

Practical Mempool Privacy via One-time Setup Batched Threshold Encryption

Arka Rai Choudhuri^{*}
Nexus

Sanjam Garg[†]
UC Berkeley

Guru-Vamsi Policharla[‡]
UC Berkeley

Mingyuan Wang[§]
NYU Shanghai

Abstract

An important consideration with the growth of the DeFi ecosystem is the protection of clients who submit transactions to the system. As it currently stands, the public visibility of these transactions in the memory pool (mempool) makes them susceptible to market manipulations such as frontrunning and backrunning. More broadly, for various reasons—ranging from avoiding market manipulation to including time-sensitive information in their transactions—clients may want the contents of their transactions to remain private until they are executed, i.e. they have *pending transaction privacy*. Therefore, *mempool privacy* is becoming an increasingly important feature as DeFi applications continue to spread.

We construct the first *practical* mempool privacy scheme that uses a *one-time* DKG setup for n decryption servers. Our scheme ensures the strong privacy requirement by not only hiding the transactions until they are decrypted but also guaranteeing privacy for transactions that were not selected in the epoch (*pending transaction privacy*). For each epoch (or block), clients can encrypt their transactions so that, once B (encrypted) transactions are selected for the epoch, they can be decrypted by each decryption server while communicating only $O(1)$ information.

Our result improves upon the best-known prior works, which either: (i) require an expensive initial setup involving a (special purpose) multiparty computation protocol executed by the n decryption servers, along with an additional *per-epoch* setup; (ii) require each decryption server to communicate $O(B)$ information; or (iii) do not guarantee pending transaction privacy.

We implement our scheme and find that transactions can be encrypted in approximately 8.5 ms, independent of committee

size, and the communication required to decrypt an entire batch of transactions is 48 bytes per party, independent of the number of transactions. If deployed on Ethereum, which processes close to 500 transactions per block, it takes close to 3.2 s for each committee member to compute a partial decryption and 3.0 s to decrypt all transactions for a block in single-threaded mode. Compared to prior work, which had an expensive setup phase per epoch, we incur $< 2\times$ overhead in the worst case. On some metrics such as partial decryption size, we actually fare better.

1 Introduction

Decentralized Finance or DeFi is a financial ecosystem that leverages blockchains to provide various financial services. The rapid growth of DeFi is best illustrated by the amount of capital that has made its way into the ecosystem, with a peak value of over \$250 billion in 2021. As the ecosystem evolves and grows, it is important that the system implements protections for clients of this system to ensure they are not taken advantage of. Unfortunately, as we illustrate below, there are several instances where such systems fall short.

Consider a scenario of a client interacting with the system and simply wants to submit a transaction. For this transaction to be considered complete by the system, it must be “executed” by the blockchain. To facilitate this, clients submit their transaction to a common *memory pool*, commonly known as *mempool*. It is then the job of the miners (or builders) not only to select from the mempool the transactions that get executed, but also the order in which the transactions are executed. In an ideal world, one would like every transaction to be treated “fairly”, i.e. the details of the transaction would not influence whether it is selected or the order in which it is executed. Ideally, miners/builders should be “blind” to the transaction details when processing them.

Unfortunately, the large financial incentives in these systems make reality far from the desired ideal world. Miners exploit the information asymmetry available to them to frontrun, backrun, or sandwich client transactions to extract addi-

^{*}arkarai.choudhuri@gmail.com. Part of this work was done while the author was a postdoc at NTT Research.

[†]sanjamg@berkeley.edu

[‡]guruvamsip@berkeley.edu

[§]mingyuan.wang@nyu.edu. Part of this work was done while the author was a postdoc at UC Berkeley.

tional value from them - termed MEV or *Miner Extractable Value* [22]. While the miners benefit financially from this asymmetry of information, it is at the expense of the clients. This has been well documented both from a miner perspective in ways to benefit the most from this system (perhaps rationally) [2, 29, 65, 68, 70], and also about the broad problems caused by the ability of the miners to exploit such a system [16, 38, 43, 58, 63]. In the most extreme case, miners can effectively censor transactions from certain clients, an undesirable outcome for the system.

Beyond simple transactions, clients may also submit updates to their existing smart contracts within their “transaction.” Such updates may be crucial if, for instance, a client discovers a bug in their smart contract that causes financial losses. In such a case, they would want to update the smart contract before others in the system become aware of the issue. However, such an update might expose the bug, allowing miners to exploit it before the transaction containing the update is executed. These vulnerabilities have led to active efforts [7, 47, 52, 57, 59, 62] aimed at addressing privacy concerns surrounding transactions in the mempool.¹

A natural cryptographic approach to reinstate the “blindness” of miners to the contents of the transactions is by having clients encrypt the contents of their transactions before they submit the transactions to the mempool. In this scenario, when the miners have to pick and order transactions, they must do so from an “encrypted” mempool, and thereby without knowledge of the contents of the transaction.² But for such a system to be of continued use, the transactions need to be executed *in the clear*. Requiring clients to return and decrypt their own transactions is undesirable for several reasons: (i) it requires the clients to be online more frequently for the system to function; and (ii) a misbehaving client may simply refuse to decrypt their transaction after it has learned the contents of the other transactions.

Instead, we want clients to be able to encrypt directly to the system, allowing the system to decrypt selected transactions without requiring further input from the clients. To distribute trust across the system, we can envision a set of decryption servers. An ideal system implementing this approach would proceed as follows: (1) The decryption servers run a one-time lightweight setup to generate a public key for the system; (2) Clients encrypt their transactions using the public key; and (3) Once B transactions for the block have been selected, each decryption server broadcasts a *partial decryption* of the ciphertexts, enabling anyone in the system to decrypt the B selected transactions. Decryption should be possible as long as a sufficient number of decryption servers broadcast their respective *partial decryptions*.

¹We note that if the cost of transaction privacy is deemed too high, our system can be modified to support both encrypted and unencrypted transactions. But for this work, we will stick to all transactions that are encrypted.

²Note that we are not demanding any “fair” ordering of the encrypted transactions, only that the ordering does not consider the transaction contents.

For this scheme to be practical, it must be concretely efficient: the setup should be lightweight, and the size of each server’s partial decryption should be constant, *independent of the number of ciphertexts B* . Intuitively, the ciphertexts should: (1) provide no information about their underlying contents until they are decrypted; and (2) continue to reveal no information if the ciphertext was *not included in the block*. The latter security property was formalized recently in [19] as *pending transaction privacy*. It is not hard to see that in the context of our earlier examples, it is critical that pending transaction privacy is maintained.

Given the importance of the problem, several works have made progress toward constructing a solution that achieves both mempool privacy and the stated efficiency goals [7, 15, 19, 27, 28, 62]. However, as we shall discuss shortly, none of these works are able to simultaneously achieve all of the desired properties – some fail to meet certain efficiency goals, while others crucially fail to provide pending transaction privacy.

1.1 Our Contribution

In this work, we present the first *concretely efficient* construction that guarantees mempool privacy and simultaneously achieves the following efficiency requirement.

1. Only requires a *single* execution of a distributed key generation (DKG) protocol.
2. Uses $O(1)$ communication per decryption server to decrypt a batch of B transactions.
3. Ensures *pending transaction privacy*.
4. Has $O(1)$ sized ciphertexts.

We created a Rust crate for our scheme and benchmarked it. It takes < 8.5 ms to encrypt a ciphertext and adds ≈ 466 bytes to the transaction size. For an Ethereum deployment, which processes ≈ 500 transactions per block, it would take close to 3.2 s to compute a partial decryption and 3.0 s to reconstruct all messages in single-threaded mode. We refer the readers to Section 6 for more details on the concrete efficiency of our scheme.

1.2 Related Works

The de facto approach to the mempool privacy problem is by the use of *threshold public key encryption* [24, 25, 32]. In such an encryption scheme there is a single public key, and the decryption servers hold shares of the decryption key such that a large fraction of them are required to decrypt any ciphertext encrypted to the public key. In the context of mempool privacy, clients can encrypt to the public key, and once the selected ciphertexts are determined, the decryption servers run the threshold decryption algorithm to only decrypt the selected ciphertexts. While the decryption servers only need to execute

a DKG as a part of the setup, the total communication sent by the n decryption servers in order to decrypt B encrypted transactions is $O(nB)$.

In order to reduce the total communication during decryption, works such as [27,55] and deployed systems such as [62] introduce fine-grained encryption schemes where the clients can encrypt to the epoch (or block) they want their transaction to be added to, and the total communication to decrypt *all* such ciphertexts is $O(n)$. In particular, this means that this information can be used to also decrypt ciphertexts that were not selected to be in the block, thereby violating the crucial requirement of *pending transaction privacy*.

Most closely related to our work is a recent scheme [19], which formalizes the security properties needed to obtain a meaningful notion of mempool privacy. They further define, and construct, a cryptographic primitive they term *batched threshold decryption* that suffices to obtain mempool privacy. While the scheme in [19] has total decryption communication $O(n)$ and guarantees pending transaction privacy, the biggest limitation is the (epoch) setup. In particular, the n decryption servers must run an initial MPC for setup, and then for every epoch (block), they must additionally run an epoch setup. Although one could design an optimized, special purpose MPC protocol, [19] only provided estimates using generic techniques for honest majority with abort MPC protocols. In contrast our work only needs a *one-time* DKG setup at the beginning of the protocol which has been extensively studied [17,37,40,44,46,60], simpler and concretely much more efficient than running a full-fledged MPC protocol. The above described comparisons are summarized in Table 1.

We note that while heavy cryptographic machinery like general purpose witness encryption [34] or indistinguishability obfuscation [4,5] can be used to derive a solution to mempool privacy, the inefficiency of known constructions limits their applicability towards constructing concretely efficient schemes. Likewise, timelock encryption [10] is also considered as a possible solution to mempool privacy, but does not satisfy pending transaction privacy in addition to requiring wasteful computation to perform decryption of the encrypted transactions.

Finally, while our focus in this work is more broadly mempool privacy, we note that for the specific case of reducing MEV, non-cryptographic approaches have also been considered [6,21,41,48,49,51,54,69]. These broadly aim to either limit how transactions can be re-ordered, or aim to provide game-theoretic defenses with both approaches resulting in reduced MEV. We refer the reader to [19] for a more detailed exposition of these non-cryptographic approaches.

Concurrent Works. Shortly after our work appeared online, two concurrent works [1,13] addressing similar problems also appeared online. We compare these works below.

In [13], the authors consider the same problem of a one-time setup batch (threshold) decryption. However, their work

uses a different approach based on *key-homomorphic puncturable PRFs*. As an exciting feature, their construction avoids the dependence on epoch id, circumventing the need for users to have to re-encrypt their message when there is a change in the epoch. On the downside, their decryption time is quadratic in the batch size.³

In [1] the authors introduce a related problem of batched-identity-based encryption (Batched IBE). In this primitive, messages can be encrypted to specific identities and batch labels. The master secret key can then be used to issue a *constant-sized* key (for a specific batch label) that enables the decryption of ciphertexts corresponding to an arbitrarily chosen subset of identities. A threshold version of this primitive implies the batch (threshold) decryption primitive considered in our work. At a technical level, their construction is essentially the same as this work, i.e., relying on witness encryption for PPE, KZG polynomial commitment scheme, and BLS signatures.

Scheme	Comm.	Pending Tx Privacy	Setup
[7,28,57]	$O(nB)$	✓	DKG
[27,55,62]	$O(n)$	✗	DKG
[19]	$O(n)$	✓	MPC + Epoch Setup
This Work	$O(n)$	✓	DKG

Table 1: A comparison of this work against prior work in terms of communication to decrypt a batch of B transactions and whether the scheme preserves the privacy of transactions that were submitted but may not have been included in the batch. DKG refers to a distributed key generation protocol and MPC refers to a (more expensive) multi-party computation protocol used during the setup phase and Epoch Setup refers to a per-epoch setup that can be carried out well before the actual epoch takes place.

2 Technical Overview

This section presents the main ideas behind our construction. We start by recalling our objective.

Our Goal. Assuming a lightweight one-time setup (e.g., standard DKG setup) among a committee, we wish to achieve the following. For each epoch, users should be able to independently (from other users) sample a ciphertext encrypting

³Their work proposes different constructions offering trade-offs between decryption time and batch decryption size. When one insists on *constant-size* batch decryption, the decryption time grows quadratic in the batch size. In comparison, our work simultaneously achieves constant-size decryption and quasi-linear decryption time.

their message. Among the ciphertexts from all users (say, ct_1, \dots, ct_m) in a given epoch, once the committee agrees on a subset $\mathcal{S} \subseteq [m]$ of size $\leq B$ to decrypt, the committee sends a constant-size partial decryption. Given more than a threshold T number of partial decryptions, one may decrypt all ciphertexts in \mathcal{S} . However, the ciphertexts outside \mathcal{S} should remain hidden. As shown in [19], constructing such a primitive suffices to obtain mempool privacy scheme. Thus, for the rest of this overview, we focus on the underlying ideas for building such a primitive. Additionally, as referenced in our discussion previously (Section 1.2), all prior works fail to achieve this goal in one way or another.

Witness Encryption for PPEs. To facilitate presenting our ideas, let us first introduce some minimal notation and a generic way of building witness encryption schemes [34] from pairing product equations (PPEs).⁴

Suppose we have a pairing-friendly group with pairing operation $e(\cdot, \cdot)$ and generators g, h, g_τ . Consider the following set of pairing product equations

$$e(g^A, h^x) = g_\tau^b, \quad (1)$$

where A is a $u \times v$ matrix of field elements, and b and x are vectors of dimension u and v , respectively. Here, we think of g^A and g_τ^b as *public* group elements and h^x as the *secret* group elements. Importantly this facilitates building a *witness encryption* scheme for the statement,

I hold the secret group elements h^x that will pass the pairing product equations $e(g^A, \cdot) = g_\tau^b$.

Specifically, to encrypt a message m for the above linear system, one samples a u -dimension random vector of field elements α and computes the ciphertext as

$$ct = (ct_1, ct_2) = \left(g^{\alpha A}, H \left(g_\tau^{(\alpha, b)} \right) \oplus m \right). \quad (6)$$

Here $H(\cdot)$ is a random oracle. Note that the public group elements g^A and g_τ^b are treated as the statement, and anyone in possession of the witness (secret group elements) h^x can decrypt by computing m as $ct_2 \oplus H(e(ct_1, h^x))$. For security, one can prove that the ciphertext is computationally hidden as long as the witness h^x is hard to compute (see Theorem 3).

Several works (e.g., [19, 35, 50]) show how WE for PPEs can be used to realize powerful applications. In these works,

⁴This general tool was observed in the literature under different names. For example, [9] observed this in the hash proof system context; [35] observed this from the witness encryption perspective (which is what we present here). Moreover, this tool is implicitly used in a less general manner in many existing works (e.g., [8, 11, 36, 50]).

⁵For a matrix/vector A of field elements, we use g^A for a matrix/vector of group elements defined in a natural way. Moreover, for two matrices g^A, g^B of group elements, $e(g^A, g^B)$ stands for the standard matrix multiplication with multiplication operation replaced by pairing operation.

⁶Note that ct_1 is a *vector* of group elements.

as in ours, the core contribution lies in casting the task at hand in the form of appropriate PPEs satisfying the constraints above. This requires “special gymnastics” to get the problem in the right format. We next explain the challenges involved in this for our work.

First Attempt. Let us start with a simple idea, by applying WE for PPEs to BLS signature [12]. Recall that, in BLS signature, the public key is $pk = h^{sk}$ for a signing key sk , and the signature σ for a tag tg is computed as $\sigma = H(tg)^{sk}$. The signature verification check is

$$e(\sigma, h) = e(H(tg), pk)$$

conforming to a PPE system as described in (1). Hence, one may apply WE for PPEs to build an encryption scheme.⁷ To extend this to multiple users while ensuring succinct batch decryption, one may ask each user to sample their respective ciphertext under the *same tag* tg .⁸ In this way, the committee could compute a (threshold) signature σ of tg as the witness to enable batch decryption of *all* ciphertext encrypted under tg . This is essentially the scheme of McFly [27].

However, this simple approach above does not provide *pending transaction privacy*. Indeed, the (threshold) BLS allows for decryption of all ciphertext encrypted under tg , not necessarily the B selected transaction. Therefore, the challenge remains how one designs a PPE system that allows succinct decryption of *only the selected transactions*.

Polynomial Encoding of a Set. To get around this challenge, it was proposed in [19] to use a *KZG polynomial commitment* [45]. Let us recall KZG polynomial commitment first. To enable committing to a degree $(B - 1)$ polynomial, the KZG common reference string consists of the tuple $(g, g^\tau, g^{\tau^2}, \dots, g^{\tau^{B-1}})$, and the commitment to a polynomial p is $com = g^{p(\tau)}$.⁹ Given a pair $(x, tg) \in \mathbb{F} \times \mathbb{F}$ such that $tg = p(x)$, the proof of opening $\pi = g^{q(\tau)}$ where $q(X)$ is a polynomial such that $q(X) = (p(X) - tg)/(X - x)$. The verification equation is thus given by

$$e(\pi, h^{\tau-x}) = e(com/g^{tg}, h).$$

An observant reader may note that the above verification equation is again compatible with (1).¹⁰ Indeed, in [19] the authors were able to leverage this to build a batch threshold decryption in the following way. For each epoch, a random group element $com = g^y$ will be sampled as the commitment and the committee shall hold secret shares of y . Now, each user will

⁷This is exactly the Boneh-Franklin IBE scheme [11].

⁸For example, this tag could be the epoch ID.

⁹This can be computed from the common reference string without knowledge of τ , which is a secret trapdoor for the scheme.

¹⁰Indeed, this property of the KZG commitment verification has been used recently in many works to construct witness encryption for polynomial opening [19, 31].

sample its own (x_i, tg_i) and compute a witness encryption for the KZG opening verification check for the aforementioned commitment com ,

$$e(\pi_i, h^{\tau-x_i}) = e(\text{com}/g^{\text{tg}_i}, h).$$

Now, suppose a set S of transactions is selected to be decrypted. For succinct batch decryption, instead of asking the committee to compute an opening proof π_i for each ciphertext in S , the committee shall leverage the *equivocal* property of the commitment com ¹¹ to explain it as the polynomial commitment of the polynomial $f(x)$ that interpolates all $\{(x_i, \text{tg}_i)\}_{i \in S}$. Note that if $f(x)$ becomes public, any one can generate the opening proof and, henceforth, decrypt the ciphertexts in S . The crucial point here is that to communicate $f(x)$, the committee does not need to send the entire $f(x)$, but only to publish an arbitrary point $(x^*, f(x^*))$ on $f(x)$.¹²

The construction we presented above is indeed the construction proposed by [19]. While this is an elegant solution, it unfortunately comes with several undesirable features.

- It necessitates a *per-epoch* setup — the committee needs to collectively sample a random group element com for every epoch. In other words, com is not a *reusable* setup.
- Equivocating $\text{com} = g^y$ to a chosen polynomial requires the committee to hold the trapdoor τ . In fact, the computation required for equivocation is a high-degree computation in τ . To enable a light-weight decryption process, [19] asks for an expensive initial setup such that committee holds secret shares of not only τ , but also τ^2, τ^3, \dots . This setup requires an expensive tailor-made MPC protocol. For instance, a simple DKG would not have sufficed.

Therefore, the question remains

Can we rely on the idea of polynomial commitment for batch decryption, while only requiring a simple one-time setup?

First Core Idea: Commitment as Witness. Upon close inspection, one may realize that the undesirable features of [19] stem from the sampling of the (equivocal) commitment com . To eliminate these drawbacks, we must allow the users to do witness encryption *without such a com that needs to be sampled afresh for each epoch*. Towards this, we first observe that the KZG opening check still conforms to the PPEs system even if *the commitment is part of the witness*. Namely,

$$e(\pi_i, h^{\tau-x_i}) = e(\text{com}/g^{\text{tg}_i}, h).$$

¹¹In [19], it was shown that to achieve equivocation, one must use a randomized version of the KZG commitment scheme. For the purposes of keeping this overview simple, we do not explain here the randomized version.

¹²Indeed, this enables a public reconstruction of $f(x)$ by applying Lagrange interpolation on $\{(x_i, \text{tg}_i)\}_{i \in S} \cup \{(x^*, f(x^*))\}$.

In other words, if the users witness encrypts with respect to the above PPE system, the committee no longer needs to sample a random commitment com per epoch — effectively removing the per-epoch setup. While this idea works nicely, putting the commitment as a part of the witness, however, comes with significant challenges as we explain next.

The “Intractability” Challenge. The core issue is the following. Since the committee does not sample the commitment, it is now *computationally easy* to find a satisfying witness. Indeed, the decryptor can pick any polynomial that interpolates (x_i, tg_i) as its witness and breaks the semantics security of the i -th ciphertext. Therefore, the challenge becomes

How can we restore the critical property that the witness is intractable to find?

One may naturally wonder if we can fix this issue by asking the committee to produce a signature on com to authenticate the commitment com . That is, each user computes a witness encryption to the following PPEs system,

$$\begin{aligned} e(\sigma, h) \cdot e(\text{com}, \text{pk})^{-1} &= g_{\tau}^0, \\ e(\pi, h^{\tau-x_i}) \cdot e(\text{com}, h)^{-1} &= e(g^{\text{tg}_i}, h)^{-1}. \end{aligned}$$

Although this idea appears to make the witness hard to compute, subtle attacks still persist. For instance, consider the adversary that wants to break the i -th ciphertext. They may still use a *constant polynomial* as the polynomial that interpolates (x_i, tg_i) . In this case, the signature on com is still easy to compute since the dLog of com is picked by the adversary.

Second Core Idea: Shifted BLS. We resolve the “intractability” challenge by using what we call a *shifted BLS* signature. Namely, we ask the committee to produce a signature, not on the commitment com itself, but on com *shifted by a public random group element*. Note that, unlike [19], the committee does not need to hold any trapdoor related to this random group element and, hence, this random group element will not require a per-epoch setup.

To elaborate, users will now encrypt to an epoch id eid ¹³ that is publicly known ahead of time. We require the committee to produce a signature on $H(\text{eid})/\text{com}$. Specifically, the final PPEs that the users apply witness encryption to is,

$$\begin{aligned} e(\sigma, h) \cdot e(\text{com}, \text{pk}) &= e(H(\text{eid}), \text{pk}), \\ e(\pi, h^{\tau-x_i}) \cdot e(\text{com}, h)^{-1} &= e(g^{\text{tg}_i}, h)^{-1}. \end{aligned} \quad (2)$$

In other words, the witness that decrypts the above ciphertext (from user i) states

I have a signature σ on a $H(\text{eid})/\text{com}$, where com has a proof of opening π at x_i to tg_i .

¹³Each epoch will select B ciphertexts to decrypt, and the users whose ciphertexts are not chosen can encrypt to the next epoch, and so on.

Intuitively, an adversary either needs to break the BLS signature scheme to generate a signature of its choice for any commitment com' of its choosing, or it needs to break the KZG polynomial commitment opening to open a fixed com to any (x', tg') of its choosing. We formalize this intuition with a new assumption (Definition 2). In the technical section, we show our new assumption reduces to more standard assumptions in the algebraic group model [33]. Intuitively, we show that the aforementioned two cases are the only possible ways that an adversary can break the system. This allows us to argue that ciphertexts not included in the B ciphertexts will never be decrypted.

Security Aspect. Finally, we note that we have ignored malleability issues so far for our ciphertext. These issues arise since an adversarial user may maul an honestly generated ciphertext such that if the mauled ciphertext is among the B ciphertexts selected, the witness for such a ciphertext can be converted to a witness for the honest ciphertext that *was not* among the B selected. For instance, if an adversary were to simply “copy” (x, tg) from an honest ciphertext ct when generating its own ciphertext ct' , it is easy to see that the witness to decrypt ct' also serves as a witness to decrypt ct . Our construction handles such non-malleability issues in an identical manner to [19]. Specifically, the users must provide a non-interactive proof of knowledge of the correct computation of the ciphertext. Further, to avoid the aforementioned copy attack, tg is generated by each user as $\text{tg} = H(g^s)$, where the user outputs S and uses the non-interactive proof of knowledge to prove knowledge of s such that $S = g^s$. We refer the readers to the technical section for details on this part.

Put Everything Together. To conclude, when a user wants to encrypt a message for an epoch eid , it applies the witness encryption scheme (see (1)) to the PPEs system (2) to generate the ciphertext. Additionally, it sends a non-interactive (zero-knowledge) proof of knowledge to attest to the correct computation of the ciphertext. Once the B ciphertexts are chosen, the committee uses the (x_i, tg_i) from the selected ciphertexts to interpolate and generate a polynomial commitment com . The committee then signs $H(\text{eid})/\text{com}$ and broadcasts the signature in a threshold manner. Any user can now decrypt the B selected ciphertexts by redoing the computation of com , and, in turn, computing the required polynomial commitment openings. They finally use the broadcast signature to obtain the witnesses needed for decryption. Since the committee only sends a single signature, it satisfies our efficiency requirements from our goal. Further, the committee only needs to run a DKG at the start of the protocol to get Shamir’s sharing of the BLS signing key.

Comparison to [19]. As we have explained above, our work builds on the ideas presented in [19]. While there are similar-

ities, our approach crucially avoids the need for a per-epoch setup. In particular, the approach in [19] uses the KZG polynomial commitment as *part of the statement* in the underlying PPE system. However, since the KZG commitment is not reusable in their scheme, the committee must generate a new commitment for each epoch, necessitating a per-epoch setup. In contrast, our key insight is that the KZG commitment can also serve as *part of the witness* in the PPE system, removing the inherent dependency on per-epoch setups. To enable this key insight, our work introduces several additional ideas such as signature authentication and shifted BLS.

3 Preliminaries

We use $[n]$ to denote the set $\{1, 2, \dots, n\}$. The security parameter is denoted by $\lambda \in \mathbb{N}$. A function $f : \mathbb{N} \rightarrow \mathbb{N}$ is said to be polynomial if there exists a constant c such that $f(n) \leq n^c$ for all $n \in \mathbb{N}$, and we write $\text{poly}(\cdot)$ to denote such a function. A function $f : \mathbb{N} \rightarrow [0, 1]$ is said to be negligible if for every $c \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$, and we write $\text{negl}(\cdot)$ to denote such a function. A probability is *noticeable* if it is not negligible, and *overwhelming* if it is equal to $1 - \text{negl}(\lambda)$ for some negligible function $\text{negl}(\lambda)$. For a set \mathcal{S} , we write $s \leftarrow \mathcal{S}$ to indicate that s is sampled uniformly at random from \mathcal{S} . An algorithm \mathcal{A} is PPT (probabilistic polynomial-time) if its running time is bounded by some polynomial in the size of its input.

Bilinear Groups. We use e to denote an efficiently computable non-degenerate bilinear pairing from the groups $(\mathbb{G}_1, \mathbb{G}_2)$ to a target group \mathbb{G}_T and use \mathbb{F} to denote the associated field. The generators of \mathbb{G}_1 and \mathbb{G}_2 are fixed to be g and h , respectively. We will represent the Shamir secret sharing of an element x as $[x]$, with i -th share represented as $[x]_i$. We note that while we overload the previously discussed notation $[n]$, the notation for secret shares will be clear from the context.

The Random Oracle Model. In the random oracle model (ROM), parties are given oracle access to some function H that is sampled uniformly at random from the space of all functions $H : \mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{X} and \mathcal{Y} are finite non-empty sets. That is, parties can query their oracle on an input $x \in \mathcal{X}$ and receive in return $H(x) \in \mathcal{Y}$. When proving the security of a protocol in the random oracle model, the simulator is able to “control” the oracle, observing queries made by the adversary and simulating responses.

Non-interactive Zero-Knowledge proofs. Let \mathcal{L} be an NP-language and \mathcal{R} the corresponding NP-relation. A Simulation Extractable-NIZK for \mathcal{R} consists of the following algorithms:

- $\text{Setup}(1^\lambda) \rightarrow (\text{crs}, \text{td})$: Takes as input a security parameter, and outputs a common reference string crs and trapdoor td .
- $\text{Prove}(\text{crs}, x, w) \rightarrow \pi$: Takes as input crs and any pair $(x, w) \in \mathcal{R}$ and outputs a proof π for the statement $x \in \mathcal{L}$.

- $\text{Verify}(\text{crs}, x, \pi) \rightarrow b$: Takes as input crs , statement x and proof π and outputs a bit b indicating whether verification has passed or failed.
- $\text{SimProve}(\text{crs}, \text{td}, x) \rightarrow \pi$: Takes as input crs , trapdoor td and statement x and outputs a *simulated* proof π .

We will drop the crs term wherever implicit. A SE-NIZK satisfies the following properties:

- *Perfect Completeness*. An SE-NIZK satisfies perfect completeness if for any $(x, w) \in \mathcal{R}$, we have

$$\Pr \left[\begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \pi \leftarrow \text{Prove}(\text{crs}, x, w) \end{array} : \text{Verify}(x, \pi) = 1 \right] = 1.$$

- *Proof of knowledge (and soundness)*. For every PPT adversary \mathcal{A} , there is a PPT extractor Ext such that

$$\Pr \left[\begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}(\text{crs}) \\ w \leftarrow \text{Ext}(\text{crs}) \end{array} : \begin{array}{l} \text{Verify}(x, \pi) = 1 \\ (x, w) \notin \mathcal{R} \end{array} \right] \leq \text{negl}(\lambda).$$

- *Computational Zero-knowledge*. There exists an algorithm SimProve such that for all PPT adversary \mathcal{A} :

$$\Pr \left[\begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ (x, w) \leftarrow \mathcal{A}(\text{crs}) \\ \pi \leftarrow \text{Prove}(\text{crs}, x, w) \end{array} : \begin{array}{l} (x, w) \in \mathcal{R} \\ \mathcal{A}(\pi) = 1 \end{array} \right] - \Pr \left[\begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ (x, w) \leftarrow \mathcal{A}(\text{crs}) \\ \pi \leftarrow \text{SimProve}(\text{crs}, \text{td}, x) \end{array} : \begin{array}{l} (x, w) \in \mathcal{R} \\ \mathcal{A}(\pi) = 1 \end{array} \right] \leq \text{negl}(\lambda).$$

- *Weak Simulation-Extractability*. For every PPT adversary \mathcal{A} , there exists a PPT extractor Ext such that

$$\Pr \left[\begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}^{\text{SimProve}} \\ w \leftarrow \text{Ext} \end{array} : \begin{array}{l} \text{Verify}(x, \pi) = 1 \\ \wedge (x, w) \notin \mathcal{R} \\ \wedge x \notin Q \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{A} has oracle access to $\text{SimProve}(\text{crs}, \text{td}, \cdot)$, and Q is a list of queries made by the adversary.

We will also demand that the proof is straight-line extractable. This means that the extractor can, on input a valid proof and (potentially) the list of random-oracle queries made by the adversary (prover), extract a witness with overwhelming probability without rewinding the prover.

Algebraic Group Model. We prove security in the algebraic group model [33]. In this model, adversaries are restricted to performing algebraic manipulation of group elements that they have seen so far. When an adversary outputs a value $Y \in \mathbb{G}$, it also produces a *linear combination* (e_1, \dots, e_m) of elements (X_1, \dots, X_m) that it has seen so far, such that $Y = \prod X_i^{e_i}$.

4 Model and Definitions

We use essentially the same security and efficiency requirements of batched-threshold encryption as defined in [19], with one main difference — our definition does not involve an epoch setup.

- $\text{Setup}(1^\lambda, n, t, B) \rightarrow \{\text{pk}, (\text{sk}_1, \dots, \text{sk}_n)\}$: On input the threshold t for n parties and a batch size B , outputs the public key pk along with secret keys for each party.
- $\text{Enc}(\text{pk}, \text{eid}, m) \rightarrow \text{ct}$: On input a message m , an epoch id eid , and a public key pk , outputs a ciphertext ct .
- $\text{BatchDec}((\text{ct}_1, \dots, \text{ct}_B), \text{eid}, \text{sk}_i) \rightarrow \text{pd}_i$: On input B ciphertexts $(\text{ct}_1, \dots, \text{ct}_B)$, the epoch id eid , and a secret key shares sk_i , outputs a partial decryption pd_i or \perp .
- $\text{Combine}(\text{pk}, \text{eid}, (\text{ct}_1, \dots, \text{ct}_B), \{\text{pd}_i\}_{i \in S}) \rightarrow (m_1, \dots, m_B)$: Takes as input the public key pk , epoch id eid and $t + 1$ partial decryptions where $S \subseteq [n]$ and $|S| = t + 1$, and outputs messages (m_1, \dots, m_B) or \perp .

Efficiency. We impose various efficiency constraints on the above algorithms. Setup must run in time $\text{poly}(n, t, B, \lambda)$. Enc must run in $O(1)$ time. BatchDec must be non-interactive i.e., each party computes its partial decryption independently. It must also run in time $\text{poly}(B, \lambda)$, and the size of each partial decryption must be sub-linear in the batch size ($o(B)$) and *independent* of the number of parties. Finally, Combine must run in time $\text{poly}(B, n, \lambda)$.

The ideal functionality. We now describe the ideal functionality (Fig. 1). Note that even though there is no interactive epoch setup, ciphertexts are still encrypted to a particular epoch and the parties must agree on a batch of ciphertexts to decrypt in each epoch. Again, our ideal functionality is identical to that of [19] except for the omission of EpochSetup . However, the encrypt and decrypt methods still make use of an epoch id eid .

5 Construction

We now present our construction of batched-threshold encryption. Note that we present an optimized version of the protocol which uses a roots of unity domain to interpolate tags across ciphertexts. This allows us to compute openings proofs of a KZG commitment at all points in the domain in $O(B \log B)$ time. However, this also means that if two clients choose the same \hat{x} , then only one of these ciphertexts can be included in the block. This can be avoided by letting clients choose an arbitrary $\hat{x} \in \mathbb{F}$ at the cost of a more expensive combine step. We prove that the latter securely emulates the ideal functionality Fig. 1.

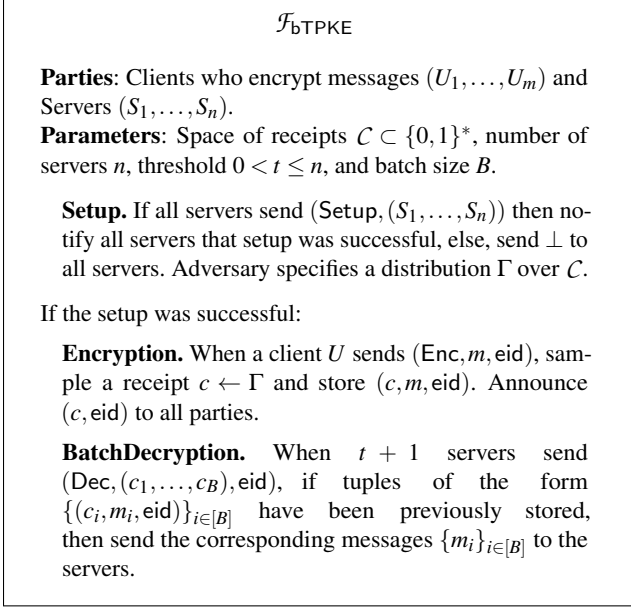


Figure 1: Ideal Functionality for a Batched Threshold Public-Key Encryption scheme.

Setup: Our setup is significantly simpler than that of [19]. We only require a *single* execution of a distributed key generation protocol to establish shamir shares of a secret key $sk \in \mathbb{F}$ among n parties. The public key is $pk = h^{sk}$ along with $\{h^{[sk]_j}\}_{j \in [n]}$. We also require a powers-of-tau CRS – $\{g, g^\tau, g^{\tau^2}, \dots, g^{\tau^{B-1}}\}$ – which can be sampled in a distributed manner as outlined in [56].¹⁴ Finally, the parties also run the setup for a NIZK proof system (in a distributed manner if necessary) and publish the public parameters (if any). In our concrete implementation, we will use sigma protocols, which are secure in the random oracle model and do not require an interactive setup.

Encrypt: Clients sample a random value $\hat{x} \in \Omega = \{1, \omega, \omega^2, \dots, \omega^{B-1}\}$, where ω is a root of unity in \mathbb{F} of order B . Informally, they create a witness encryption to the following relation, using the compiler from [9, 35]:

I know an opening proof π for some KZG commitment [45] com , at the point \hat{x} to the value tg , and I know a BLS signature σ on $\delta = H(\text{eid})/\text{com}$,

where variables in **red** are part of the witness and those in **blue** are part of the statement. More formally the set of PPE’s that need to be satisfied can be written as:

$$e(H(\text{eid}) \cdot g^{-\text{tg}}, h) = e(\pi, h^{(\tau-\hat{x})}) \cdot e(\delta, h)$$

¹⁴We note that one can also use a powers of tau CRS that has been previously setup in a distributed manner. For instance, <https://ceremony.ethereum.org> has 141,416 contributions. Soundness holds as long as at least one party deleted the randomness they used as part of their contribution.

$$e(\delta, pk) = e(\sigma, h)$$

Or equivalently,

$$e(\text{com} \cdot g^{-\text{tg}}, h) = e(\pi, h^{(\tau-\hat{x})})$$

$$e(H(\text{eid}) \cdot \text{com}^{-1}, pk) = e(\sigma, h)$$

where terms in red denote the witness.

Adding CCA2 Security. To achieve CCA2 security for our encryption scheme, we must a) prove that ciphertext is well-formed and b) prevent an adversary from choosing the same tg as an honest party.

We require the former to ensure correctness for ciphertexts created by the adversary, which is crucially used by the simulator to extract the underlying message as part of the security proof. The latter is necessary to prevent a concrete attack that can be launched by the adversary. Namely, the adversary first looks at an honest party’s ciphertext encrypted to some string tg^* and creates another “dummy” ciphertext which is encrypted to the same statement – $(\hat{x}, \text{eid}, \text{tg}^*)$. The adversary then gets the dummy ciphertext decrypted by including it in the next batch, and the partial decryptions will allow the adversary to not only decrypt the dummy ciphertext but also the honest party’s ciphertext as they require the *same* witness. To prevent this, we will choose tg^* as the output of a one-way function f on a randomly chosen input s and demand a proof of knowledge of s . Concretely, we use $f(s) := g^s$ as it allows us to use very efficient sigma protocols with small proofs.

Batch Decryption: When choosing a batch of ciphertexts to decrypt, the committee must ensure that the \hat{x} values across ciphertexts are distinct. The committee members first compute $\text{com} = g^{p(\tau)}$, where $p(X)$ is the degree- $(B-1)$ polynomial interpolating the tags across the B ciphertexts. The j -th party then publishes $\sigma_j = (H(\text{eid})/\text{com})^{[sk]_j}$ as their partial decryption.

Combine: We first recover the commitment $\text{com} = g^{p(\tau)}$, as done in the Batch Decryption step, and compute opening proofs for all points in Ω using [30]. Given t partial signatures on $H(\text{eid})/\text{com}$, we can interpolate in the exponent to recover σ such that $e(H(\text{eid})/\text{com}, pk) = e(\sigma, h)$. Finally, we simply run the witness encryption’s decrypt algorithm from [9, 35] to obtain the underlying messages.

Theorem 1. *The protocol in Fig. 2 securely emulates the $\mathcal{F}_{\text{bTRKE}}$ ideal functionality (Fig. 1) in the dealer model against any static PPT adversary A corrupting up to $t < n/2$ parties. This holds under the assumption that a weakly simulation-extractable NIZK is available and that the i -kzg game (Definition 2) is hard in the programmable random oracle model.*

Batched-Threshold Encryption

Parameters: A pairing friendly group $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$, and random oracles $H : \{0, 1\}^* \rightarrow \{0, 1\}^*$, $H_F : \{0, 1\}^* \rightarrow \mathbb{F}$, and $H_G : \{0, 1\}^* \rightarrow \mathbb{G}_1$.

• **Setup** $(1^\lambda, n, t, B)$: A trusted dealer provides the committee with (t, n) -shares of a random field element $sk \leftarrow \mathbb{F}$ and publishes the committee's public key $pk = h^{sk}$ along with $\{h^{[sk]_j}\}_{j \in [n]}$, and $crs_{KZG} = (g, g^\tau, \dots, g^{\tau^{B-1}}, h, h^\tau)$. They also run the setup for NIZK and publish crs_{NIZK} .

• **Enc** (pk, eid, m) : Clients first sample a random value $\hat{x} \in \Omega = \{1, \omega, \omega^2, \dots, \omega^{B-1}\}$, where ω is a root of unity in \mathbb{F} of order B . They also sample $(\alpha, \beta, s) \leftarrow \mathbb{F}^3$, and compute $S \leftarrow g^s$, and $tg \leftarrow H_F(S)$. The ciphertext is set to be:

- $ct^{(1)} := H(e(H_G(eid) \cdot g^{-tg}, h)^\alpha) \oplus M$
- $ct^{(2)} := h^{\alpha(\tau - \hat{x})}$
- $ct^{(3)} := h^\alpha pk^\beta$
- $ct^{(4)} := h^\beta$
- S
- \hat{x}
- $\Phi = \text{NIZK.Prove}\{\alpha, \beta, s \mid S = g^s \wedge ct^{(1)} \wedge ct^{(2)} = h^{\alpha(\tau - \hat{x})} \wedge ct^{(3)} = h^\alpha pk^\beta \wedge ct^{(4)} = h^\beta\}$

• **BatchDec** $((ct_0, \dots, ct_{B-1}), eid, sk_i)$: A threshold number of parties agree on a batch of ciphertexts $\{ct_0, \dots, ct_{B-1}\}$ such that every ciphertext has a valid proof Φ_i and $\{ct_0 \cdot \hat{x}, \dots, ct_{B-1} \cdot \hat{x}\} = \Omega$. Compute $com = g^{p(\tau)}$ using crs_{KZG} , where $p(X)$ is a degree- $(B-1)$ polynomial such that $\{p(ct_i \cdot \hat{x}) = tg_i\}_{i \in [B]}$, where $tg_i = H(S_i)$. The j -th party then publishes $\sigma_j = (H_G(eid)/com)^{[sk]_j}$ as the partial decryption.

• **Combine** $(pk, (ct_1, \dots, ct_B), \{\sigma_i\}_{i \in \mathcal{S}})$: Recover the commitment $com = g^{p(\tau)}$, where $p(X)$ is the degree- $(B-1)$ polynomial such that $\{p(ct_i \cdot \hat{x}) = tg_i\}_{i \in [B]}$, where $tg_i = H_F(ct_i \cdot S)$. Compute $\delta = H_G(eid)/com$ and opening proofs $\{\pi_i = g^{q_i(\tau)}\}_{i \in [B]}$, where $q_i(X) = (p(X) - tg_i)/(X - ct_i \cdot \hat{x})$. Given t partial decryptions $\{\sigma_j \mid e(H_G(eid)/com, pk_j) = e(\sigma_j, h)\}_{j \in \mathcal{S}}$, where $\mathcal{S} \subseteq [n]$ and $|\mathcal{S}| > t$, interpolate in the exponent to recover σ such that $e(H_G(eid)/com, pk) = e(\sigma, h)$. Output

$$\{ct_i^{(1)} \oplus H(e(\pi_i, ct_i^{(2)}) \cdot e(\delta, ct_i^{(3)}) \cdot e(\sigma^{-1}, ct_i^{(4)}))\}_{i \in [B]}.$$

Figure 2: Protocol for Batched-Threshold Encryption

6 Implementation and Evaluation

To evaluate the concrete performance of our scheme, and to ensure a fair comparison against prior work, we implemented a performant version of our scheme in Rust using the same libraries. We use `arkworks` [3] for implementations of pairing-friendly curves and their associated algebra, and the `merlin` [23] library to handle the Fiat-Shamir transform. Our implementation is available at <https://github.com/guruvamsi-policharla/batched-threshold-pp>

Setup. All of our experiments were run on a 2019 MacBook Pro with a 2.4 GHz Intel Core i9 processor and 16 GB of DDR4 RAM in single-threaded mode.

Algebra. We use BLS12-381 as our pairing friendly curve and the BLAKE3 hash function as our random oracle.

Dealer. We assume a trusted dealer for the distributed key generation during the setup phase but this can be emulated using any verifiable secret sharing scheme [17, 37, 40, 44, 46, 60].

Evaluation. We now evaluate the performance of our scheme. We answer the following questions:

- How long does it take to encrypt a message? What is the size of the corresponding ciphertexts?
- How long does it take for a committee member to compute partial decryptions, and how does it vary with batch/committee size?
- How long does it take to recover all messages given partial decryptions, and how does it vary with batch size?

Batch size	Partial decrypt (ms)	Reconstruct (ms)
8	49.2	60.5
32	199.8	168.1
128	809.9	646.1
512	3203.9	3026.5

Table 2: Time taken to compute partial decryptions and to reconstruct messages with varying batch size.

Encryption. The ciphertext contains one \mathbb{G}_1 element, three \mathbb{G}_2 elements, 4 \mathbb{F} elements, a two byte description of \hat{x} , and a string proportional to the message length (32 bytes), resulting in a total size of 498 bytes. Creating a ciphertext takes approximately 8.5ms, and is independent of committee/batch size. Our encryption time is $\approx 40\%$ slower and ciphertexts are $\approx 35\%$ larger than prior work [19]. In contrast, we only need a very simple *one-time* setup, whereas [19] required a complicated multi-party computation-based setup and per-epoch

Parameter	[7]	[55, 62]	[19]	This Work
Increase in Ciphertext size	$ \mathbb{G}_1 + \mathbb{G}_2 $	$ \mathbb{G}_2 $	$2 \mathbb{G}_2 + \mathbb{G}_1 + 3 \mathbb{F} + 2$	$3 \mathbb{G}_2 + \mathbb{G}_1 + 4 \mathbb{F} + 2$
Partial Decryptions	$nB \mathbb{G}_1 $ (3 MB)	$n \mathbb{G}_1 $ (6 KB)	$n(\mathbb{G}_1 + \mathbb{F})$ (10 KB)	$n(\mathbb{G}_1)$ (6 KB)
Partial Decryptions / Block Size	500%	1%	2%	1%

Table 3: Ciphertext size and total decryption broadcast size for n committee members, and B transactions per block. Concrete numbers were obtained with the BLS12-381 curve, using the average transaction and block size on Ethereum in 2023 for a committee size of 128 with a block of 512 transactions. We note that the numbers of [55, 62] are for CPA-secure ciphertexts. Adding CCA security would increase ciphertext size as shown in [27]. We compare against the most efficient schemes with weakest security guarantees to highlight that despite the handicap given to prior work, our scheme is competitive.

setup. We view this as a clearly favorable trade-off for the setting of mempool privacy.

Batch Decryption. Before issuing a partial decryption, each party must check that all ciphertexts in the batch are valid by verifying the accompanying zero-knowledge proof. We find that this is the most expensive step in the protocol, encompassing $\approx 99\%$ of the times reported in Table 2. Each partial decryption is one \mathbb{G}_1 element (48 bytes with compression), and 40% smaller than the partial decryptions in [19] but computing them is $\approx 20\%$ slower.

Reconstruct. Given sufficiently many partial decryptions, any party can recover the messages from the entire batch of ciphertexts that were decrypted in that epoch. In the process, we need to produce KZG opening proofs at all points in the domain Ω . Using the techniques from [30], we can produce all openings using $O(B \log B)$ group operations and $O(B)$ pairings. In Table 2, we provide numbers assuming that all parties respond with partial decryptions, and each partial decryption is verified to be correct. We note here an important distinction when comparing numbers reported in [19] with those in Table 2. [19] reports “optimistic” timings and size, where they assume that the entire committee responds and that all partial decryptions are correct. If an inconsistency is detected, they demand proofs that the partial decryption was computed correctly. Hence, they exclude the time taken to generate these proofs and sizes of these proofs from the reported numbers. Our partial decryptions are publicly testable using a pairing check, and hence malformed partial decryptions can be detected without incurring any additional communication. At a batch size of 512, we are faster than [19] as we use an optimized implementation to compute pairing product equation where the miller loop is evaluated only once instead of for every pairing. We also note that increasing committee size does not have a noticeable impact on the reconstruction time as computing KZG opening proofs takes a majority of the time – which is independent of committee size.

Committee Churn. Parties may leave/join the protocol for a variety of reasons. In such situations, we can ensure the protocol can continue without interruption using techniques from the line of work on Proactive Secret Sharing [14, 26, 39,

42, 53, 61, 64, 66, 67]. From the benchmarks in [19], this could be completed in under 10 seconds for a committee size of 64 in a simulated WAN on the Google Cloud Platform using the `NetEm` package with a delay time of 200 ms, jitter time of 20 ms and a rate of 10 mbit per second.

7 Acknowledgements

This work is supported in part by AFOSR Award FA9550-24-1-0156, and research grants by Bakar Fund, J. P. Morgan Faculty Research Award, Supra Inc., Visa Inc, and the Stellar Development Foundation. Guru-Vamsi Policharla is supported in part by a Berkeley Center for Responsible, Decentralized Intelligence (RDI) Fellowship. Mingyuan Wang is supported in part by an NYU startup fund.

8 Ethics Considerations

Following the guidelines laid out for ethical considerations, we believe that our work does not have any potentially negative outcomes from the research and publication process. Our primary contribution is a cryptographic scheme that enhances the privacy of user transactions in the context of decentralized finance (DeFi). We believe that such a system, if ever deployed would enable users to choose to protect their transactions. We view this as a positive outcome for the users of such a system, especially given that the transactions that are completed are eventually revealed to all users of the system.

9 Compliance with Open Science Policy

In addition to making the full version of our paper available on the *Cryptology ePrint Archive* [20], we have also made the associated research artifacts public at <https://zenodo.org/records/14642466>. For the latest version we refer the reader to <https://github.com/guruvamsi-policharla/batched-threshold-pp>.

References

- [1] Amit Agarwal, Rex Fernando, and Benny Pinkas. Efficiently-thresholdizable batched identity based encryption, with applications. Cryptology ePrint Archive, Paper 2024/1575, 2024.
- [2] Guillermo Angeris, Alex Evans, and Tarun Chitra. A note on bundle profit maximization. *Stanford University*, 2021.
- [3] arkworks contributors. arkworks zksnark ecosystem. <https://arkworks.rs>, 2022.
- [4] Gennaro Avitabile, Nico Döttling, Bernardo Magri, Christos Sakas, and Stella Wohnig. Signature-based witness encryption with compact ciphertext. Cryptology ePrint Archive, Paper 2024/1477, 2024.
- [5] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Berlin, Heidelberg, August 2001.
- [6] Carsten Baum, Bernardo David, and Tore Kasper Frederiksen. P2DEX: Privacy-preserving decentralized cryptocurrency exchange. In Kazuo Sako and Nils Ole Tippenhauer, editors, *ACNS 21 International Conference on Applied Cryptography and Network Security, Part I*, volume 12726 of *LNCS*, pages 163–194. Springer, Cham, June 2021.
- [7] Joseph Bebel and Dev Ojha. Ferveo: Threshold decryption for mempool privacy in BFT networks. Cryptology ePrint Archive, Report 2022/898, 2022.
- [8] Fabrice Benhamouda and Huijia Lin. k-round multi-party computation from k-round oblivious transfer via garbled interactive circuits. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 500–532. Springer, Cham, April / May 2018.
- [9] Olivier Blazy and Céline Chevalier. Structure-preserving smooth projective hashing. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 339–369. Springer, Berlin, Heidelberg, December 2016.
- [10] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 757–788. Springer, Cham, August 2018.
- [11] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer, Berlin, Heidelberg, August 2001.
- [12] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Berlin, Heidelberg, December 2001.
- [13] Jan Bormet, Sebastian Faust, Hussien Othman, and Ziyang Qu. BEAT-MEV: Epochless approach to batched threshold encryption for MEV prevention. Cryptology ePrint Archive, Paper 2024/1533, 2024.
- [14] Christian Cachin, Klaus Kursawe, Anna Lysyanskaya, and Reto Strohbl. Asynchronous verifiable secret sharing and proactive cryptosystems. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 88–97. ACM Press, November 2002.
- [15] Ran Canetti and Shafi Goldwasser. An efficient threshold public key cryptosystem secure against adaptive chosen ciphertext attack. In Jacques Stern, editor, *EUROCRYPT’99*, volume 1592 of *LNCS*, pages 90–106. Springer, Berlin, Heidelberg, May 1999.
- [16] Agostino Capponi, Ruizhe Jia, and Ye Wang. The evolution of blockchain: from lit to dark. *arXiv preprint*, 2022.
- [17] Ignacio Cascudo and Bernardo David. Publicly verifiable secret sharing over class groups and applications to DKG and YOSO. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part V*, volume 14655 of *LNCS*, pages 216–248. Springer, Cham, May 2024.
- [18] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. Marlin: Pre-processing zkSNARKs with universal and updatable SRS. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Cham, May 2020.
- [19] Arka Rai Choudhuri, Sanjam Garg, Julien Piet, and Guru-Vamsi Policharla. Mempool privacy via batched threshold encryption: Attacks and defenses. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [20] Arka Rai Choudhuri, Sanjam Garg, Guru-Vamsi Policharla, and Mingyuan Wang. Practical mempool privacy via one-time setup batched threshold encryption. Cryptology ePrint Archive, Paper 2024/1516, 2024.
- [21] Michele Ciampi, Muhammad Ishaq, Malik Magdon-Ismail, Rafail Ostrovsky, and Vassilis Zikas. Fairmm: A fast and frontrunning-resistant crypto market-maker. *IACR Cryptology ePrint Archive*, 2021.

- [22] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy*, pages 910–927. IEEE Computer Society Press, May 2020.
- [23] dalek-cryptography contributors. merlin. <https://github.com/dalek-cryptography/merlin>, 2020.
- [24] Alfredo De Santis, Yvo Desmedt, Yair Frankel, and Moti Yung. How to share a function securely. In *26th ACM STOC*, pages 522–533. ACM Press, May 1994.
- [25] Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In Gilles Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 307–315. Springer, New York, August 1990.
- [26] Yvo Desmedt and Sushil Jajodia. Redistributing secret shares to new access structures and its applications. Technical report, Citeseer, 1997.
- [27] Nico Döttling, Lucjan Hanzlik, Bernardo Magri, and Stella Wahnig. McFly: Verifiable encryption to the future made practical. In Foteini Baldimtsi and Christian Cachin, editors, *FC 2023, Part I*, volume 13950 of *LNCS*, pages 252–269. Springer, Cham, May 2023.
- [28] Taher ElGamal. On computing logarithms over finite fields. In Hugh C. Williams, editor, *CRYPTO’85*, volume 218 of *LNCS*, pages 396–402. Springer, Berlin, Heidelberg, August 1986.
- [29] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. Sok: Transparent dishonesty: front-running attacks on blockchain. In *International Conference on Financial Cryptography and Data Security*, 2019.
- [30] Dankrad Feist and Dmitry Khovratovich. Fast amortized KZG proofs. Cryptology ePrint Archive, Report 2023/033, 2023.
- [31] Nils Fleischhacker, Mathias Hall-Andersen, and Mark Simkin. Extractable witness encryption for KZG commitments and efficient laconic OT. Cryptology ePrint Archive, Report 2024/264, 2024.
- [32] Yair Frankel. A practical protocol for large group oriented networks. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *EUROCRYPT’89*, volume 434 of *LNCS*, pages 56–61. Springer, Berlin, Heidelberg, April 1990.
- [33] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Cham, August 2018.
- [34] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013.
- [35] Sanjam Garg, Dimitris Kolonelos, Guru-Vamsi Policharla, and Mingyuan Wang. Threshold encryption with silent setup. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 352–386, Cham, 2024. Springer Nature Switzerland.
- [36] Sanjam Garg and Akshayaram Srinivasan. Two-round multiparty secure computation from minimal assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 468–499. Springer, Cham, April / May 2018.
- [37] Craig Gentry, Shai Halevi, and Vadim Lyubashevsky. Practical non-interactive publicly verifiable secret sharing with thousands of parties. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 458–487. Springer, Cham, May / June 2022.
- [38] Arthur Gervais, Ghassan O Karame, Karl Wüst, Vasileios Glykantzis, Hubert Ritzdorf, and Srdjan Capkun. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016.
- [39] Vipul Goyal, Abhiram Kothapalli, Elisaweta Masserova, Bryan Parno, and Yifan Song. Storing and retrieving secrets on a blockchain. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part I*, volume 13177 of *LNCS*, pages 252–282. Springer, Cham, March 2022.
- [40] Jens Groth. Non-interactive distributed key generation and key resharing. Cryptology ePrint Archive, Report 2021/339, 2021.
- [41] Lioba Heimbach and Roger Wattenhofer. Eliminating sandwich attacks with the help of game theory. *arXiv preprint*, 2022.
- [42] Amir Herzberg, Stanislaw Jarecki, Hugo Krawczyk, and Moti Yung. Proactive secret sharing or: How to cope with perpetual leakage. In Don Coppersmith, editor, *CRYPTO’95*, volume 963 of *LNCS*, pages 339–352. Springer, Berlin, Heidelberg, August 1995.

- [43] Aljosha Judmayer, Nicholas Stifter, Philipp Schindler, and Edgar R. Weippl. Estimating (miner) extractable value is hard, let's go shopping! *IACR Cryptology ePrint Archive*, 2021.
- [44] Aniket Kate, Easwar Vivek Mangipudi, Pratyay Mukherjee, Hamza Saleem, and Sri Aravinda Krishnan Thyagarajan. Non-interactive VSS using class groups and application to DKG. *Cryptology ePrint Archive*, Report 2023/451, 2023.
- [45] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Berlin, Heidelberg, December 2010.
- [46] Jonathan Katz. Round-optimal, fully secure distributed key generation. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 285–316, Cham, 2024. Springer Nature Switzerland.
- [47] Alireza Kavousi, Duc V. Le, Philipp Jovanovic, and George Danezis. Blindperm: Efficient mev mitigation with an encrypted mempool and permutation. *Cryptology ePrint Archive*, Paper 2023/1061, 2023. <https://eprint.iacr.org/2023/1061>.
- [48] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. *IACR Cryptology ePrint Archive*, 2021.
- [49] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In *International Cryptology Conference*, 2020.
- [50] Dimitris Kolonelos, Giulio Malavolta, and Hoeteck Wee. Distributed broadcast encryption from bilinear groups. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 407–441. Springer, Singapore, December 2023.
- [51] Klaus Kursawe. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 2020.
- [52] Dahlia Malkhi and Pawel Szalachowski. Maximal extractable value (mev) protection on a dag, 2022.
- [53] Sai Krishna Deepak Maram, Fan Zhang, Lun Wang, Andrew Low, Yupeng Zhang, Ari Juels, and Dawn Song. CHURP: Dynamic-committee proactive secret sharing. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2369–2386. ACM Press, November 2019.
- [54] Conor McMenamin, Vanesa Daza, and Matthias Fitzl. Fairtradex: A decentralised exchange preventing value extraction. *arXiv preprint*, 2022.
- [55] Peyman Momeni, Sergey Gorbunov, and Bohan Zhang. Fairblock: Preventing blockchain front-running with minimal overheads. In *SecureComm*, volume 462 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 250–271. Springer, 2022.
- [56] Valeria Nikolaenko, Sam Ragsdale, Joseph Bonneau, and Dan Boneh. Powers-of-tau to the people: Decentralizing setup ceremonies. In Christina Pöpper and Lejla Batina, editors, *Applied Cryptography and Network Security*, pages 105–134, Cham, 2024. Springer Nature Switzerland.
- [57] Julien Piet, Vivek Nair, and Sanjay Subramanian. Mevade: An mev-resistant blockchain design. In *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9, 2023.
- [58] Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? *arXiv preprint*, 2021.
- [59] Antoine Rondelet and Quintus Kilbourn. Threshold encrypted mempools: Limitations and considerations, 2023.
- [60] Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 148–164. Springer, Berlin, Heidelberg, August 1999.
- [61] David A. Schultz, Barbara Liskov, and Moses Liskov. Mobile proactive secret sharing. In Rida A. Bazzi and Boaz Patt-Shamir, editors, *27th ACM PODC*, page 458. ACM, August 2008.
- [62] Shutter Network contributors. The shutter network. <https://shutter.network>, 2021.
- [63] Christof Ferreira Torres, Ramiro Camino, et al. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In *30th USENIX Security Symposium*, 2021.
- [64] Robin Vasantlall, Eduardo Alchieri, Bernardo Ferreira, and Alysson Bessani. COBRA: Dynamic proactive secret sharing for confidential BFT services. In *2022 IEEE Symposium on Security and Privacy*, pages 1335–1353. IEEE Computer Society Press, May 2022.
- [65] Ye Wang, Yan Chen, Shuiguang Deng, and Roger Wattenhofer. Cyclic arbitrage in decentralized exchange markets. *Available at SSRN 3834535*, 2021.

- [66] T.M. Wong, Chenxi Wang, and J.M. Wing. Verifiable secret redistribution for archive systems. In *First International IEEE Security in Storage Workshop, 2002. Proceedings.*, pages 94–105, 2002.
- [67] Lidong Zhou, Fred B. Schneider, and Robbert Van Renesse. Apss: proactive secret sharing in asynchronous systems. *ACM Trans. Inf. Syst. Secur.*, 8(3):259–286, aug 2005.
- [68] Liyi Zhou, Kaihua Qin, Antoine Cully, Benjamin Livshits, and Arthur Gervais. On the just-in-time discovery of profit-generating transactions in defi protocols. In *2021 IEEE Symposium on Security and Privacy (SP)*, 2021.
- [69] Liyi Zhou, Kaihua Qin, and Arthur Gervais. A2mm: Mitigating frontrunning, transaction reordering and consensus instability in decentralized exchanges. *arXiv preprint*, 2021.
- [70] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. High-frequency trading on decentralized on-chain exchanges. In *2021 IEEE Symposium on Security and Privacy (SP)*, 2021.

A Security Proof

We prove security with guaranteed output delivery against a fully malicious adversary corrupting up to $t < n/2$ members of the committee and any number of users creating ciphertexts. In the process, we introduce a new interactive assumption (Definition 2) and justify its hardness. We then prove Theorem 1 which shows that the protocol in Fig. 2 securely emulates the ideal functionality in Fig. 1.

We start by describing our new interactive assumption, that we term i-kzg, below.

Definition 2 (i-kzg Game). *Let g and h be the generators of \mathbb{G}_1 , and \mathbb{G}_2 , respectively. We define an oracle $\mathcal{O}^{a,t,\tilde{x},\{b_i\}_{i \in [m]}}(j, f(X))$ parametrized by $t, a, \tilde{x}, \{b_i\}_{i \in [m]} \in \mathbb{F}$. It takes as input a degree- $(B-1)$ polynomial $f(X)$ and an index j , and if $f(\tilde{x}) \neq 0$ it returns $(g^{b_j} - g^{f(t+\tilde{x})})^a$, else it returns \perp . Consider the following game for $B, m = \text{poly}(\lambda)$, where \mathcal{A} is a non-uniform probabilistic polynomial time algorithm:*

1. Sample a, t, u, v, \tilde{x} , and $\{b_i\}_{i \in [m]}$ from \mathbb{F} and send the following to \mathcal{A} :

$$(\tilde{x}, g, g^t, g^{t^2}, \dots, g^{t^{B-1}}, \{g^{b_i}\}_{i \in [m]}, h, h^t, h^a, h^v, h^u, h^{u+av})$$

2. For each $j \in [m]$, \mathcal{A} can make one query of the form

$$\mathcal{O}^{a,t,\tilde{x},\{b_i\}_{i \in [m]}}(j, f_j(X)),$$

for any degree- $(B-1)$ polynomial $f_j(X)$ of its choice.

3. \mathcal{A} wins if it outputs $e(g, h)^{b_j u}$ for some $j \in [m]$.

The i -kzg game is said to be hard if for all non-uniform probabilistic polynomial time algorithms \mathcal{A} , the probability that \mathcal{A} wins the game is negligible.

While the above interactive assumption closely resembles the security properties required from our scheme, we show below that the above assumption can be reduced to the standard (non-interactive) k -discrete log assumption in the algebraic group model (AGM). To do so, we follow the analysis of algebraic adversaries as presented in [9, 35].

We first prove a master theorem for the compilers in [9, 35] that will help simplify our analysis. We use the following notation for convenience. Let $\mathbf{X} \in \mathbb{G}^{u \times v}$, for some group \mathbb{G} and $e \in \mathbb{F}^n$. We will denote the i -th row by \mathbf{X}_i and $X_{i,j}$ to denote the entry in the j -th column of the i -th row. We will write \mathbf{X}_i^e to denote $\prod_{j \in [u]} X_{i,j}^{e_j}$. We will also use \circ to denote matrix multiplication between two matrices with elements in \mathbb{G}_1 and \mathbb{G}_2 , respectively, where the multiplication operation is replaced with a pairing and addition is replaced with the group operation. Similarly, we will use \cdot to denote multiplication between a matrix of field elements and a matrix of group elements where the multiplication operation is replaced by group exponentiation and addition is replaced by the group operation.

Theorem 3 (Master Theorem). *Let $\mathbf{L} \in \mathbb{G}_1^{u \times v_1}$, $\mathbf{R} \in \mathbb{G}_2^{v_2 \times u}$ and $\mathbf{b} \in \mathbb{G}_2^u$, define a system of pairing product equations, that are satisfied by $\mathbf{w}_L \in \mathbb{G}_2^{v_1}$, and $\mathbf{w}_R \in \mathbb{G}_1^{v_2}$ iff $\mathbf{L}_j \circ \mathbf{w}_L \cdot (\mathbf{R}^\top)_j \circ \mathbf{w}_R = \mathbf{b}_j$ for all $j \in [u]$. For all non-uniform PPT algebraic adversaries \mathcal{A} such that*

$$\Pr[\mathcal{A}(\mathbf{L}, \mathbf{R}, \mathbf{b}, \{(\mathbf{L}^\top)_i^r\}_{i \in [v_1]}, \{\mathbf{R}_i^r\}_{i \in [v_2]}) \rightarrow \mathbf{b}^r \mid r \leftarrow \mathbb{F}^u] = \epsilon,$$

there exists a uniform PPT algebraic adversary $\mathcal{A}'^{\mathcal{A}}$ which outputs $\mathbf{w}_L \in \mathbb{G}_2^{v_1}$, and $\mathbf{w}_R \in \mathbb{G}_1^{v_2}$ such that $\mathbf{L}_j \circ \mathbf{w}_L \cdot (\mathbf{R}^\top)_j \circ \mathbf{w}_R = \mathbf{b}_j$ for all $j \in [u]$, along with corresponding representations, with probability ϵ .

Proof. Let \mathcal{T} denote the completion set of target group elements viz. all possible pairing combinations of elements in \mathbb{G}_1 and \mathbb{G}_2 seen by \mathcal{A} so far. Since \mathcal{A} is an algebraic adversary, for any group element in the target group that it outputs, it also attaches the specific linear combination $\alpha \in \mathbb{F}^{|\mathcal{T}|}$ of elements in \mathcal{T} that it used to produce the output. The terms in \mathcal{T} can be categorized into three different types:

- Degree-0: Terms obtained by pairing an element from \mathbf{L} with an element from \mathbf{R} or a term from \mathbf{b} .
- Degree-1: Terms obtained by pairing an element from \mathbf{L} with an element from $\{\mathbf{R}_i^r\}_{i \in [v_2]}$ or an element from $\{(\mathbf{L}^\top)_i^r\}_{i \in [v_1]}$ with an element from \mathbf{R} .
- Degree-2: Terms obtained by pairing an element from $\{(\mathbf{L}^\top)_i^r\}_{i \in [v_1]}$ with an element from $\{\mathbf{R}_i^r\}_{i \in [v_2]}$.

Observe that the coefficients of Degree-0 and Degree-2 terms must be zero in the linear combination used to produce \mathbf{b}^r . If not, observe that we can solve for at least one of the r_i 's and \mathcal{A} can then be used to solve for r^* given $(g^{r^*}, h^{r^*}) \in \mathbb{G}_1 \times \mathbb{G}_2$ – a variant of the discrete logarithm problem.

We now describe the strategy of $\mathcal{A}'^{\mathcal{A}}$. Let \mathbf{R}_j denote the j -th row of \mathbf{R} . We can then define $S_j = \{s \mid e(s, \mathbf{R}_j^r) \in \mathcal{T}\}$, to be the terms where some element in \mathbb{G}_1 is paired with \mathbf{R}_j^r . We also define $\alpha_j \subseteq \alpha$ to be the corresponding coefficients of the elements in the set $\{e(s, \mathbf{R}_j^r) \mid s \in S_j\}$, in the linear combination used to produce \mathbf{b}^r . It is now easy to see that $\mathcal{A}'^{\mathcal{A}}$ can produce a satisfying \mathbf{w}_R by setting the j -th element of \mathbf{w}_R to be $\mathbf{w}_{R,j} \leftarrow S_j^{\alpha_j}$. Analogously for \mathbf{w}_L . \square

Definition 4 ((k) -dlog Assumption). *Let g , and h be the generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively. Then for all non-uniform PPT algorithm \mathcal{A} , it holds that*

$$\Pr[\mathcal{A}(g, g^a, g^{a^2}, \dots, g^{a^k}, h, h^k) \rightarrow k \mid a \leftarrow \mathbb{F}] \leq \text{negl}(\lambda).$$

Lemma 5. *The i -kzg game is hard in the algebraic group model provided the $(B-1)$ -dlog assumption holds.*

Proof. First observe that we can take any algebraic adversary \mathcal{A} which produces $e(g, h)^{b_j u}$ for some $j \in [m]$ and construct an adversary \mathcal{A}' which produces δ, π , and σ (and corresponding linear combinations of elements in \mathbb{G}_1) such that:

$$\begin{aligned} e(g^{b_j}, h) &= e(\pi, h^t) \cdot e(\delta, h) \\ e(\delta, h^a) &= e(\sigma, h) \end{aligned}$$

This can be achieved by applying the master theorem Theorem 3 with

$$\mathbf{R} = \begin{bmatrix} h & h^a \\ h^t & 0 \\ 0 & h \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 0 \\ e(g, h)^{b_j u} \end{bmatrix},$$

where a satisfying witness is of the form $\mathbf{w}_R = (\delta, \pi, \sigma^{-1})$. We can then redefine $\text{com} = g^{b_j} - \delta$ to obtain the triple $(\text{com}, \pi, \sigma)$ which satisfies:

$$\begin{aligned} e(\text{com}, h) &= e(\pi, h^t) \\ e(g^{b_j}, \text{pk}) &= e(\sigma, h) \cdot e(\text{com}, \text{pk}) \end{aligned}$$

Suppose there existed an adversary \mathcal{A} which succeeded in the i -kzg game with non-negligible probability. Then we will construct an adversary \mathcal{A}' which can solve the $(B-1)$ -dlog problem with non-negligible probability. We begin by describing $m+2$ types of strategies, and \mathcal{A}' will run one of these, chosen uniformly at random. We will then show that for all non-uniform PPT machines \mathcal{A} , at least one of these strategies will succeed, and hence \mathcal{A}' will solve the $(B-1)$ -dlog problem with non-negligible probability. Let the $(B-1)$ -dlog challenge be $(g, g^z, g^{z^2}, \dots, g^{z^{B-1}}, h, h^z)$.

1. Set $t := z$ and sample a, \tilde{x}, u, v , and b_i 's uniformly at random. \mathcal{A}' then runs \mathcal{A} with the challenge instance

$$(\tilde{x}, g, g^z, g^{z^2}, \dots, g^{z^{B-1}}, \{g^{b_i}\}_{i \in [m]}, h, h^z, h^a, h^v, h^{zu}, h^{u+av})$$

and can simulate responses from the oracle using knowledge of a .

2. Similarly, Set $a := z$, and sample remaining variables uniformly at random. Responses to oracle queries can be simulated using knowledge of b_i 's and t .
3. We describe a strategy for each $j \in [m]$: Set $b_j := z$ and sample remaining variables uniformly at random. Responses to oracle queries can be simulated using knowledge of a .

Since all witness variables are in \mathbb{G}_1 and hence expressed as a linear combination over

$$(g, g^t, g^{t^{B-1}}, \{g^{b_i}, (g^{b_i} - g^{f_i(t+\tilde{x})})^a\}_{i \in [m]}).$$

Let $Q = \{(i, f_i, (g^{b_i} - g^{f_i(t+\tilde{x})})^a)\}_{i \in [m]}$ denote the query-response pairs for queries made by the adversary to the oracle in the i -kzg game and \mathcal{A} outputs $(\text{com}, \pi, \sigma)$, along with their representation. Consider the following events:

1. We denote by E_1 , the event that com is expressed as a linear combination of powers-of- t $(g, g^t, \dots, g^{t^{B-1}})$ – and no other elements.
2. We denote by E_2 , the event that there exists $(j, f_j, \sigma) \in Q$ such that $\text{com} = g^{f_j(t)}$.

Case $\neg E_1$: Suppose \mathcal{A} returned a representation of com where the coefficients of $\{g^{b_i}, (g^{b_i} - g^{f_i(t+\tilde{x})})^a\}_{i \in [m]}$ are non-zero. Then we have coefficients $p_i, \beta_i, \tilde{\beta}_i, \gamma_i, \tilde{\gamma}_i$ such that:

$$\begin{aligned} \sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i + \sum_{i=1}^m \tilde{\beta}_i \cdot b_i t + \\ \sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t+\tilde{x}))a + \sum_{i=1}^m \tilde{\gamma}_i \cdot (b_i - f_i(t+\tilde{x}))at = 0. \end{aligned}$$

We will now use these coefficients to solve the $(B-1)$ -dlog problem. Suppose

$$\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t+\tilde{x}))a + \sum_{i=1}^m \tilde{\gamma}_i \cdot (b_i - f_i(t+\tilde{x}))at \neq 0$$

and. In the case where we set $a := z$, we can recover z as

$$z = -\frac{\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i + \sum_{i=1}^m \tilde{\beta}_i \cdot b_i t}{\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t+\tilde{x})) + \sum_{i=1}^m \tilde{\gamma}_i \cdot (b_i - f_i(t+\tilde{x}))t}$$

Thus, the following equations must hold:

$$\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m (\beta_i + \tilde{\beta}_i \cdot t) \cdot b_i = 0$$

$$\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x}))a + \sum_{i=1}^m \tilde{\gamma}_i \cdot (b_i - f_i(t + \tilde{x}))at = 0$$

Note that since E_1 is false, at least one of $\beta_i, \tilde{\beta}_i, \gamma_i, \tilde{\gamma}_i$ must be non-zero. There are now two cases to analyze:

1. If one of $\beta_i, \tilde{\beta}_i$ is non-zero, and $\sum_{i=1}^m (\beta_i + \tilde{\beta}_i \cdot t) \cdot b_i \neq 0$, then there must exist some $i^* \in [m]$ such that $(\beta_{i^*} + \tilde{\beta}_{i^*} \cdot t) \cdot b_{i^*} \neq 0$. We can then guess this i^* , uniformly at random, and set $b_{i^*} := z$. With non-negligible probability we will be correct, and we can solve for z as:

$$z = -\frac{\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i \neq i^*} (\beta_i + \tilde{\beta}_i \cdot t) \cdot b_i}{(\beta_{i^*} + \tilde{\beta}_{i^*} \cdot t)}$$

If instead $\sum_{i=1}^m (\beta_i + \tilde{\beta}_i \cdot t) \cdot b_i = 0$, but one of the $\beta_i, \tilde{\beta}_i$ is non-zero, then we can set $t := z$, and recover z as

$$z = -\frac{\sum_{i=1}^m \beta_i \cdot b_i}{\sum_{i=1}^m \tilde{\beta}_i \cdot b_i}$$

Hence, it must be true that all the $\beta_i, \tilde{\beta}_i$ are zero.

2. If one of $\gamma_i, \tilde{\gamma}_i$ is non-zero, and

$$\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x}))a + \sum_{i=1}^m \tilde{\gamma}_i \cdot (b_i - f_i(t + \tilde{x}))at = 0,$$

then we set $t := z$ and solve for z by finding the roots of

$$\sum_{i=1}^m (\gamma_i + \tilde{\gamma}_i \cdot z) \cdot (b_i - f_i(z + \tilde{x})) = 0.$$

We can then test all roots to see which one is consistent with the challenge instance. Hence, it must be true that all the $\gamma_i, \tilde{\gamma}_i$ are zero.

Thus, E_1 must be true.

Case $E_1 \wedge \neg E_2$: Since E_1 is true, com is a linear combination of $(g, g^t, \dots, g^{t^{B-1}})$. Let f' be the degree- $(B-1)$ polynomial that \mathcal{A} provided as a representation of com in the powers-of- t basis viz. $\text{com} = g^{f'(t)}$. We then define f to be the degree- $(B-1)$ polynomial such that $f'(t) = f(t + \tilde{x})$, effectively giving us a commitment in the powers-of- $(t + \tilde{x})$ basis. This implies $\sigma = g^{(b_j - f(t + \tilde{x}))a}$. Furthermore, since E_2 is false, $(j, f, \cdot) \notin \mathcal{Q}$.

Suppose the j -th query was never made by the adversary viz. $(j, \cdot, \cdot) \notin \mathcal{Q}$, we then have a representation of σ in

the basis of $(g, g^t, \dots, g^{t^{B-1}}, \{g^{b_i}\}_{i \in [m]}, \{g^{(b_i - f_i(t + \tilde{x}))}\}_{i \in [m]; i \neq j})$ which implies we have p_i 's, β_i 's and γ_i 's such that

$$\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i + \sum_{i=1, i \neq j}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x}))a = (b_j - f(t + \tilde{x}))a.$$

If $\sum_{i \neq j}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x})) - (b_j - f(t + \tilde{x})) \neq 0$, then in the case where we set $a := z$, we can solve for z as

$$z = \frac{\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i}{(b_j - f(t + \tilde{x})) - \sum_{i=1, i \neq j}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x}))}$$

We are then left with the following two equations:

$$\sum_{i \neq j}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x})) - (b_j - f(t + \tilde{x})) = 0$$

$$\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i = 0$$

In this case, if we had set $b_i := z$, then we can solve for z as

$$z = \sum_{i \neq j}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x})) + f(t + \tilde{x}).$$

Therefore, \mathcal{A} cannot win if it didn't make the j -th query and there exists $(j, f_j, \sigma_j) \in \mathcal{Q}$. There are now two cases:

1. $f(t + \tilde{x}) = f_j(t + \tilde{x})$. Corresponding to the case where $\sigma = \sigma_j$, but recall that there does not exist $(j, f, \cdot) \in \mathcal{Q}$, and hence $f \neq f_j$. In the case where we set $t := z$, we can solve for z by finding the roots of

$$f(z + \tilde{x}) - f_j(z + \tilde{x}) = 0.$$

2. $f(t + \tilde{x}) \neq f_j(t + \tilde{x})$. Corresponding to the case $\sigma \neq \sigma_j$. We have a representation of $g^{(b_j - f(t + \tilde{x}))a}$ in the basis of $(g, g^t, \dots, g^{t^{B-1}}, \{g^{b_i}\}_{i \in [m]}, \{g^{(b_i - f_i(t + \tilde{x}))a}\}_{i \in [m]})$.

$$\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i +$$

$$\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x}))a = (b_j - f(t + \tilde{x}))a.$$

If $\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x})) - (b_j - f(t + \tilde{x})) \neq 0$, then in the case where we set $a := z$, we can solve for z as

$$z = \frac{\sum_{i=0}^{B-1} p_i \cdot t^i + \sum_{i=1}^m \beta_i \cdot b_i}{(b_j - f(t + \tilde{x})) - \sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x}))}$$

On the other hand, suppose $\sum_{i=1}^m \gamma_i \cdot (b_i - f_i(t + \tilde{x})) - (b_j - f(t + \tilde{x})) = 0$. Note that the case where all γ_i 's are zero except for $\gamma_j = 1$ cannot occur as this would imply $f(t + \tilde{x}) = f_j(t + \tilde{x})$. In all other cases we could have set some $b_{i^*} := z$, where $\gamma_{i^*} \neq 0$, solved for z .

Hence, E_2 must be true if E_1 is true $\implies E_2$ is true.

Case $E_1 \wedge E_2$: Here \mathcal{A} outputs a $\text{com} = g^{f_j(t+\tilde{x})}$ that has the representation of a polynomial $f_j(X)$ that was queried, and used the same σ that it received. If we redefine $t = \tau + \tilde{x}$, we see that com is a KZG [45] polynomial commitment to f_j and the adversary output a proof π that corresponds to the opening of the commitment com at the point \tilde{x} to the value 0. Note that querying such a polynomial $f_j(X)$ is expressly forbidden by the oracle. Hence, \mathcal{A} violated the KZG binding property, which can then be reduced to solving $(B-1)$ -dlog as shown in [18]. \square

Theorem 1. *The protocol in Fig. 2 securely emulates the $\mathcal{F}_{\text{BTRKE}}$ ideal functionality (Fig. 1) in the dealer model against any static PPT adversary A corrupting up to $t < n/2$ parties. This holds under the assumption that a weakly simulation-extractable NIZK is available and that the i -kzg game (Definition 2) is hard in the programmable random oracle model.*

We begin by describing the simulator Sim , who simulated all interactions between the adversary and the honest parties using only the ideal functionality.

Sim remembers all queries made to the Random Oracle. Unless specified, if the query has not been previously made, it lazily samples the response and remembers it. If the query has been made previously, Sim returns the same response.

During the Setup phase, Sim plays the role of the trusted dealer as (i) uses the NIZK simulator to sample crs_{NIZK} ; (ii) samples elements of crs_{KZG} honestly; and (iii) samples random field elements as the corrupt parties' shares of the secret key and then sets sk to be a random value. This fully determines the honest parties' shares of the secret key, and hence $h_{j \in [n]}^{[\text{sk}]}$, which Sim publishes. Here, Sim also specifies Γ to be the distribution of an encryption of a random message, under this public key.

For convenience, we will maintain a list of previously seen receipts for honest party ciphertexts \mathcal{L} . Every time the simulator sees a receipt ct from an honest party, it updates the set as $\mathcal{L} \leftarrow \mathcal{L} \cup \{\text{ct}\}$. When simulating the Encryption phase, there are two cases:

- Every time an honest party encrypts a message, Sim receives a receipt ct (say) from the ideal functionality. If $\text{ct} \in \mathcal{L}$, Sim aborts. Else, Sim adds ct to \mathcal{L} and forwards ct to the adversary.
- For any ciphertext $\text{ct} = (\text{ct}^{(1)}, \text{ct}^{(2)}, \text{ct}^{(3)}, \text{ct}^{(4)}, S, \hat{x}, \Phi)$ created by corrupt parties, Sim verifies the NIZK proof Φ and discards the ciphertext if it fails. Next, it checks whether there exists some $\text{ct}' \in \mathcal{L}$ such that $H(\text{ct}.S) = H(\text{ct}'.S)$ and aborts if so. Otherwise, Sim runs the extractor and receives the randomness (α, β) used by the adversary to create the ciphertext. Sim then extracts the message as $M = \text{ct}^{(1)} \oplus H(e(H(\text{eid}) \cdot g^{-\text{tg}}, h)^\alpha)$.

Sim receives the entire batch of messages $\{M_1, \dots, M_B\}$ corresponding to the batch of ciphertexts $\{\text{ct}_1, \dots, \text{ct}_B\}$ that

are decrypted in epoch number eid . Now, Sim computes the degree- $(B-1)$ polynomial $p(X)$ such that $\{p(i) = \text{tg}_i\}_{i \in [B]}$, where $\text{tg}_i := H_F(\text{ct}_i.S)$. It then computes $\text{com} = g^{p(\tau)}$ using crs_{kzg} and sets $\delta = H(\text{eid})/\text{com}$. Note that the partial decryptions are fully determined as $\{\sigma_i := \delta^{[\text{sk}]_i}\}_{i \in [B]}$. Sim can compute partial decryptions that should have been sent by honest parties using the shares of the secret key it create during Setup. It then recovers $\sigma = \delta^{\text{sk}}$.

Finally, the simulator must program the random oracle to ensure that the simulated honest party ciphertexts remain consistent with the messages obtained from the ideal functionality. For every decrypted ciphertext ct_i that came as the receipt for an honest party's ciphertext, Sim computes the KZG opening proof π_i as described in Fig. 2 and computes $e(\pi, \text{ct}_i^{(2)}) \cdot e(\delta, \text{ct}_i^{(3)}) \cdot e(\sigma^{-1}, \text{ct}_i^{(4)})$. If this point has been previously queried to the random oracle, Sim aborts. Otherwise, it programs the random oracle output to be $M_i \oplus \text{ct}_i^{(1)}$.

We now describe a sequence of hybrids that are indistinguishable from each other starting from the batch-decryption protocol in Fig. 2 in the real world and ending in the ideal world where a simulator handles all interaction between the parties and the ideal functionality.

- H₀:** The real world where the adversary interacts with the honest parties who execute the protocol as described in Fig. 2.
 - H₁:** The simulator emulates the trusted dealer and executes the Setup sub protocol faithfully. For all Random Oracle queries made by the adversary, Sim queries the random oracle defined in the protocol and forwards all responses. This hybrid has an identical distribution to the previous hybrid.
 - H₂:** This hybrid is identical to the previous hybrid, except for one change. When a corrupt party creates a ciphertext, Sim first verifies the NIZK proof and discards the ciphertext if it fails. Additionally, it maintains a list \mathcal{L} of ciphertexts created by honest parties and checks if there exists some $\text{ct}' \in \mathcal{L}$ such that $H(\text{ct}.S) = H(\text{ct}'.S)$ and aborts if so. If both the above checks pass, then Sim runs the extractor to obtain the randomness (α, β) used by the adversary to create the ciphertext. Sim can then extract the message as $M = \text{ct}^{(1)} \oplus H(e(H_G(\text{eid}) \cdot g^{-\text{tg}}, h)^\alpha)$.
- We will now argue that Sim aborts with negligible probability. From the knowledge soundness property of the NIZK proof, if the proof verifies, then the statement must be true, and the adversary knows s such that $\text{ct}.S = g^s$, except with negligible probability. Furthermore, we have that $H_F(\text{ct}.S) = H_F(\text{ct}'.S)$. This implies that either (i) $\text{ct}.S = \text{ct}'.S$, or (ii) the adversary has found a collision for the random oracle, which occurs with negligible probability by definition. In the former case, we can use this adversary to solve the discrete logarithm problem.

The reduction embeds the challenge instance from the discrete logarithm game as the S component of some ciphertext created by honest parties. To generate the proof for this ciphertext, we can run SimProve . Since S is a randomly chosen group element in the batch-threshold encryption protocol, and from the weak-simulation extractability of the NIZK the distributions are computationally indistinguishable. With noticeable probability, the adversary will create a ciphertext with the same S as the challenge instance. By running the extractor, we can compute $s = \text{dlog}(S)$ and hence solve the discrete logarithm problem. Hence, $\text{ct}.S \neq \text{ct}'.S$ with overwhelming probability which implies Sim aborts with negligible probability. Thus, this hybrid is computationally indistinguishable from the previous hybrid.

H₃: In this hybrid, Sim begins to program the output of the random oracle as follows. For each eid , Sim samples a random field element b_{eid} and sets $H_G(\text{eid}) = g^{b_{\text{eid}}}$. For all other queries, Sim lazily samples a response to the random oracle query. This hybrid has an identical distribution to the previous hybrid.

H₄: This hybrid is identical to the previous hybrid, except Sim aborts when \mathcal{A} queries the random oracle on a *bad* point which allows it to break the semantic security of honest party ciphertexts.

More formally, for each ciphertext ct created by an honest party for epoch eid , let (α, β) denote the randomness used to create the ciphertext and define $\mathcal{P}_{\text{ct}} = e(H_G(\text{eid}) \cdot g^{-\text{tg}}, h)^\alpha$, where $\text{tg} = H_F(\text{ct}.S)$. Let $\mathcal{B}_{\text{eid}} = \{\text{ct}_1, \dots, \text{ct}_B\}$ denote the batch of ciphertexts that are decrypted in epoch number eid . If a batch decryption for that particular epoch has not occurred yet, then we set $\mathcal{B}_{\text{eid}} = \emptyset$. We can then define the set $\mathcal{P}_{\text{eid}} = \{\mathcal{P}_{\text{ct}} \mid \text{ct} \in \mathcal{L}_{\text{eid}} \setminus \mathcal{B}_{\text{eid}}\}$, where \mathcal{L}_{eid} is the set of honest party ciphertexts encrypted to epoch eid . If \mathcal{A} queries the random oracle at any point in \mathcal{P}_{eid} , Sim aborts. We will now argue that Sim aborts with negligible probability.

Claim 6. *If the i -kzg game (Definition 2) is hard, then for all PPT adversaries \mathcal{A} , and any epoch eid , $\Pr[\mathcal{A} \rightarrow p \mid p \in \mathcal{P}_{\text{eid}}] \leq \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} corrupting up to $t < n/2$ committee members and an arbitrary number of users. If \mathcal{A} queries a point $p \in \mathcal{P}_{\text{eid}}$, for any epoch eid with non-negligible probability, then we can use \mathcal{A} to construct Sim' that wins the i -kzg game with non-negligible probability. The high-level idea is to embed the instance from the i -kzg game in a randomly chosen ciphertext $\tilde{\text{ct}}$ created by an honest party. With some non-negligible probability, \mathcal{A} will query the point $\mathcal{P}_{\tilde{\text{ct}}}$ corresponding to $\tilde{\text{ct}}$, as there are a polynomial (in security parameter) number of epochs and ciphertexts. We

will then use $\mathcal{P}_{\tilde{\text{ct}}}$ to produce a winning output in the i -kzg game. We now describe Sim' .

Let the challenge instance in the i -kzg game be

$$(\tilde{x}, g, g^t, g^{t^2}, \dots, g^{t^B}, \{g_i^b\}_{i \in [m]}, h, h^t, h^a, h^v, h^{tu}, h^{u+av}),$$

with access to an oracle $\mathcal{O}^{a,t,\tilde{x},\{b_i\}_{i \in [m]}}$ and let \mathcal{C} be the set of committee members corrupted by the adversary.

During setup, Sim' sets $\tau = t + \tilde{x}$. It can then compute $\text{crs}_{\text{kzg}} = (g, g^\tau, \dots, g^{\tau^B}, h, h^\tau)$ using linear operations over $(g, g^t, g^{t^2}, \dots, g^{t^B}, h, h^t)$. Sim' sets the public key of the committee to be h^a , and samples random values $\{[a]_i\}_{i \in \mathcal{C}}$ for the shares of corrupt parties. Also note that h^a and $\{h^{[a]_i}\}_{i \in \mathcal{C}}$, can be used to compute $\{h^{[a]_i}\}_{i \in [n] \setminus \mathcal{C}}$ using linear interpolation such that $\{[a]_i\}_{i \in [n]}$ forms a (t, n) -secret sharing of a . Finally, Sim' runs the setup for NIZK and publishes $\{\text{crs}_{\text{kzg}}, \text{crs}_{\text{NIZK}}, \{h^{[a]_i}\}_{i \in [n]}\}$ and sends $\{[a]_i\}_{i \in \mathcal{C}}$ to the corrupt parties.

Let \mathcal{M} be the space of messages. Sim' samples $s \leftarrow \mathbb{F}$ and computes $S \leftarrow g^s$. It then sets $\tilde{\text{ct}}$ to be:

- $\tilde{\text{ct}}^{(1)} \leftarrow \mathbb{F}^{\mathcal{M}}$
- $\tilde{\text{ct}}^{(2)} := h^{tu}$
- $\tilde{\text{ct}}^{(3)} := h^{u+av}$
- $\tilde{\text{ct}}^{(4)} := h^v$
- S
- \tilde{x}
- $\tilde{\Phi} = \text{NIZK.SimProve}\{\alpha, \beta, s \mid S = g^s \wedge \text{ct}^{(1)} \wedge \text{ct}^{(2)} = h^{\alpha(\tau - \tilde{x})} \wedge \text{ct}^{(3)} = h^{\alpha \text{pk}^\beta} \wedge \text{ct}^{(4)} = h^\beta\}$

By defining $\alpha := u$, $\beta := v$, it is easy to see that the statement $\tilde{\Phi}$ attests to is indeed in the language, but Sim' does not know the witness values α, β and hence simulates the proof.

Let $C(\lambda)$ be an upper bound on the number of honest party ciphertexts the adversary sees, which depends on the run time of the adversary. Sim' samples $j \leftarrow [C(\lambda)]$ and replaces the j -th ciphertext created by an honest party with $\tilde{\text{ct}}$.

For random oracle queries of the form $H_G(\text{eid})$, where eid is the epoch id, Sim' sets $H_G(\text{eid}) = g^{b_{\text{eid}} + \tilde{\text{tg}}}$, where $\tilde{\text{tg}} \leftarrow H_F(\tilde{\text{ct}}.S)$. We choose m large enough, based on the run time of the adversary, to ensure Sim' can respond to all queries.

For each batch decryption, Sim' simulates partial decryptions as follows. Let $\mathcal{B}_{\text{eid}} = \{\text{ct}_1, \text{ct}_2, \dots, \text{ct}_B\}$ be the set of ciphertexts that are decrypted in epoch eid . If $\tilde{\text{ct}} \in \mathcal{B}_{\text{eid}}$, then Sim' aborts and restarts \mathcal{A} . Now, Sim' queries the oracle as $\mathcal{O}^{a,t,\tilde{x},\{b_i\}_{i \in [m]}}(\text{eid}, f_{\text{eid}}(X) - \tilde{\text{tg}})$, where $f_{\text{eid}}(X)$ is the degree- $(B-1)$ polynomial such

that $f_{\text{eid}}(\text{ct}_i.\hat{x}) = \text{tg}_i$, where $\text{tg}_i = H_F(\text{ct}_i.S)$. In response, it receives $\sigma := (H(\text{eid})/g^{f_{\text{eid}}(t+\hat{x})})^a = (g^{b_{\text{eid}}+\text{tg}_i-f_{\text{eid}}(\tau)})^a$. Note that it could have also received \perp if $f_{\text{eid}}(\hat{x}) = \text{tg}$, but this can only happen in two cases i) $\tilde{\text{ct}}$ was included in \mathcal{B}_{eid} or ii) the adversary created a ciphertext ct^* such that $H_F(\text{ct}^*.S) = H_F(\tilde{\text{ct}}.S)$. In the former case, Sim' would have already aborted and restarted \mathcal{A} . The latter case occurs with negligible probability, as argued in earlier hybrids. Hence, Sim' does not receive \perp , except with a negligible probability. Sim' can now simulate all partial decryptions in a straightforward manner using $\{[a]_i\}_{i \in C}$.

If \mathcal{A} terminates before Sim' has replaced a ciphertext, then Sim' restarts \mathcal{A} and tries again. If Sim' manages to replace a ciphertext before \mathcal{A} terminates, then Sim' randomly samples one of the queries made by \mathcal{A} to the random oracle and outputs it as the response to the i-kzg game. This is a winning output with non-negligible probability, as the adversary has queried a point in \mathcal{P}_{eid} with non-negligible probability. Thus, Sim' wins the i-kzg game with non-negligible probability, which contradicts the hardness of the i-kzg game. \square

From Claim 6 Sim aborts with negligible probability provided the i-kzg game is hard. Thus, this hybrid is computationally indistinguishable from the previous hybrid.

- H₅:** This hybrid is identical to the previous hybrid, except Sim now simulates *all* honest party ciphertexts as described in the previous hybrid. When an honest party's ciphertext $\tilde{\text{ct}}$ is included in the batch of ciphertexts decrypted in that epoch, Sim programs the random oracle output to be $\tilde{M} \oplus \tilde{\text{ct}}.\text{ct}^{(1)}$, where \tilde{M} is the message the honest party received as input. The distribution of *simulated* honest party ciphertexts $\tilde{\text{ct}}$ is identical to the distribution of *honestly* generated ciphertexts, provided the point $\mathcal{P}_{\tilde{\text{ct}}}$ is never queried by the adversary. Using a similar argument as in the previous hybrid, \mathcal{A} queries such points with negligible probability. Thus, this hybrid is computationally indistinguishable from the previous hybrid.
- H₆:** This hybrid is identical to the previous hybrid, except Sim computes the point to program using a different strategy. Recall that in previous hybrids Sim controlled honest parties and hence has access to the randomness used to create ciphertexts. In this hybrid, Sim uses its knowledge of the secret key to simulate partial decryptions by faithfully following the protocol. This hybrid has an identical distribution to the previous hybrid.
- H₇:** In the final hybrid, Sim does not control any honest parties and only handles interaction between the ideal functionality and the adversary. During setup Sim runs the

setup sub-protocol from Fig. 2 faithfully to, and specifies the distribution Γ to be the encryption of a randomly chosen message under the resulting public key. Instead of simulating honest party ciphertexts as in previous hybrids, it simply forwards the receipts from the ideal functionality to the adversary. During batch decryption, Sim receives the messages corresponding to the batch of ciphertexts. Here it programs the random oracle *after* receiving the messages as described in previous hybrids.