Efficiently-Thresholdizable Batched Identity Based Encryption, with Applications

Amit Agarwal*
University of Illinois
Urbana-Champaign (UIUC)
amita2@illinois.edu

Rex Fernando
Aptos Labs
rex1fernando@gmail.com

Benny Pinkas
Bar Ilan University
benny@pinkas.net

Abstract

We propose a new cryptographic primitive, "batched identity-based encryption" (Batched IBE), and its thresholdized version. The new primitive allows encrypting messages with specific identities and batch labels, where the latter can represent, for example, a block number on a blockchain. Given an arbitrary subset of identities for a particular batch, our primitive enables efficient issuance of a single decryption key that can be used to decrypt all ciphertexts having identities that are included in the subset while preserving the privacy of all other ciphertexts. At the heart of our construction is a new technique that enables public aggregation (i.e. without knowledge of any secrets) of any subset of identities, into a succinct digest. This digest is used to derive, via a master secret key, a single *succinct* decryption key for all identities that were digested in this batch. In a threshold system, where the master key is distributed as secret shares among multiple authorities, our method significantly reduces the communication (and in some cases, computation) of the authorities. It achieves this by making their costs for key issuance independent of the batch size.

We present a concrete instantiation of a Batched IBE scheme based on the KZG polynomial commitment scheme by Kate et al. (Asiacrypt'10) and a modified form of the BLS signature scheme by Boneh et al. (Asiacrypt'01). The construction is proven secure in the generic group model (GGM).

In a blockchain setting, the new construction can be used for achieving mempool privacy by encrypting transactions to a block, opening only the transactions included in a given block and hiding the transactions that are not included in it. With the thresholdized version, multiple authorities (validators) can collaboratively manage the decryption process. Other possible applications include scalable support via blockchain for fairness of dishonest majority MPC, and conditional batched threshold decryption that can be used for implementing secure Dutch auctions and privacy preserving options trading.

*Part of this work was done while the author was an intern at Aptos Labs.

Contents

1 Introduction

This paper studies problems related to efficient batch decryption of identity-based encryption (IBE), and particularly threshold batch decryption in a blockchain setting. The solutions we propose support batched threshold decryption of arbitrary subsets of ciphertexts, and can be used to provide a scalable support via blockchain for achieving *fairness* in dishonest majority MPC.

Both standard threshold encryption and threshold identity-based encryption have received much attention in the blockchain setting. For example, (standard) threshold encryption has been used to achieve mempool privacy [BO22] (further discussed below), and threshold IBE has been used to achieve "encryption to the future" [Cam+21; GMR23; Döt+23; Cer+23]. With encryption to the future, the idea is that if the validators of a particular chain (namely, the nodes responsible for running the chain) have shared an IBE master secret key, then during a block they can release decryption keys for the ID which is that block's number or the time on which that block is published. Anyone who wants to encrypt a message to be decrypted at a specific block in the future, or at a specific time, can use the IBE public key with the appropriate ID. However, a major drawback of this technique is that it decrypts *all* messages encrypted to a specific block, and does not enable to dynamically choose which of these messages to decrypt.

A possible solution to the aforementioned all-or-nothing decryption problem is to encrypt each message using a separate key. In that case, if there are B ciphertexts to be decrypted, and n holders of the secret key shares (i.e., the validators, in the blockchain setting), with a decryption threshold of $\Omega(n)$, standard threshold decryption requires these parties to do $O(nB)$ computation and communication (per party that needs to decrypt the message).

On particular setting where this overhead is especially problematic is that of achieving mempool privacy. The goal of mempool privacy stems from the way that transactions are submitted to blockchains. Before transactions are finalized, they are held in a *memory pool (mempool)*, which is publicly readable. That a mempool is public is inherent: it must be readable by all the validators so that they are able to build the next block, and the design of blockchains as permissionless networks allows for anyone to run a validator. The fact that the mempool is public and contains information on what transactions will be in future blocks makes it ripe for exploitation. Such exploitation is widespread, and has been termed *miner-extractable value*, or MEV.[1] The main technique for combating this type of abuse is to encrypt the transactions in the mempool, and to only decrypt the transactions in a block after the block has been finalized (see, e.g., [Kav+23] and references within). That way, although the mempool is still public, it is opaque. In addition, since

[1]This concept was introduced in [Dai+20] which explored transaction reordering and front-running in decentralized exchanges (DEXs) on the Ethereum blockchain. The term MEV refers to the additional profits validators can extract by reordering, censoring, or including transactions in a specific way within a block. This result was groundbreaking because it highlighted a significant and underexplored vulnerability in decentralized finance (DeFi) systems. This result has since been highly influential in research on blockchain economics, consensus protocols, and DeFi security.

a block has limited capacity and might not include all transactions submitted to it, it is crucial that the decryption process can choose to decrypt any subset of the transactions encrypted to the block, while keeping hidden the transactions outside of this subset. On the other hand, independent threshold decryption of each transaction incurs a total overhead of $O(nB)$ communication and computation, which will likely be too high for modern blockchains that are built to achieve high throughput and very low latency. Several recent works have studied this problem and have proposed solutions which we discuss in Section 1.3 and **??**.

1.1 Our results

We introduce and construct a new primitive, which we call **Batched Identity Based Encryption** (Batched IBE). This new notion solves the efficiency problem described above for threshold decryption, and also has several other interesting applications, in both the threshold and the non-threshold setting. Our new notion works as follows.

- As with standard IBE, encryptions are done with respect to some ID. In addition, an encryption also specifies a *batch label*.
- Any set of IDs, up to a pre-specified maximum batch size, can be *publicly aggregated* to produce a succinct digest.
- A succinct decryption key can be computed from this digest, a specified batch label, and a master secret key. This computation is done in *constant time* relative to the batch size. The key can then be used to decrypt any ciphertext that was encrypted with respect to any ID in the digest and the matching batch label.
- Optionally, the decryption key computation can be thresholdized.

We have the following contributions in this work.

- In Section 5, we formally define the notion of Batched Identity Based Encryption (Batched IBE).
- In Section 6, we provide an efficient construction for Batched IBE based on Type-3 pairings and prove its security in the Generic Group Model (GGM).
- In Section 7, we discuss different optimizations which can be used to speed up the digestion and decryption process in our basic construction and an efficient way to add non-malleability for mempool privacy application.
- In Section A, we define a threshold verison of Batched IBE, called Thresholdizable Batched IBE , and efficiently extend our non-threshold construction to the threshold version.
- In Section 8, we discuss the implementation of our scheme and analyze its concrete performance, both in the threshold and non-threshold setting.
- In Section 1.2 and Section 3, we provide many interesting applications of our primitive.

1.2 Cryptographic Applications

Our construction opens up a wide range of applications. We briefly describe a few of them here and provide more detailed descriptions in Section 3. The mempool privacy application was already discussed in the beginning of Section 1. Another application is for fair dishonest-majority MPC via blockchain [Cho+17]. In this context, we demonstrate how Batched IBE can efficiently implement the theoretical construction presented in [Cho+17] and replace general-purpose Witness Encryption. Notably, the decryption overhead of the servers in this construction involves only a few exponentiations and is independent of the number of MPC computations they support. More broadly, in Section 3 we outline a general framework for using our Batched IBE scheme as a means for "conditional" batched threshold decryption, with applications such as secure Dutch auction and options trading.

1.3 Related Works

Prior works have studied the aforementioned efficiency problem for threshold decryption, specifically in the context of mempool privacy, and have proposed solutions. We succinctly summarize these works and our distinction in Table 1. We now proceed to briefly describe these related works.

		Computation		Per batch
Scheme	Communication	Public	Private	setup needed ?
Threshold Encryption [ElG86; CG99; BO22]	$O(nB)$	$O(nB)$	$O(B)$	No
Choudhuri et. al. [Cho+24a]	$O(n)$	$O(B \log B) + O(n)$	$O(B \log B)$	Yes
This work	$O(n)$	$O(B \log B) + O(n)$	$O(1)$	No

Table 1: Comparison of the per-server communication and computation costs required for performing batched threshold decryption for a batch size B and n servers, represented in terms of the number of group/field elements and group/field operations respectively w.r.t a group/field size of $O(\lambda)$ bits where λ is the security parameter. For computation, we separately show the public and private computation cost where private refers to any local computation that is performed using a local secret state (e.g. secret key share). We also show whether the scheme requires a batch specific setup phase. This table only lists the approaches whose security guarantee matches with that of our scheme.

Until recently, there were two standard approaches. The first approach [BO22] independently encrypts and decrypts each transaction using a standard threshold encryption [CG99; ElG86]. As discussed in the introduction, this approach has the disadvantage of incurring $O(nB)$ communication cost between the validators (holding shares of the decryption key), i.e. the communication cost scales linearly with the batch size B.

The second approach [Cam+21; GMR23; Döt+23; Cer+23] is based on using threshold IBE [BF01] as a mechanism for "encryption to the future." That is, in each round, the validators collabora-

tively compute and publish a *single decryption key* corresponding to the current block number b, by setting the IBE ID to be b. This single key can be used to decrypt an arbitrary number of ciphertexts that have been encrypted towards this block. Thus, this solution has the benefit of incurring just $O(n)$ communication per validator (instead of $O(nB)$ as in the first approach). Notice, however, that this decryption is *all-or-nothing* and opens all ciphertexts with the corresponding ID b, whether or not they have actually been included in the block. As explained earlier, blocks have limited capacity, and if the mempool is large enough then many transactions in the mempool will get left out of any given next block. With this IBE approach, these transactions which get left out of a block completely lose their privacy. As such, threshold IBE solves the efficiency problem but yields a fundamentally weaker privacy guarantee in terms of mempool privacy.

Recently, [Cho+24a] attempts to provide a solution to resolve the shortcomings associated with the two standard approaches described above. Specifically, they end up with an *online* phase which has $O(n)$ communication per validator, i.e. independent of the batch size B, and allows arbitrary subsets of ciphertexts to be decrypted, but relies on an *expensive offline per-block interactive setup phase*. The first version of [SAS24] also achieves similar *syntactic* parameters, namely sublinear communication per validator and an expensive per-block setup phase, but does so with a weaker "time-lock" security definition: a ciphertext is only guaranteed to be hidden before the corresponding block decryption key is released, but there are no guarantees about its privacy afterwards, even if the block did not contain the ciphertext.

1.4 Concurrent Works

Concurrent and independent to our work, two other works, namely [Cho+24b] and [Bor+24], achieve similar results to ours. In addition, the work of [SAS24] has also been updated, replacing its original construction with one that is substantially more similar to ours.[2] In terms of presentation, all these works present their construction as a *batched threshold encryption* scheme, whereas we present it as a *batched identity-based encryption (IBE)* scheme (which is optionally thresholdizable). The works of [Cho+24b] and [Bor+24] focus on the specific application of mempool privacy, whereas we discuss a wider range of applications, such as fairness in dishonest majority MPC, secure dutch auctions, and privacy preserving options trading. The fact that our scheme is presented as an IBE aids in exploring these applications.

Comparison to [Cho+24b] and [SAS24]: The technical ideas in [Cho+24b] and the updated version of [SAS24] are quite similar to ours, except that those works provide a security proof in the Algebraic Group Model (AGM) whereas we use the Generic Group Model (GGM). In the GGM, we are able to prove a more adaptive version of security than the works of [Cho+24b; Sue24]. Both of these works rely on non-interactive zero-knowledge (NIZKs) to achieve ciphertext non-malleability (which is required for the mempool privacy application), whereas we are able to leverage our stronger security definition along with a standard signature scheme to achieve the same property without NIZKs (see Section 7.4). As discussed in Section 8 when comparing to the

[2] To avoid confusion, we note that the authors of [SAS24] have significantly revised their paper with a new title and a new construction, but have posted this revision to the same ePrint record (2024/762). The old title is "Extractable Witness Encryption for Signed Vector Digests from Pairings and Trust-Scalable One-Time Programs" and the new title is "Constant-Cost Batched Partial Decryption in Threshold Encryption."

earlier work of [Cho+24a] (which also uses NIZKs to achieve non-malleability), avoiding general-purpose NIZKs allows for nontrivial performance improvements when compared to constructions that require them.

Comparison to [Bor+24]: The construction in [Bor+24] is markedly different from ours (and [Cho+24b; SAS24]), and thus it offers some tradeoffs w.r.t our construction. We discuss these tradeoffs here. First, [Bor+24] allows the ciphertexts to use identities from only a polynomial sized identity space, whereas we can support an exponential sized identity space. In applications where the identities are selected randomly by different encryptors, it is essential that there is no collision among the identities. This is not an issue when the identity space is exponential and [Cho+24b; SAS24]) but does become an issue when the identity space is of polynomial size. To reduce the identity collision probability, [Bor+24] discusses two possible solutions: (1) Each ciphertext can be created w.r.t $O(\log \kappa)$ different random identities, where κ is a statistical security parameter (e.g. $\log \kappa = 40$), which ensures that there exists a perfect matching between the ciphertexts and identity space with overwhelming probability. (On the other hand, this solution does not address collisions caused by malicious clients who deliberately collide identities.) (2) Each ciphertext is created w.r.t a single random identity but decryption of ciphertexts is performed in smaller sub-batches (by splitting the batch of size B into α smaller disjoint sub-batches each of size B/α). In the threshold setting, this leads to a communication cost of $O(\alpha)$ between the authorities (holders of the shares of the master secret key) and the probability of collision reduces exponentially as α increases. Note that in our work, due to the exponential identity space, there is no such collision problem and each ciphertext is always created w.r.t a single identity and thus the communication cost between the authorities in the threshold setting is always $O(1)$ group elements.

Second, the size of the public parameter in [Bor+24] grows quadratically with the size of identity space whereas in our construction grows linearly with the batch size (which is typically much smaller than the size of identity space). Third, each ciphertext in our construction (and equivalently in [Cho+24b; SAS24]) is tied to not just an identity but also a batch label. This means that if $\mathsf{ct}_{\mathsf{id},t}$ is a ciphertext created w.r.t identity id and batch label t, then $\mathsf{ct}_{\mathsf{id},t}$ can only be decrypted during batch t and not in other batches. This can be seen as a shortcoming in some applications, for example mempool privacy, where one could hope to include the ciphertexts (representing encrypted transactions) in a future batch t' (representing a future block) if they were not included in the current batch/block t. In our construction, this would require generating two different independent ciphertexts, $\mathsf{ct}_{\mathsf{id},t}$ and $\mathsf{ct}_{\mathsf{id},t'}$, tied to batch label t and t' respectively. Interestingly, the construction in [Bor+24] does not have this caveat; each ciphertext ct is tied *only to* the identity id and not to the batch label. This means that a ciphertext $\mathsf{ct}_{\mathsf{id}}$, generated w.r.t identity id, can be decrypted in any batch which need not be decided at the time of creating the ciphertext.

2 Technical Overview

Standard IBE versus Batched IBE. In a standard Identity-Based Encryption (IBE) scheme, we have a universe $\mathcal{I}$ of public identities (for example, these could be email ids) and users can create a ciphertext ct w.r.t any $\mathsf{id} \in \mathcal{I}$ so that ct can be decrypted given the identity specific-secret key $\mathsf{sk}_{\mathsf{id}}$. These identity-specific secret keys $\mathsf{sk}_{\mathsf{id}}$ are typically issued by an authority (resp. a set of authorities) holding a master secret key msk (resp. shares of msk) which is used to derive the

identity-specific secret key $\mathsf{sk}_{\mathsf{id}}$.

We would like to extend this standard IBE to a batched setting where we have a pool of ciphertexts $\{\mathsf{ct}_1, \ldots, \mathsf{ct}_n\}$ and each ct_j is a ciphertext w.r.t. some identity $\mathsf{id}_j \in \mathcal{I}$. Given a *dynamically selected*[3] public batch of identities $S \subseteq \mathcal{I}$, we would like the authority to release a secret key sk_S which is *succinct* (i.e., independent of the size of set S), and which enables the decryption of all ciphertexts ct_j where $\mathsf{id}_j \in S$ while ensuring that all ciphertexts ct_j corresponding to $\mathsf{id}_j \notin S$ remain hidden.

A naive solution to achieve this would be to simply have sk_S be the set of standard IBE secret keys for the individual identities in the batch S, i.e, $\mathsf{sk}_S = \{\mathsf{sk}_j | \mathsf{id}_j \in S\}$. While this works, the secret key sk_S here has size proportional to $|S|$ which is not succinct.

Boneh-Franklin IBE. Our starting point is the (standard) IBE scheme of Boneh-Franklin [BF01], hereafter referred to as BF-IBE. Here, the master secret key msk is the signing key of a BLS signature scheme [BLS01] and $\mathsf{sk}_{\mathsf{id}}$ is simply a BLS signature on id using the signing key msk. Recall that a BLS signature scheme is defined on a pairing-friendly group $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ of prime order p with group operation $+ : \mathbb{G}_i \times \mathbb{G}_i \to \mathbb{G}_i$ and a pairing operation $\circ : \mathbb{G}_2 \times \mathbb{G}_1 \to \mathbb{G}_T$. For a group $\mathbb{G}_i$ with generator g_i, we will use the notation $[x]_i$ to represent the group element $x \cdot g_i$ in the group $\mathbb{G}_i$ where $x \in \mathbb{Z}_p$. The signer holds a signing key $\mathsf{msk} \in \mathbb{Z}_p$ and publishes a verification key $\mathsf{vk} = [\mathsf{msk}]_2$. The signature on a message $m \in \{0,1\}^*$ is simply $\sigma_m = \mathsf{msk} \cdot H(m) \in \mathbb{G}_1$ where $H : \{0,1\}^* \to \mathbb{G}_1$ is a hash function modeled as Random Oracle. To verify a claimed signature σ on a message m, the following pairing check is performed.

$$[1]_2 \circ \sigma = \mathsf{vk} \circ H(m)$$

Coming back to the BF-IBE construction, the authority holds a BLS signing key msk and uses it to derive an identity-specific secret key $\mathsf{sk}_{\mathsf{id}} := \sigma_{\mathsf{id}} = \mathsf{msk} \cdot H(\mathsf{id})$. To encrypt a message $m \in \mathbb{G}_T$ w.r.t. identity id, the encryptor samples a random $r \leftarrow \mathbb{Z}_p$ and produces the following ciphertext

$$\mathsf{ct} = (\mathsf{ct}_1, \mathsf{ct}_2) = (r \cdot [1]_2, r \cdot (\mathsf{vk} \circ H(\mathsf{id})) + m)$$

Given an identity secret key $\mathsf{sk}_{\mathsf{id}} := \sigma_{\mathsf{id}}$, the message can be recovered as $\mathsf{ct}_2 - (\mathsf{ct}_1 \circ \mathsf{sk}_{\mathsf{id}})$.

Extension to the batched setting using an accumulator. We would like to extend the above basic BF-IBE construction to the batched setting. As mentioned earlier, simply concatenating the individual $\mathsf{sk}_{\mathsf{id}}$ values of all ids in the batch $S \subseteq \mathcal{I}$ doesn't lead to a succinct key for the batch. To remedy this, we use a cryptographic accumulator scheme. Such a scheme enables compressing a set $S = \{s_1, \ldots, s_n\}$ of items into a *succinct* public digest d. Using the set S and digest d, it is possible to compute a short cryptographic proof π_s proving that a specific element s is contained in S. The verification algorithm, given the digest d, claimed element s and proof π_s, outputs a bit indicating either accept or reject. The completeness of the scheme ensures that correctly generated proofs π_s for $s \in S$ always pass the verification check, whereas soundness ensures that it is hard for a computationally bounded adversary to compute valid proofs π_s for $s \notin S$.

Given such an accumulator scheme, a natural approach is to create a succinct digest d for the public batch of identities $S = \{\mathsf{id}_1, \ldots, \mathsf{id}_n\}$, and then compute a succinct secret key for the batch by setting sk_S to be a BLS signature on the digest d, i.e. $\mathsf{sk}_S := \sigma_S = \mathsf{msk} \cdot H(d) \in \mathbb{G}_1$. Now,

[3] By "dynamic", we mean that the subset S of identities can be selected *after* the ciphertexts have been created.

one could hope to create an encryption scheme where a ciphertext ct, generated w.r.t. a specific identity id, is decryptable given a triple $(d, \pi_{\mathsf{id}}, \mathsf{sk}_S)$ as a witness, and if and only if the following two conditions hold: 1) π_{id} is a valid membership proof of id w.r.t. the digest d, 2) sk_S is a valid signature on the digest d.

Challenges and next steps. Although the above template conceptually works, it is not clear how to build an *efficient* encryption scheme satisfying the required properties and, in general, it seems to require a general purpose witness encryption [Gar+13; Gar+16; Tsa22; VWW22] which requires strong cryptographic assumptions and/or is often inefficient in practice.

We utilize the observation in [Gar+24] that the BF-IBE construction can be seen as a special case of a general technique which transforms a public linear constraint system, defined over some cryptographically hard, pairing-friendly groups, into an efficient witness encryption scheme. In such a constraint system, each constraint involves some public group elements (which are part of the statement) and some private group elements (which are part of the witness). The "linearity" condition requires that the witness group elements are never paired with each other in the constraint. To be specific, we can consider the following linear constraint system containing n constraints and m witness elements.

$$\begin{aligned} a_{1,1} \circ w_1 + \ldots + a_{1,m} \circ w_m &= b_1 \\ a_{2,1} \circ w_1 + \ldots + a_{2,m} \circ w_m &= b_2 \\ \ldots &= \ldots \\ a_{n,1} \circ w_1 + \ldots + a_{n,m} \circ w_m &= b_n \end{aligned}$$

where $\{a_{i,j}\}_{i\in[n],j\in[m]}$ and $\{b_i\}_{i\in[n]}$ are public group elements in $\mathbb{G}_2$ and $\mathbb{G}_T$ respectively whereas $\{w_i\}_{i\in[m]}$ are private witness elements in $\mathbb{G}_1$ (highlighted in grey). The above system of constraints can be succinctly expressed as

$$\mathbf{A} \circ \mathbf{w} = \mathbf{b}$$

where

$$\mathbf{A} = \begin{pmatrix} a_{1,1} & \ldots & a_{1,m} \\ \vdots & \vdots & \vdots \\ a_{n,1} & \ldots & a_{n,m} \end{pmatrix}, \qquad \mathbf{w} = \begin{pmatrix} w_1 \\ \vdots \\ w_m \end{pmatrix}, \qquad \mathbf{b} = \begin{pmatrix} b_1 \\ \vdots \\ b_n \end{pmatrix}$$

Given such a constraint system, a message $m \in \mathbb{G}_T$ can be witness-encrypted in the following way (analogous to BF-IBE encryption).

$$\mathsf{ct} = (\mathsf{ct}_{\mathbf{1}}, \mathsf{ct}_2) = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b} + m)$$

where $\mathbf{r} \leftarrow \mathbb{Z}_p^n$ is a randomly sampled column vector and $\mathbf{r}^T$ denotes its transpose. Intuitively, one could think of the term $\mathbf{r}^T \cdot \mathbf{b}$ as a one-time-pad that is applied to the message m.

Given a witness $\mathbf{w}$, the message can be recovered from ct by simply computing $\mathsf{ct}_2 - \mathsf{ct}_1 \circ \mathbf{w}$ (again analogous to BF-IBE decryption).

Recall from our template discussed earlier that in the context of batched IBE, we want to create an encryption scheme where a ciphertext ct, generated w.r.t. a specific identity id, is decryptable given a witness triple $\mathbf{w} = (d, \pi_{\mathsf{id}}, \mathsf{sk}_S)$, iff the following two conditions hold: 1) π_{id} is a valid membership proof of id w.r.t. the digest d, 2) sk_S is a valid signature on the digest d. Each of these two conditions will induce a constraint on the witness $\mathbf{w}$. Therefore, in order for us to utilize the aforementioned general technique of building an encryption scheme from a constraint system, we need to select our cryptographic ingredients, namely the accumulator and signature scheme, in a careful way so that the induced constraints are linear w.r.t. the witness $\mathbf{w}$.

Observation 1. Our first obervation is that the well-known KZG commitment scheme [KZG10] satisfies the required property of linear verification. Let d denote the digest created out of a set $S \subseteq \mathcal{I} = \mathbb{Z}_p$. Without going too much into the details of the KZG scheme, we note that the digest d is created by interpolating a univariate polynomial f whose roots are all the elements in set S, and evaluating it at a secret point τ "in the exponent". In other words, d is simply $[f(\tau)]_1$. The verification of a membership proof $\pi_{\mathsf{id}} \in \mathbb{G}_1$ w.r.t. $\mathsf{id} \in \mathcal{I}$ involves checking the following constraint.

$$[1]_2 \circ d = ([\tau]_2 - [\mathsf{id}]_2) \circ \pi_{\mathsf{id}}$$

where $[\tau]_2$ is a public group element generated during a one-time setup phase.

Observation 2. Unfortunately, the BLS signature scheme [BLS01], which we have been discussing so far in the context of BF-IBE, does not satisfy the desired property of linear verification in our context. To be more specific, while the BLS verification constraint is linear w.r.t. a signature (which is the sk_S part of our witness $\mathbf{w}$), it is not linear w.r.t. the message being signed (which is the digest d part of our witness $\mathbf{w}$). In our context, the unmodified BLS verification constraint would look as follows,

$$[1]_2 \circ \mathsf{sk}_S = \mathsf{vk} \circ H(d)$$

As we can see, this constraint involves applying the hash function H on the digest d, which is a highly *non-linear* operation! It turns out that by adding a slight modification to the BLS signature scheme, we can restore linearity to the constraint. Let $\alpha \leftarrow \mathbb{G}_1$ be a random group element whose discrete logarithm is unknown to any party. Then, the modified BLS scheme would sign a message $m \in \mathbb{G}_1$ as $\sigma_m := \mathsf{msk} \cdot (m + \alpha)$. The verification constraint for a claimed signature σ on a message $m \in \mathbb{G}_1$ would simply check whether $[1]_2 \circ \sigma = \mathsf{vk} \circ (m + \alpha)$ holds. Translating the notations to our context, where we need to sign the digest d to produce sk_S, the verification constraint would be,

$$[1]_2 \circ \mathsf{sk}_S = \mathsf{vk} \circ (d + \alpha)$$

We note that this modification to the BLS scheme does not come for free. Firstly, it requires α to be generated as part of the setup. Secondly, if α is reused for signing more than once, then it can be used to forge signatures[4]. Thirdly, there is concrete forgery attack against this scheme in

[4]Given signatures $\sigma_{m_1} := \mathsf{msk} \cdot (m_1 + \alpha)$ and $\sigma_{m_2} := \mathsf{msk} \cdot (m_2 + \alpha)$ on messages m_1 and m_2 respectively, one can easily get a valid signature $\sigma_{m_3} := 2\sigma_{m_1} - \sigma_{m_2} = \mathsf{msk} \cdot (2m_1 - m_2 + \alpha)$ on message $m_3 := (2m_1 - m_2)$

Type-1 and Type-2 pairing group (even when α is used only once)[5]. The first two limitations can be easily addressed by generating a fresh α as the output of a random oracle on a nonce (which could be the batch label in our context). To address the third limitation, we restrict ourselves to Type-3 pairing groups and prove the security of our overall scheme in the GGM.

Putting things together. Given these two ingredients (KZG commitments and modified BLS), we can now express our verification constraints in a form that is linear w.r.t. the witness $\mathbf{w} = (d, \pi_{\mathsf{id}}, \mathsf{sk}_S)$.

$$[1]_2 \circ d = ([\tau]_2 - [\mathsf{id}]_2) \circ \pi_{\mathsf{id}}$$
$$[1]_2 \circ \mathsf{sk}_S = \mathsf{vk} \circ (d + \alpha)$$

Rearranging the above constraint system, we get the following matrix form.

$$\underbrace{\begin{pmatrix} [1]_2 & [\mathsf{id}]_2 - [\tau]_2 & 0 \\ \mathsf{vk} & 0 & -[1]_2 \end{pmatrix}}_{\mathbf{A}} \circ \underbrace{\begin{pmatrix} d \\ \pi_{\mathsf{id}} \\ \mathsf{sk}_S \end{pmatrix}}_{\mathbf{w}} = \underbrace{\begin{pmatrix} [0]_T \\ -(\mathsf{vk} \circ \alpha) \end{pmatrix}}_{\mathbf{b}}$$

Given such a linear constraint system, a message m can be encrypted by sampling $\mathbf{r} \leftarrow \mathbb{Z}_p^2$ and computing

$$\mathsf{ct} = (\mathsf{ct}_{\mathbf{1}}, \mathsf{ct}_2) = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b} + m)$$

as described earlier. Given a witness $\mathbf{w}$, the message can be recovered from ct by simply computing $\mathsf{ct}_2 - \mathsf{ct}_1 \circ \mathbf{w}$.

3 Applications

3.1 Mempool privacy

As was described in Section 1, our result enables users to submit transactions to a specific block, where the details of each transaction are encrypted using its unique ID and a batch label equal to the block number. Once the validators agree on the transactions that will be included in the block, they aggregate the IDs of these transactions to produce a succinct digest. Given this digest, the validators compute a succinct decryption key for this block. The computation of this decryption key is the only procedure requiring access to a secret (namely, the master secret key), and is thus the only operation that must be computed by a threshold computation. Finally, given the succinct decryption key of this block, it is possible to decrypt all transactions that were included in the digest. All other transactions that were submitted to the block but were not included in it, remain hidden.

[5] In such groups, one can use $\mathsf{vk} = [\mathsf{msk}]_2$ to get $[\mathsf{msk}]_1 = \mathsf{msk} \cdot [1]_1$. Given a signature $\sigma_{m_1} := \mathsf{msk} \cdot (m_1 + \alpha)$ on message m_1, we can derive a signature $\sigma_{m_2} := \sigma_{m_1} + \mathsf{msk} \cdot [1]_1 = \mathsf{msk} \cdot (m_1 + 1 + \alpha)$ on message $m_2 = m_1 + 1$

As mentioned in [Cho+24a], the application of mempool privacy has a specific requirement that ciphertexts must satisfy a form of non-malleability. To get the required non-malleability property, they rely on the generic technique of using NIZK proofs for achieving CCA2-security. We observe that our Batched IBE scheme already satisfies a form of adaptive security and this can be leveraged to achieve non-malleability significantly more cheaply than [Cho+24a] by replacing NIZK proofs with standard digital signatures. We refer the readers to Section 7.4 for the formal details.

3.2 Fair Dishonest majority MPC via Blockchain

Secure Multi-party Computation (MPC) allows two or more parties to compute any public function over their privately-held inputs, without revealing any information beyond the result of the computation. An extremely desirable property of such a MPC protocol is *fairness*, namely ensuring that either all parties learn the output or no one does. Unfortunately, in a dishonest majority setting (where the adversary can actively corrupt more than half the number of parties), achieving fairness is impossible [Cle86] in general in the standard MPC model [Gol09]. Yet this property is crucial for applications like sealed-bid auctions and contract signing, where information asymmetry can be exploited by a malicious party.

An intriguing method introduced in [Cho+17] suggests achieving fairness through witness encryption (WE) and public bulletin boards such as blockchains. The parties execute an unfair MPC protocol, completely off-chain, to compute a WE ciphertext of the output, rather than the output itself. The construction ingeniously employs a "release token" as a witness for decrypting the WE ciphertext. First, each party generates a secret share of the release token, ensuring that the token value is shared between all parties. The statement used for encrypting the output with WE is designed so that its valid witness is a *proof* of publishing the full release token on a public bulletin board. This proof can correspond, for example, for a signature of the blockchain on all shares of the release token. In that case, to decrypt the output a party must have all shares of the release token be *published* on the blockchain. As a result, all shares are available to all parties, not only to the last party to provide its share.

While the result in [Cho+17] is theoretical and requires general-purpose WE which is prohibitively inefficient, we observe that a slightly modified version of their idea can be efficiently instantiated using our Batched IBE primitive.

- Instead of computing a WE ciphertext, the MPC computes an encryption to a specific ID that identifies this MPC session, and a specific block label.
- Additionally, before executing the MPC, the MPC participants deploy a smart contract that stores the public keys of all MPC participants. This smart contract will order the blockchain to release the output of the MPC computation only given a signature on the ID from each of the participants in this computation. (Essentially, the signature of participant i corresponds to the agreement of this participant to the publication of the output.)
- The blockchain follows the smart contract, and only if all participants provided their signatures, it includes this instance ID in the digest of IDs whose decryption is enabled.

This construction has several nice properties. First, it enables a blockchain to support fairness for an arbitrary number of MPC instances, while running between its validators a *single* threshold

computation whose overhead is independent of the number of MPC instances. This property is crucial for scaling, since blockchain validators typically already have a high load related to executing transactions and achieving consensus. Second, the validators do not need to decrypt the results of each MPC session. They merely compute (in constant time) the decryption key that enables the decryption of every MPC session for which all release approvals were given. Third, our specific construction of Batched IBE based on pairing-friendly groups is compatible with the popular SPDZ MPC framework [Dam+12] which is one of the most efficient MPC protocol for performing general-purpose (unfair) secure computation in a dishonest majority setting. In [SA19; Dal+20], it is shown that the SPDZ framework (which natively works over a field) can be efficiently extended to perform secure computation over groups. In the context of fair MPC application, the parties would be required to securely emulate via MPC the encryption procedure of the Batched IBE construction, which requires just 6 group exponentiations. Based on the results provided in [Dal+20], we estimate that the overhead of this computation would be less than 20 ms (resp. 500 ms) when using SPDZ MPC protocol in a LAN (resp. continent-wide WAN) setting. (These results correspond to a two-party computation, or to a 3-party computation with an honest majority. The addition to the MPC computation includes sampling a random group element k, encrypting it using our Batched IBE scheme, and encrypting the actual output y using k as a symmetric key. This final encryption can use k as a one-time pad, if the length of y is not longer than the length of k, or otherwise use AES encryption with k as the key. The overhead of computing an AES encryption should be marginal compared to the IBE encryption.)

3.3 Conditional Batched Threshold Decryption

More generally, our (thresholdizable) Batched IBE construction enables a form of conditional batch threshold decryption of ciphertexts.

- The fact the decryption is *conditional* enables to decide at the last minute which ciphertexts will be decrypted.
- The *threshold* property enables to distribute trust between multiple servers or validators. These servers check if the condition is met and, if so, enable decryption.
- The *batch* property enables scalability, since the threshold computation of the servers is independent of the number of ciphertexts that need to be decrypted.

Let us elaborate more on the "batch" and the "conditional decryption" properties.

The "batch label" notion. In its simplest form, the batch label can correspond to an event that progresses monotonically along a single dimension. The most obvious examples are batch labels corresponding to a block id or to the time. In particular, the latter option implements *time-lock encryption* (see [RSW96] and followup work). This implementation of time-lock encryption is more scalable than the notion based on moderately-hard computation, while trust becomes dependent on the assumption that the server (or a large enough number of the servers in a threshold setting) are not malicious.

A more complex batch label can correspond to a combination of multiple events in different dimensions. For example, it could correspond to the event "(the USD/EUR exchange rate is above 1 OR the GBP/JPY rate is below 200) AND the date is January 1, 2026". The servers responsible

for producing the decryption key for batches will only process the ciphertexts having a batch label that matches the specified event of interest.

Conditional decryption. The server, or group of servers sharing the master secret key, can decide to enable decryption for any *arbitrary subset* of the ciphertexts that have the same batch label. There are many examples where this subset cannot be predicted in advance. For instance, the fair-MPC application, or a Dutch auction application described below. As another example, consider the mempool privacy application where each encrypted transaction submitted to a block has a public maximal fee that its sender is willing to pay, and a public upper bound on the amount of gas that the transaction might consume. A block has limited capacity in terms of gas, and therefore the validators need to solve a knapsack problem in order to find which subset of submitted transactions will maximize their revenues. Their decision on which transactions to decrypt will be based on the solution to this problem.

Sample applications Let us further explore some applications that our construction can support.

- **Secure Dutch auctions on the blockchain**: A *Dutch auction* is an auction where the price starts high and gradually decreases, and the first bidder to accept the current price wins. This type of auction helps determine the market value by finding the highest acceptable price, with a process that is transparent and visible to all participants. We would like to implement this type of auction in a non-interactive manner, while hiding all bids except for the highest one(s).[6]

  Suppose that the price of the good for sale, i.e. the item being auctioned, has a price resolution of m values (say, $m = 1000$), and that bidder i wants to bid a price of $b_i \in [1, m]$. This bidder submits m encryptions with batch labels $1, \ldots, m$, where the encryption with label b_i is of a 1 value, and all encryptions with labels greater than b_i are of 0. (The encryptions of labels smaller than b_i can be arbitrary.) All bidders write a smart contract which begins with decrypting all encryptions with label m. If all these encryptions are of 0, then in the next block the encryptions with label $m - 1$ are decrypted, etc. This process stops when one of the decrypted values is 1. When this happens, the bidder (or bidders) who encrypted a 1 value for the current label are declared the winners and have to pay a price equal to the current label. No further values are decrypted. It is easy to verify that this process implements the Dutch auction and hides all bids except for the winning ones. In terms of latency, each price point is decrypted in a separate block, but given the existence of low-latency blockchains, such as as Aptos, Sui or Solana, with sub-second block latency, the overall run time of the auction can be sufficiently fast for many applications. As for efficiency, a single blockchain can support a very large number of auctions, since a single threshold computation by the validators enables to decrypt the bids of all relevant auctions, and the decryption itself can be done publicly and does not require threshold computation.

- **Options trading**: European options are a type of financial derivative that grants the holder the right, but not the obligation, to buy or sell an underlying asset at a predetermined price on a specified date, known as the expiration date. The defining feature of European options is that

[6] A Dutch auction is roughly the interactive equivalent of the (non-interactive) first-price sealed-bid auction. So a non-interactive implementation of a Dutch auction is also an implementation of a sealed-bid first-price auction.

they can only be exercised on the expiration date, not before. Our construction can be used to hide the terms of such options until the expiration date, and only reveal and execute an option if its holder chooses to execute it. More specifically, every option is encrypted with a batch label that is equal to its expiration date, and with an ID that identifies its holder. Before or at the expiration date, the holder of the option must send a signed execution command to authorize it. A digest of all options which were authorized to be executed is computed. Afterwards, the server (resp. set of servers) that has the master key (resp. shares of master secret key) publishes the corresponding decryption key that enables to decrypt and execute the authorized options while maintaining the privacy of remaining options.

4 Preliminaries

Notation We use λ to denote a computational security parameter, $[n]$ to represent the set of integers $\{1, \ldots, n\}$, $x \leftarrow S$ to denote that x is an element sampled uniformly at random from set S. We use bold-letters to indicate vectors and matrices. For a vector $\mathbf{v}$ of length n, we use the notation v_i to indicate the i^{th} co-ordinate of $\mathbf{v}$ where $i \in [n]$. By $\mathsf{poly}(\lambda)$ and $\mathsf{negl}(\lambda)$, we mean the class $\lambda^{O(1)}$ and $\frac{1}{\lambda^{\omega(1)}}$. Given a security parameter λ, we use PPT to denote probabilitic $\mathsf{poly}(\lambda)$-time Turing Machines with $\mathsf{poly}(\lambda)$-sized advice.

4.1 Bilinear Groups

We follow the notation used in [Gar+24], Section 3.1. A bilinear group, denoted as $\mathcal{BG}$, is a set of three groups $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ of prime order p, with a (non-degenerate) bilinear map or pairing, denoted as e. This map takes one element from $\mathbb{G}_1$ and one element from $\mathbb{G}_2$ and produces an element in $\mathbb{G}_T$. The groups $\mathbb{G}_1$ and $\mathbb{G}_2$ are referred to as source groups, while $\mathbb{G}_T$ is the target group. The groups $\mathbb{G}_1, \mathbb{G}_2$ have random generators g_1, g_2, and we use the notation $[x]_s$ to represent $x \cdot g_s$ in the group $\mathbb{G}_s$, for $s \in \{1, 2, T\}$, where $x \in \mathbb{Z}_p$ and the generator of $\mathbb{G}_T$ is $g_T = e(g_1, g_2)$. The group operation is additive, and therefore $[x]_s + [y]_s = [x + y]_s$.

We can represent the pairing operation $e([x]_1, [y]_2)$ as $[x]_1 \circ [y]_2 = [y]_2 \circ [x]_1 = [x \cdot y]_T$. (This notation makes it easier to write expressions which compose pairings with linear operations.) As a result, the operation $\circ$ is commutative and can be applied to vectors of equal length. For example, for $\mathbf{u} \in (\mathbb{G}_1)^n, \mathbf{v} \in (\mathbb{G}_2)^n$, where $\mathbf{u} = ([u_1]_1, \ldots, [u_n]_1)$ and $\mathbf{v} = ([v_1]_2, \ldots, [v_n]_2)$, we have that $\mathbf{u}^T \circ \mathbf{v} = [u_1 v_1 + \cdots u_n v_n]_T$. It is further possible to use this notation for matrix-vector multiplication. If $\mathbf{A} \in (\mathbb{G}_1)^{n \times m}$ and $\mathbf{b} \in (\mathbb{G}_2)^m$, then $\mathbf{A} \circ \mathbf{b}$ is the vector in $(\mathbb{G}_T)^n$ with the coordinates $(\mathbf{A}_1 \circ \mathbf{b}, \ldots, \mathbf{A}_n \circ \mathbf{b})$ where $\mathbf{A}_i$ is the i^{th} row vector of the matrix $\mathbf{A}$.

4.2 Generic group model (GGM)

At a high level, this model captures the class of 'generic' adversaries - adversaries that don't exploit concrete representations of the elements of the group and only perform generic group operations. This model is aimed to capture the possible *algebraic* attacks that an adversary can perform. The following description of Shoup's GGM [Sho97] is taken from [Zha22]. Let $p \in \mathbb{Z}$ be a positive integer, and let $S \subseteq \{0,1\}^*$ be a set of strings of cardinality at least p. We will assume an upper bound is known on the length of strings in S. An arbitrary group $\mathbb{G}$ of prime order p is generically modeled by sampling a *random* injection $L : \mathbb{Z}_p \to S$, which we will call the *labeling function*. We

will think of $L(x)$ as corresponding to (the string representation) of g^x, where $g \in \mathbb{G}$ is a fixed generator of the group $\mathbb{G}$. All parties are able to make the following queries:

- **Labeling queries.** The party submits $x \in \mathbb{Z}_p$, and receives $L(x)$.
- **Group operations.** The party submits $(\ell_1, \ell_2, a_1, a_2) \in S^2 \times \mathbb{Z}_p^2$. If there exists $x_1, x_2 \in \mathbb{Z}_p$ such that $L(x_1) = \ell_1$ and $L(x_2) = \ell_2$, then the party receives $L(a_1 x_1 + a_2 x_2)$. Otherwise, the party receives $\perp$.

We note that the labeling map L need not be explicitly materialized all at once. In fact, it is typical in security proofs to sample such a map in a lazy fashion as the queries are made. When performing such a lazy sampling, we will think of L as a "dictionary" and use $x \in L$ to mean the action of checking whether the element x exists as a "key" in L.

In a bilinear group $\mathcal{BG}$, the parties are equipped with labeling queries and group operations in $\mathbb{G}_i$, with a random labeling function $L_i : \mathbb{Z}_p \to S_i$ for $S_i \subseteq \{0,1\}^*$ and $i \in \{1, 2, T\}$. Additionally, the parties can make the following query for pairing operation [BBG05].

- **Pairing operation.** The party submits $(\ell_1, \ell_2) \in S_1 \times S_2$. If there exists $x_1, x_2 \in \mathbb{Z}_p$ such that $L_1(x_1) = \ell_1$ and $L_2(x_2) = \ell_2$, then the party receives $L_T(x_1 \cdot x_2)$. Otherwise, the party receives $\perp$.

The typical GGM does not capture an adversary's ability to "hash" arbitrary strings into the group To model this, we can extend the GGM with an appropriate hashing oracle as discussed in [LPS23; Bau+23]. Here, a hash function $H : \{0,1\}^* \to \mathbb{Z}_p$ is sampled at random and parties can make the following query.

- **Hash query.** The party submits $x \in \{0,1\}^*$ and receives $L(H(x))$.

Security proofs in the GGM. The security of a cryptographic scheme is usually defined by a game where two parties, namely the challenger and the adversary, interact as per some specification and the challenger finally outputs a bit b indicating the verdict of the game ($b = 1$ for win and $b = 0$ for lose). When writing security proofs in the GGM, the challenger is supposed to simulate the GGM oracle towards the adversary. Let $n \in \mathbb{Z}$ be a positive integer. The game usually involves the challenger sampling some secrets $x_1, \ldots, x_n \in \mathbb{Z}_p$ and leaking $L(f(x_1, \ldots, x_n))$ to the adversary for some f.

A standard technique for security proofs in the GGM, introduced in [Sho97], is the following. Instead of explicitly sampling secret $x_i \in \mathbb{Z}_p$, the challenger replaces it with an indeterminate X_i. Additionally, the domain of the labeling function is expanded from $\mathbb{Z}_p$ to $\mathbb{Z}_p[X_1, \ldots, X_n]$, the set of all n variate polynomials with coefficients in $\mathbb{Z}_p$. Now, instead of leaking $L(f(x_1, \ldots, x_n))$ to the adversary, the challenger leaks $L(f(X_1, \ldots, X_n))$. This model, where the challenger substitutes all the secrets with their corresponding indeterminates, is known as the "symbolic model". Proving security in this model boils down to showing independence between a target polynomial and the polynomials whose labels/encodings are present in the view of the adversary. We will use this technique in proving the security of our scheme.

4.3 Polynomial Commitment

A polynomial commitment enables the computation of a compact value, denoted as *com*, for a polynomial f that may have a high degree over a finite field $\mathbb{F}$. Subsequently, one can compute concise openings to prove that the polynomial committed to by *com* evaluates to some value β at a specific point α. The polynomial commitment must be binding, meaning it should be infeasible to open the same point to two distinct values. We refer the reader to [KZG10] for the detailed definition of polynomial commitments.

The KZG polynomial commitment scheme [KZG10] uses as public parameters powers of a secret point τ in the exponent of a group generator. In the group $\mathbb{G}_1$ used in our construction these would be the values $[\tau]_1, \ldots, [\tau^d]_1$, where d is an upper bound on the degree of the polynomial. These parameters are computed in a setup phase, and the value of τ is kept secret. Committing to a polynomial is done by evaluating it in the exponent at point τ, namely computing $[f(\tau)]_1$. A crucial property is that this computation can be done using the public powers of τ, but without knowledge of τ itself.

Theorem 4.1 ([KZG10; Chi+20]). *If the d-DLOG assumption holds with respect to parameter generation algorithm of the KZG commitment scheme described in [KZG10], then that commitment scheme is a correct and binding polynomial commitment scheme in the AGM, according to the definitions of [KZG10; Chi+20].*

4.4 Digital Signatures

A signature scheme consists of three algorithms $(\mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify})$:

- $\mathsf{Setup}(1^\lambda) \to (\mathsf{sk}, \mathsf{vk})$ is a randomized algorithm that takes as input a security parameter 1^λ and outputs a keypair $(\mathsf{sk}, \mathsf{vk})$.
- $\mathsf{Sign}(\mathsf{sk}, m) \to \sigma$ is a (potentially randomized) algorithm which takes as input a signing key sk and a message m, and outputs a signature σ.
- $\mathsf{Verify}(\mathsf{vk}, m, \sigma) \to 0$ or 1 is a deterministic algorithm which takes as input a verification key vk, a message m, and a signature σ, and either accepts or rejects.

In Section 7.4, we will require a signature scheme that satisfied strong existential unforgeability, which we define below:

Definition 4.2 (Strong existential unforgeability). A signature scheme $(\mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify})$ satisfies strong existential unforgeability if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Verify}(\mathsf{vk}, m^*, \sigma^*) = 1 \wedge (m^*, \sigma^*) \notin Q : \begin{array}{c}(\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{Setup}(1^\lambda) \\ (m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathsf{Sign}(\mathsf{vk}, \cdot)}(1^\lambda, \mathsf{vk})\end{array}\right] < \mathsf{negl}(\lambda),$$

where Q is the set of query-response pairs (m_i, σ_i) generated by the interaction between $\mathcal{A}$ and $\mathsf{Sign}(\mathsf{vk}, \cdot)$.

4.5 Collision-Resistant Hash Function Families

We define collision-resistance for a hash function family $\mathcal{H}$:

Definition 4.3 (Collision-resistance)**.** A hash function family $\mathcal{H}$: is collision-resistant if for all PPT adversaries $\mathcal{A}$ the following holds:

$$\Pr\left[H(x_1) = H(x_2) : \begin{array}{c} H \xleftarrow{\$} \mathcal{H}(1^\lambda) \\ (x_1, x_2) \leftarrow \mathcal{A}(1^\lambda, H) \end{array}\right] < \mathsf{negl}(\lambda).$$

5 Defining Batched Identity Based Encryption

5.1 Syntax

As in a standard IBE [BF01], a Batched IBE will include a Setup algorithm (with an additional parameter for batch size) which generates public parameters and KeyGen algorithm that generates a public key pk and a master secret key msk. Typically, KeyGen will be executed by a central authority which privately stores msk and publishes pk publicly. Using the pk, anyone can encrypt a message m w.r.t. a specific id, along with a batch label t, to produce a ciphertext. In contrast to a standard IBE where the authority issues identity specific secret keys, our Batched IBE will enables secret key issuance for a *batch* of identities. To capture this, we split the secret key derivation process into two parts: 1) A Digest algorithm that, given a batch of identities, produces a digest, 2) A ComputeKey algorithm that, given a digest and the batch label, produces a secret key sk. We note that only the ComputeKey algorithm uses the master secret key. The Digest algorithm only uses public information that is accessible to all participants. Finally, anyone holding the digest-batch label-specific key secret key sk can use the Decrypt algorithm to decrypt all ciphertexts whose identities were part of the digest.

Definition 5.1 (Batched IBE Syntax)**.** A Batched IBE scheme BIBE is specified by six algorithms: Setup, KeyGen, Encrypt, Decrypt, Digest, ComputeKey.

- $\mathsf{Setup}(1^\lambda, 1^B) \to \mathsf{params}$: A randomized algorithm that takes as input a security parameter $\lambda \in \mathbb{N}$ and a batch size $B = B(\lambda)$. It outputs params (system parameters) which includes a description of the message space $\mathcal{M}$, identity space $\mathcal{I}$, batch label space $\mathcal{T}$ and ciphertext space $\mathcal{C}$.
- $\mathsf{KeyGen}(\mathsf{params}) \to (\mathsf{msk}, \mathsf{pk})$: a randomized algorithm that takes as input params and outputs msk (master secret key) and pk (public key).
- $\mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t) \to c$: a randomized algorithm that takes as input a message $m \in \mathcal{M}$, an identity $\mathsf{id} \in \mathcal{I}$, a batch label $t \in \mathcal{T}$, public key pk and outputs a ciphertext $c \in \mathcal{C}$.
- $\mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}) \to d$: a deterministic algorithm that takes as input the public key pk and a list of identities $\mathsf{id}_1, \ldots, \mathsf{id}_B$ where each $\mathsf{id}_i \in \mathcal{I}$. It outputs a digest d.
- $\mathsf{ComputeKey}(\mathsf{msk}, d, t) \to \mathsf{sk}$: a deterministic algorithm that takes as input the master secret key msk, digest d, batch label t and outputs a digest-batch label-specific secret key sk.

- $\mathsf{Decrypt}(c, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, \mathsf{id}, t) \to m$: a deterministic algorithm that takes as input a ciphertext c, secret key sk, digest d, a list of identities $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and an identity-batch label pair (id, t). It outputs a message $m \in \mathcal{M}$.

Remark. We note that the syntax doesn't enforce any restrictions on how the batch list $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\}$ is selected as input for the Digest algorithm. We leave this choice to higher-level applications that use the BIBE primitive. One such application, namely mempool privacy, is discussed in Section 7.4. Other applications are discussed in Section 1.2.

Remark. Optionally, a Batched IBE scheme might include a batched version of Decrypt procedure, $\mathsf{BatchDecrypt}(\mathsf{CT}, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, t) \to \{m_1, \ldots, m_B\}$, which will be able to decrypt a set of ciphertexts CT, encrypted w.r.t ids from $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\}$, with efficiency better than naively invoking Decrypt on each ciphertext in the set CT. Our construction in Section 6 will have this feature as described in Section 7.3.

5.2 Correctness, Non-triviality and Security

The above algorithms should satisfy the following requirements.

Definition 5.2 (Batched IBE Correctness)**.** For all $\lambda \in \mathbb{N}, B \in \mathbb{N}, m \in \mathcal{M}, t \in \mathcal{T}, \mathsf{id} \in \mathcal{I}, S \subseteq \mathcal{I}$ s.t. $|S| = B$ and $\mathsf{id} \in S$, the following should hold:

$$\Pr\left[\mathsf{Decrypt}(c, \mathsf{sk}, d, S, \mathsf{id}, t) = m \;\middle|\; \begin{array}{c} \mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B) \\ (\mathsf{pk}, \mathsf{msk}) \leftarrow \mathsf{KeyGen}(\mathsf{params}) \\ c \leftarrow \mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t) \\ d \leftarrow \mathsf{Digest}(\mathsf{pk}, S) \\ \mathsf{sk} \leftarrow \mathsf{ComputeKey}(\mathsf{msk}, d, t) \end{array}\right] = 1$$

Definition 5.3 (Batched IBE Non-triviality/Efficiency)**.** We require that the running time of ComputeKey be independent of the batch size B (which implies that the digest d and sk are also independent of B)

Remark. We enforce the above requirement for the following reasons: 1) Without this requirement, one could come up with a trivial scheme using standard IBE where the secret key sk for a set of ids $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\}$ and batch label t is simply the standard IBE secret keys for identities $\{\mathsf{id}_1||t, \ldots, \mathsf{id}_B||t\}$. Here, the running time of ComputeKey and its output sk will be $O(B)$. 2) This feature, while already useful in a non-threshold setting, becomes even more useful in a threshold setting where the msk is split among multiple authority members (using a secret sharing scheme) and they securely emulate the execution of ComputeKey procedure using their share of msk to produce sk. In such a setting, the above requirement will ensure that the running time and communication cost of the secure emulation is independent of the batch size (which can be a huge cost saving in practice). We refer the readers to Section A for more details regarding the threshold setting.

Our definition of security is an adaptation of the standard definition of IBE by Boneh et. al. [BF01] and captures the fact that any ciphertext c created w.r.t. an identity id^* and batch label t^* remains hidden as long as at least one of the following two conditions hold: 1) The sk for batch with label t^* has not been released, 2) id^* is not included in the batch with label t^*. Moreover, we allow each batch label to be used only once as this is what our construction achieves and

suffices for many of the discussed applications (where batch label can be considered as a "round number")

Definition 5.4 (Batched Identity Based Encryption Security)**.** We define a security game $\mathsf{Expt}^{\mathsf{BIBE}}_{\mathcal{A},b}(1^\lambda, B)$ with respect to adversary $\mathcal{A}$ in the box below.

We say that a batched IBE scheme is secure if for all $B \in \mathbb{N}$, for all PPT adversaries $\mathcal{A}$ there exists some negligible function ϵ_A such that the following holds:

$$\left|\Pr[\mathsf{Expt}^{\mathsf{BIBE}}_{\mathcal{A},0}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{BIBE}}_{\mathcal{A},1}(1^\lambda, B) = 1]\right| < \epsilon_A(\lambda).$$

The security game $\mathsf{Expt}^{\mathsf{BIBE}}_{\mathcal{A},b}(1^\lambda, B)$.

Setup: The challenger takes as input the security parameter λ and the batch size B. It runs $\mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B)$, and then runs $(\mathsf{msk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(\mathsf{params})$. Finally, it sends $(\mathsf{params}, \mathsf{pk})$ to $\mathcal{A}$.

The rest of the game proceeds in rounds, as follows.

Pre-challenge queries: $\mathcal{A}$ may issue an arbitrary number of *key computation queries:*

- A sends a list ids of B identities along with a batch label t to the challenger.
- If a key computation query has already been made with batch label t, the challenger halts the game.
- Otherwise, the challenger does the following:
  - Compute a digest $d \leftarrow \mathsf{Digest}(\mathsf{pk}, \mathsf{ids})$ of the ids in ids using public key pk.
  - Compute a secret key $\mathsf{sk} \leftarrow \mathsf{ComputeKey}(\mathsf{msk}, d, t)$, using the digest d computed from the previous step.
  - Send sk to $\mathcal{A}$.

Challenge round: Once during the game, $\mathcal{A}$ may decide that the current round is the *challenge round.* The challenge round proceeds as follows:

- $\mathcal{A}$ sends two messages $m_0, m_1 \in \mathcal{M}$ and an identity-batch label pair (id^*, t^*) on which it wishes to be challenged.
- If key computation query (ids, t) has already been made with batch label $t = t^*$ and where $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger computes $c \leftarrow \mathsf{Encrypt}(\mathsf{pk}, m_b, \mathsf{id}^*, t^*)$ and sends c to $\mathcal{A}$.

Post-challenge queries: After the challenger round, $\mathcal{A}$ may again issue an arbitrary number of *key computation queries,* with the additional restriction that $\mathcal{A}$ cannot query (t^*, ids) with $\mathsf{id}^* \in \mathsf{ids}$:

- $\mathcal{A}$ sends a list ids of B identities along with a batch label t to the challenger.
- If a key computation query has already been made with batch label t **or if** $t = t^*$ **and** $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger does the following:
  - Compute a digest $d \leftarrow \mathsf{Digest}(\mathsf{pk}, \mathsf{ids})$ of the ids in ids using public key pk.
  - Compute a secret key $\mathsf{sk} \leftarrow \mathsf{ComputeKey}(\mathsf{msk}, d, t)$, using the digest d computed from the previous step.
  - Send sk to $\mathcal{A}$.

Output: At any point in time, $\mathcal{A}$ can decide to halt and output a bit $b' \in \{0, 1\}$. The game then halts with the same output b'.

6 Our Batched Identity Based Encryption construction

6.1 Construction

In this section, we will provide a construction of BIBE scheme based on Kate et. al. [KZG10] polynomial commitment scheme and the modified BLS signature scheme that we discussed in Section 2. We show the formal construction in Fig. 1. The key details of the construction are as follows. The construction is in a Type-3 asymmetric pairing setting, with groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of order p, and a pairing operation $e(\mathbb{G}_1, \mathbb{G}_2) \to \mathbb{G}_T$. The message space is $\mathbb{G}_T$ and the identity space is $\mathbb{Z}_p$. The construction also uses a hash function H whose range is $\mathbb{G}_1$.

Key generation includes choosing a random master secret key $\mathsf{msk} \leftarrow \mathbb{Z}_p$ and a random secret parameter $\tau \leftarrow \mathbb{Z}_p$. The secret key is msk, while the public key includes a powers-of-Tau setting of order B in $\mathbb{G}_1$, namely $([\tau]_1, \ldots, [\tau^B]_1$, as well as the following two elements in $\mathbb{G}_2$, $[\tau]_2$ and $[\mathsf{msk}]_2$, corresponding to raising the generator $g_2 \in \mathbb{G}_2$ to the powers of τ and msk.

The encryption of a message m with identity id to a batch label t, includes computing a matrix $\mathbf{A} \in (\mathbb{G}_2)^{2\times 3}$ and a vector $\mathbf{b} \in (\mathbb{G}_T)^2$. In addition to being dependent on the public key, $\mathbf{A}$ depends on the identity id, whereas $\mathbf{b}$ depends on the batch label (and therefore $\mathbf{b}$ is identical for all messages encrypted to this batch label). The encryption is $c = (c_1, c_2) = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b} + m)$, where $\mathbf{r}$ is a random column vector in $(\mathbb{Z}_p)^2$. Note that c_1 is a vector of dimension 3 in $\mathbb{G}_2$, and c_2 is an element of $\mathbb{G}_T$.

The digest algorithm takes B identities and then interpolates a polynomial of degree B in $\mathbb{Z}_p$ whose roots are these identities, and whose leading coefficient is 1. The digest d is the KZG commitment of this polynomial, namely the value of this polynomial at the point $[\tau]_1$.

The compute-key algorithm outputs the secret key $\mathsf{sk} := \mathsf{msk} \cdot (d + H(t)) \in \mathbb{G}_1$. This is the only algorithm that uses the master secret key msk.

Decryption is computed independently for every ciphertext c. The decryption of a message with identity id involves interpolating a polynomial which has roots in all identities in the digest, *except* for id, and computing a KZG opening proof π using this polynomial. The ciphertext is parsed as (c_1, c_2), and decrypted by computing message $c_2 - c_1 \circ (d, \pi, \mathsf{sk})^T$.

Theorem 6.1. *Assuming Type-3 pairing group $\mathcal{BG}$, there exists a construction (Fig. 1) for Batched IBE*

BIBE **construction**

- $\mathsf{Setup}(1^\lambda, 1^B)$: Output three groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of order p, where p is a λ-bit prime, equipped with generators g_1, g_2, g_T, respectively, and an efficiently computable pairing operation $\circ : \mathbb{G}_2 \times \mathbb{G}_1 \to \mathbb{G}_T$. Set the message space $\mathcal{M} := \mathbb{G}_T$, identity space $\mathcal{I} := \{0, \ldots, p-1\}$, and batch label space $\mathcal{T} := \{0,1\}^\lambda$. Also output a randomly sampled hash function $H : \mathcal{T} \to \mathbb{G}_1$.

- $\mathsf{KeyGen}(\mathsf{params})$: Sample $\mathsf{msk} \leftarrow \mathbb{Z}_p$ and $\tau \leftarrow \mathbb{Z}_p$. Output $\mathsf{msk}, \mathsf{pk} := ([\tau]_1, \ldots, [\tau^B]_1, [\tau]_2, [\mathsf{msk}]_2)$.

- $\mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t)$: Let $\mathbf{A}$ be a matrix in $(\mathbb{G}_2)^{2\times 3}$ and $\mathbf{b}$ be a vector in $(\mathbb{G}_T)^2$, defined as follows.

$$\mathbf{A} := \begin{pmatrix} [1]_2 & [\mathsf{id}]_2 - [\tau]_2 & 0 \\ [\mathsf{msk}]_2 & 0 & -[1]_2 \end{pmatrix}$$

$$\mathbf{b} := \begin{pmatrix} [0]_T \\ -([\mathsf{msk}]_2 \circ H(t)) \end{pmatrix}$$

  Sample a (column) vector $\mathbf{r} = (r_1, r_2) \leftarrow (\mathbb{Z}_p)^2$ and output the ciphertext c where

$$c = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b} + m)$$

- $\mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\})$: Let $f(X) = \sum_{i=0}^{B} f_i \cdot X^i$ be a univariate polynomial of degree B over $\mathbb{Z}_p$ with roots at $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and leading coefficient 1. Output digest $d := \sum_{i=0}^{B} f_i \cdot [\tau^i]_1$.

- $\mathsf{ComputeKey}(\mathsf{msk}, d, t)$: Output the secret key $\mathsf{sk} := \mathsf{msk} \cdot (d + H(t))$.

- $\mathsf{Decrypt}(c, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, \mathsf{id}, t)$: Let $q(X) = \sum_{i=0}^{B-1} q_i \cdot X^i$ be a univariate polynomial of degree $B-1$ with roots at $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\} \setminus \{\mathsf{id}\}$ and leading coefficient 1. Set $\pi := \sum_{i=0}^{B-1} q_i \cdot [\tau^i]_1$ and set $\mathbf{w}$ to be the following vector.

$$\mathbf{w} = \begin{pmatrix} d \\ \pi \\ \mathsf{sk} \end{pmatrix}$$

  Finally, parse c as (c_1, c_2) and output the decrypted message $m^* := c_2 - c_1 \circ \mathbf{w}$.

Figure 1: Our construction for Batched Identity Based Encryption (BIBE)

which is secure in the Generic group model.

6.2 Analysis

6.2.1 Efficiency

Parameter	Size
Public parameters size	$B\|\mathbb{G}_1\| + 2\|\mathbb{G}_2\|$
Ciphertext size	$3\|\mathbb{G}_2\| + \|\mathbb{G}_T\|$
Digest size	$\|\mathbb{G}_1\|$
Decryption key size	$\|\mathbb{G}_1\|$

Table 2: Parameter sizes for our Batched IBE construction (Fig. 1).

We discuss the computation cost of each of the algorithms in Fig. 1. Setup requires $O(\lambda)$ operations. KeyGen requires $O(B)$ group exponentiations. Encrypt requires $O(1)$ group exponentiations. Digest and Decrypt require $O(B \log B)$ field multiplications (via DFT) and $O(B)$ group exponentiations. In addition, Decrypt requires a single multipairing of size 3. ComputeKey requires $O(1)$ group operations (independent of batch size B). We summarize the concrete parameter sizes below in Table 2.

6.2.2 Correctness

The correctness of the scheme is straight forward. Given a ciphertext $c = (c_1, c_2)$, and a witness $\mathbf{w}$ generated as per the specified algorithms, we have the following decrypted message.

$$
\begin{aligned}
m^* &= c_2 - c_1 \circ \mathbf{w} \\
&= \mathbf{r}^T \cdot \mathbf{b} + m - (\mathbf{r^T} \cdot \mathbf{A}) \circ \mathbf{w} \\
&= -r_2([\mathsf{msk}]_2 \circ H(t)) + m - \Big([r_1 + r_2 \cdot \mathsf{msk}]_2, [r_1 \cdot \mathsf{id}]_2 - [r_1 \cdot \tau]_2, -[r_2]_2\Big) \circ \begin{pmatrix} d \\ \pi \\ \mathsf{sk} \end{pmatrix} \\
&= -r_2([\mathsf{msk}]_2 \circ H(t)) + m - \Big([r_1 + r_2 \cdot \mathsf{msk}]_2, [r_1 \cdot \mathsf{id}]_2 - [r_1 \cdot \tau]_2, -[r_2]_2\Big) \\
&\circ \begin{pmatrix} [\Pi_{i=1}^{B}(\tau - \mathsf{id}_i)]_1 \\ [\Pi_{i=1, \mathsf{id}_i \neq \mathsf{id}}^{B}(\tau - \mathsf{id}_i)]_1 \\ \mathsf{msk} \cdot ([\Pi_{i=1}^{B}(\tau - \mathsf{id}_i)]_1 + H(t)) \end{pmatrix} \\
&= -r_2([\mathsf{msk}]_2 \circ H(t)) + m - [r_1 \cdot \Pi_{i=1}^{B}(\tau - \mathsf{id}_i) + r_2 \cdot \mathsf{msk} \cdot \Pi_{i=1}^{B}(\tau - \mathsf{id}_i) \\
&- r_1 \cdot \Pi_{i=1}^{B}(\tau - \mathsf{id}_i) - r_2 \cdot \mathsf{msk} \cdot \Pi_{i=1}^{B}(\tau - \mathsf{id}_i)]_T + [r_2]_2 \circ \mathsf{msk} \cdot H(t) \\
&= m
\end{aligned}
$$

6.2.3 Security

We will prove the security of our scheme in the GGM model equipped with hash queries. In this model, the challenger will implement the group oracle and the hash oracle (along with an oracle for key computation queries as defined in the security game $\mathsf{Expt}^{\mathsf{BIBE}}$ earlier).

Theorem 6.2. *For all $B \in \mathbb{N}$ and all unbounded adversaries $\mathcal{A}$ making at most q queries (including queries to the group oracle, hash oracle, and key computation queries), we have:*

$$\left|\Pr[\mathsf{Expt}^{\mathsf{BIBE,GGM}}_{\mathcal{A},0}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{BIBE,GGM}}_{\mathcal{A},1}(1^\lambda, B) = 1]\right| \leq 2\binom{q+B+5}{2}\frac{(B+2)}{p}$$

where $\mathsf{Expt}^{\mathsf{BIBE,GGM}}$ refers to the same experiment $\mathsf{Expt}^{\mathsf{BIBE}}$ as defined in Thm. 5.4 except that we specialize it for our specific Construction (Fig. 1) and model it in the GGM, i.e., all the group and hash operations performed by the adversary are simulated by the challenger as defined in the GGM model (Section 4.2). We formally define the experiment below.

The security game $\mathsf{Expt}^{\mathsf{BIBE,GGM}}_{\mathcal{A},b}(1^\lambda, B)$.

Setup: The challenger takes as input the security parameter λ and the batch size B. It runs $\mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B)$. Let $S_1 = S_2 = S_T = \{0,1\}^{p'}$ where $2^{p'} \geq p$ and p is λ bit prime (denoting the group order) contained in params. Then it initializes the maps including the labeling function $L_i : \mathbb{Z}_p \to S$ for each $i \in \{1, 2, T\}$ and the hash function $H : \{0,1\}^* \to \mathbb{Z}_p$. During the KeyGen phase, it performs the following steps:

- Sample $e_{\mathsf{msk}} \leftarrow S_2$, $\mathsf{msk} \leftarrow \mathbb{Z}_p$. Set $L_2(\mathsf{msk}) := e_{\mathsf{msk}}$ and $S_2 := S_2 \setminus \{e_{\mathsf{msk}}\}$.
- Sample $\tau \leftarrow \mathbb{Z}_p$.
- For $i \in [B-1]$, sample $e_{\tau^i} \leftarrow S_1$ and set $L_1(\tau^i) := e_{\tau^i}$ and $S_1 := S_1 \setminus \{e_{\tau^i}\}$.
- Sample $e' \leftarrow S_2$ and set $L_2(\tau) = e'$ and $S_2 := S_2 \setminus \{e'\}$.

It sets $\mathsf{pk} = (L_1(\tau), \ldots, L_1(\tau^B), L_2(\tau), L_2(\mathsf{msk}))$ and sends $(\mathsf{params}, \mathsf{pk})$ to $\mathcal{A}$.

The rest of the game proceeds in rounds, as follows.

Pre-challenge queries: $\mathcal{A}$ may issue an arbitrary number of the following types of queries:

Labeling query: For each labeling query for group $\mathbb{G}_i$, the challenger receives a value $v \in \mathbb{Z}_p$ from $\mathcal{A}$. If $v \notin L_i$, it samples $z \leftarrow S_i$, sets $L_i(v) := z$ and $S_i := S_i \setminus \{z\}$. It sends $L_i(v)$ to $\mathcal{A}$.

Group operation: For each group operation query for group $\mathbb{G}_i$, the challenger receives $(\ell_1, \ell_2, a_1, a_2) \in (\{0,1\}^{p'})^2 \times \mathbb{Z}_p^2$. If there doesn't exists $x_1, x_2 \in \mathbb{Z}_p$ such that $L_i(x_1) = \ell_1$ and $L_i(x_2) = \ell_2$, then send $\perp$ to $\mathcal{A}$. Otherwise, execute an internal labeling query step on group $\mathbb{G}_i$ with input $x_3 := a_1 x_1 + a_2 x_2$ and send $L_i(x_3)$ to $\mathcal{A}$.

Pairing operation: The challenger receives $(\ell_1, \ell_2) \in \{0,1\}^{p'} \times \{0,1\}^{p'}$. If there doesn't exist $x_1, x_2 \in \mathbb{Z}_p$ such that $L_1(x_1) = \ell_1$ and $L_2(x_2) = \ell_2$, then the challenger sends $\perp$ to $\mathcal{A}$. Otherwise, it computes $x_3 = x_1 \cdot x_2$. It executes an internal labeling query step for group $\mathbb{G}_T$ with input x_3 and sends $L_T(x_3)$ to $\mathcal{A}$.

Hash query: For each hash query, the challenger receives a string $s \in \{0,1\}^*$ from $\mathcal{A}$. If $s \notin H$, the challenger samples $z \leftarrow \mathbb{Z}_p$. It sets $H(s) := z$ and executes an internal labeling query step for group $\mathbb{G}_1$ with input z. It sends $L_1(z)$ to $\mathcal{A}$.

Key computation query:

- $\mathcal{A}$ sends a list ids of B identities, $S = \{\mathsf{id}_1, \ldots, \mathsf{id}_B\} \subseteq \mathcal{I}$, along with a batch label $t \in \mathcal{T}$ to the challenger.
- If a key computation query has already been made with batch label t, the challenger halts the game.
- Otherwise, the challenger does the following:
  - Let $f(X) = \sum_{i=0}^{B} f_i \cdot X^i$ be a univariate polynomial of degree B over $\mathbb{Z}_p$ with roots at $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and leading coefficient 1. If $f(\tau) \notin L_1$, then execute an internal labeling query step for group $\mathbb{G}_1$ with input $f(\tau)$.
  - If $t \notin H$, execute an internal hash query step with input t.
  - Let $P := (f(\tau) + H(t)) \cdot \mathsf{msk}$. If $P \notin L_1$, then execute an internal labeling query step for group $\mathbb{G}_1$ with input P.
  - Finally, send the secret key $\mathsf{sk} := L_1(P)$ to $\mathcal{A}$.

Challenge round: Once during the game, $\mathcal{A}$ may decide that the current round is the *challenge round.* The challenge round proceeds as follows:

- $\mathcal{A}$ sends two messages $m_0, m_1 \in \{0,1\}^{p'}$ and an identity-batch label pair (id^*, t^*) on which it wishes to be challenged.
- If key computation query (ids, t) has already been made with batch label $t = t^*$ and where $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger executes the following steps.
  - If $t^* \notin H$, execute a hash query step internally using input t^*.
  - Let $\mathbf{A}$ be a matrix and $\mathbf{b}$ be a vector defined as follows.

$$\mathbf{A} := \begin{pmatrix} 1 & \mathsf{id}^* - \tau & 0 \\ \mathsf{msk} & 0 & -1 \end{pmatrix}$$

$$\mathbf{b} := \begin{pmatrix} 0 \\ -H(t^*) \cdot \mathsf{msk} \end{pmatrix}$$

– Sample a (column) vector $\mathbf{r} \leftarrow (\mathbb{Z}_p^*)^2$. Compute a list $\mathbf{y}$ of four $\mathbb{Z}_p$ values where

$$\mathbf{y} = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b})$$

– Parse $\mathbf{y}$ as (y_0, y_1, y_2, y_3). For each $i \in \{0, 1, 2, 3\}$, execute an internal labeling query step for group $\mathbb{G}_2$ with input y_i.

– Send the ciphertext $\mathbf{c} = (L_2(y_0), L_2(y_1), L_2(y_2), L_2(y_3 + L_2^{-1}(m_b)))$ to $\mathcal{A}$

Post-challenge queries: After the challenger round, $\mathcal{A}$ may again issue an arbitrary number of all the query types that were part of Pre-challenge queries with the following additional restriction on *key computation queries*: $\mathcal{A}$ cannot query (t^*, ids) with $\mathsf{id}^* \in \mathsf{ids}$. More formally, we have the following:

Key computation query:

- $\mathcal{A}$ sends a list ids of B identities along with a batch label t to the challenger.
- If a key computation query has already been made with batch label t **or if** $t = t^*$ **and** $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger executes the same steps as described in the Pre-challenge queries.

Output: At any point in time, $\mathcal{A}$ can decide to halt and output a bit $b' \in \{0, 1\}$. The game then halts with the same output b'.

Proof. To prove Thm. 6.2, we will proceed in two steps. In the first step, we will show that the difference between symbolic versions of the experiments $\mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,GGM}}$ and $\mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,GGM}}$, denoted by $\mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,SM}}$ and $\mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,SM}}$ respectively, is zero (Thm. 6.3). In the second step, we will show that for $b \in \{0, 1\}$, the difference between $\mathsf{Expt}_{\mathcal{A},b}^{\mathsf{BIBE,GGM}}$ and $\mathsf{Expt}_{\mathcal{A},b}^{\mathsf{BIBE,SM}}$ is at most $\epsilon = \binom{q+B+5}{2} \frac{(B+2)}{p}$ (Thm. 6.4). Combining these two steps implies the following:

$$\mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,GGM}} \approx_\epsilon \mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,SM}} \equiv \mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,SM}} \approx_\epsilon \mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,GGM}}$$

which implies that $\mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,GGM}} \approx_{2\epsilon} \mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,GGM}}$ and completes the proof of Thm. 6.2. □

Corollary 6.2.1. *For $B = \mathsf{poly}(\lambda)$, $q = \mathsf{poly}(\lambda)$ and all unbounded adversaries $\mathcal{A}$ making at most q queries (including queries to the group oracle, hash oracle, and key computation queries), we have:*

$$\left|\Pr[\mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,GGM}}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,GGM}}(1^\lambda, B) = 1]\right| \leq \mathsf{negl}(\lambda)$$

Lemma 6.3. *For all $B \in \mathbb{N}$ and all unbounded adversaries $\mathcal{A}$, we have:*

$$\left|\Pr[\mathsf{Expt}_{\mathcal{A},0}^{\mathsf{BIBE,SM}}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}_{\mathcal{A},1}^{\mathsf{BIBE,SM}}(1^\lambda, B) = 1]\right| = 0$$

The security game $\mathsf{Expt}^{\mathsf{BIBE,SM}}_{\mathcal{A},b}(1^\lambda, B)$.

Setup: The challenger takes as input the security parameter λ and the batch size B. It runs $\mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B)$. Let $S_1 = S_2 = S_T = \{0,1\}^{p'}$ where $2^{p'} \geq p$ and p is a prime (denoting the group order) contained in params. Then it initializes the maps including the labeling function $L_i : \mathbb{Z}_p \to S$ for each $i \in \{1, 2, T\}$ and the hash function $H : \{0,1\}^* \to \mathbb{Z}_p$. During the KeyGen phase, it performs the following steps:

- Sample $e_{\mathsf{msk}} \leftarrow S_2$, sets $L_2(X_{\mathsf{msk}}) := e_{\mathsf{msk}}$ where X_{msk} is an indeterminate corresponding to the master secret key msk. Set $S_2 := S_2 \setminus \{e_{\mathsf{msk}}\}$.
- Let X_τ be an indeterminate.
- For $i \in [B-1]$, sample $e_{\tau^i} \leftarrow S_1$ and set $L_1(X_\tau^i) := e_{\tau^i}$ and $S_1 := S_1 \setminus \{e_{\tau^i}\}$.
- Also, sample $e' \leftarrow S_2$ and set $L_2(X_\tau) := e'$ and $S_2 := S_2 \setminus \{e'\}$.

It sets $\mathsf{pk} = (L_1(X_\tau), \ldots, L_1(X_\tau^B), L_2(X_\tau), L_2(X_{\mathsf{msk}}))$ and sends $(\mathsf{params}, \mathsf{pk})$ to $\mathcal{A}$.

The rest of the game proceeds in rounds, as follows.

Pre-challenge queries: $\mathcal{A}$ may issue an arbitrary number of the following types of queries:

Labeling query: For each labeling query for group $\mathbb{G}_i$, the challenger receives a value $v \in \mathbb{Z}_p$ from $\mathcal{A}$. If $v \notin L_i$, it samples $z \leftarrow S_i$, sets $L_i(v) := z$ and $S_i := S_i \setminus \{z\}$. It sends $L_i(v)$ to $\mathcal{A}$.

Group operation: For each group operation query for group $\mathbb{G}_i$, the challenger receives $(\ell_1, \ell_2, a_1, a_2) \in \{0,1\}^{{p'}^2} \times \mathbb{Z}_p^2$. If there doesn't exist polynomials $x_1, x_2 \in \mathbb{Z}_p[*]$ such that $L(x_1) = \ell_1$ and $L(x_2) = \ell_2$, then send $\perp$ to $\mathcal{A}$. Otherwise, execute an internal labeling query step for group $\mathbb{G}_i$ with polynomial $x_3 = a_1 x_1 + a_2 x_2$ and send $L_i(x_3)$ to $\mathcal{A}$.

Pairing operation: The challenger receives $(\ell_1, \ell_2) \in \{0,1\}^{p'} \times \{0,1\}^{p'}$. If there doesn't exist polynomials $x_1, x_2 \in \mathbb{Z}_p[*]$ such that $L_1(x_1) = \ell_1$ and $L_2(x_2) = \ell_2$, then the challenger sends $\perp$ to $\mathcal{A}$. Otherwise, it computes polynomial $x_3 = x_1 \cdot x_2$. It executes an internal labeling query step for group $\mathbb{G}_T$ with input x_3 and sends $L_T(x_3)$ to $\mathcal{A}$.

Hash query: For each hash query, the challenger receives a string $s \in \{0,1\}^*$ from $\mathcal{A}$. If $s \notin H$, then the challenger sets $H(s) := X_s$, where X_s is an indeterminate, and executes an internal labeling query step for group $\mathbb{G}_1$ with input X_s. It sends $L_1(H(s))$ to $\mathcal{A}$.

Key computation query:

- $\mathcal{A}$ sends a list ids of B identities, $S = \{\mathsf{id}_1, \ldots, \mathsf{id}_B\} \subseteq \mathcal{I}$, along with a batch label $t \in \mathcal{T}$ to the challenger.

- If a key computation query has already been made with batch label t, the challenger halts the game.
- Otherwise, the challenger does the following:
  - Let $f(X_\tau) = \sum_{i=0}^{B} f_i \cdot X_\tau^i$ be a univariate polynomial of degree B over $\mathbb{Z}_p$ with roots at $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and leading coefficient 1. If $f(X_\tau) \notin L_1$, the challenger samples $z \leftarrow \{0,1\}^*$ and sets $L_1(f(X_\tau)) := z$.
  - If $t \notin H$, it executes a hash query step with input t.
  - Let $P := (f(X_\tau) + H(t)) \cdot X_{\mathsf{msk}}$ be a polynomial where the indeterminates are X_τ, $H(t)$ and X_{msk}. If $P \notin L_1$, the challenger samples $z \leftarrow \{0,1\}^*$ and sets $L_1(P) := z$.
  - Finally, it sends the secret key $\mathsf{sk} := L_1(P)$ to $\mathcal{A}$.

Challenge round: Once during the game, $\mathcal{A}$ may decide that the current round is the *challenge round.* The challenge round proceeds as follows:

- $\mathcal{A}$ sends two messages $m_0, m_1 \in \{0,1\}^*$ and an identity-batch label pair (id^*, t^*) on which it wishes to be challenged.
- If key computation query (ids, t) has already been made with batch label $t = t^*$ and where $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger executes the following steps.
  - If $t^* \notin H$, execute a hash query step internally using input t^*.
  - Let $\mathbf{A}$ be a matrix and $\mathbf{b}$ be a vector defined as follows.

$$\mathbf{A} := \begin{pmatrix} 1 & \mathsf{id}^* - X_\tau & 0 \\ X_{\mathsf{msk}} & 0 & -1 \end{pmatrix}$$

$$\mathbf{b} := \begin{pmatrix} 0 \\ -H(t^*) \cdot X_{\mathsf{msk}} \end{pmatrix}$$

  - Let $\mathbf{X_r} := \begin{pmatrix} X_{r_1} \\ X_{r_2} \end{pmatrix}$ be a vector of two indeterminates. Compute a list $\mathbf{y}$ of four polynomials where

$$\mathbf{y} = (\mathbf{X_r}^T \cdot \mathbf{A}, \mathbf{X_r}^T \cdot \mathbf{b})$$

  - Parse $\mathbf{y}$ as (y_0, y_1, y_2, y_3). For each $i \in \{0,1,2,3\}$, set $L_2(y_i) \leftarrow \{0,1\}^*$.
  - Send the ciphertext $\mathbf{c} = (L_2(y_0), L_2(y_1), L_2(y_2), L_2(y_3 + L_2^{-1}(m_b)))$ to $\mathcal{A}$

Post-challenge queries: After the challenger round, $\mathcal{A}$ may again issue an arbitrary number of *key computation queries,* with the additional restriction that A cannot query (t^*, ids) with $\mathsf{id}^* \in \mathsf{ids}$:

- $\mathcal{A}$ sends a list ids of B identities along with a batch label t to the challenger.
- If a key computation query has already been made with batch label t **or if** $t = t^*$ **and** $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger executes the same steps as described in the Pre-challenge queries.

Output: At any point in time, $\mathcal{A}$ can decide to halt and output a bit $b' \in \{0, 1\}$. The game then halts with the same output b'.

Proof. To prove the above lemma, we will introduce an intermediate hybrid experiment $\mathsf{Expt}^{\mathsf{Hyb},\mathsf{SM}}_{\mathcal{A}}(1^\lambda, B)$ and show that for $b \in \{0, 1\}$, the following holds:

$$\mathsf{Expt}^{\mathsf{BIBE},\mathsf{SM}}_{\mathcal{A},b}(1^\lambda, B) \equiv \mathsf{Expt}^{\mathsf{Hyb},\mathsf{SM}}_{\mathcal{A}}(1^\lambda, B) \tag{1}$$

$\mathsf{Expt}^{\mathsf{Hyb},\mathsf{SM}}_{\mathcal{A}}(1^\lambda, B)$ is same as $\mathsf{Expt}^{\mathsf{BIBE},\mathsf{SM}}_{\mathcal{A},b}(1^\lambda, B)$ except that the fourth component of the ciphertext c in the challenge round is computed as follows:

- The challenger defines an indeterminate u and sets $L_2(u) \leftarrow S_2$ and updates $S_2 := S_2 \setminus L_2(u)$.
- The challenger sets the fourth component of ciphertext $\mathbf{c}$ to be $L_2(u)$.

To prove Eq. (1), we need to show that the polynomial y_3 involved in the ciphertext of challenge round of $\mathsf{Expt}^{\mathsf{BIBE},\mathsf{SM}}_{\mathcal{A},b}(1^\lambda, B)$ is independent of the polynomials corresponding to all the other group element encodings which are in the view of adversary.

Without loss of generality, we will assume that the adversary makes n key computation queries with batch labels $t_1, \ldots, t_n$ and requests a challenge on $t_{i^*} = t^*$ for some $i^* \in [n]$ and identity id^*. We will use f^i to denote the degree B univariate polynomial having as roots the ids used in the i^{th} key computation query. Without loss of generality, we will assume that the hash queries are made on all the batch labels $t_1, \ldots, t_n$.

We will now list down the polynomials corresponding to the encodings held by the adversary.

$$L_1 = \left\{ 1, \{X^i_\tau\}_{i\in[B]}, \{H(t_i)\}_{i\in[n]}, \{(f^i(X_\tau) + H(t_i)) \cdot X_{\mathsf{msk}})\}_{i\in[n]} \right\}$$

$$L_2 = \left\{ 1, X_\tau, X_{\mathsf{msk}}, \underbrace{X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}}}_{\mathbf{r}^T\cdot a_1}, \underbrace{X_{r_1}(\mathsf{id}^* - X_\tau)}_{\mathbf{r}^T\cdot a_2}, \underbrace{-X_{r_2}}_{\mathbf{r}^T\cdot a_3} \right\}$$

The polynomial y_3 involved in $\mathsf{Expt}^{\mathsf{BIBE},\mathsf{SM}}_{\mathcal{A},b}$ is $(-X_{r_2} \cdot H(t_{i^*}) \cdot X_{\mathsf{msk}})$ and we wish to show that it is outside the span of $L_1 \otimes L_2$, i.e., it is linearly independent of the list of polynomials obtained by multiplying polynomials in L_1 with polynomials in L_2. Let's assume, for the sake of contradiction, that this is not the case, i.e., y_3 happens to be a linear combination of polynomials in $L_1 \otimes L_2$. Then, by inspection, we have the following observations about the coefficients of the polynomials involved in the linear combination:

- Let $S_1 \subseteq L_1$ and $S_2 \subset L_2$ be the following sets:

$$S_1 = \left\{ 1, \{X_\tau^i\}_{i\in[B]}, \{H(t_i)\}_{i\in[n]}, \{(f^i(X_\tau) + H(t_i)) \cdot X_{\mathsf{msk}})\}_{i\in[n]} \right\}$$

$$S_2 = \left\{ 1, X_\tau, X_{\mathsf{msk}} \right\}$$

  The coefficients of the polynomials in $S_1 \otimes S_2$ will be zero because the monomials in such polynomials do not occur in the target polynomial and are not present as monomials in other polynomials in$L_1 \otimes L_2$ (therefore they cannot be cancelled out).

- The coefficients of terms generated by the following completion will be zero for the same reason as above.

$$\left\{ 1, \{X_\tau^i\}_{i\in[B]}, \{H(t_i)\}_{i\in[n]} \right\} \otimes \left\{ \underbrace{-X_{r_2}}_{\mathbf{r}^T \cdot a_3} \right\}$$

- Similarly, the coefficients of the polynomials generated by the following completion will be zero as they contain monomials of the form $H(t_i) \cdot X_{r_2} \cdot X_{\mathsf{msk}}^2$ and $H(t_i) \cdot X_{\mathsf{msk}} \cdot X_{r_1} \cdot X_\tau$ which are neither present in the target polynomial nor in other polynomials in $L_1 \otimes L_2$.

$$\left\{ (f^i(X_\tau) + H(t_i)) \cdot X_{\mathsf{msk}})\}_{i\in[n]} \right\} \otimes \left\{ \underbrace{X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}}}_{\mathbf{r}^T \cdot a_1}, \underbrace{X_{r_1}(\mathsf{id}^* - X_\tau)}_{\mathbf{r}^T \cdot a_2} \right\}$$

- The coefficients of terms generated by the following completion will be zero.

$$\left\{ \{H(t_i)\}_{i\in[n]} \right\} \otimes \left\{ \underbrace{X_{r_1}(\mathsf{id}^* - X_\tau)}_{\mathbf{r^T} \cdot a_2} \right\}$$

  The reason is that it generates polynomials having monomials of the form $H(t_i) \cdot X_{r_1} \cdot X_\tau$. Since these monomials occur neither in the target polynomial, nor as monomials in the polynomials generated by other terms in $L_1 \otimes L_2$, their coefficients will be zero.

- The coefficients of terms generated by the following completion will be zero.

$$\left\{ \{H(t_i)\}_{i\in[n]} \right\} \otimes \left\{ \underbrace{X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}}}_{\mathbf{r^T} \cdot a_1} \right\}$$

  The polynomials generated by this completion contains monomials of the form $H(t_i) \cdot X_{r_1}$.
  The only other completion which generates this polynomial is $\left\{ \{H(t_i)\}_{i\in[n]} \right\} \otimes \left\{ \underbrace{X_{r_1}(\mathsf{id}^* - X_\tau)}_{\mathbf{r^T} \cdot a_2} \right\}$.

However, by previous observation, the coefficient of all terms of those completion are zero which, in turn, forces the coefficient of all terms in this completion to be also zero.

Finally, we are left with the following completion terms.

$$S_1 := \left\{ \{X_\tau^i\}_{i\in[B]} \right\} \otimes \left\{ \underbrace{X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}}}_{\mathbf{r}^T \cdot a_1}, \underbrace{X_{r_1}(\mathsf{id}^* - X_\tau)}_{\mathbf{r}^T \cdot a_2} \right\}$$

$$S_2 := \left\{ \{(f^i(X_\tau) + H(t_i)) \cdot X_{\mathsf{msk}})\}_{i\in[n]} \right\} \otimes \left\{ \underbrace{-X_{r_2}}_{\mathbf{r}^T \cdot a_3} \right\}$$

We note that each polynomial in S_2 has a monomial of the form $H(t_i) \cdot X_{\mathsf{msk}} \cdot X_{r_2}$. For all $i \in [n], i \neq i^*$, the coefficients of such polynomials would be zero as the monomial is neither present in the target polynomial nor in the other polynomials generated by the remaining completion. Hence, we are left with the following polynomials in the completion.

$$\left(\left\{ \{X_\tau^i\}_{i\in[B]} \right\} \otimes \left\{ \underbrace{X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}}}_{\mathbf{r}^T \cdot a_1}, \underbrace{X_{r_1}(\mathsf{id}^* - X_\tau)}_{\mathbf{r}^T \cdot a_2} \right\} \right) \bigcup \left\{ -(f^{i^*}(X_\tau) + H(t_{i^*})) \cdot X_{\mathsf{msk}}) \cdot X_{r_2} \right\}$$

$$= \left\{ \{X_\tau^i \cdot (X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}})\}_{i\in[B]}, \{X_\tau^i \cdot X_{r_1}(\mathsf{id}^* - X_\tau)\}_{i\in[B]}, -(f^{i^*}(X_\tau) + H(t_{i^*})) \cdot X_{\mathsf{msk}} \cdot X_{r_2} \right\}$$

Recall that the target polynomial y_3 is $(-X_{r_2} \cdot H(t_{i^*}) \cdot X_{\mathsf{msk}})$. Let $\{c_i, d_i\}_{i\in[0,B]}, e \in \mathbb{Z}_p$ be coefficients s.t.

$$\begin{aligned} -X_{r_2} \cdot H(t_{i^*}) \cdot X_{\mathsf{msk}} = &\sum_{i=0}^{B} c_i \cdot X_\tau^i \cdot (X_{r_1} + X_{r_2} \cdot X_{\mathsf{msk}}) \\ &+ \sum_{i=0}^{B} d_i \cdot X_\tau^i \cdot X_{r_1}(\mathsf{id}^* - X_\tau) \\ &- e(f^{i^*}(X_\tau) + H(t_{i^*})) \cdot X_{\mathsf{msk}} \cdot X_{r_2} \end{aligned}$$

From the above, it is clear that $e = 1$ to get the $(-X_{r_2} \cdot H(t_{i^*}) \cdot X_{\mathsf{msk}})$ monomial on the R.H.S. This implies that $\sum_{i=0}^{B} c_i \cdot X_\tau^i = f^{i^*}(X_\tau)$ so that the monomials involving $X_{r_2}, X_{\mathsf{msk}}$ vanish on the R.H.S. Hence, we have the following remaining constraint.

$$f^{i^*}(X_\tau) \cdot X_{r_1} = \sum_{i=0}^{B} d_i \cdot X_\tau^i \cdot X_{r_1} \cdot (X_\tau - \mathsf{id}^*)$$

The above constraint implies that id^* is a root of the polynomial $f^{i^*}(X_\tau)$ which is a contradiction as per the rules of the security game.

□

Lemma 6.4. *For all $B \in \mathbb{N}$ and all unbounded adversaries $\mathcal{A}$ making at most q queries (including queries to the group oracle, hash oracle, and key computation queries), for $b \in \{0,1\}$, we have:*

$$\left|\Pr[\mathsf{Expt}^{\mathsf{BIBE,GGM}}_{\mathcal{A},b}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{BIBE,SM}}_{\mathcal{A},b}(1^\lambda, B) = 1]\right| \leq \binom{q+B+5}{2}\frac{(B+2)}{p}$$

Proof. To prove this, we consider the following experiment[7]: At the end of $\mathsf{Expt}^{\mathsf{BIBE,SM}}_{\mathcal{A},b}$, the challenger samples uniformly random values from $\mathbb{Z}_p$ for all the indeterminates involved in $\mathsf{Expt}^{\mathsf{BIBE,SM}}_{\mathcal{A},b}$ and replaces all the polynomials with their evaluations. This is a perfect simulation of $\mathsf{Expt}^{\mathsf{BIBE,GGM}}_{\mathcal{A},b}$ unless the following bad event happens: The sampled values result in a an identical evaluation of polynomials which are not identical. For any pair of polynomials (f_1, f_2), of total degree (d_1, d_2) respectively, we can bound the probability of this bad event happening to be at most $\frac{\max(d_1,d_2)}{p}$ using Schwartz-Zippel lemma [Sch80; Zip79]. There are $q+B+5$ polynomials in total across all three groups (q polynomials from queries and $B+5$ polynomials from the setup phase) with maximum total degree at most $B+2$. Therefore, union bounding across all possible pairs of polynomials gives us a maximum failure probability of $\binom{q+B+5}{2}\frac{(B+2)}{p}$. □

7 Extensions and Optimizations

7.1 Thresholdizing the scheme

In our Batched IBE scheme (Thm. 5.1), the master secret-key msk is held by a central authority which runs $\mathsf{ComputeKey}$ procedure to derive the batch label specific secret keys sk and distribute them as necessary. In many applications, including the ones we discussed in the introduction, it is favorable to distribute this trust among multiple authorities. Specifically, instead of a single authority holding the complete msk, we would like to have multiple, let's say n, authorities where each authority $i \in [n]$ holds a "partial" master secret key msk_i (e.g., in the form of secret shares of msk) and the adversary can corrupt at most a threshold f number of authorities. In such a distributed setting, it is highly desirable to construct a scheme where multiple authorities can securely issue the batch label specific secret keys sk without leaking their "partial" master secret key msk_i. Our construction in Fig. 1 readily admits such an efficient threshold version. This is due to the fact that the $\mathsf{ComputeKey}$ procedure in our construction simply computes a BLS-like signature which can be efficiently thresholdized as observed in prior works [Bol03]. For completeness, we define Thresholdizable Batched Identity Based Encryption in Section A and show how our non-threshold construction can be modified to obtain a threshold version.

7.2 Outsourcing the digest computation

In our construction, the digest $d = \mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\})$ is computed as a KZG commitment to the polynomial $f(X) = \sum_{i=0}^{B} f_i \cdot X^i$ that has roots at $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and leading coefficient 1. The digest is $d := \sum_{i=0}^{B} f_i \cdot [\tau^i]_1$.

[7] We note that the master theorem proved in [BBG05] cannot be directly applied here as the adversary can make oracle queries other than standard group operations such as key computation queries. To account for this difference, we redo the analysis here which essentially follows the same style that is used in proving the master theorem.

We note that the work of computing the digest $d = \mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\})$ can be outsourced to a single server, such that the result can be efficiently verified by everyone. The main observation is that, since the polynomial $f(X)$ can be represented as $f(X) = \prod_{i=1}^{B}(X - \mathsf{id}_i)$ then for any value z in the field, the value $f(z)$ can be efficiently computed using only B multiplications, which is much faster than interpolating the polynomial and computing the digest.

Verification can thus be implemented using the Fiat-Shamir paradigm:

- The server which computed the digest d computes a random field point $z = H(d)$, and computes $y = f(z)$.
- To provide an evaluation proof for the KZG commitment d, it computes the quotient polynomial $q(X) = \frac{f(X)-y}{X-z}$ and then computes and publishes a KZG evaluation proof $\pi = [q(\tau)]_1$.
- Given d and π, anyone can compute $z = H(d)$ and then efficiently compute $y = f(z)$. Then the proof can be verified by checking that $[1]_2 \circ (d - y) = ([\tau]_2 - [z]_2) \circ \pi$.

One can use the well-known Schwartz-Zippel lemma [Sch80; Zip79] to show that the above procedure satisifies soundness.

The work of computing the group multiplications in $d := \sum_{i=0}^{B} f_i \cdot [\tau^i]_1$ can also be distributed among m servers. The result can be efficiently verified by everyone, given that they have access to the polynomial $f(X)$. (Computing $f(X)$ itself is non-trivial, as it is essentially computing an interpolation.)

Assume that this work is distributed between m servers such that server $k = 1, \ldots, m$ computes $d_k := \sum_{i=(k-1)\cdot B/m}^{k\cdot B/m-1} f_i \cdot [\tau^i]_1$. (To simplify the notation we assume that B/m is an integer.) Then, given these results from the servers it is easy to compute $d = \sum_{k=1}^{m} d_k$. The work of each server is roughly $1/m$ the work of computing d. The main remaining issue is how to efficiently verify that each d_k is computed correctly (or, in the case of outsourcing to a single server, verifying that it computed d correctly). We note that server k actually computes a KZG commitment d_k for the polynomial $f_k = \sum_{i=(k-1)B/m}^{kB/m-1} f_i \cdot X^i$. This polynomial is of degree $\frac{kB}{m} - 1 \leq B$. Therefore d_k can be verified using the same procedure outlined above for verifying d.[8]

7.3 Batching the Decryption Procedure

Some applications, such as mempool privacy, require multiple ciphertexts to be decrypted by a single party (e.g. blockchain validator in the mempool privacy application) once the secret key sk for a particular batch (or block in the mempool privacy application) is released. Assuming the number of ciphertexts is $O(B)$, a naive application of Decrypt will result in a total decryption time of $O(B^2 \log B)$ for our Batched IBE construction as per the analysis in Section 6.2.1. We discuss

[8] It is possible to do an efficient *optimistic* verification of the work of the servers, in the sense that if all servers are honest then the result can be efficiently verified without interpolating $f(X)$. Otherwise, after interpolating $f(X)$ it is possible to identify which server k provided an incorrect d_k value. The proof process is as follows. First, a value z is computed as $H(d_1, \ldots, d_m)$. Then each server provides $y_k = f_k(z)$ and a proof that the polynomial committed to by d_k has the output y_k at the point z. The verifier computes $f(z)$ and verifies that it is equal to $\sum_{k=1}^{m} y_k$. If $d = \sum_{k=1}^{m} d_k$ is not a commitment to $f(X)$, the by the Schwartz-Zippel theorem this check fails with all but negligible probability. (Note that this is not a check that each d_k is computed correctly, but rather that $d = \sum_{k=1}^{m} d_k$ is correct, which is the property that we need.

The batch decryption procedure BatchDecrypt**.**

$\mathsf{BatchDecrypt}(\{c_1, \ldots, c_B\}, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, \mathsf{id}, t)$**:**

- **(KZG opening proof computation)** Use [FK23] to derive openings $\pi_1, \ldots, \pi_B$ for the roots $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\}$ of the polynomial $f(X) := (X - \mathsf{id}_1) \ldots (X - \mathsf{id}_B)$, with respect to digest d.
- **(Decryption)** For each $i \in [B]$:
  1. Set
$$\mathbf{w}_i = \begin{pmatrix} d \\ \pi_i \\ \mathsf{sk} \end{pmatrix}.$$
  2. Parse c_i as $(c_{i,1}, c_{i,2})$, and set $m_i = c_{i,2} - c_{i,1} \circ \mathbf{w}$.
- Output $\{m_1, \ldots, m_B\}$.

Figure 2: Batching the decryption procedure in our Batched IBE construction.

below how to improve this time. Note that in this section, we refer to "time" as the total number of group and field operations required to perform a task.

Note that the computation that dominates the running time for decryption is the computation of KZG commitment opening proofs. The time for computing all the B proofs naively is $O(B^2 \log B)$, whereas the time for the remaining work after the opening proofs are finished is simply $O(B)$. The work of [FK23] shows how to compute a set of opening proofs for a KZG commitment much more efficiently than the naive approach. Specifically, assume as in our decryption procedure we have a KZG commitment to a polynomial of degree B, and we need opening proofs for each of its B roots. [FK23] show that if the polynomial has been constructed so the roots are all part of some set of roots of unity Ω, then it is possible to compute all openings in time $O(|\Omega| \log |\Omega|)$. Specifically, if we want to support a maximum batch size B, we can set Ω such that $|\Omega| = B$, and then as long as we choose the IDs from Ω, we can batch-decrypt in time $O(B \log B + B) = O(B \log B)$. In addition, [FK23] show that even if the roots of the polynomial are chosen arbitrarily, it is still possible to compute the B openings in time $O(|\Omega| \log^2 |\Omega|)$. So if the particular application requires more flexibility in choosing IDs, it is still possible to batch-decrypt in time $O(B \log^2 B + B) = O(B \log^2 B)$. We formally a batch decryption procedure BatchDecrypt in Fig. 2.

7.4 Non-Malleability for Mempool Privacy Application

In the mempool privacy application, the ciphertext encrypts transaction details and the mempool contains contains many such ciphertexts (which represent pending transactions in an encrypted form). Note that such a mempool is *public*, i.e. all the ciphertexts are publicly visible. By some arbitrary (possibly adversarial) process, a subset of the ciphertexts lying in the mempool are selected

whose underlying transaction would be included in the next block[9]. Once the subset is finalized, all the ciphertexts in this subset are decrypted, executed and recorded in the next block.

Combining the fact that the mempool is *public* and the subset selection process is *adversarial*, we want to prevent the following malleability attack: The adversary observes a ciphertext ct submitted by an honest user in the mempool, creates a new ciphertext ct^* by mauling ct in a way so that the plaintext m^* underlying ct^* is related in some way to the plaintext m underlying ct, and then includes ct^* in the subset while excluding ct. Once ct^* is decrypted, it would reveal m^* and hence leak information about m even though ct was never selected in the subset. This implies that the adversary managed to learn information about the transaction m while preventing m from being executed in the next block, something which is not desirable in the mempool privacy application. A naive application of Batched IBE to mempool privacy does not prevent this: imagine clients choose a random ID id and submit $(\mathsf{id}, \mathsf{ct})$ to the mempool. An adversary can simply submit a different $(\mathsf{id}, \mathsf{ct}^*)$, where ct^* was computed using the same ID as as an honest party, and can include $(\mathsf{id}, \mathsf{ct}^*)$ in the subset while excluding $(\mathsf{id}, \mathsf{ct})$. Once the Batched IBE secret key is released for the subset, it will enables decryption of all ciphertexts that have an identity which is the subset, including ct, even though the ciphertext ct itself is not part of the subset.

As mentioned in [Cho+24a], the application of mempool privacy has a specific requirement that ciphertexts must satisfy a form of non-malleability. Specifically, the authors of [Cho+24a] state that "adding non-malleability to ciphertexts corresponds closely to securing [the] encryption scheme against chosen ciphertexts." To get CCA2-security, they rely on generic technique of using NIZK proofs. We observe that our Batched IBE scheme already satisfies a form of adaptive security where the adversary is allowed to make key computation queries even after observing the challenge ciphertext. We can use this fact to achieve non-malleability significantly more cheaply than [Cho+24a]. Specifically, instead of using NIZK proofs, we will leverage a standard signature scheme to prevent the attack described before. The high-level idea is the following: To encrypt a message m (which represents transaction details in the mempool privacy application), an honest user would sample a signing key pair $(\mathsf{vk}^{\mathsf{Sign}}, \mathsf{sk}^{\mathsf{Sign}})$, set $\mathsf{id} = H(\mathsf{vk}^{\mathsf{Sign}})$ where H is a hash function, and then encrypt m under identity id and batch label t (which represents the next block number in the mempool privacy application) using our Batched IBE scheme to create a ciphertext $\mathsf{ct}^{\mathsf{BIBE}}$. Instead of sending $\mathsf{ct}^{\mathsf{BIBE}}$ to the mempool, the user would be required to send an expanded ciphertext $\mathsf{ct} = (\mathsf{vk}^{\mathsf{Sign}}, \mathsf{ct}^{\mathsf{BIBE}}, \sigma)$ where σ is a signature on ct using signing key $\mathsf{sk}^{\mathsf{Sign}}$.

Assuming the block capacity is B, once a subset of B such expanded ciphertexts, $\{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}$, are selected, a set S of approved identities is computed. To do so, each expanded ciphertext ct_i is parsed as $(\mathsf{vk}_i^{\mathsf{Sign}}, \mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i)$ and the identity $\mathsf{id}_i = H(\mathsf{vk}_i^{\mathsf{Sign}})$ corresponding to $\mathsf{ct}_i^{\mathsf{BIBE}}$ is added to the set S only if σ_i is a valid signature on $\mathsf{ct}_i^{\mathsf{BIBE}}$ w.r.t $\mathsf{vk}_i^{\mathsf{Sign}}$. Once the set S is computed, the decryption process begins by invoking the $\mathsf{Digest}^{\mathsf{BIBE}}$ procedure on set S followed by $\mathsf{ComputeKey}^{\mathsf{BIBE}}$ and finally $\mathsf{Decrypt}^{\mathsf{BIBE}}$. The reason why using signatures in this fashion suffices to counter the malleability attacks discussed before is the following. Suppose an adversary observes an honest ciphertext $\mathsf{ct} = (\mathsf{vk}^{\mathsf{Sign}}, \mathsf{ct}^{\mathsf{BIBE}}, \sigma)$ and would like to create a new mauled ciphertext $\widetilde{\mathsf{ct}} = (\widetilde{\mathsf{vk}^{\mathsf{Sign}}}, \widetilde{\mathsf{ct}^{\mathsf{BIBE}}}, \widetilde{\sigma})$, where $\widetilde{\mathsf{ct}^{\mathsf{BIBE}}} \neq \mathsf{ct}^{\mathsf{BIBE}}$, in a way so that by having just $\widetilde{\mathsf{ct}^{\mathsf{BIBE}}}$ decrypted, it would learn information about the plaintext underlying the honestly generated $\mathsf{ct}^{\mathsf{BIBE}}$. To do so, it needs to be the case that the identity $\widetilde{\mathsf{id}} = H(\widetilde{\mathsf{vk}^{\mathsf{Sign}}})$ corresponding to $\widetilde{\mathsf{ct}^{\mathsf{BIBE}}}$ matches with the identity

[9]This subset selection needs to be done because the size of mempool can be unbounded but the size of a block is bounded.

$\mathsf{id} = H(\mathsf{vk}^{\mathsf{Sign}})$ corresponding to $\mathsf{ct}^{\mathsf{BIBE}}$ (otherwise the plaintext underlying $\mathsf{ct}^{\mathsf{BIBE}}$ is hidden via the security of BIBE scheme). Now there are two cases: either $\widetilde{\mathsf{vk}^{\mathsf{Sign}}} \neq \mathsf{vk}^{\mathsf{Sign}}$ or $\widetilde{\mathsf{vk}^{\mathsf{Sign}}} = \mathsf{vk}^{\mathsf{Sign}}$. The former case means that the adversary found a hash collision (which would not be possible for a polynomial time adversary) and the latter case means that the adversary was able to forge a signature (which would not be possible by the unforgeability of the signature scheme).

In the remainder of this section, we will formalize the above ideas by defining the syntax and semantics of a *Batched Encryption* scheme (BE) aimed towards capturing the mempool privacy application. We note that the difference between a Batched Encryption scheme and Batched IBE scheme is that the former doesn't have any explicit notion of identities and resembles more like a standard public-key encryption scheme. We then provide a construction for BE scheme (Fig. 3) by combining a BIBE scheme along with a signature scheme Sign.

Syntax:

- $\mathsf{Setup}(1^\lambda, 1^B) \to \mathsf{params}$: A randomized algorithm that takes as input a security parameter $\lambda \in \mathbb{N}$ and a batch size $B = B(\lambda)$. It outputs params (system parameters) which includes a description of the message space $\mathcal{M}$, batch label space $\mathcal{T}$ and ciphertext space $\mathcal{C}$.
- $\mathsf{KeyGen}(\mathsf{params}) \to (\mathsf{pk}, \mathsf{sk})$: a randomized algorithm that takes as input params and outputs sk (secret key) and pk (public key).
- $\mathsf{Encrypt}(\mathsf{pk}, t, m) \to \mathsf{ct}$: a randomized algorithm that takes as input a public key pk along with the batch label $t \in \mathcal{T}$ and message $m \in \mathcal{M}$. It outputs a ciphertext $\mathsf{ct} \in \mathcal{C}$.
- $\mathsf{Decrypt}(\mathsf{sk}, \{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}, t) \to \{m_1, \ldots m_B\}$: a deterministic algorithm that takes as input B ciphertexts, $\{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}$, along with a batch label $t \in \mathcal{T}$ and outputs B plaintexts $\{m_1, \ldots m_B\}$.

Definition 7.1 (Batched Encryption Correctness)**.** For all $\lambda \in \mathbb{N}, B \in \mathbb{N}, \{m_1, \ldots, m_B\} \in \mathcal{M}^B, t \in \mathcal{T}$, the following should hold:

$$\Pr\left[\begin{array}{c|c} \mathsf{Decrypt}(\mathsf{sk}, \{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}, t) & \mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B) \\ = & (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(\mathsf{params}) \\ \{m_1, \ldots, m_B\} & \forall i \in [B] : \mathsf{ct}_i \leftarrow \mathsf{Encrypt}(\mathsf{pk}, m_i, t) \end{array}\right] = 1$$

Definition 7.2 (Batched Encryption Non-triviality/Efficiency)**.** We require that Decrypt can be split into two parts: 1) A *public* part, meaning computation *without using* the secret key sk, whose running time depends on the batch size B, 2) A *private* part, meaning computation *using* sk, whose running time is independent of the batch size B.

Remark. Similar to the motivation for defining the non-triviality/efficiency for Batched IBE scheme, we enforce the above requirement for the following reasons: 1) Without this requirement, one could come up with a trivial Batched Encryption scheme using a standard public key encryption scheme where the running time of Decrypt that depends on the secret key sk will be $O(B)$. 2) This feature is useful in a threshold setting where the sk is split among multiple parties (using a secret sharing scheme) and these parties need to securely emulate the execution of Decrypt procedure using their share of sk. In such a setting, the above requirement will ensure that the running time

and communication cost of the secure emulation is independent of the batch size B (which can be a huge cost saving in practice).

Definition 7.3 (Batched Encryption Security)**.** We say that a Batched Enc. scheme is secure if for all $B \in \mathbb{N}$, for all PPT adversaries $\mathcal{A}$ there exists some negligible function $\epsilon_{\mathcal{A}}$ such that the following holds:

$$\left|\Pr[\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},0}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},1}(1^\lambda, B) = 1]\right| < \epsilon_{\mathcal{A}}(\lambda)$$

with respect to the security game $\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},b}(1^\lambda, B)$ defined below.

The security game $\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},b}(1^\lambda, B)$**.**

Setup: The challenger takes as input the security parameter λ and the batch size B. It runs $\mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B)$, and then runs $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(\mathsf{params})$. Finally, it sends $(\mathsf{params}, \mathsf{pk})$ to $\mathcal{A}$.

The rest of the game proceeds in rounds, as follows.

Pre-challenge queries: $\mathcal{A}$ may issue an arbitrary number of *batch decryption queries*:

- $\mathcal{A}$ sends $\langle\{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}, t\rangle$ to the challenger.
- If a decryption query has already been made with batch label t, the challenger halts the game.
- The challenger runs $\mathsf{Decrypt}(\mathsf{sk}, \{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}, t) \to \{m_1, \ldots m_B\}$.
- The challenger sends $\{m_1, \ldots m_B\}$ to $\mathcal{A}$.

Challenge round: Once during the game, $\mathcal{A}$ may decide that the current round is the *challenge round*. The challenge round proceeds as follows:

- $\mathcal{A}$ sends t^* and $\{m_0, m_1\}$ to the challenger.
- The challenger runs $\mathsf{Encrypt}(\mathsf{pk}, t^*, m_b) \to \mathsf{ct}^*$.
- The challenger sends ct^* to $\mathcal{A}$.

Post-challenge queries: After the challenge rounds, $\mathcal{A}$ may again issue an arbitrary number of *batch decryption queries*, with the additional restriction that $\mathcal{A}$ cannot query (CT, t^*) with $\mathsf{ct}^* \in \mathsf{CT}$:

- $\mathcal{A}$ sends $\langle\{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}, t\rangle$ to the challenger.
- If a decryption query has already been made with batch label t **or if** $\mathsf{ct}^* \in \{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}$, the challenger halts the game.

- The challenger runs $\mathsf{Decrypt}(\mathsf{sk}, \{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}, t) \rightarrow \{m_1, \ldots m_B\}$.
- The challenger sends $\{m_1, \ldots m_B\}$ to $\mathcal{A}$.

Output: At any time, $\mathcal{A}$ can halt and output a bit $b' \in \{0, 1\}$. The experiment then ends with the same output b'.

Remark (Thresholdizing the scheme). One can define a thresholdizable version of Batched Encryption scheme (TBE) where the secret key sk is split among multiple, let's say n, parties (using a secret sharing scheme) in a manner similar to how we extend the definition of Batched IBE scheme to the threshold setting (see Section A) where correctness and security are required to hold against an adversary which corrupts at-most $\lfloor \frac{n-1}{2} \rfloor$ of the parties. In such a TBE scheme, the Decrypt procedure is an interactive protocol and the non-triviality requirement is that the communication cost of such a protocol should be independent of the batch size B.

Theorem 7.4. *Assuming the existence of a Batched Identity Based Encryption scheme* BIBE*, a collision-resistant hash function family, and a strongly existential-unforgeable signature scheme* Sign *, there exists (Fig. 3) a Batched Encryption scheme* BE*.*

The correctness of our BE scheme follows directly from the correctness of BIBE scheme and Sign scheme. Also, the non-triviality of our BE scheme follows from the non-triviality of BIBE scheme.

We now proceed to prove the security of our BE scheme. Let $\mathcal{A}$ be a PPT adversary against the game $\mathsf{Expt}^{\mathsf{mempool}}$. We build a sequence of hybrids to show that $\left|\Pr[\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},0}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},1}(1^\lambda, B) = 1]\right| < \epsilon_{\mathcal{A}}(\lambda)$.

- $\mathsf{Expt}^{\mathsf{mempool},0}_{\mathcal{A},b}$: This is same as $\mathsf{Expt}^{\mathsf{mempool}}_{\mathcal{A},b}$, instantiated with the scheme from Fig. 3.
- $\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}$: This is same as $\mathsf{Expt}^{\mathsf{mempool},0}_{\mathcal{A},b}$, except that:
  - During the challenge round, the experiment parses $\mathsf{ct}^* = (\mathsf{vk}^{*,\mathsf{Sign}}, \mathsf{ct}^{*,\mathsf{BIBE}}, \sigma^*)$, and saves $\mathsf{vk}^{*,\mathsf{Sign}}$. If $H(vk^{*,\mathsf{Sign}})$ is equal to the hash of some verification key given by the adversary during a pre-challenge query, the experiment halts and outputs "FAIL 1".
  - during each post-challenge query, before decryption, the experiment parses each $\mathsf{ct}_i = (\mathsf{vk}_i^{\mathsf{Sign}}, \mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i)$, and if $\mathsf{vk}_i^{\mathsf{Sign}} \neq \mathsf{vk}^{*,\mathsf{Sign}}$ but $H(\mathsf{vk}_i^{\mathsf{Sign}}) = H(\mathsf{vk}^{*,\mathsf{Sign}})$, it halts and outputs "FAIL 1".
- $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},b}$: This is same as $\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}$, except that:
  - during each post-challenge query, after verifying that after checking that $\mathsf{ct}^* \notin \{\mathsf{ct}_1, \ldots, \mathsf{ct}_B\}$ and after verifying the signatures included in each ciphertext, the experiment parses each $\mathsf{ct}_i = (\mathsf{vk}_i^{\mathsf{Sign}}, \mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i)$, and if $\mathsf{vk}_i^{\mathsf{Sign}} = \mathsf{vk}^{*,\mathsf{Sign}}$, it halts and outputs "FAIL 2".

To finish the proof of security, we state and prove the following claims of indistinguishability of adjacent pairs of hybrids.

Batched Encryption (BE) construction.

- $\mathsf{Setup}(1^\lambda, 1^B)$: Execute $\mathsf{Setup}^{\mathsf{BIBE}}(1^\lambda, 1^B) \to \mathsf{params}$. Choose a hash function $H : \{0,1\}^* \to \mathcal{I}$ where $\mathcal{I}$ is identity space of BIBE scheme. Output params and H.

- $\mathsf{KeyGen}(\mathsf{params})$: Execute $\mathsf{KeyGen}^{\mathsf{BIBE}}(\mathsf{params}) \to (\mathsf{msk}^{\mathsf{BIBE}}, \mathsf{pk}^{\mathsf{BIBE}})$. Set $(\mathsf{pk}^{\mathsf{BE}}, \mathsf{sk}^{\mathsf{BE}}) := (\mathsf{pk}^{\mathsf{BIBE}}, \mathsf{msk}^{\mathsf{BIBE}})$ and output $(\mathsf{pk}^{\mathsf{BE}}, \mathsf{sk}^{\mathsf{BE}})$

- $\mathsf{Encrypt}(\mathsf{pk}^{\mathsf{BE}}, t, m)$:

  1. Sample random signing keypair $(\mathsf{sk}^{\mathsf{Sign}}, \mathsf{vk}^{\mathsf{Sign}})$.
  2. Set $\mathsf{id} = H(\mathsf{vk}^{\mathsf{Sign}})$.
  3. Use BIBE to encrypt message m: set $\mathsf{ct}^{\mathsf{BIBE}} = \mathsf{BIBE.Encrypt}(\mathsf{pk}^{\mathsf{BE}}, m, \mathsf{id}, t)$.
  4. Use signature scheme to compute signature $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}^{\mathsf{Sign}}, \mathsf{ct})$.
  5. Output $\mathsf{ct}^{\mathsf{BE}} := (\mathsf{vk}^{\mathsf{Sign}}, \mathsf{ct}^{\mathsf{BIBE}}, \sigma)$.

- $\mathsf{Decrypt}(\mathsf{sk}^{\mathsf{BE}}, \{\mathsf{ct}_1^{\mathsf{BE}}, \ldots, \mathsf{ct}_B^{\mathsf{BE}}\}, t)$:

  1. Let S be a set representing the set of "approved" ids, initialized to be empty.
  2. For each i, perform the following steps:
     (a) Parse $\mathsf{ct}_i^{\mathsf{BE}} = (\mathsf{vk}_i^{\mathsf{Sign}}, \mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i)$.
     (b) If σ_i is not a valid signature on $\mathsf{ct}_i^{\mathsf{BIBE}}$ w.r.t $\mathsf{vk}_i^{\mathsf{Sign}}$, then set $m_i := \perp$.
     (c) Otherwise, add $\mathsf{id}_i = H(\mathsf{vk}_i^{\mathsf{Sign}})$ to S.
  3. $\mathsf{Digest}^{\mathsf{BIBE}}(\mathsf{pk}^{\mathsf{BE}}, S) \to d$
  4. $\mathsf{ComputeKey}(\mathsf{sk}^{\mathsf{BE}}, d, t) \to \mathsf{sk}^{\mathsf{BIBE}}$
  5. For each $\mathsf{id}_i \in S$, $\mathsf{Decrypt}(\mathsf{ct}_i^{\mathsf{BIBE}}, \mathsf{sk}^{\mathsf{BIBE}}, d, S, \mathsf{id}_i, t) \to m_i$.
  6. Output $\{m_1, \ldots, m_B\}$.

Figure 3: Our construction for Batched Encryption scheme BE using the Batched Identity Based Encryption scheme BIBE (Fig. 1) and an arbitrary signature scheme Sign.

Claim. *Let $\mathcal{A}$ be a PPT adversary against the game $\mathsf{Expt}^{\mathsf{mempool}}$. Assuming H is chosen from a collision-resistant hash function family, for all $b \in \{0,1\}$,*

$$\left|\Pr[\mathsf{Expt}^{\mathsf{mempool},0}_{\mathcal{A},b}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}(1^\lambda, B) = 1]\right| < \epsilon_{\mathcal{A}}(\lambda).$$

Proof. Assume that the claim does not hold. This implies that the probability that $\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}(1^\lambda, B)$ outputs "FAIL 1" is non-negligible. We can then build a reduction to collision-resistance of the hash function. The reduction receives H from the challenger, and then runs $\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}(1^\lambda, B)$, outputting H as part of the setup; if the output of the experiment is "FAIL 1", then the reduction sends the $\mathsf{vk}_i^{\mathsf{Sign}}$ and $\mathsf{vk}^{*,\mathsf{Sign}}$ which caused failure to the collision-resistance challenger.

Since $\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}(1^\lambda, B)$ outputs "FAIL 1" exactly when $\mathsf{vk}_i^{\mathsf{Sign}}$ and $\mathsf{vk}^{*,\mathsf{Sign}}$ form a collision over H, and since the probability of "FAIL 1" is non-negligible by assumption, we have contradicted collision-resistance of H. □

Claim. *Let $\mathcal{A}$ be a PPT adversary against the game $\mathsf{Expt}^{\mathsf{mempool}}$. Assuming the signature scheme satisfies strong existential unforgeability, for all $b \in \{0,1\}$,*

$$\left|\Pr[\mathsf{Expt}^{\mathsf{mempool},1}_{\mathcal{A},b}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},b}(1^\lambda, B) = 1]\right| < \epsilon_{\mathcal{A}}(\lambda).$$

Proof. Assume that the claim does not hold. This implies that the probability that $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},b}(1^\lambda, B)$ outputs "FAIL 2" is non-negligible. We can then build a reduction to existential unforgeability of the signature scheme.

The reduction works as follows. It runs $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},b}(1^\lambda, B)$ until the challenge round. During that round, it asks for a verification key from the existential unforgeability challenger, and uses this as $\mathsf{vk}^{*,\mathsf{Sign}}$, hashing it to get the ID that it uses to generate the ciphertext $\mathsf{ct}^{*,\mathsf{BIBE}}$. It then queries the signature oracle of the existential unforgeability game to get a signature σ^* over the message $\mathsf{ct}^{*,\mathsf{BIBE}}$ before sending $(\mathsf{vk}^{*,\mathsf{Sign}}, \mathsf{ct}^{*,\mathsf{BIBE}}, \sigma^*)$ to $\mathcal{A}$. After this, if at any point $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},b}(1^\lambda, B)$ halts with output "FAIL 2", then the reduction sends message $\mathsf{ct}_i^{\mathsf{BIBE}}$ and σ_i for which $\mathsf{vk}_i^{\mathsf{Sign}} = \mathsf{vk}^{*,\mathsf{Sign}}$ to the existential unforgeability challenger.

The experiment outputs "FAIL 2" only if the following holds, with respect to some i:

- σ_i verifies correctly under vk_i^{Sign},
- $\mathsf{vk}_i^{\mathsf{Sign}} = \mathsf{vk}^{*,\mathsf{Sign}}$, the verification key given by the existential forgeability challenger, and
- $\mathsf{ct}^* = (\mathsf{vk}^{*,\mathsf{Sign}}, \mathsf{ct}^{*,\mathsf{BIBE}}, \sigma^*) \neq (\mathsf{vk}_i^{\mathsf{Sign}}, \mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i) = \mathsf{ct}_i$.

Notice that the first and last condition together imply that $(\mathsf{ct}^{*,\mathsf{BIBE}}, \sigma^*) \neq (\mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i)$. Thus, a "FAIL 2" output implies that the reduction was able to produce a valid message-signature pair that that was not queried to the signature oracle, breaking strong existential unforgeability of the signature scheme.

□

Claim. *Let $\mathcal{A}$ be a PPT adversary against the game $\mathsf{Expt}^{\mathsf{mempool}}$. Assuming the BIBE scheme satisfies the security definition in Thm. 5.4,*

$$\left|\Pr[\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},1}(1^\lambda, B) = 1]\right| < \epsilon_{\mathcal{A}}(\lambda).$$

Proof. Assume that the claim does not hold. We build a reduction to the security game of the BIBE. The reduction runs $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B)$, with the following two changes:

- During the challenge round, instead of running $\mathsf{Encrypt}^{\mathsf{BIBE}}$ directly, it sends a query $(m_0, m_1, \mathsf{id}^*, t^*)$ to the BIBE challenger, where m_0, m_1, and t^* are the plaintexts and batch label chosen by $\mathcal{A}$, and $id^* = H(\mathsf{vk}^{\mathsf{Sign},*})$ is the hash of the verification key chosen by $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B)$. After receiving the response ciphertext $\mathsf{ct}^{\mathsf{BIBE},*}$ from the BIBE challenger, it computes the final challenge round response in the same way as $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B)$.
- During the pre-challenge and post-challenge queries, **after running all the verification steps in** $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B)$, it parses each ct_i as $(\mathsf{vk}_i^{\mathsf{Sign}}, \mathsf{ct}_i^{\mathsf{BIBE}}, \sigma_i)$, sets $\mathsf{id}_i = H(\mathsf{vk}_i^{\mathsf{Sign}})$, and sends the set $\{\mathsf{id}_1, \ldots, id_n\}$ and t to the BIBE challenger. It then gets a secret key back and uses this to compute the plaintexts to return to $\mathcal{A}$.

Observe that when $b = 0$, the reduction's behavior differs from the behavior of $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B)$ if and only if it sends a query to the BIBE challenger which results in the BIBE challenger aborting. A BIBE abort happens in three cases:

- During a pre-challenge query, a batch label t is given that has already been used. *The reduction, following the logic of* $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}$*, checks this is not the case before issuing a BIBE query, so this will never happen.*
- During the challenge round, the t^* and id^* given to the BIBE challenger are such that a key computation query has already been made w.r.t. t^* and including id^*. *The reduction, following the logic of* $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}$*, aborts if* $\mathsf{id}^* = \mathsf{id}_i = H(\mathsf{vk}_i^{\mathsf{Sign}})$ *for some* $\mathsf{vk}_i^{\mathsf{Sign}}$ *given by the adversary during the pre-challenge queries, so this will never happen.*
- During a post-challenge query, a batch label t is given that has already been used, or $t = t^*$ and the challenge id id^* is in the list ids sent to the BIBE challenger. *The reduction, following the logic of* $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}$*, would have output "FAIL 1" or "FAIL 2" before querying the BIBE oracle if this were the case, so this will never happen.*

Because of this, the reduction has identical behavior to $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},0}(1^\lambda, B)$ when $b = 0$. In the same way, the reduction has identical behavior to $\mathsf{Expt}^{\mathsf{mempool},2}_{\mathcal{A},1}(1^\lambda, B)$ when $b = 1$. This means that an adversary who can distinguish these two experiments causes the reduction to contradict security of the BIBE scheme.

□

Corollary 7.4.1. *Assuming the existence of a Thresholdizable Batched Identity Based Encryption scheme* TBIBE *(Fig. 5), a collision-resistant hash function family, and a secure signature scheme* Sign*, there exists a Thresholdizable Batched Encryption scheme* TBE*.*

7.5 Combining non-malleability and batched decryption

When combining the extensions described in Sections 7.3 and 7.4, we need to modify our Batched IBE construction slightly to make these extensions mutually compatible. This is because one of

the batch decryption extensions described in Section 7.3 requires the identity space to be B^{th} roots of unity, where $B = \mathsf{poly}(\lambda)$ is the batch size. On the other hand, the non-malleability extension described in Section 7.4 requires the identity space to be the output space of a hash function, which must be of super-polynomial (in λ) size to guarantee security.

The modification is simple: Instead of encoding the identities $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\}$ as the roots of polynomial when computing the digest, we encode them as evaluation points (which is in fact the standard way of encoding a vector when computing the KZG [KZG10] polynomial commitment). Specifically, each identity id will now be a pair of field elements $(\mathsf{id}_x, \mathsf{id}_y)$ where we interpret the first element as the x co-ordinate and the second element as the y co-ordinate. We describe this alternative construction in Fig. 4. The security proof for this modified construction is similar to the security proof of our main construction.

This modification allows us to simultaneously apply the extensions described in Sections 7.3 and 7.4. To enable the batch decryption extension Section 7.3, we can restrict the space of id_x to be B^{th} roots of unity. To enable the non-malleability extension Section 7.4, we can set the space of id_y to be the output space of a hash function H and derive $\mathsf{id}_y := H(\mathsf{vk}^{\mathsf{Sign}})$ as described in Fig. 3.

8 Concrete Performance

In this section, we discuss the concrete performance of our scheme. We have implemented our Batched IBE construction (Fig. 1) in rust, using the `arkworks` framework [con22] and the BLS12-381 curve. The benchmarks were run using a Google Cloud VM of type `t2d-standard-4`, with a four-core AMD EPYC Milan CPU and 16GB of RAM. All benchmarks were run with parallelism enabled. Table 3 shows the time taken to compute a digest, to compute a decryption key, to encrypt, and to decrypt, with respect to several different batch sizes. We have omitted benchmarking the setup time, since our setup is simply a KZG powers-of-tau setup plus a BLS key-generation, both of which are standard and have been well-studied.

Batch Size	Digest	ComputeKey	Encrypt	Decrypt
100	11.5 ms	720 µs	6.4 ms	9.1 ms
1,000	104.4 ms	643 µs	6.5 ms	88.7 ms
10,000	877 ms	681 µs	6.4 ms	778.5 ms
100,000	8.6 s	759 µs	6.4 ms	8.6 s

Table 3: Running times for a *single* invocation of different procedures in our Batched IBE scheme for varying batch sizes.

Mempool Privacy: Comparison with [Cho+24a] and related works. In addition to implementing the vanilla version of our scheme, we also implemented the extensions in Sections 7.1 and 7.3 to 7.5, in order to compare the performance of our scheme with that of [Cho+24a]. That is, we implemented a version of our scheme with threshold decryption key computation, with batched decryption using [FK23] (as described in Fig. 2), and with signature verification over the submitted ciphertexts (as described in Fig. 3). We used the `ed25519-dalek` library for signatures.

BIBE **construction**

- $\mathsf{Setup}(1^\lambda, 1^B)$: Output three groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of order p, where p is a λ-bit prime, equipped with generators g_1, g_2, g_T, respectively, and an efficiently computable pairing operation $\circ : \mathbb{G}_2 \times \mathbb{G}_1 \to \mathbb{G}_T$. Set the message space $\mathcal{M} := \mathbb{G}_T$, identity space $\mathcal{I} := \mathcal{I}_1 \times \mathcal{I}_2$ where $\mathcal{I}_1$ is the B sized subset of $\mathbb{Z}_p$ containing B^{th} roots of unity and $\mathcal{I}_2 := \{0, \ldots, p-1\}$, and batch label space $\mathcal{T} := \{0,1\}^\lambda$. Also output a randomly sampled hash function $H : \mathcal{T} \to \mathbb{G}_1$.
- $\mathsf{KeyGen}(\mathsf{params})$: Sample $\mathsf{msk} \leftarrow \mathbb{Z}_p$ and $\tau \leftarrow \mathbb{Z}_p$. Output $\mathsf{msk}, \mathsf{pk} := ([\tau]_1, \ldots, [\tau^{B-1}]_1, [\tau]_2, [\mathsf{msk}]_2)$.
- $\mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t)$: Parse id as $(\mathsf{id}_x, \mathsf{id}_y)$. Let $\mathbf{A}$ be a matrix in $(\mathbb{G}_2)^{2\times 3}$ and $\mathbf{b}$ be a vector in $(\mathbb{G}_T)^2$, defined as follows.

$$\mathbf{A} := \begin{pmatrix} [1]_2 & [\mathsf{id}_x]_2 - [\tau]_2 & 0 \\ [\mathsf{msk}]_2 & 0 & -[1]_2 \end{pmatrix}$$

$$\mathbf{b} := \begin{pmatrix} [1]_2 \circ [\mathsf{id}_y]_1 \\ -([\mathsf{msk}]_2 \circ H(t)) \end{pmatrix}$$

  Sample a (column) vector $\mathbf{r} = (r_1, r_2) \leftarrow (\mathbb{Z}_p)^2$ and output the ciphertext c where

$$c = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b} + m)$$

- $\mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\})$: Let $f(X) = \sum_{i=0}^{B-1} f_i \cdot X^i$ be a univariate polynomial of degree $B-1$ over $\mathbb{Z}_p$ with evaluation points as $\mathsf{id}_1, \ldots, \mathsf{id}_B$. Output digest $d := \sum_{i=0}^{B} f_i \cdot [\tau^i]_1$.
- $\mathsf{ComputeKey}(\mathsf{msk}, d, t)$: Output the secret key $\mathsf{sk} := \mathsf{msk} \cdot (d + H(t))$.
- $\mathsf{Decrypt}(c, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, \mathsf{id}, t)$: Parse each id_i as $(\mathsf{id}_{i_x}, \mathsf{id}_{i_y})$ and id as $(\mathsf{id}_x, \mathsf{id}_y)$. Let $S = \{(x,y) | (\mathsf{id}_{i_x}, \mathsf{id}_{i_y}) \in \{\mathsf{id}_1, \ldots, \mathsf{id}_B\} \setminus \{\mathsf{id}\} \wedge x := \mathsf{id}_{i_x} \wedge y := (\mathsf{id}_{i_y} - \mathsf{id}_y) \cdot (\mathsf{id}_{i_x} - \mathsf{id}_x)^{-1}$. Let $q(X) = \sum_{i=0}^{B-2} q_i \cdot X^i$ be a univariate polynomial of degree $B-2$ with evaluation points being elements of S. Set $\pi := \sum_{i=0}^{B-2} q_i \cdot [\tau^i]_1$ and set $\mathbf{w}$ to be the following vector.

$$\mathbf{w} = \begin{pmatrix} d \\ \pi \\ \mathsf{sk} \end{pmatrix}$$

  Finally, parse c as (c_1, c_2) and output the decrypted message $m^* := c_2 - c_1 \circ \mathbf{w}$.

Figure 4: Alternative construction for Batched Identity Based Encryption by encoding identities as evaluation points (instead of roots) of the digest polynomial with the difference from the original construction (Fig. 1) highlighted in blue. This construction supports the extension described in Section 7.5.

Batch Size	[Cho+24a] w/o Setup	[Cho+24a] w/ Setup	Ours
8	83.218 ms	18.08 s	30.96 ms
32	337.5 ms	18.34 s	114.7 ms
128	1.422 s	19.42 s	462.1 ms
512	6.02 s	24.02 s	1.92 s

Table 4: Total time to decrypt an entire batch of ciphertexts for the mempool privacy application for varying batch sizes. We compare with [Cho+24a] where "Setup" refers to the per-batch setup phase needed in the construction of [Cho+24a]. Our scheme doesn't have such a per-batch setup phase.

In order to get an accurate comparison, we reran the benchmarks from [Cho+24a] on the same Google Cloud VM of type `t2d-standard-4` which we used for benchmarking our scheme, using their publicly-available source code.[10]. We instantiated our scheme with the same threshold parameters (number of servers $n = 16$, corruption threshold $f = 8$) as in their benchmark, and with IDs chosen as roots of unity to enable the fast version of [FK23] during decryption. In Table 4 we show, for both our scheme and [Cho+24a], the total *computation* time for a single batch decryption (i.e. total time to decrypt a batch of B ciphertexts). Recall that [Cho+24a] requires an expensive, per-batch-decryption setup phase. The MPC protocol for computing this setup phase was not implemented by them, as far as we know, but [Cho+24a] estimates in their paper that it would take around 18 seconds. We have included the time for their scheme both with and without this setup cost.

9 Conclusion and Open Problems

We proposed a new cryptographic primitive called "batched identity-based encryption" (Batched IBE) and its thresholdized version. We provided an efficient construction for this primitive using pairing friendly groups along with useful extensions, applications and implementation benchmarks. We now discuss some interesting future directions.

- **Proving security in non-idealized models.** We prove the security of our construction in the Generic Group Model (GGM). Currently, it is unclear whether the security of this scheme can be reduced to a simple "DLOG-style" computational assumption without resorting to idealized models.

- **Removing the need for batch label.** In our Batched IBE scheme, each ciphertext is tied not only to an identity but also to a batch label. This means that if $\mathsf{ct}_{\mathsf{id},t}$ is a ciphertext created w.r.t identity id and batch label t, then $\mathsf{ct}_{\mathsf{id},t}$ can only be decrypted in batch t and not in other batches. This can be seen as a shortcoming in some applications, for example mempool privacy, where one could hope to include the ciphertexts (representing encrypted transactions) in a future batch t' (representing a future block) if they were not included in

[10]URL is `https://github.com/guruvamsi-policharla/batched-threshold-encryption`.

the current batch/block t. In our construction, this would require generating two different independent ciphertexts, $\mathsf{ct}_{\mathsf{id},t}$ and $\mathsf{ct}_{\mathsf{id},t'}$, tied to batch labels t and t', respectively.

Ideally, a scheme where the ciphertexts are not linked to the batch or block number would be the best way to handle this issue. This is an interesting problem that we currently lack a solution for our specific construction approach. Having said that, we note that in our construction users can proactively submit several ciphertexts with several sequential block numbers at once (rather than waiting for the non-inclusion event and then re-encrypting). The exact number of these ciphertexts can be set based on the probability of block inclusion for a transaction (which depends on how in-demand block space is). Given that the encryption time is quite fast (≈ 6.5 ms) and ciphertext sizes are relatively small (≈ 864 bytes), we expect the performance overhead to be reasonable unless the block inclusion probability is very low.

- **Support for weighted thresholds.** As described in Section 7.1 and Section A, our construction can be extended to a threshold setting where the trust of key issuance is distributed across n authorities out of which f can be corrupt. It would be interesting to explore whether we can efficiently extend the construction to a *weighted* threshold setting where each authority $i \in [n]$ has a prescribed weight $w_i \in \mathbb{N}$ and the usual threshold f is replaced with a threshold in terms of total weight value. This modeling makes sense for a Proof of Stake (PoS) blockchain where each validator has a specific stake which represents its "weight".

- **Reducing/Eliminating the setup phase.** Our construction requires a one-time setup phase which generates "powers-of-tau" and a public-private key pair. In the threshold setting, this setup would have to be performed in a distributed fashion. While there are many known efficient protocols for performing this task [Gen+99; FS01; Das+22; NBBR16; Nik+24; WCB25; DXR24], and there is also a possibility of reusing existing setups such as [Eth; Zfn], it would be interesting to explore whether the setup phase can be reduced/eliminated by leveraging the silent-setup techniques such as in [Gar+24].

Acknowledgements

We would like to thank Alin Tomescu and Andrei Tonkikh for many useful discussions related to this project.

References

[Bau+23] Balthazar Bauer, Pooya Farshim, Patrick Harasser, and Markulf Kohlweiss. "The uber-knowledge assumption: A bridge to the AGM". In: *Cryptology ePrint Archive* (2023).

[BBG05] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. "Hierarchical identity based encryption with constant size ciphertext". In: *Annual international conference on the theory and applications of cryptographic techniques*. Springer. 2005, pp. 440–456.

[BF01] Dan Boneh and Matt Franklin. "Identity-based encryption from the Weil pairing". In: *Annual international cryptology conference*. Springer. 2001, pp. 213–229.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. "Short signatures from the Weil pairing". In: *International conference on the theory and application of cryptology and information security*. Springer. 2001, pp. 514–532.

[BO22] Joseph Bebel and Dev Ojha. "Ferveo: Threshold decryption for mempool privacy in bft networks". In: *Cryptology ePrint Archive* (2022).

[Bol03] Alexandra Boldyreva. *Efficient threshold signature, multisignature and blind signature schemes based on the Gap-Diffie-Hellman-group signature scheme, PKC 2003, LNCS 2139*. 2003.

[Bor+24] Jan Bormet, Sebastian Faust, Hussien Othman, and Ziyan Qu. "BEAT-MEV: Epochless Approach to Batched Threshold Encryption for MEV Prevention". In: *Cryptology ePrint Archive* (2024).

[Cam+21] Matteo Campanelli, Bernardo David, Hamidreza Khoshakhlagh, Anders Konring, and Jesper Buus Nielsen. *Encryption to the Future: A Paradigm for Sending Secret Messages to Future (Anonymous) Committees*. Cryptology ePrint Archive, Paper 2021/1423. 2021. URL: https://eprint.iacr.org/2021/1423.

[Cer+23] Andrea Cerulli, Aisling Connolly, Gregory Neven, Franz-Stefan Preiss, and Victor Shoup. "Vetkeys: How a blockchain can keep many secrets". In: *Cryptology ePrint Archive* (2023).

[CG99] Ran Canetti and Shafi Goldwasser. "An efficient threshold public key cryptosystem secure against adaptive chosen ciphertext attack". In: *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 1999, pp. 90–106.

[Chi+20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. "Marlin: Preprocessing zkSNARKs with Universal and Updatable SRS". In: *Advances in Cryptology - EUROCRYPT 2020*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12105. Lecture Notes in Computer Science. Springer, 2020, pp. 738–768.

[Cho+17] Arka Rai Choudhuri, Matthew Green, Abhishek Jain, Gabriel Kaptchuk, and Ian Miers. "Fairness in an unfair world: Fair multiparty computation from public bulletin boards". In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 2017, pp. 719–728.

[Cho+24a] Arka Rai Choudhuri, Sanjam Garg, Julien Piet, and Guru-Vamsi Policharla. "Mempool Privacy via Batched Threshold Encryption: Attacks and Defenses". In: *Cryptology ePrint Archive* (2024).

[Cho+24b] Arka Rai Choudhuri, Sanjam Garg, Guru-Vamsi Policharla, and Mingyuan Wang. "Practical Mempool Privacy via One-time Setup Batched Threshold Encryption". In: *Cryptology ePrint Archive* (2024).

[Cle86] Richard Cleve. "Limits on the security of coin flips when half the processors are faulty". In: *Proceedings of the eighteenth annual ACM symposium on Theory of computing*. 1986, pp. 364–369.

[con22] Arkworks contributors. *arkworks zkSNARK ecosystem*. 2022. URL: https://arkworks.rs.

[Dai+20] 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, SP 2020, San Francisco, CA, USA, May 18-21, 2020*. IEEE, 2020, pp. 910–927. DOI: 10.1109/SP40000.2020.00040. URL: https://doi.org/10.1109/SP40000.2020.00040.

[Dal+20] Anders Dalskov, Claudio Orlandi, Marcel Keller, Kris Shrishak, and Haya Shulman. "Securing DNSSEC keys via threshold ECDSA from generic MPC". In: *Computer Security–ESORICS 2020: 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14–18, 2020, Proceedings, Part II 25*. Springer. 2020, pp. 654–673.

[Dam+12] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. "Multiparty computation from somewhat homomorphic encryption". In: *Annual Cryptology Conference*. Springer. 2012, pp. 643–662.

[Das+22] Sourav Das, Thomas Yurek, Zhuolun Xiang, Andrew Miller, Lefteris Kokoris-Kogias, and Ling Ren. "Practical asynchronous distributed key generation". In: *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2022, pp. 2518–2534.

[Döt+23] Nico Döttling, Lucjan Hanzlik, Bernardo Magri, and Stella Wohnig. "McFly: verifiable encryption to the future made practical". In: *International Conference on Financial Cryptography and Data Security*. Springer. 2023, pp. 252–269.

[DXR24] Sourav Das, Zhuolun Xiang, and Ling Ren. "Powers of Tau in Asynchrony". In: *NDSS*. 2024.

[ElG86] Taher ElGamal. "On computing logarithms over finite fields". In: *Advances in Cryptology—CRYPTO'85 Proceedings 5*. Springer. 1986, pp. 396–402.

[Eth] *Ethereum Ceremony*. URL: https://ceremony.ethereum.org/.

[FK23] Dankrad Feist and Dmitry Khovratovich. "Fast amortized KZG proofs". In: *IACR Cryptol. ePrint Arch.* (2023), p. 33. URL: https://eprint.iacr.org/2023/033.

[FS01] Pierre-Alain Fouque and Jacques Stern. "One round threshold discrete-log key generation without private channels". In: *Public Key Cryptography: 4th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2001 Cheju Island, Korea, February 13–15, 2001 Proceedings 4*. Springer. 2001, pp. 300–316.

[Gar+13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. "Witness encryption and its applications". In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 2013, pp. 467–476.

[Gar+16] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. "Candidate indistinguishability obfuscation and functional encryption for all circuits". In: *SIAM Journal on Computing* 45.3 (2016), pp. 882–929.

[Gar+24] Sanjam Garg, Dimitris Kolonelos, Guru-Vamsi Policharla, and Mingyuan Wang. "Threshold Encryption with Silent Setup". In: *Annual International Cryptology Conference*. Springer. 2024, pp. 352–386.

[Gen+99] Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. "Secure distributed key generation for discrete-log based cryptosystems". In: *Advances in Cryptology—EUROCRYPT'99: International Conference on the Theory and Application of Cryptographic Techniques Prague, Czech Republic, May 2–6, 1999 Proceedings 18*. Springer. 1999, pp. 295–310.

[GMR23] Nicolas Gailly, Kelsey Melissaris, and Yolan Romailler. "tlock: Practical timelock encryption from threshold bls". In: *Cryptology ePrint Archive* (2023).

[Gol09] Oded Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, 2009.

[Kav+23] 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. URL: https://eprint.iacr.org/2023/1061.

[KZG10] Aniket Kate, Gregory M Zaverucha, and Ian Goldberg. "Constant-size commitments to polynomials and their applications". In: *Advances in Cryptology-ASIACRYPT 2010: 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings 16*. Springer. 2010, pp. 177–194.

[LPS23] Helger Lipmaa, Roberto Parisella, and Janno Siim. "Algebraic group model with oblivious sampling". In: *Theory of Cryptography Conference*. Springer. 2023, pp. 363–392.

[NBBR16] Wafa Neji, Kaouther Blibech, and Narjes Ben Rajeb. "Distributed key generation protocol with a new complaint management strategy". In: *Security and communication networks* 9.17 (2016), pp. 4585–4595.

[Nik+24] Valeria Nikolaenko, Sam Ragsdale, Joseph Bonneau, and Dan Boneh. "Powers-of-tau to the people: Decentralizing setup ceremonies". In: *International Conference on Applied Cryptography and Network Security*. Springer. 2024, pp. 105–134.

[RSW96] Ronald L Rivest, Adi Shamir, and David A Wagner. *Time-lock puzzles and timed-release Crypto*. Technical Report MIT-LCS-TR-684. Massachusetts Institute of Technology, 1996. URL: https://people.csail.mit.edu/rivest/pubs/RSW96.pdf.

[SA19] Nigel P Smart and Younes Talibi Alaoui. "Distributing any Elliptic Curve Based Protocol: With an Application to MixNets." In: *IACR Cryptol. ePrint Arch.* 2019 (2019), p. 768.

[SAS24] Sora Suegami, Shinsaku Ashizawa, and Kyohei Shibano. *Constant-Cost Batched Partial Decryption in Threshold Encryption*. Cryptology ePrint Archive, Paper 2024/762. 2024. URL: https://eprint.iacr.org/2024/762.

[Sch80] Jacob T Schwartz. "Fast probabilistic algorithms for verification of polynomial identities". In: *Journal of the ACM (JACM)* 27.4 (1980), pp. 701–717.

[Sho97] Victor Shoup. "Lower bounds for discrete logarithms and related problems". In: *Advances in Cryptology—EUROCRYPT'97: International Conference on the Theory and Application of Cryptographic Techniques Konstanz, Germany, May 11–15, 1997 Proceedings 16*. Springer. 1997, pp. 256–266.

[Sue24] Sora Suegami. "Extractable Witness Encryption for Signed Vector Digests from Pairings and Trust-Scalable One-Time Programs". In: *Cryptology ePrint Archive* (2024).

[Tsa22] Rotem Tsabary. "Candidate witness encryption from lattice techniques". In: *Annual International Cryptology Conference*. Springer. 2022, pp. 535–559.

[VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. "Witness encryption and null-IO from evasive LWE". In: *International Conference on the Theory and Application of Cryptology and Information Security*. Springer. 2022, pp. 195–221.

[WCB25] Faxing Wang, Shaanan Cohney, and Joseph Bonneau. "SoK: Trusted setups for powers-of-tau strings". In: *Cryptology ePrint Archive* (2025).

[Zfn] *Conclusion of the Powers of Tau Ceremony*. URL: `https://zfnd.org/conclusion-of-the-powers-of-tau-ceremony/`.

[Zha22] Mark Zhandry. "To label, or not to label (in generic groups)". In: *Annual International Cryptology Conference*. Springer. 2022, pp. 66–96.

[Zip79] Richard Zippel. "Probabilistic algorithms for sparse polynomials". In: *International symposium on symbolic and algebraic manipulation*. Springer. 1979, pp. 216–226.

A Thresholdizable Batched Identity Based Encryption

In this section, we describe a threshold version of the Batched IBE scheme which we defined in Thm. 5.1. At a high-level, this threshold version is aimed to capture settings where one would like to distribute the trust of key issuance across multiple authorities among which the adversary can corrupt at most some threshold number of authorities. Specifically, instead of a single authority holding the complete msk, we would like to have multiple, let's say n, authorities where each authority $i \in [n]$ holds a "partial" master secret key msk_i (e.g., in the form of secret shares of msk). Accordingly, the ComputeKey procedure (as defined in Batched IBE) will be split into two parts: 1) ComputeKeyShare algorithm which will be used by each authority to produce a *partial* secret key w.r.t a specific batch using its private partial master secret key msk_i, 2) ComputeKeyAggregate algorithm which can be used to combine the partial secret key w.r.t a specific batch into a full secret key.

In the following sections, we will formally define the syntax and semantics of Thresholdizable Batched IBE . We will then present a construction of Thresholdizable Batched IBE (which is a straightforward adaptation of our Batched IBE construction) and analyze it. For the ease of readability, we will highlight all the differences between Thresholdizable Batched IBE and Batched IBE in blue.

A.1 Syntax

Definition A.1 (Thresholdizable Batched IBE Syntax)**.** A Thresholdizable Batched IBE scheme TBIBE is specified by seven algorithms: Setup, KeyGen, Encrypt, Decrypt, Digest, ComputeKeyShare, ComputeKeyAggregate.

- $\mathsf{Setup}(1^\lambda, 1^B, 1^n, 1^f) \to \mathsf{params}$: A randomized algorithm that takes as input a security parameter $\lambda \in \mathbb{N}$, a batch size $B = B(\lambda)$, number of authorities $n = n(\lambda)$ and corruption threshold $f = f(\lambda)$. It outputs params (system parameters) which includes a description of the message space $\mathcal{M}$, identity space $\mathcal{I}$, batch label space $\mathcal{T}$ and ciphertext space $\mathcal{C}$.
- $\mathsf{KeyGen}(\mathsf{params}) \to (\{\mathsf{msk}_i\}_{i\in[n]}, \{\mathsf{pk}_i\}_{i\in[n]}, \mathsf{pk})$: a randomized algorithm that takes as input params and outputs n many msk_i (partial master secret key), n many pk_i (partial public key) and a single pk (global public key).
- $\mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t) \to c$: a randomized algorithm that takes as input a message $m \in \mathcal{M}$, an identity $\mathsf{id} \in \mathcal{I}$, a batch label $t \in \mathcal{T}$, global public key pk and outputs a ciphertext $c \in \mathcal{C}$.
- $\mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}) \to d$: a deterministic algorithm that takes as input the global public key pk and a list of identities $\mathsf{id}_1, \ldots, \mathsf{id}_B$ where each $\mathsf{id}_i \in \mathcal{I}$. It outputs a digest d.
- $\mathsf{ComputeKeyShare}(\mathsf{msk}_i, d, t) \to \mathsf{sk}_i$: a deterministic algorithm that takes as input the partial master secret key msk_i, digest d, batch label t and outputs a partial digest-batch label-specific secret key sk_i.
- $\mathsf{ComputeKeyAggregate}(\{\mathsf{pk}_i\}_{i\in[n]}, \{\mathsf{sk}_i\}_{i\in[n]}, d, t) \to \mathsf{sk}$: a deterministic algorithm that takes as input all the partial public keys $\{\mathsf{pk}_i\}_{i\in[n]}$ and all the partial digest-batch label-specific secret key $\{\mathsf{sk}_i\}_{i\in[n]}$ and outputs a digest-batch label-specific secret key sk.

- $\mathsf{Decrypt}(c, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, \mathsf{id}, t) \to m$: a deterministic algorithm that takes as input a ciphertext c, secret key sk, digest d, a list of identities $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and an identity-batch label pair (id, t). It outputs a message $m \in \mathcal{M}$.

A.2 Correctness, Non-triviality and Security

The above algorithms should satisfy the following requirements.

For correctness, we generalize the correctness requirement of the (non-threshold) Batched IBE to allow the adversary to (statically) corrupt at most f out of n authorities, where $f \leq \lfloor \frac{n-1}{2} \rfloor$ is the corruption threshold, and get their partial master secret-keys, and observe the partial secret-keys issued by uncorrupted authorities for arbitrary inputs of the adversary's choice.

Definition A.2 (Thresholdizable Batched IBE Correctness)**.** For all $\lambda \in \mathbb{N}, B \in \mathbb{N}, n \in \mathbb{N}, f \leq \lfloor \frac{n-1}{2} \rfloor, m \in \mathcal{M}, t \in \mathcal{T}, \mathsf{id} \in \mathcal{I}, S \subseteq \mathcal{I}$ s.t. $|S| = B$ and $\mathsf{id} \in S$, $\mathsf{Cor} \subset [n]$ s.t. $|\mathsf{Cor}| \leq f$ and for any unbounded adversary $\mathcal{A}$, the following should hold:

$$\Pr\left[\mathsf{Decrypt}(c, \mathsf{sk}, d, S, \mathsf{id}, t) = m \;\middle|\; \begin{array}{c} \mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B, 1^n) \\ (\{\mathsf{msk}_i\}_{i\in[n]}, \{\mathsf{pk}_i\}_{i\in[n]}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(\mathsf{params}) \\ c \leftarrow \mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t) \\ d \leftarrow \mathsf{Digest}(\mathsf{pk}, S) \\ \forall i \in [n] \setminus \mathsf{Cor}, \mathsf{sk}_i \leftarrow \mathsf{ComputeKeyShare}(\mathsf{msk}_i, d, t) \\ \forall i \in \mathsf{Cor}, \mathsf{sk}_i \leftarrow \mathcal{A}^{\{\mathsf{ComputeKeyShare}(\mathsf{msk}_k, \cdot, \cdot)\}_{k\in[n]\setminus\mathsf{Cor}}}(\{\mathsf{msk}_j\}_{j\in\mathsf{Cor}}, \\ \{\mathsf{pk}_j\}_{j\in[n]\setminus\mathsf{Cor}}, \mathsf{pk}, m, \mathsf{id}, t, c, S) \\ \mathsf{sk} \leftarrow \mathsf{ComputeKeyAggregate}(\{\mathsf{pk}_i\}_{i\in[n]}, \{\mathsf{sk}_i\}_{i\in[n]}, d, t) \end{array}\right] = 1$$

Definition A.3 (Thresholdizable Batched IBE Non-triviality/Efficiency)**.** We require that the running time of $\mathsf{ComputeKeyShare}$ and $\mathsf{ComputeKeyAggregate}$ be independent of the batch size B (which implies that the digest d and sk are also independent of B). We also require that the running time of $\mathsf{ComputeKeyShare}$ and the size of sk be independent of the number of authorities n.

For security, we generalize the security requirement of the (non-threshold) Batched IBE to allow the adversary to (statically) corrupt at most f out of n authorities, where $f \leq \lfloor \frac{n-1}{2} \rfloor$, get their partial master secret-keys, and observe the partial secret-keys issued by uncorrupted authorities for arbitrary inputs of the adversary's choice while constrained to the same rules as defined earlier in the non-threshold version.

Definition A.4 (Thresholdizable Batched IBE Security)**.** We define a security game $\mathsf{Expt}^{\mathsf{TBIBE}}_{\mathcal{A},b}(1^\lambda, B, n)$ with respect to adversary $\mathcal{A}$ in the box below.

We say that a Thresholdizable Batched IBE scheme is secure if for all $n \in \mathbb{N}, f \leq \lfloor \frac{n-1}{2} \rfloor, B \in \mathbb{N}$, for all PPT adversaries $\mathcal{A}$, for all $\mathsf{Cor} \subset [n]$ s.t. $|\mathsf{Cor}| \leq f$, there exists some negligible function $\epsilon_{\mathcal{A}}$ such that the following holds:

$$\left|\Pr[\mathsf{Expt}^{\mathsf{TBIBE}}_{\mathcal{A},0}(1^\lambda, B, n, f, \mathsf{Cor}) = 1] - \Pr[\mathsf{Expt}^{\mathsf{TBIBE}}_{\mathcal{A},1}(1^\lambda, B, n, f, \mathsf{Cor}) = 1]\right| < \epsilon_{\mathcal{A}}(\lambda).$$

The security game $\mathsf{Expt}^{\mathsf{TBIBE}}_{\mathcal{A},b}(1^\lambda, B, n, f, \mathsf{Cor})$.

Setup: The challenger takes as input the security parameter λ and the batch size B. It runs $\mathsf{params} \leftarrow \mathsf{Setup}(1^\lambda, 1^B, 1^n, 1^f)$, and then runs
$(\{\mathsf{msk}_i\}_{i\in[n]}, \{\mathsf{pk}_i\}_{i\in[n]}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(\mathsf{params})$.
Finally, it sends $(\mathsf{params}, \mathsf{pk}, \{\mathsf{msk}_i\}_{i\in\mathsf{Cor}}, \{\mathsf{pk}_i\}_{i\in[n]})$ to $\mathcal{A}$.

The rest of the game proceeds in rounds, as follows.
Pre-challenge queries: $\mathcal{A}$ may issue an arbitrary number of *key computation queries:*

- $\mathcal{A}$ sends a list ids of B identities along with a batch label t to the challenger.
- If a key computation query has already been made with batch label t, the challenger halts the game.
- Otherwise, the challenger does the following:
  - Compute a digest $d \leftarrow \mathsf{Digest}(\mathsf{pk}, \mathsf{ids})$ of the ids in ids using public key pk.
  - For all $i \in [n] \setminus \mathsf{Cor}$, compute a partial secret key $\mathsf{sk}_i \leftarrow \mathsf{ComputeKeyShare}(\mathsf{msk}_i, d, t)$, using the digest d computed from the previous step.
  - Send $\{\mathsf{sk}_i\}_{i\in[n]\setminus\mathsf{Cor}}$ to $\mathcal{A}$.

Challenge round: Once during the game, $\mathcal{A}$ may decide that the current round is the *challenge round*. The challenge round proceeds as follows:

- $\mathcal{A}$ sends two messages $m_0, m_1 \in \mathcal{M}$ and an identity-batch label pair (id^*, t^*) on which it wishes to be challenged.
- If key computation query (ids, t) has already been made with batch label $t = t^*$ and where $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger computes $c \leftarrow \mathsf{Encrypt}(\mathsf{pk}, m_b, \mathsf{id}^*, t^*)$ and sends c to $\mathcal{A}$.

Post-challenge queries: After the challenger round, $\mathcal{A}$ may again issue an arbitrary number of *key computation queries*, with the additional restriction that A cannot query (t^*, ids) with $\mathsf{id}^* \in \mathsf{ids}$:

- $\mathcal{A}$ sends a list ids of B identities along with a batch label t to the challenger.
- If a key computation query has already been made with batch label t **or if** $t = t^*$ **and** $\mathsf{id}^* \in \mathsf{ids}$, the challenger halts the game.
- Otherwise, the challenger does the following:
  - Compute a digest $d \leftarrow \mathsf{Digest}(\mathsf{pk}, \mathsf{ids})$ of the ids in ids using public key pk.

– For all $i \in [n] \setminus \mathsf{Cor}$, compute a partial secret key $\mathsf{sk}_i \leftarrow \mathsf{ComputeKeyShare}(\mathsf{msk}_i, d, t)$, using the digest d computed from the previous step.

– Send $\{\mathsf{sk}_i\}_{i \in [n] \setminus \mathsf{Cor}}$ to $\mathcal{A}$.

Output: At any point in time, $\mathcal{A}$ can decide to halt and output a bit $b' \in \{0, 1\}$. The game then halts with the same output b'.

A.3 Construction

Theorem A.5. *Assuming Type-3 pairing group $\mathcal{BG}$, there exists a construction (Fig. 5) for Thresholdizable Batched IBE which is secure in the Generic group model.*

A.4 Correctness

The correctness follows directly from the correctness of our non-threshold version of the scheme along with the robustness property of threshold BLS signatures [Bol03] which ensures that any $f + 1$ valid partial signatures (which corresponds to partial digest-batch label-specific secret keys sk_i in our scheme) suffice to reconstruct the aggregate signature (which corresponds to the digest-batch label-specific secret key sk in our scheme). Since $f \leq \lfloor \frac{n-1}{2} \rfloor$ in our setting, we will always have a set of $n - f \geq f + 1$ valid partial sk_i's to reconstruct the final sk.

A.5 Security

We will prove the security of our construction in the GGM model equipped with oblivious sampling. In this model, the challenger will implement the group oracle and the hash oracle (along with an oracle for key computation queries as defined in the security game).

Theorem A.6. *For all $B = \mathsf{poly}(\lambda), n = \mathsf{poly}(\lambda), f \leq \lfloor \frac{n-1}{2} \rfloor$, for all $\mathsf{Cor} \subset [n]$ s.t. $|\mathsf{Cor}| \leq f$, and all unbounded adversaries $\mathcal{A}$ making at most $q = \mathsf{poly}(\lambda)$ queries (including queries to the group oracle, hash oracle, and key computation queries), we have*

$$\left|\Pr[\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}_{\mathcal{A},0}(1^\lambda, B, n, f, \mathsf{Cor}) = 1] - \Pr[\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}_{\mathcal{A},1}(1^\lambda, B, n, f, \mathsf{Cor}) = 1]\right| \leq \mathsf{negl}(\lambda)$$

where $\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}$ refers to the same experiment $\mathsf{Expt}^{\mathsf{TBIBE}}$ as defined in Thm. A.4 except that we specialize it for our specific Construction (Fig. 5) and model it in the GGM, i.e., all the group and hash operations performed by the adversary are simulated by the challenger as defined in the GGM model (Section 4.2). This is done in a manner similar to $\mathsf{Expt}^{\mathsf{BIBE},\mathsf{GGM}}$ in Thm. 6.2.

Proof. To prove the above theorem, we will show that the security of the threshold version of our construction w.r.t $\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}$ can be reduced to the security of the non-threshold version of our construction w.r.t. $\mathsf{Expt}^{\mathsf{BIBE},\mathsf{GGM}}$.

Assume, for the sake of contradiction, that Thm. A.6 is false. Then, there exists $B = \mathsf{poly}(\lambda), n = \mathsf{poly}(\lambda), f \leq \lfloor \frac{n-1}{2} \rfloor$, $\mathsf{Cor} \subset [n]$ s.t. $|\mathsf{Cor}| \leq \lfloor \frac{n-1}{2} \rfloor$, there exists an adversary $\mathcal{A}$ making at most q queries s.t. there exists a polynomial $\mathsf{poly}_{\mathcal{A}}$ where

TBIBE **construction**

- $\mathsf{Setup}(1^\lambda, 1^B, 1^n, 1^f)$: Output three groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of order p, where p is a λ-bit prime, equipped with generators g_1, g_2, g_T, respectively, and an efficiently computable pairing operation $\circ : \mathbb{G}_2 \times \mathbb{G}_1 \to \mathbb{G}_T$. Set the message space $\mathcal{M} := \mathbb{G}_T$, identity space $\mathcal{I} := \{0, \ldots, p-1\}$, and tag space $\mathcal{T} := \{0,1\}^\lambda$. Also output a randomly sampled hash function $H : \mathcal{T} \to \mathbb{G}_1$.

- $\mathsf{KeyGen}(\mathsf{params})$: Sample $\mathsf{msk} \leftarrow \mathbb{Z}_p$ and $\tau \leftarrow \mathbb{Z}_p$. Set $\{\mathsf{msk}_i\}_{i\in[n]} \leftarrow \mathsf{ShamirShare}(\mathsf{msk}, f, n)$, i.e, n shamir shares of msk using a degree f polynomial. For all $i \in [n]$, set $\mathsf{pk}_i := [\mathsf{msk}_i]_2$. Output $\{\mathsf{msk}_i\}_{i\in[n]}, \{\mathsf{pk}_i\}_{i\in[n]}, \mathsf{pk} := ([\tau]_1, \ldots, [\tau^B]_1, [\tau]_2, [\mathsf{msk}]_2)$.

- $\mathsf{Encrypt}(\mathsf{pk}, m, \mathsf{id}, t)$: Let $\mathbf{A}$ be a matrix in $(\mathbb{G}_2)^{2\times 3}$ and $\mathbf{b}$ be a vector in $(\mathbb{G}_T)^2$, defined as follows.

$$\mathbf{A} := \begin{pmatrix} [1]_2 & [\mathsf{id}]_2 - [\tau]_2 & 0 \\ [\mathsf{msk}]_2 & 0 & -[1]_2 \end{pmatrix}$$

$$\mathbf{b} := \begin{pmatrix} [0]_T \\ -([\mathsf{msk}]_2 \circ H(t)) \end{pmatrix}$$

  Sample a (column) vector $\mathbf{r} = (r_1, r_2) \leftarrow (\mathbb{Z}_p)^2$ and output the ciphertext c where

$$c = (\mathbf{r}^T \cdot \mathbf{A}, \mathbf{r}^T \cdot \mathbf{b} + m)$$

- $\mathsf{Digest}(\mathsf{pk}, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\})$: Let $h(X) = \sum_{i=0}^{B} h_i \cdot X^i$ be a univariate polynomial of degree B over $\mathbb{Z}_p$ with roots at $\mathsf{id}_1, \ldots, \mathsf{id}_B$ and leading coefficient 1. Output digest $d := \sum_{i=0}^{B} h_i \cdot [\tau^i]_1$.

- $\mathsf{ComputeKeyShare}(\mathsf{msk}_i, d, t)$: Output the partial secret key $\mathsf{sk}_i := \mathsf{msk}_i \cdot (d + H(t))$.

- $\mathsf{ComputeKeyAggregate}(\{\mathsf{pk}_i\}_{i\in[n]}, \{\mathsf{sk}_i\}_{i\in[n]}, d, t)$: Let $U = \{i | i \in [n], [1]_2 \circ \mathsf{sk}_i = \mathsf{pk}_i \circ (d + H(t))\}$ which represents the set of valid partial secret keys. Let $V = \{v_1, \ldots, v_k\} \subseteq U$ be any arbitrary subset of U of size $k = f + 1$ and let $L_i(0) = \Pi_{j\neq i} \frac{(-v_j)}{(v_i - v_j)}$ be the i^{th} Lagrange coefficient for all $i \in [k]$. Output $\mathsf{sk} = L_1(0) \cdot \mathsf{sk}_{v_1} + \ldots + L_k(0) \cdot \mathsf{sk}_{v_k}$.

- $\mathsf{Decrypt}(c, \mathsf{sk}, d, \{\mathsf{id}_1, \ldots, \mathsf{id}_B\}, \mathsf{id}, t)$: Let $q(X) = \sum_{i=0}^{B-1} q_i \cdot X^i$ be a univariate polynomial of degree $B-1$ with roots at $\{\mathsf{id}_1, \ldots, \mathsf{id}_B\} \setminus \{\mathsf{id}\}$ and leading coefficient 1. Set $\pi := \sum_{i=0}^{B-1} q_i \cdot [\tau^i]_1$ and set $\mathbf{w}$ to be the following vector.

$$\mathbf{w} = \begin{pmatrix} d \\ \pi \\ \mathsf{sk} \end{pmatrix}$$

  Finally, parse c as (c_1, c_2) and output the decrypted message $m^* := c_2 - c_1 \circ \mathbf{w}$.

Figure 5: Our construction for Thresholdizable Batched IBE (TBIBE) where the differences from the non-threshold Batched IBE construction (Fig. 1) are highlighted in blue.

$$\left|\Pr[\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}_{\mathcal{A},0}(1^\lambda, B, n, f, \mathsf{Cor}) = 1] - \Pr[\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}_{\mathcal{A},1}(1^\lambda, B, n, f, \mathsf{Cor}) = 1]\right| > \frac{1}{\mathsf{poly}_{\mathcal{A}}(\lambda)}$$

We will now construct an adversary $\mathcal{B}$ which will contradict Corollary 6.2.1. The adversary $\mathcal{B}$ works as follows:

- Let $\mathsf{Cor}' = \mathsf{Cor} \cup S$ where $S \subseteq [n] \setminus \mathsf{Cor}$ is an arbitrary subset s.t. $|S| = f - |\mathsf{Cor}|$.
- During the Setup and KeyGen phase, $\mathcal{B}$ performs the following steps:
  - Receive $(\mathsf{params}, \mathsf{pk})$ from the challenger.
  - For all $i \in \mathsf{Cor}'$, sample $\mathsf{msk}_i \leftarrow \mathbb{Z}_p$ and perform a $\mathbb{G}_2$ labeling query step using input msk_i to obtain $\mathsf{pk}_i = [\mathsf{msk}_i]_2$.
  - Let $U = \mathsf{Cor}' \cup \{0\}$. For all $u \in U$, let L_u be the lagrange polynomial of degree f.
  - Set $\mathsf{pk}_0 := \mathsf{pk}$.
  - For all $i \in [n] \setminus \mathsf{Cor}'$, perform $\mathbb{G}_2$ group operation queries to obtain $\mathsf{pk}_i := \sum_{\forall u \in U} L_u(i) \cdot \mathsf{pk}_u$.
  - Send $(\mathsf{params}, \mathsf{pk}, \{\mathsf{msk}_i\}_{i \in \mathsf{Cor}}, \{\mathsf{pk}_i\}_{i \in [n]})$ to $\mathcal{A}$.
- During the *key computation queries* (during pre-challenge and post-challenge phases), $\mathcal{B}$ performs the following steps.
  - It receives a list ids of B identities along with a batch label t and forwards it to the challenger.
  - If the challenger halts the game, then $\mathcal{B}$ also halts.
  - Otherwise, $\mathcal{B}$ receives sk from the challenger and performs the following steps.
    - Compute a digest $d \leftarrow \mathsf{Digest}(\mathsf{pk}, \mathsf{ids})$ of the ids in ids using public key pk and $\mathbb{G}_1$ group operation queries.
    - Perform a hash query step to obtain $H(t)$.
    - For $i \in \mathsf{Cor}'$, perform a $\mathbb{G}_2$ group operation query to obtain $\mathsf{sk}_i := \mathsf{msk}_i \cdot (d + H(t))$.
    - Let $U = \mathsf{Cor}' \cup \{0\}$. For all $u \in U$, let L_u be the lagrange polynomial of degree f.
    - Set $\mathsf{sk}_0 := \mathsf{sk}$.
    - For all $i \in [n] \setminus \mathsf{Cor}'$, use $\mathbb{G}_2$ group operation queries to get $\mathsf{sk}_i := \sum_{\forall u \in U} L_u(i) \cdot \mathsf{sk}_u$.
    - Send $\{\mathsf{sk}_i\}_{i \in [n] \setminus \mathsf{Cor}}$ to $\mathcal{A}$.
- During the challenge round, $\mathcal{B}$ forwards transparently forwards the messages received from $\mathcal{A}$ to the challenger and vice-versa.
- For all labeling queries, group operation queries, hash queries and pairing operation queries received from $\mathcal{A}$, forward it to the challenger, and then forward the response received back to $\mathcal{A}$.

By construction of $\mathcal{B}$, the following holds:

$$\Pr[\mathsf{Expt}^{\mathsf{BIBE},\mathsf{GGM}}_{\mathcal{B},0}(1^\lambda, B) = 1] = \Pr[\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}_{\mathcal{A},0}(1^\lambda, B, n, f, \mathsf{Cor}) = 1]$$

$$\Pr[\mathsf{Expt}^{\mathsf{BIBE},\mathsf{GGM}}_{\mathcal{B},1}(1^\lambda, B) = 1] = \Pr[\mathsf{Expt}^{\mathsf{TBIBE},\mathsf{GGM}}_{\mathcal{A},1}(1^\lambda, B, n, f, \mathsf{Cor}) = 1]$$

Moreover, the adversary $\mathcal{B}$ makes at most $q' = \mathsf{poly}(q, n) = \mathsf{poly}(\lambda)$ queries to the challenger. Hence, we get that,

$$\left|\Pr[\mathsf{Expt}^{\mathsf{BIBE},\mathsf{GGM}}_{\mathcal{B},0}(1^\lambda, B) = 1] - \Pr[\mathsf{Expt}^{\mathsf{BIBE},\mathsf{GGM}}_{\mathcal{B},1}(1^\lambda, B) = 1]\right| > \frac{1}{\mathsf{poly}_{\mathcal{A}}(\lambda)}$$

which contradicts Corollary 6.2.1. □