions that preserve the security requirements of threshold receipt-freeness while avoiding direct decryption of all individual votes, whether they are valid or not.

Tallying processes based on the homomorphic aggregation of ballots instead of a mixnet can offer a solution to these problems in the context of approval voting (and in other cases), while tally-hiding protocols offer more general solutions: in these protocols, individual votes are never decrypted.

Contributions. We propose a new paradigm for threshold receipt-free voting that addresses a core limitation in [9]: the lack of vote validity enforcement before tallying. Instead of requiring voters to generate complex zero-knowledge proofs at ballot submission time, we defer validity checking to a *pre-tally phase*, executed after combined ciphertexts are reconstructed but before they are tallied.

In particular, each voter secret-shares their vote and encrypts the shares under a traceable RF encryption scheme (TREnc) [6], then submits the resulting ciphertexts (ballot shares) to a set of processing servers. These servers rerandomize the ballot shares and post them to a public bulletin board. After voting concludes, any party (e.g., the talliers) reconstructs each final ciphertext by combining a threshold of rerandomized ballot shares. These ciphertexts are then validated in a *pre-tally phase* via a multi-party computation (MPC) protocol that checks whether each encrypts a vote in the prescribed domain, without

revealing the vote or auxiliary information. Valid ballots are then forwarded to a standard tallying procedure. This paradigm shift decouples vote-domain enforcement from the voter's submission step and moves it to a joint server-side protocol, overcoming limitations of threshold RF settings, where secret sharing and rerandomization rendered traditional ZK-based validity checks infeasible. We realize this paradigm through two constructions:

- MiniMix: mirrors the architecture of Doan et al. [9], with voters submitting encrypted shares of their vote under TREnc, and performs MPC-based pre-tally validation over the reconstructed ciphertext.
- HomoRand: allows voters to submit auxiliary information that facilitates more efficient validity checks in certain settings, while preserving RF.

These constructions extend the threshold RF model of [9] beyond mixnet-based designs to support homomorphic tallying over binary and general vote domains. HomoRand achieves asymptotically superior pre-tally performance (logarithmic in the bit-length required to represent the domain span), but with greater voter-side computation; MiniMix, by contrast, is optimized for low client overhead, making it well-suited to moderate domain sizes or when client efficiency is critical.

Unlike general-purpose frameworks for collaborative zero-knowledge, such as that of Ozdemir and Boneh [15], which adapt ZK-SNARKs to distributed-witness settings, our protocols are tailored to voting-specific constraints. Here, the witness originates from a potentially malicious voter, and correctness and privacy must be enforced jointly, without compromising RF. Notably, the RF definition we target does not cover scenarios where an adversary prevents voting, such as forcing abstention or invalid ballots.

Remark. Although our protocols are motivated by [9], our pre-tally validation technique is broadly applicable. For instance, one could remove the 0/1 proofs from the CGS97 protocol [5] and insert a round of MiniMix to enforce ballot validity. Such a substitution would functionally ensure well-formed ballots, shifting the computational efforts from voting clients to servers.

2 Building Blocks

Secret Sharing. We use Shamir's t-out-of-n threshold secret sharing scheme [16], consisting of two algorithms. The sharing algorithm $\mathsf{Share}(n, t, m)$ splits a secret m into n shares $(m_1, \ldots, m_n)$ such that any subset of at least t shares can reconstruct m, while any set of fewer than t shares reveals no information about m. The reconstruction algorithm $\mathsf{Combine}(n, t, m_1, \ldots, m_t)$ uses Lagrange interpolation to recover m from any t shares.

Traceable Receipt-Free Encryption (TREnc). TREnc [6] is a public key encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, augmented with a 5-tuple of algorithms: LGen, on input a security parameter λ and a public encryption key pk, outputs a link key lk; LEnc encrypts a message m using (pk, lk) and outputs a ciphertext CT. Trace outputs the trace τ of CT. Rand randomizes CT to output another ciphertext. Ver checks if a ciphertext is valid and outputs 1 if true, and 0 otherwise.

Informally, a TREnc scheme is *traceable* if no efficient adversary can produce a second ciphertext with the same trace as a given ciphertext, yet which decrypts to a different message. It is *TCCA-secure* (Traceable Chosen-Ciphertext Attack secure) if re-randomized ciphertexts are indistinguishable from fresh ciphertexts with the same trace, even in the presence of a restricted decryption oracle. We assume it is straightforward to extract specific parts of the TREnc's ciphertext CT. In particular:

Strip(pk, CT): Returns the CPA-encryption component of CT. For instance, in TREnc [6], this component is $\boldsymbol{c} = (c_0, c_1, c_2) = (G^m f^\theta, g^\theta, h^\theta)$, where m is the message, $\theta \leftarrow\!\!\$\, \mathbb{Z}_p$, and $(G, f, g, h) \in \mathsf{pk}$.

CPA(pk, m; r): Computes the CPA component of a TREnc ciphertext on message m with randomness r under TREnc's public key pk. We note that this algorithm can be used independently of TREnc.

We further assume that the underlying CPA encryption is additively homomorphic over the message space $\mathbb{Z}_p$ and that the encryption operates over a cyclic group of prime order p. That is, given $\boldsymbol{C}_1 = \mathsf{CPA}(\mathsf{pk}, m_1; r_1)$ and $\boldsymbol{C}_2 = \mathsf{CPA}(\mathsf{pk}, m_2; r_2)$, one can compute $\boldsymbol{C}_1 \cdot \boldsymbol{C}_2 = \mathsf{CPA}(\mathsf{pk}, m_1 + m_2; r_1 + r_2)$ (up to group notation). In many cases, we omit the randomness when it is sampled uniformly at random. In the rest of the paper, we use the prefix TREnc. to indicate that an algorithm belongs to the TREnc suite.

Non-Interactive Proofs. Informally, a proof for an NP relation R is a protocol by which a prover P convinces a probabilistic polynomial-time (PPT) verifier V that $\exists w : R(w, x) = 1$, where x is a called a statement, and w is a witness for x. In the non-interactive setting, the proof consists of a single message from P to V and proceeds as follows. A setup algorithm $\mathsf{PrfSetup}(1^\lambda)$, on input a security parameter λ, generates a common reference string (CRS) σ. The prover then computes $\mathsf{PrfProve}(\sigma, x; w)$, outputting a proof π if $R(x, w) = 1$, or $\perp$ otherwise. Finally, the verifier checks the proof via $\mathsf{PrfVerify}(\sigma, x, \pi)$ and outputs 1 if the proof is valid, and 0 otherwise.

A non-interactive zero-knowledge proof of knowledge (NIZKPoK) satisfies completeness, knowledge soundness, and zero-knowledge [11].

3 Verifiable Ciphertext Validity in MPC

We present two MPC-based constructions for verifying whether a ciphertext $\boldsymbol{C} = \mathsf{CPA}(\mathsf{PK}, m)$, under a CPA-secure encryption scheme with the corresponding decryption $\mathsf{Dec}(\mathsf{SK}, \boldsymbol{C}) = m$, encrypts a message m satisfying a given constraint (e.g., $m \in \{v_1, \ldots, v_d\} \subset \mathbb{Z}_p$). The protocol reveals only the outcome of the check and leaks no additional information about the underlying plaintext or any related value, even in the case of failure. In the following, we describe two such approaches in detail.

3.1 MPC-MiniMix Protocol

The first construction, referred to as MPC-MiniMix (Algorithm 1), employs an OR-proof mechanism to verify message validity. The core idea is to check whether

any of the ciphertexts $\boldsymbol{C}_j = \mathsf{CPA}(\mathsf{PK}, m - v_j)$ encrypts the value 0, where each $\boldsymbol{C}_j$ is derived as $\boldsymbol{C} \cdot \mathsf{CPA}(\mathsf{PK}, v_j)^{-1}$ for $j \in [d]$, exploiting the additive homomorphism of the CPA scheme. Conceptually, this construction generalizes plaintext equivalence proofs [14] to enable a form of privacy-preserving range verification.

The sender first encrypts the message m, from which all parties can compute the ciphertexts $\{\boldsymbol{C}_j\}_{j\in[d]}$. The parties then engage in a collaborative verification process, structured as a simplified mixnet operating over the d ciphertexts. Each party $\mathcal{T}_k$, *in turn*, shuffles the ciphertexts (lines 4–9), with the output of $\mathcal{T}_k$ serving as the input to $\mathcal{T}_{k+1}$, for all $k \in [T-1]$, where T denotes the number of parties. Each shuffle includes a re-randomization step (line 7), ensuring that—under honest execution—the ciphertext encrypting 0 becomes unlinkable to its initial position, thereby preserving zero-knowledge.

A significant challenge arises when the message m is a carefully crafted value. In such cases, decryption may reveal not only the success or failure of the range verification but also partial information about m, specifically $m - v_j$ for some j. This could potentially assist the adversary in extracting structural patterns and re-identify the voter. To mitigate this risk, we incorporate a masking step following the shuffle: after shuffling, each party $\mathcal{T}_k$ raises each ciphertext to a fresh random exponent $\alpha_j^{(k)}$ and proves correctness via a ZKPoK $\boldsymbol{\pi}^{(k)}$ (lines 10–14). These ciphertexts and proofs are then broadcast (line 15). The shuffle proof can be instantiated using any standard zero-knowledge proof system for shuffle correctness (e.g., [1, 17]), while the re-randomization proof can be implemented via a standard Σ-protocol.

We assume an honest majority setting. If the majority finds that a proof fails verification (line 16), the corresponding party $\mathcal{T}_k$ is excluded from the protocol, and the output of the last honest party is used to proceed. Although the shuffling and exponentiation phases are presented separately for clarity and to support later comparative analysis (Section 5.4), they can be efficiently merged into a single phase with a unified zero-knowledge proof. Finally, the parties jointly decrypt the resulting ciphertexts (line 20). If the message m is valid, at least one ciphertext will decrypt to 0 while being different of $\mathsf{CPA}(\mathsf{PK}, 0; 0)$; otherwise all decrypted values appear as random group elements, preventing any information leakage about m.

3.2 MPC-HomoRand Protocol

In this approach, we leverage pairings to ensure that the encrypted message m lies within a prescribed integer interval $[v_1, v_d] \subset \mathbb{Z}_p$. Without loss of generality, let l be the smallest number of bits such that the span of the domain satisfies $v_d - v_1 < 2^l$. To prove that $m \in [v_1, v_d]$ (inclusive), it suffices to verify that $m - v_1 \in [0, 2^l - 1]$ and $v_d - m \in [0, 2^l - 1]$. For simplicity, here we demonstrate how to prove that $m \in [0, 2^l - 1]$. This is accomplished by having the sender decompose m into an l-bit string $b_1 b_2 \dots b_l$, and separately encrypt each bit in two distinct source groups $\mathbb{G}$ and $\hat{\mathbb{G}}$. Subsequently, a MPC protocol operates in the target group $\mathbb{G}_{\mathcal{T}}$ to jointly verify that each encrypted bit b_j satisfies $b_j \in \{0, 1\}$ for all $j \in [l]$. This binary encoding allows the sender's computational effort to scale logarithmically with the bit length l, i.e., $O(\log l)$.

Algorithm 1 MPC-MiniMix: Sequential Multi-Party Randomization & Verification

procedure MPC-MiniMix($\mathsf{PK}, \mathsf{SK}, \boldsymbol{C} = \{\boldsymbol{C}_j\}_{j\in[d]}$)
  Initialize $\boldsymbol{C}^{(0)} \leftarrow \boldsymbol{C}$ ▷ i.e., $\boldsymbol{C}_j^{(0)} \leftarrow \boldsymbol{C}_j$ for $j \in [d]$
  for $k = 1$ to T **do** ▷ Each party acts sequentially
    Sample at random a permutation P_k of $\{1, \ldots, d\}$
    for $j = 1$ to d **do** ▷ Shuffle the ciphertexts
      $s_j^{(k)} \leftarrow\!\!\$\ \mathbb{Z}_p \setminus \{0\}$
      $\boldsymbol{C}'^{(k)}_{P_k(j)} \leftarrow \boldsymbol{C}_j^{(k-1)} \cdot \mathsf{CPA}(\mathsf{PK}, 1_{\mathbb{G}}; s_j^{(k)})$
    end for
    $\boldsymbol{\pi}_{sf}^{(k)} \leftarrow \mathsf{PrfProve}(\mathsf{PK}, \{\boldsymbol{C}'^{(k)}_j, \boldsymbol{C}_j^{(k-1)}\}_{j\in[d]}; P_k, \{s_j^{(k)}\}_{j\in[d]})$
    for $j = 1$ to d **do** ▷ Apply random factors
      Sample random values $\alpha_j^{(k)}, r_j^{(k)} \leftarrow\!\!\$\ \mathbb{Z}_p \setminus \{0\}$
      $\boldsymbol{C}_j^{(k)} \leftarrow (\boldsymbol{C}'^{(k)}_j)^{\alpha_j^{(k)}} \cdot \mathsf{CPA}(\mathsf{PK}, 1_{\mathbb{G}}; r_j^{(k)})$
      $\boldsymbol{\pi}_{rd,j}^{(k)} \leftarrow \mathsf{PrfProve}(\mathsf{PK}, \boldsymbol{C}'^{(k)}_j, \boldsymbol{C}_j^{(k)}; \alpha_j^{(k)}, r_j^{(k)})$
    end for
    Publish $\boldsymbol{C}^{(k)} = \{\boldsymbol{C}_j^{(k)}\}_{j\in[d]}$ and $\boldsymbol{\pi}^{(k)} = (\boldsymbol{\pi}_{sf}^{(k)}, \{\boldsymbol{\pi}_{rd,j}^{(k)}\}_{j\in[d]})$
    if $\mathsf{PrfVer}(\boldsymbol{\pi}^{(k)}) = 0$ **then return** $\perp$
    end if
  end for
  for $j = 1$ to d **do**
    $m_j \leftarrow \mathsf{Dec}(\mathsf{SK}, C_j^{(T)})$ ▷ Decrypt the final ciphertexts
    if $m_j = 0$ **then return** 1 ▷ Return 1 if any decrypted value is 0
    end if
  end for
  return 0 ▷ Return 0 if no ciphertext decrypts to 0
end procedure

More precisely, the sender encodes each bit $b_j \in \{0, 1\}$ by encrypting G^{b_j} and $\hat{G}^{b_j}$ under public keys PK_1 and PK_2, respectively, yielding ciphertexts $\boldsymbol{c}_j = \mathsf{CPA}(\mathsf{PK}_1, G^{b_j}) \in \mathbb{G}$ and $\hat{\boldsymbol{c}}_j = \mathsf{CPA}(\mathsf{PK}_2, \hat{G}^{b_j}) \in \hat{\mathbb{G}}$. The parties then jointly evaluate the expression $e(G, \hat{G})^{\sum_{j\in[l]} \gamma_j b_j (1-b_j)} \stackrel{?}{=} 1_{\mathbb{G}_\mathcal{T}}$, where $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_\mathcal{T}$ denotes a bilinear pairing, and each $\gamma_j \leftarrow\!\!\$\ \mathbb{Z}_p$ is a blinding factor. This check succeeds if and only if all bits are well-formed, i.e., $b_j \in \{0, 1\}$ for every $j \in [l]$. Crucially, the use of masking randomness ensures that the check does not reveal the index or value of any invalid bit, thereby preserving privacy even in the case of malformed inputs. In the MPC setting, the blinding factors γ_j are additively shared across parties: each party $\mathcal{T}_k$ for $k \in [T]$ locally samples $\gamma_{k,j} \leftarrow\!\!\$\ \mathbb{Z}_p$, and $\gamma_j = \sum_{k\in[T]} \gamma_{k,j}$. These values are generated *in parallel* across all parties.

We instantiate this technique in the MPC-HomoRand protocol, specified in Algorithms 2 and 3, using a concrete CPA encryption scheme. Let the $\mathsf{PK} = (G, f, g, h) \in \mathbb{G}^4$, with secret key $\mathsf{SK} = (\alpha, \beta) \in \mathbb{Z}_p^2$ and $f = g^\alpha h^\beta$. Encryption of a message $m \in \mathbb{Z}_p$ with randomness r proceeds as $\boldsymbol{c} = \mathsf{CPA}(\mathsf{PK}, m; r) = (c_0, c_1, c_2) = (G^m f^r, g^r, h^r)$, and decryption yields $\mathsf{Dec}(\mathsf{SK}, \boldsymbol{c}) = c_0 \cdot c_1^{-\alpha} \cdot c_2^{-\beta} = G^m$. In Algorithm 2, each party independently masks the ciphertexts in parallel and proves correctness via a ZKPoK $\{\boldsymbol{\pi}_{k,j}\}_{j\in[l]}$, which are broadcast alongside the resulting ciphertexts. Algorithm 3 specifies the subsequent verification phase:

Algorithm 2 MPC-HomoRand.Part1: Parallel Randomization by each $\mathcal{T}_k$ ($k \in [T]$)

procedure MPC-HomoRand.Part1($\mathsf{PK}_1, \mathsf{PK}_2, \boldsymbol{C}$) ▷ Input: $\boldsymbol{C} = \{(\boldsymbol{c}_j, \hat{\boldsymbol{c}}_j)\}_{j\in[l]}$
  for $j = 1$ to l **do** ▷ For each vote bit
    $\gamma_{k,j}, \lambda_{k,j}, r_{k,j}, s_{k,j} \leftarrow\$ \mathbb{Z}_p \setminus \{0\}$
    $\boldsymbol{c}''_{k,j} \leftarrow \boldsymbol{c}_j^{\gamma_{k,j}} \cdot \mathsf{CPA}(\mathsf{PK}_1, G^{\lambda_{k,j}}; r_{k,j})$ ▷ Apply random factors
    $\hat{\boldsymbol{c}}''_{k,j} \leftarrow (\mathsf{CPA}(\mathsf{PK}_2, \hat{G}; 0)/\hat{\boldsymbol{c}}_j)^{\lambda_{k,j}} \cdot \mathsf{CPA}(\mathsf{PK}_2, 1_{\hat{\mathbb{G}}}; s_{k,j})$ ▷ Apply random factors
    $\boldsymbol{\pi}_{k,j} \leftarrow \mathsf{PrfProve}(\mathsf{PK}_1, \mathsf{PK}_2, \boldsymbol{C}, \boldsymbol{c}''_{k,j}, \hat{\boldsymbol{c}}''_{k,j}; \gamma_{k,j}, \lambda_{k,j}, r_{k,j}, s_{k,j})$
    $\boldsymbol{C}''_{k,j} \leftarrow (\boldsymbol{c}''_{k,j}, \hat{\boldsymbol{c}}''_{k,j}, \boldsymbol{\pi}_{k,j})$
  end for
  return $\boldsymbol{C}''_k = \{\boldsymbol{C}''_{k,j}\}_{j\in[l]}$
end procedure

Algorithm 3 MPC-HomoRand.Part2: Multi-Party Verification

1: **procedure** MPC-HomoRand.Part2($\mathsf{PK}_1, \mathsf{PK}_2, \mathsf{SK}_1, \mathsf{SK}_2, \boldsymbol{C}, \{\boldsymbol{C}''_k\}_{k\in[T]}$)
2: **for** $j = 1$ to l **do** ▷ For each vote bit
3: **for** $k = 1$ to T **do**
4: **if** $\mathsf{PrfVer}(\mathsf{PK}, \boldsymbol{C}''_{k,j}, \boldsymbol{\pi}_{k,j}) = 0$ **then return** $\perp$
5: **end if**
6: **end for**
7: $\boldsymbol{c}''_j = (c''_{0j}, c''_{1j}, c''_{2j}) \leftarrow \prod_{k=1}^{T} \boldsymbol{c}''_{k,j}$ ▷ Aggregate the ciphertexts
8: $\hat{\boldsymbol{c}}''_j = (\hat{c}''_{0j}, \hat{c}''_{1j}, \hat{c}''_{2j}) \leftarrow \prod_{k=1}^{T} \hat{\boldsymbol{c}}''_{k,j}$
9: $X_j \leftarrow \mathsf{Dec}(\mathsf{SK}_1, \boldsymbol{c}''_j)$
10: $\bar{\boldsymbol{c}}'_j = (\bar{c}'_{0j}, \bar{c}'_{1j}, \bar{c}'_{2j}) \leftarrow \mathsf{CPA}(\mathsf{PK}_2, \hat{G}; 0)/\hat{\boldsymbol{c}}_j$
11: $x_j, y_j, z_j \leftarrow e(X_j, \bar{c}'_{0j})/e(G, \hat{c}''_{0j}), e(X_j, \bar{c}'_{1j})/e(G, \hat{c}''_{1j}), e(X_j, \bar{c}'_{2j})/e(G, \hat{c}''_{2j})$
12: **end for**
13: $(x, y, z) \leftarrow (\prod_{j=1}^{l} x_j, \prod_{j=1}^{l} y_j, \prod_{j=1}^{l} z_j)$
14: $Y \leftarrow \mathsf{Dec}(\mathsf{SK}_2, (x, y, z))$ ▷ Decrypt the aggregated result
15: **return** $Y = 1_{\mathbb{G}_\mathcal{T}}$? 1 : 0
16: **end procedure**

each proof is validated (line 4), and any party that fails verification is excluded. The combination of valid ciphertexts (lines 7–8) is computed solely from the outputs of honest parties, and the final values are jointly decrypted (line 14). The values (x_j, y_j, z_j) computed in line 11 are used to securely mask the term $b_j(1 - b_j)$ which would otherwise leak information if $b_j \notin \{0, 1\}$. Although the description employs a specific CPA encryption scheme for concreteness, our construction generalizes to any additive homomorphic CPA scheme. The full version, which includes complete proofs and additional technical details, is available in [10]. The core idea, securely verifying bit decomposition without leakage, remains applicable in broader cryptographic contexts.

4 Threshold Receipt-Free Voting

We adopt the voting system model of Doan et al. [9], which supports multiple ballot processing servers and ensures both threshold receipt-freeness and threshold correctness. We then review the first threshold receipt-free scheme from [9], whose limitations motivate our new constructions in Section 5.

Definition 4.1 (Voting System [9]). *A voting system with n ballot processing servers consists of algorithms:* (SetupElection, Vote, ProcessBallot, TraceBallot,

$\mathsf{Valid}, \mathsf{Append}, \mathsf{Publish}, \mathsf{VerifyVote}, \mathsf{Tally}, \mathsf{VerifyResult}$), *and a result function* $\rho_m : \mathbb{V}^m \cup \{\perp\} \to \mathbb{R}$*, where* $\mathbb{V}$ *is the vote domain and* $\mathbb{R}$ *is the result space.*

The election is initialized via SetupElection. Each voter runs Vote to generate a ballot consisting of n shares, one per ballot processing server. Each share is independently randomized by a distinct server using ProcessBallot. A tracking code is derived via TraceBallot, enabling voters to later trace their ballots on the public bulletin board PBB. Validated shares are collected and appended to the ballot box BB using Append, then made publicly available on PBB via Publish. Voters verify correct recording using VerifyVote and their tracking codes. The tally is computed over valid ballots, checked by Valid, using Tally, and its correctness is publicly verifiable via VerifyResult.

Threshold Receipt-Freeness ($\mathsf{t_{rf}}$). The definition of threshold RF proceeds in two parts. The first ensures that no adversary–despite corrupting up to $\mathsf{t_{rf}}$ out of n ballot processing servers–can coerce a voter into producing a receipt that convincingly proves how they voted. This guarantee holds even if the adversary learns the voter's randomness or attempts to bias the ballot construction. The second part, extends the indistinguishability-based receipt-freeness notion of Devillez et al. [6] to the threshold setting. It formalizes the requirement that even if a voter colludes with up to $\mathsf{t_{rf}}$ servers and deviates arbitrarily from the honest ballot distribution, a third party, given full knowledge of how the ballot was constructed and access to the public bulletin board, should not be able to determine whether the voter submitted that ballot or a different one encoding another vote.

Threshold Correctness ($\mathsf{t_{corr}}$). Threshold correctness ensures that the choices of honest voters are faithfully reflected in the final tally, even in the presence of limited adversarial control over the ballot processing servers. Concretely, it guarantees that as long as at least $n - \mathsf{t_{corr}}$ processed shares of each honestly generated ballot appear on the public bulletin board–regardless of whether they were handled by malicious servers–the election outcome remains uniquely determined. That is, the adversary should not be able to construct two sets of valid-looking ballot boxes that lead to different outcomes.

The Doan et al.'s Voting Scheme. The first protocol to simultaneously achieve threshold receipt-freeness and correctness was proposed by Doan et al. [9] (E-Vote-ID 2024). Their construction, illustrated in Figure 1, introduces a threshold receipt-free voting system that significantly reduces trust assumptions on ballot randomizers.

In their scheme, each voter secret-shares their vote, encrypts each share using a TREnc (Section 2), and submits the resulting ballot shares $\{\mathsf{b}_i\}_{i\in[n]}$ to a set of independent randomizers. Each randomizer then rerandomizes one ballot share and output b'_i made available on the public board BB. After the voting phase, the talliers (or any observer) extract standard CPA components from the TREnc outputs, and use Lagrange interpolation to reconstruct an encryption of the original vote. This process results in a final ciphertext $\boldsymbol{C}$ that is then submitted into a mixnet-based tallying process. Receipt-freeness is preserved as long as at least one honest rerandomized share b'_i is used in the reconstruction of $\boldsymbol{C}$.

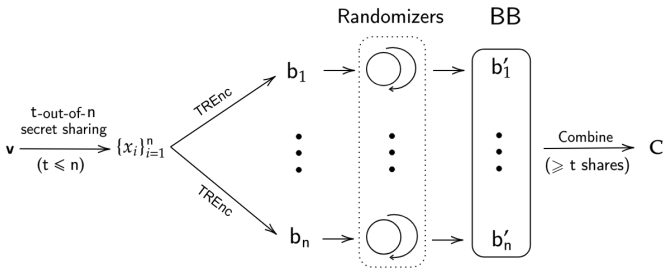

Fig. 1. Doan et al.'s voting scheme.

Limitations. The protocol does not ensure that $\boldsymbol{C}$ is a valid vote encryption, meaning a malicious $\boldsymbol{C}$ could encode an invalid vote, that, upon decryption, leaks unintended information and potentially violates violating RF. While TREnc provides strong privacy guarantees, it offers limited support when it comes to verifying election-specific constraints on the encrypted vote, which typically require non-linear proofs. For instance, proving a vote v is a bit requires a quadratic relation, $\mathsf{v}(1-\mathsf{v})=0$, requiring a non-linear proof (see, e.g., [8]).

Although NIZK proofs can, in principle, support such constraints, integrating them into this setting with the secret sharing and rerandomization is non-trivial. In particular, a non-linear proof constructed before randomization would not survive the transformation, and constructing such a proof after randomization is infeasible for the voter, who lacks control over the ciphertext randomness. In a multi-randomizer setting, where multiple servers independently rerandomize different shares of a ballot, adapting a non-linear proof locally is also challenging. The transformation applied by each server generally depends on the randomness chosen by others, rendering any isolated adaptation incorrect or unverifiable. As a result, constructing such proofs typically requires interaction between the voter and randomizing parties, which could introduce complexity and security risks, as malicious servers could collude with voters. To address these issues, we introduce a technique in the next section to enforce non-linear constraints on encrypted votes without requiring any voter-randomizer coordination.

5 Voting Schemes

We propose a new approach that avoids requiring voters to construct complex zero-knowledge proofs at ballot submission time. Instead, we shift the validity checking to a dedicated *pre-tally phase*, which takes place after the final combined ciphertexts $\boldsymbol{C}$ have been reconstructed (see Figure 1), but before tallying begins.

At a high level, our scheme preserves the overall voting flow of Doan et al. from the voter's perspective. The key difference lies in how the reconstructed ciphertexts $\boldsymbol{C}$ are handled. Rather than passing them directly to a mixnet, they are first processed in a multi-party computation (MPC) protocol that verifies whether each ciphertext encodes a valid vote. This verification is performed collaboratively by the talliers or those who hold shares of the decryption key without revealing the vote or any auxiliary information. Only ciphertexts that pass this validation step are forwarded to the standard tallying process.

In the following, we present two concrete instantiations of this paradigm: MiniMix and HomoRand. These schemes differ in the inputs provided by voters to the pre-tally phase and the way the underlying MPC protocol operates. Before delving into the designs of MiniMix (Section 5.2) and HomoRand (Section 5.3), we first outline their shared structure in Section 5.1, following the general voting system definition in Definition 4.1. We then give the correctness and the security theorem statements of the constructions in Section 6.

5.1 Overview

Key Generation. The election authority EA initiates the setup by running the SetupElection algorithm, generating a public/secret key pair $(\mathsf{PK}, \mathsf{SK})$. The public key PK is posted on a public bulletin board PBB, while the private key SK is shared among T talliers using a threshold decryption scheme. Key generation can be distributed securely in prime-order groups through standard techniques.

Voting Phase. To cast a vote v, a voter executes the Vote algorithm. In principle, this procedure begins by secret-sharing v into n shares using a t-out-of-n threshold secret sharing scheme (Section 2). Each share is then independently encrypted under a TREnc (Section 2), producing ballot shares $\{\mathsf{b}_i\}_{i\in[n]}$. The complete ballot b consists of this set of encrypted shares.

Each b_i is then submitted to a distinct randomizing server, while the corresponding trace $\tau = \mathsf{TraceBallot}(\mathsf{b})$ is simultaneously transmitted to PBB. Upon receiving a b_i, a server performs a local verification using the validity checks defined by TREnc, optionally including a zero-knowledge proof verification. It further ensures that no duplicate traces are present with respect to all traces on PBB. Valid ballot shares are re-randomized using $\mathsf{ProcessBallot}(\mathsf{b}_i)$ and made available on the bulletin board. Voters can later verify inclusion of their vote by querying VerifyVote on the trace τ and confirming its appearance on PBB.

Tallying Phase. After voting concludes, the tallying phase begins. Talliers aggregate ballot shares by comparing the traces on PBB with the complete traces τ posted by each voter. The Valid function checks whether at least t valid shares are present for each ballot. If the condition is met, the CPA components are extracted from the corresponding shares and combined using Lagrange interpolation. Owing to the linearity of the interpolation, this process can be performed directly in the encrypted domain, yielding a final ciphertext that encrypts the original vote. Ballots failing the Valid check are discarded.

As noted by Doan et al. [9], this reconstruction step is critical for mitigating adversarial influence. A malicious voter may attempt to to deviate from the protocol by submitting malformed or inconsistent shares. However, since all valid ballot shares on the bulletin board (possibly more than t) are combined, only one message can be reconstructed. As the adversary cannot anticipate which subset of shares will be successfully posted (e.g., due to the random failure of non-operational randomizers), they are incentivized to submit a consistent and correct ballot to ensure that the tally reflects their intended vote.

Finally, the reconstructed ciphertexts corresponding to valid ballots are submitted to the PreTally function, which checks that the encrypted vote is within

the intended range. Depending on its inputs, the PreTally function behaves differently. In the MiniMix construction (Figure 2), which utilizes the MPC-MiniMix protocol from Algorithm 1, the voter submits only the encrypted vote. In contrast, the HomoRand construction uses the MPC-HomoRand protocol from Algorithms 2 and 3, where the voter also sends additional values to assist the participating parties in verifying the validity of the encrypted vote. Depending on the deployment scenario, the participating parties may include the talliers or, alternatively, external authorities. For simplicity, we assume that the talliers are responsible for performing this validation in the current setting. Invalid ballots are discarded based on the outcome of the PreTally check. The talliers then run the Tally function to compute the election result r based on the valid ones according to the election rules. A proof of correctness Π is generated and can be verified by anyone using the VerifyResult algorithm. In the instantiations below, the number of servers n and the threshold t are implicit inputs for all algorithms.

Relationship between correctness and receipt-freeness thresholds. Our constructions rely on a secret sharing scheme where at least t valid ballot shares are needed to reconstruct a vote. This implies that the system can tolerate up to $n-t$ missing or invalid shares without affecting correctness. On the privacy side, receipt-freeness holds as long as the reconstruction includes at least one honestly rerandomized share. So even if up to $t-1$ ballot shares are maliciously processed, it is infeasible for the adversary to compromise voter privacy. In other words, the system remains secure against up to $t-1$ compromised processing servers, which matches the tight bound shown in prior work by Doan et al. [9].

5.2 The MiniMix Construction

In this scheme, the input to the PreTally is a ballot $\mathsf{b} = \{\mathsf{b}_i\}_{i\in[n]}$ randomized by randomizers, where each b_i is a TREnc's encryption of a share x_i of the voter's intended message v.

To initiate the PreTally procedure, the talliers (and optionally the public) verify each ballot share b_i by applying the TREnc.Ver function (see Section 2). A ballot b is deemed valid by $\mathsf{Valid}(\mathsf{BB}, \mathsf{b})$ (see Figure 2) if at least t valid shares are posted on the public bulletin board. Upon successful validation, the combination algorithm Combine is publicly executed by any of the talliers. It takes as input the homomorphic components of all available valid ballot shares–up to n shares–and outputs a combined ciphertext $\boldsymbol{C}_0$, representing the encrypted vote v. Due to the properties of Lagrange interpolation, reconstructing with more than t valid shares preserves the correctness of the resulting ciphertext.

Subsequently, T talliers jointly executes the MPC-MiniMix as described in Algorithm 1 with the input $\boldsymbol{C} = \{\boldsymbol{C}_j\}_{j\in[d]}$, where $\boldsymbol{C}_j = \boldsymbol{C}_0 \cdot \mathsf{CPA}(\mathsf{PK}, -v_j)$. In particular, each tallier $\mathcal{T}_k$ for $k \in [T]$ shuffles the d ciphertexts in $\boldsymbol{C}$, applying an independent random factor. It also produces a publicly verifiable ZKPoK proof, attesting to correctness. Misbehaving parties are identified through proof verification. After the last tallier has completed their respective round, all the talliers jointly decrypt the final output $\boldsymbol{C}^{(T)} = \{\boldsymbol{C}_j^{(T)}\}_{j\in[d]}$. As $\boldsymbol{C}^{(T)}$ is encrypted under the TREnc's PK, decryption proceeds using $\mathsf{TREnc.Dec}(\mathsf{SK}, \cdot)$ (Algorithm

SetupElection (λ)

$(\mathsf{SK}, \mathsf{PK}) \leftarrow \mathsf{TREnc.Gen}(1^\lambda)$
return PK

Vote(id, v[, aux])

$\{x_1, \ldots, x_n\} \leftarrow \mathsf{Share}(n, t, \mathsf{v})$
if aux is empty **for** $i = 1$ to n do
  $\mathsf{lk}_i \leftarrow\!\!\$\ \mathsf{TREnc.LGen}(\mathsf{PK})$
else $\{\mathsf{lk}_i\}_{i=1}^n \leftarrow$ aux
for $i = 1$ to n do
  $\mathsf{b}_i \leftarrow \mathsf{TREnc.LEnc}(\mathsf{PK}, \mathsf{lk}_i, x_i)$
return $\mathsf{b} = \{\mathsf{b}_i\}_{i=1}^n$

Valid(BB, b)

if $\exists \mathsf{b}' \in \mathsf{BB} \wedge \exists \tau_i' \subset \mathsf{TraceBallot}(\mathsf{b}')$:
$\tau_i' \subset \mathsf{TraceBallot}(\mathsf{b})$ **then return** $\perp$
$k = 0$
for $i = 1$ to $|\mathsf{b}|$ do
  if $\mathsf{TREnc.Ver}(\mathsf{PK}, \mathsf{b}_i) = 1$
  then $k \leftarrow k + 1$
if $k \geq t$ **then return** 1 **else return** 0

ProcessBallot(b_i)

return $\mathsf{TREnc.Rand}(\mathsf{PK}, \mathsf{b}_i)$

TraceBallot(b)

$\tau_i \leftarrow \mathsf{TREnc.Trace}(\mathsf{b}_i)$
return $\tau = \{\tau_i\}_{i=1}^n$

PreTally (BB, SK, b)

if $\mathsf{Valid}(\mathsf{BB}, \mathsf{b}) = 0$ **then return** 0
for $i = 1$ to $|\mathsf{b}|$ do
  $\boldsymbol{c}_i \leftarrow \mathsf{TREnc.Strip}(\mathsf{PK}, \mathsf{b}_i)$
$\boldsymbol{C}_0 \leftarrow \mathsf{Combine}(n, t, \{\boldsymbol{c}_i\}_{i=1}^{\leq n})$
for $j = 1$ to d do
  $\boldsymbol{C}_j \leftarrow \boldsymbol{C}_0 \cdot \mathsf{CPA}(\mathsf{PK}, -v_j)$
return $\mathrm{MPC\text{-}MiniMix}(\mathsf{PK}, \mathsf{SK}, \boldsymbol{C} = \{\boldsymbol{C}_j\}_{j=1}^d)$

VerifyVote(PBB, τ)

if $\exists \mathsf{b} \in \mathsf{PBB} : \mathsf{Valid}(\mathsf{b}) \wedge \tau == \mathsf{TraceBallot}(\mathsf{b})$
then return 1 **else return** 0

Fig. 2. MiniMix instantiation of our voting scheme.

1, line 20). The PreTally procedure outputs 1 as soon as one of the decrypted ciphertexts equals 0, confirming that the original vote v belongs to the designated valid range; otherwise, it outputs 0. Invalid votes are discarded, and the Tally function is applied to the combined ciphertext $\boldsymbol{C}_0$ of each valid ballot to compute the final outcome according to the election rules. A formal correctness proof of MPC-MiniMix is provided in the full version [10].

Discussion. It is worth noting that the MiniMix construction can be adapted to minimize online-phase computation. In the presented variant, the shuffle is performed after a ballot is available on the PBB, which may lead to inefficiencies such as increased latency and synchronization delays.

As an alternative, the authorities may jointly shuffle the encrypted vote domain $\{\boldsymbol{V}_j = \mathsf{CPA}(v_j)\}_{j\in[d]}$ under a secret permutation π, yielding a randomized sequence $\{\boldsymbol{V}'_{\pi(j)}\}_{j\in[d]}$. Given a combined ciphertext $\boldsymbol{C}_0$ (as defined earlier), they compute $\boldsymbol{C}'_j = \boldsymbol{C}_0/\boldsymbol{V}'_{\pi(j)}$ for each $j \in [d]$, resulting in ciphertexts encrypting $(\mathsf{v} - v_{\pi(j)})$. These can be tested for equality to zero to verify whether v belongs to the valid domain. As π remains hidden, the test leaks no information about the actual vote. To avoid linkability, each ballot must be verified against a freshly shuffled domain encryption. Otherwise, repeated shuffles could reveal identical vote positions, enabling correlation across ballots. Hence, the number of domain shuffles must scale with the number of ballots or eligible voters.

5.3 The HomoRand Construction

Due to space constraints, the full specification of HomoRand is provided in [10] (Appendix A); here we give a high-level overview. The input to PreTally is a randomized ballot $\mathsf{b} = \{\mathsf{b}_i\}_{i\in[n]}$, where each b_i now consists of l components b_{ji} for $j \in [l]$. To this end, the vote v is first decomposed into an l-bit string $\{b_j\}_{j\in[l]}$,

and each bit b_j is shared using t-out-of-n secret sharing scheme to obtain n share $\{b_{ji}\}_{i\in[n]}$. Each share b_{ji} is then encrypted separately under two TREnc public keys, yielding a pair of ciphertexts in $\mathbb{G}$ and $\hat{\mathbb{G}}$. A non-interactive randomizable proof (e.g., Groth-Sahai proofs [12]) accompanies each pair, proving consistency across the two encryptions, i.e., demonstrating that they indeed encrypt the same value. As a result, each ballot share contains l components, where each includes two TREnc ciphertexts and a consistency proof. Each ballot share is re-randomized, and the resulting share is posted to the public board.

The $\mathsf{PreTally}$ protocol begins by validating each ballot b posted on the public bulletin board. This involves applying $\mathsf{TREnc.Ver}$ to the ciphertexts and invokes $\mathsf{PrfVerify}$ on the associated proof. A ballot is deemed valid if, for every bit index $j \in [l]$, there exist at least t valid shares $\{\mathsf{b}_{ji}\}_i$ posted on the board. Once a ballot passes verification, the associated proofs are discarded. Next, the $\mathsf{Combine}$ algorithm aggregates the homomorphic components of the accepted ballot shares (up to n) to reconstruct two CPA encryptions of $\{b_j\}_{j\in[l]}$ in $\mathbb{G}$ and $\hat{\mathbb{G}}$. The protocol then invokes the MPC-$\mathsf{HomoRand}$ procedure from Algorithms 2 and 3 to jointly verify that each b_j is a bit.

5.4 Efficiency Discussion

Table 1 compares the computational costs incurred during the casting of a single ballot under the two proposed constructions. All costs reported are per voter, per randomizer, or per tallier, as appropriate.

Computational Costs. On the voter's side, the Vote algorithm in $\mathsf{MiniMix}$ requires n invocations of $\mathsf{TREnc.Enc}$, as each ballot share is encrypted independently. In contrast, $\mathsf{HomoRand}$ imposes a higher load: $2ln$ encryptions (two per bit per share) and $16ln$ exponentiations for Groth-Sahai consistency proofs across all $i \in [n]$ and $j \in [l]$ (see [10] for details on proof computation). Each randomizer, executing $\mathsf{ProcessBallot}$, performs 1 TREnc rerandomization in $\mathsf{MiniMix}$, while in $\mathsf{HomoRand}$, this increases to $2l$, alongside rerandomizing the associated proofs. Talliers executing $\mathsf{PreTally}$ incur costs primarily from the underlying MPC protocols. In MPC-$\mathsf{MiniMix}$ (Algorithm 1), each tallier (i) shuffles d ciphertexts with verifiable proofs at a cost denoted $\mathsf{shuffle}(d)$ (lines 4–9); (ii) applies random factors to shuffled ciphertexts with proofs (lines 11–13), amounting to d rand operations; and (iii) participates in the joint decryption of d ciphertexts (line 19–24), contributing d dec operations. In MPC-$\mathsf{HomoRand}$, each tallier performs $2l$ rand operations (Algorithm 2, lines 2–8) and engages in $l{+}1$ dec operations (Algorithm 3, lines 9 and 14).

Verification Costs. Anyone, including external auditors or voters, can verify the $\mathsf{PreTally}$ result by checking the posted proofs. In $\mathsf{HomoRand}$, the verification requires $2l \cdot \mathsf{rand} + (l+1) \cdot \mathsf{dec}$ while in $\mathsf{MiniMix}$ it involves $d \cdot (\mathsf{rand} + \mathsf{dec})$, plus $\mathsf{shuffle}(d)$ proofs from each of the T talliers. Since $l = \log_2(d)$, the relative verification cost ratio scales as roughly $d/(2\log_2 d)$ in favor of $\mathsf{HomoRand}$ for large domains d, though $\mathsf{MiniMix}$ incurs additional overhead from shuffle proofs, which scale with both d and T.

	SetupElection	Vote	ProcessBallot	PreTally
MiniMix	$1 \cdot \mathsf{TREnc.Gen} \in \mathbb{G}$	$n \cdot \mathsf{TREnc.Enc}$	$1 \cdot \mathsf{TREnc.Rand}$	$\mathsf{shuffle}(d) + d \cdot \mathsf{rand} + d \cdot \mathsf{dec}$
HomoRand	$1 \cdot \mathsf{TREnc.Gen} \in \mathbb{G}$ $1 \cdot \mathsf{TREnc.Gen} \in \hat{\mathbb{G}}$	$2ln \cdot \mathsf{TREnc.Enc}$ $\mathbb{G}^{16ln} \times \hat{\mathbb{G}}^{16ln}$	$2l \cdot \mathsf{TREnc.Rand}$ $\mathbb{G}^{14l} \times \hat{\mathbb{G}}^{14l}$	$2l \cdot \mathsf{rand} + (l+1) \cdot \mathsf{dec}$

Table 1. Computational cost per ballot under each construction.

While HomoRand achieves PreTally costs that scale with $l = \log_2(d)$, asymptotically outperforming the linear-in-d cost of MiniMix, this efficiency comes at the price of substantially higher computational overhead for voters and randomizers. Specifically, the cost of Vote in HomoRand scales with both n and l, due to bitwise encryption and the generation of Groth-Sahai proofs. In contrast, MiniMix imposes minimal and domain-independent client-side costs, making it an attractive choice in settings where lightweight voting is essential and the number of valid vote options is modest. Conversely, HomoRand may be preferable in settings where the verification of PreTally is a critical concern–such as public audits or large-scale elections–particularly when voter-side computation is not a bottleneck and the domain size d is large enough to outweigh its setup and voting costs.

6 Security of the Voting Schemes

The proposed voting schemes achieve both *threshold receipt-freeness* and *threshold correctness* (Section 4). Due to space constraints, we provide proof sketches here. Formal proofs appear in the full version [10].

Theorem 6.1 (Threshold Receipt-Freeness). *Let* TREnc *be verifiable and* TCCA*-secure, and let the proof systems employed for the pre-tally and for verifying tally correctness be zero-knowledge. Then both constructions achieve threshold receipt-freeness under a t-out-of-n sharing scheme with threshold $\mathsf{t_{rf}} = t-1$. More precisely, for the* MiniMix *scheme,* $\Pr[\mathsf{Exp}^{\mathsf{deceive}}_{\mathcal{A},\mathcal{V},\mathsf{t_{rf}}}(\lambda) = 1] \leq \varepsilon_{\mathrm{verif}}$ *and* $\mathsf{Adv}^{\mathsf{rf},\mathsf{t_{rf}},\beta}_{\mathcal{A},\mathcal{V}}(1^\lambda) \leq \varepsilon_{\mathrm{ZK}} + q(n - \mathsf{t_{rf}})\varepsilon_{\mathrm{tcca}}$*. For the* HomoRand *scheme,* $\Pr[\mathsf{Exp}^{\mathsf{deceive}}_{\mathcal{A},\mathcal{V},\mathsf{t_{rf}}}(\lambda) = 1] \leq \varepsilon_{\mathrm{verif}}$ *and* $\mathsf{Adv}^{\mathsf{rf},\mathsf{t_{rf}},\beta}_{\mathcal{A},\mathcal{V}}(1^\lambda) \leq \varepsilon_{\mathrm{ZK}} + lq(n - \mathsf{t_{rf}})(2\varepsilon_{\mathrm{tcca}} + \varepsilon_{\mathsf{sxdh}})$*. Here, l is the bit-length of the vote domain, and $\varepsilon_{\mathrm{ZK}}$, $\varepsilon_{\mathrm{sxdh}}$, $\varepsilon_{\mathrm{verif}}$, $\varepsilon_{\mathrm{tcca}}$ bound the adversarial advantage against the ZK proof systems, the SXDH (Symmetric eXternal Diffie-Hellman) assumption, verifiability, and* TCCA*-security of* TREnc*, respectively; q is the number of ballot-append queries.*

Proof. We sketch the proof for MiniMix; the case of HomoRand is analogous. Threshold receipt-freeness is defined via two experiments.

In the first, $\mathsf{Exp}^{\mathsf{deceive}}_{\mathcal{A},\mathcal{V},\mathsf{t_{rf}}}(\lambda)$ [9], security follows from TREnc's TCCA security, verifiability, and the correctness of the secret sharing scheme, as shown in [9].

In the second, $\mathsf{Exp}^{\mathsf{rf},\mathsf{t_{rf}},\beta}_{\mathcal{A},\mathcal{V}}(\lambda)$, the adversary must provide two valid ballots with matching traces. We define a sequence of hybrid games starting from $\beta = 0$ (honest setup) and ending at $\beta = 1$ (simulated setup). Let $\mathcal{A}$ query ballot pairs $\mathsf{B}_0, \mathsf{B}_1$ with identical traces to BB_0 and BB_1 respectively. We progressively replace the processed ballot shares of B_0 with those of B_1 in BB_0, relying on the TCCA security of TREnc to preserve indistinguishability at each step, incurring at most $\varepsilon_{\mathrm{tcca}}$ advantage per swap. This may introduce inconsistencies in the tally since this changes the underlying plaintext. However, since the vote intent of B_0 is

known from TREnc.Dec's correctness, the zero-knowledge simulator can emulate the decryption step (Algorithm 1, line 20) to match the expected result of pre-tally phase. Thus, $\mathcal{A}$ cannot distinguish whether B_0 or B_1 was processed, even when B_0 encodes an invalid vote, except with an error bounded by $\varepsilon_{\mathrm{ZK}}$. A hybrid argument over q queries yields $\mathsf{Adv}^{\mathsf{rf},\mathsf{t}_{\mathsf{rf}},\beta}_{\mathcal{A},\mathcal{V}}(1^\lambda) \leq \varepsilon_{\mathrm{ZK}} + q(n - \mathsf{t}_{\mathsf{rf}})\varepsilon_{\mathrm{tcca}}$.

Theorem 6.2 (Threshold Correctness). *Let* TREnc *be traceable and verifiable, and assume that the underlying NIZK proof systems are correct. Then both constructions achieve threshold correctness with* $\mathsf{t}_{\mathsf{corr}} = n - t$ *under a t-out-of-n secret sharing scheme. More precisely, for any efficient adversary* $\mathcal{A}$ *making* q *ballot-append queries,* $\Pr[\mathsf{Exp}^{\mathsf{corr},\mathsf{t}_{\mathsf{corr}}}_{\mathcal{A},\mathcal{V}}(\lambda) = 1] \leq qn\varepsilon_{\mathrm{trace}} + \varepsilon_{\mathrm{corr}}$ *for* MiniMix *and* $\Pr[\mathsf{Exp}^{\mathsf{corr},\mathsf{t}_{\mathsf{corr}}}_{\mathcal{A},\mathcal{V}}(\lambda) = 1] \leq 2qln\varepsilon_{\mathrm{trace}} + \varepsilon_{\mathrm{corr}}$ *for* HomoRand*, where* l *is the vote bit-length, and* $\varepsilon_{\mathrm{trace}}$*,* $\varepsilon_{\mathrm{corr}}$ *bound* $\mathcal{A}$*'s advantage against the traceability of* TREnc *and the correctness error of the underlying MPC protocol, respectively.*

Proof. In the threshold correctness experiment, denoted $\mathsf{Exp}^{\mathsf{corr},\mathsf{t}_{\mathsf{corr}}}_{\mathcal{A},\mathcal{V}}(\lambda)$ [9], $\mathcal{A}$ submits q pairs of valid ballots to two election views, each omitting at most $\mathsf{t}_{\mathsf{corr}}$ ballot shares. As all ballots derive from honestly generated ones encoded as TREnc ciphertexts, traceability ensures $\mathcal{A}$ cannot alter the vote intent without changing trace values, except with negligible probability. This yields a bound of $qn\varepsilon_{\mathrm{trace}}$ for MiniMix, and $2qln\varepsilon_{\mathrm{trace}}$ for HomoRand. Once ballots are posted to the PBB, the correctness of the secret sharing scheme guarantees that the inputs to the pre-tally phase correctly reflect the original vote intent. The correctness of the underlying NIZK proofs in MPC-MiniMix and MPC-HomoRand then ensures that both views yield identical outputs for valid votes.

Verifiability. Our schemes ensure both *individual* and *universal verifiability*. Since the voter's voting process in MiniMix follows the protocol of [9], it inherits individual verifiability via the traceability of TREnc. In HomoRand, voters additionally provide consistency proofs across the two TREnc ciphertexts in each ballot share, without affecting traceability. Hence, no adversary can alter the vote intent without detection. For universal verifiability, both constructions employ publicly verifiable ZKPoKs in pre-tally and tally phases. The soundness of these proofs ensures that all accepted ballots encode valid votes and that decryption is performed correctly, thereby guaranteeing a trustworthy tally outcome.

7 Conclusion

We propose a new direction for validating encrypted ballots in threshold receipt-free voting systems, where ballots are non-interactively rerandomized by multiple independent ballot processing servers. Our approach introduces a pre-tally validation phase executed in a multiparty computation style, allowing authorities to verify the validity of encrypted votes without decrypting them or requiring voter-supplied validity proofs.

We develop two constructions that achieve threshold receipt-freeness and verifiability while addressing privacy risks present in prior mixnet-based systems. Both schemes support homomorphic or mixnet-based tallying, depending on vote range constraints, and crucially reveal only the validity status of a ballot, nothing more. To our knowledge, this is the first threshold receipt-free solution to achieve better privacy independently of the tallying technique.

Our efficiency analysis recommends using MiniMix when voter-side efficiency is critical and the number of valid vote options is small or moderate, while HomoRand is better suited for large or complex vote domains. In addition to our core designs, we also discuss how existing validity-check techniques can be adapted to fit our pre-tally framework. An open direction is to reduce the computational and communication complexity of the pre-tally process, or to devise alternative mechanisms such as randomizable validity proofs that can be verified in a single-pass setting, even in the presence of fully malicious randomizers.

Acknowledgments

Thomas Peters is a research associate of the Belgian Fund for Scientific Research (F.R.S.-FNRS). This work has been funded in part by the Walloon Region through the project CyberExcellence (convention number 2110186).

References

1. Bayer, S., Groth, J.: Efficient zero-knowledge argument for correctness of a shuffle. In: Advances in Cryptology–EUROCRYPT 2012. pp. 263–280. Springer (2012)
2. Benaloh, J., Tuinstra, D.: Receipt-free secret ballot elections. In: Proc. 26th ACM Symp. on Theory of Computing (STOC). pp. 544–553. ACM (1994)
3. Blazy, O., Fuchsbauer, G., Pointcheval, D., Vergnaud, D.: Signatures on randomizable ciphertexts. In: Public Key Cryptography–PKC. pp. 403–422. Springer (2011)
4. Chaidos, P., Cortier, V., Fuchsbauer, G., Galindo, D.: BeleniosRF: A non-interactive receipt-free electronic voting scheme. In: Proc. ACM CCS. pp. 1614–1625. ACM (2016)
5. Cramer, R., Gennaro, R., Schoenmakers, B.: A secure and optimally efficient multi-authority election scheme. In: EUROCRYPT '97. pp. 103–118. Springer (1997)
6. Devillez, H., Pereira, O., Peters, T.: Traceable receipt-free encryption. In: Proc. ASIACRYPT 2022. pp. 273–303. Springer (2022)
7. Devillez, H., Pereira, O., Peters, T., Yang, Q.: Can we cast a ballot as intended and be receipt free? In: Proc. IEEE S&P 2024. pp. 3440–3457. IEEE (2024)
8. Doan, T.V.T., Pereira, O., Peters, T.: Encryption mechanisms for receipt-free and perfectly private verifiable elections. In: Proc. ACNS. pp. 257–287. Springer (2024)
9. Doan, T.V.T., Pereira, O., Peters, T.: Threshold receipt-free single-pass evoting. In: Proc. E-Vote-ID 2024. pp. 20–36. Springer (2024)
10. Doan, T.V.T., Pereira, O., Peters, T.: Threshold receipt-free voting with server-side vote validation. Cryptology ePrint Archive, Paper 2025/1321 (2025), `https://eprint.iacr.org/2025/1321`
11. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for np. In: Advances in Cryptology-EUROCRYPT 2006. pp. 339–358. Springer (2006)
12. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Proc. EUROCRYPT 2008. pp. 415–432. Springer (2008)
13. Hirt, M.: Multi party computation: Efficient protocols, general adversaries, and voting (2001)
14. McMurtry, E., Pereira, O., Teague, V.: When is a test not a proof? In: Proc. ESORICS 2020. pp. 23–41. Springer (2020)
15. Ozdemir, A., Boneh, D.: Experimenting with collaborative zk-snarks: Zero-knowledge proofs for distributed secrets. In: Proc. USENIX Security 2022. pp. 4291–4308. USENIX Association (2022)
16. Shamir, A.: How to share a secret. Communications of the ACM pp. 612–613 (1979)
17. Terelius, B., Wikström, D.: Proofs of restricted shuffles. In: Africacrypt. pp. 100–113. Springer (2010)