

Time-Lock Encrypted Storage for Blockchains

Amit Agarwal, Kushal Babel, Sourav Das, and Babak Poorebrahim Gilkalaye

Category Labs

Abstract. We introduce time-lock encrypted storage (tTLES), a storage service provided by blockchains. In tTLES, clients store encrypted values towards a future decryption time τ_{tgt} (measured in block height). The security of tTLES requires that a value is decrypted only if (i) the encrypted value is included in the blockchain, and (ii) the time τ_{tgt} has passed. This is crucially different from existing schemes, which only enforce either of these conditions but not both. We formalize tTLES, and present an efficient protocol that relies on (in a black-box manner) a threshold identity-based encryption scheme, and a recent batch threshold decryption scheme. Finally, we discuss various applications that will benefit from tTLES.

1 Introduction

Modern blockchains are increasingly hosting applications ranging from finance and on-chain governance to games and contests that require storing data confidentially for a period of time before it is decrypted and processed [14,19,10].

Consider the example of an on-chain Kaggle*-like Machine Learning model training contest. Here, participants (clients) submit their ML model weights to the on-chain application before a submission deadline, and the models are evaluated after the deadline to distribute rewards to the submitters in proportion to their model performance. The on-chain submissions to the contest need to remain confidential until the submission deadline so that participants do not gain an unfair advantage by seeing the model weights of other participants before the deadline. Moreover, submissions that participants send to the mempool but are not included in the blockchain (e.g., due to congestion or denial-of-service attacks) need to remain private indefinitely, even after the submission deadline, to protect the intellectual property of the clients.

In this work, we are interested in the following two key confidentiality properties of privacy-sensitive blockchain applications that can receive encrypted values (as transactions) as inputs:

- **Time-lock indistinguishability:** Encrypted values that are included in the blockchain need to remain confidential until a specified target time τ_{tgt} (e.g., submission deadline in the above example), after which they are decrypted publicly and processed.
- **Non-inclusion privacy:** Encrypted values that are *not* included in the blockchain (e.g., due to congestion or denial-of-service attacks) need to remain confidential indefinitely, even after the target time τ_{tgt} has passed.

One naive strategy to achieve these properties is to use secret sharing. The client can secret share its storage value among the blockchain validators, with one share given to each validator, who then reconstruct the storage value at the target time τ_{tgt} if it is included in the blockchain, and do not reconstruct it otherwise. This strategy, while achieving the desired properties, is not practical as it requires communication and computation costs that scale linearly with the number of validators and the number of encrypted values, and the computation cost is incurred in an “online manner” at the target time τ_{tgt} .

Existing efficient solutions for time-lock encryption [13] fail to provide non-inclusion privacy, i.e., they do not provide confidentiality for storage values that are not included in the blockchain. At a high level, these solutions use threshold identity-based encryption (tIBE) [2] where clients encrypt storage values with the identity τ_{tgt} , and the decryption key for time τ_{tgt} is derived (in a threshold fashion) from a public parameter that is made available at time τ_{tgt} . While this gives the time-lock indistinguishability property efficiently

*<https://www.kaggle.com/>

with the release of the key material at τ_{tgt} , it does not provide non-inclusion privacy since the decryption key for time τ_{tgt} can be used to decrypt all storage values encrypted with id τ_{tgt} , including those that are not included in the blockchain.

Recent solutions that threshold decrypt a batch of ciphertexts [1,6,9,5] provide non-inclusion privacy, as any ciphertexts that are not included in the batch remain confidential. However, such solutions do not inherently provide the time-lock indistinguishability property. They can be invoked at the target time τ_{tgt} with the input batch comprised of all submitted ciphertexts encrypted towards the target time τ_{tgt} , but these solutions incur a threshold decryption cost that necessarily depends on the size of the batch, and hence are not practical for decrypting at once a large number of values submitted towards τ_{tgt} . This form of decryption cost is incurred in an “online manner” when the threshold decryption is to be performed, and hence can be prohibitive.

1.1 Our contributions

In this work, we have the following contributions:

- We define the notion of (thresholdizable) time-lock encrypted storage (tTLES) (see Section 3), which achieves both the time-lock indistinguishability property and non-inclusion privacy, without incurring an “online” cost at τ_{tgt} that depends on the number of values encrypted towards τ_{tgt} . Instead, the (threshold) decryption cost of decrypting the stored values is amortized over time during the submission of the clients’ encrypted values.
- We provide a concrete construction of time-lock encrypted storage (tTLES) primitive, that relies on a threshold identity-based encryption (tIBE) and batched threshold decryption (BTE) in a black-box manner (see Section 4).
- We discuss some interesting applications of time-lock encrypted storage (tTLES) primitive in the context of blockchains (see Section 1.4).

1.2 Our technique

At a high-level, we construct a threshold time-lock encrypted storage (tTLES) scheme by simply combining, in a black-box manner, a batch threshold decryption primitive (BTE) with a time-lock encryption primitive (tIBE). In a bit more detail, a client that wishes to encrypt a message m towards a target time τ_{tgt} samples a symmetric key s , and encrypts m using this key simply as $m \oplus \text{PRG}(s)$, where PRG is a vanilla pseudorandom generator. The client then derives two secret shares, s_1 and s_2 , of the symmetric key s , i.e., $s_1 \oplus s_2 = s$. It encrypts the first share s_1 using the batch decryption primitive towards time τ_{tgt} and the second share s_2 using the time-lock encryption primitive towards time τ_{tgt} . The final ciphertext ct consists of a tuple of four values $(\tau_{\text{tgt}}, \text{BTE.Enc}(s_1), \text{tIBE.Enc}(s_2, \tau_{\text{tgt}}), m \oplus \text{PRG}(s))$.[†]

The batch threshold decryption component $\text{BTE.Enc}(s_1)$ provides confidentiality for non-included values (since s_1 will be hidden unless ct is included in the blockchain), while the threshold time-lock encryption component $\text{tIBE.Enc}(s_2, \tau_{\text{tgt}})$ gives us the confidentiality for included values until time τ_{tgt} has reached (since s_2 will be hidden until time τ_{tgt}). For efficiency, the batch decryption component is invoked during the submission time (more precisely at the time of inclusion) so as to amortize the cost over the submission window and recover one secret share s_1 of the key s . At the target time τ_{tgt} , decryption is made efficient (independent of the number of values encrypted towards τ_{tgt}) by just releasing (in a threshold fashion) a succinct key $k_{\tau_{\text{tgt}}}$ related to τ_{tgt} in order to recover the second share of the encryption key. More precisely, the tIBE interface provides an efficient threshold key issuance protocol $\Pi\text{-KeyIss}$ which can be used to derive $k_{\tau_{\text{tgt}}} \leftarrow \Pi\text{-KeyIss}(\tau_{\text{tgt}})$. Note that the running time and communication cost of this $\Pi\text{-KeyIss}$ protocol is completely independent of the number of values encrypted towards τ_{tgt} . However, once $k_{\tau_{\text{tgt}}}$ is made publicly

[†]An alternative approach would be to have nested encryption such that $\text{ct} = (\tau_{\text{tgt}}, \text{BTE.Enc}(\text{tIBE.Enc}(s, \tau_{\text{tgt}})), m \oplus \text{PRG}(s))$. We describe the secret-sharing based construction in this paper due to its simplicity.

available, anyone (e.g., a smart contract) can publicly decrypt $\text{tIBE.Enc}(s_2, \tau_{\text{tgt}})$ to recover the second share s_2 .

Our technique maps cleanly to the operation of a blockchain. In every block, clients submit requests to store encrypted values. The blockchain validators can *partially decrypt* these values in every block (i.e. $\text{BTE.Enc}(s_1)$), and store these intermediate decryption values in the blockchain to be decrypted with a small cost at the target time of decryption. Values that are not included in any block until the target block remain confidential.

1.3 Related Work

To the best of our knowledge, no prior work has studied time-lock encrypted storage that simultaneously guarantees both time-lock indistinguishability and non-inclusion privacy.

Prior works that only provide time-lock indistinguishability guarantees can be categorized into two groups based on their underlying design approach: (i) *time-release* schemes and (ii) *time-lock puzzles*. In time-release schemes, a committee holds secret-shared state and releases decryption material at a designated time, enabling delayed access to ciphertexts [7,12,13]. In contrast, in time-lock puzzles, decryption involves solving a sequential computational puzzle whose expected solution time matches the target delay [16,11,15]. Puzzle-based methods face two practical drawbacks in decentralized settings: they require significant computational resources and are sensitive to hardware assumptions.

Numerous recent papers have studied the batched threshold encryption (BTE) primitive that can be trivially extended to build encrypted storage that only provides the non-inclusion privacy property [8,1,6,9,5,3]. Clients can submit encrypted values using a BTE scheme, a subset of which will then be included in a block in the blockchain. The blockchain validators then decrypt the subset of encrypted values included using the BTE decryption algorithm. Clearly, this approach does not provide time-lock indistinguishability. Moreover, existing BTE constructions impose batch-size-dependent costs and require $\Omega(B)$ bilinear pairing operations to decrypt a batch of size B — which effectively caps the number of ciphertexts that can be decrypted in a single batch.

1.4 Other Applications

Besides the ML contest example we discussed above, the tTLES primitive can be useful for several other kinds of blockchain applications, such as:

On-chain puzzle solving. Game hosts post challenges with submission deadlines, and participants store their encrypted solutions using a tTLES. The time-lock indistinguishability property of tTLES prevents solutions from being copied before the deadline. The non-inclusion privacy property of tTLES ensures that solutions excluded from the blockchain (e.g., due to censorship) remain hidden, preventing adversaries from learning honest participants’ work without allowing them to earn rewards.

On-chain governance. Voters submit their encrypted votes before a voting deadline using a tTLES. The time-lock indistinguishability property of tTLES prevents vote manipulation and bandwagon effects by keeping votes secret until the deadline. The non-inclusion privacy property protects the privacy of voters whose votes get censored, i.e., did not get included in the blockchain. Moreover, the non-inclusion privacy property prevents potential misinformation claims about election fraud from excluded votes.

Sealed-bid auctions. Bidders store their encrypted bids before an auction deadline using a tTLES. The time-lock indistinguishability property of the tTLES ensures the sealed nature of the auction until the deadline, while non-inclusion privacy protects bidders whose bids are not included in the blockchain, preventing sellers from learning excluded valuations that could lead to seller’s remorse or strategic manipulation in future auctions.

2 Preliminaries

Notation. Let λ denote the security parameter. We assume that all algorithms get λ in unary as input. By $\text{poly}(\lambda)$ and $\text{negl}(\lambda)$, we mean the class $\lambda^{O(1)}$ and $\frac{1}{\lambda^{\omega(1)}}$ respectively. Given a security parameter λ , we use PPT to denote probabilistic $\text{poly}(\lambda)$ -time Turing Machines with $\text{poly}(\lambda)$ -sized advice. For a pair of distribution ensembles $(\{X_\lambda\}_{\lambda \in \mathbb{N}}, \{Y_\lambda\}_{\lambda \in \mathbb{N}})$, we say the ensembles are “computationally indistinguishable”, denoted by $X_\lambda \approx_c Y_\lambda$, if $|\Pr[X_\lambda = 1] - \Pr[Y_\lambda = 1]| = \text{negl}(\lambda)$. For a finite set S , we write $s \leftarrow \$ S$ to denote that s is sampled uniformly at random from S , and we write $|S|$ to denote the size of S . For an integer $a \in \mathbb{N}$, we use $[a]$ to denote the ordered set $\{1, \dots, a\}$. Furthermore, we write “ \leftarrow ” for probabilistic assignment and “ $:=$ ” for deterministic assignment. Also, we use the notation $v \leftarrow \Pi\text{-}\star(u)$ for any algorithm to denote that it is a multi-party computation protocol among a set of parties, where each party has some private input, a common public input u and each honest party produces a common public output v (which can be \perp).

Threat model. We consider a network of n parties, denoted with the set $\{1, \dots, n\}$. We consider an adversary \mathcal{A} that can statically corrupt up to $t < n$ parties in the protocol. The corrupted parties are considered malicious and can deviate arbitrarily from the specified protocol. We also consider a trusted entity \mathcal{E} to emulate an idealized blockchain protocol run by n parties. We elaborate on our assumptions on \mathcal{E} in Section 3.

2.1 Threshold Identity-Based Encryption

Our protocol relies on the standard threshold identity-based encryption (tIBE) scheme [18], as defined below.

Definition 1 (Threshold Identity-Based Encryption). *A threshold identity-based encryption scheme tIBE has the following PPT algorithms and protocols:*

- $\text{pp} \leftarrow \text{Setup}(1^\lambda)$: *The randomized setup algorithm takes as input the security parameter λ and outputs public parameters pp , which are an implicit input to all subsequent algorithms and protocols.*
- $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$: *The randomized key-generation algorithm takes as input the total number of parties n , the corruption threshold t , and outputs a master public key mpk to all, and the i -th secret key sk_i to party i .*
- $\text{ct} \leftarrow \text{Enc}(\text{mpk}, m, \text{id})$: *The randomized encryption algorithm takes the master public key mpk , a message $m \in \mathcal{M}$, and an identity $\text{id} \in \{0, 1\}^*$, and outputs a ciphertext $\text{ct} \in \mathcal{C}$. Here, \mathcal{M} and \mathcal{C} are the message and ciphertext spaces, respectively.*
- $k_{\text{id}}/\perp \leftarrow \Pi\text{-KeyIss}(\text{id})$: *This is an interactive protocol among all parties for key issuance, where each party P_i has a private input sk_i , and all parties have a common input id . The protocol outputs a decryption key k_{id} for the identity id or \perp indicating failure.*
- $m/\perp := \text{Dec}(\text{ct}, k_{\text{id}})$: *The deterministic decryption algorithm takes a ciphertext ct and the identity key k_{id} , and outputs a message m or \perp on failure.*

We require the tIBE scheme to satisfy the standard correctness, indistinguishability under chosen-plaintext attack (IND-CPA) security and the robustness properties. Intuitively, the correctness definition ensures messages encrypted under a identity id can be decrypted correctly with the corresponding issued key k_{id} . Similarly, the IND-CPA security guarantees privacy of messages encrypted under identities whose keys are unknown to the adversary. Note that this should hold even when the adversary statically corrupts up to t parties. Finally, the robustness requirement ensures that under-honest majority assumption, i.e., with $t < n/2$, an adversary can not prevent honest parties from successfully decrypting honestly generated ciphertexts. We note that in our robustness definition, we let the adversary to learn secret keys of all n parties. We formalize these properties next and refer the reader to [4] for more details of these properties.

| <u>Game $\text{Expt}_{\mathcal{A},b}^{\text{tIBE-cpa}}(n, t, \lambda)$:</u> | <u>Game $\text{Expt}_{\mathcal{A}}^{\text{tIBE-rob}}(n, t, \lambda)$:</u> |
|---|---|
| 1: $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ 2: $\mathcal{Q}_{\text{Keylss}} \leftarrow \emptyset$ 3: $\mathcal{C} \leftarrow \mathcal{A}(\text{pp})$ 4: $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$ 5: $(m_0, m_1, \text{id}^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Keylss}}}(\text{mpk}, \{\text{sk}_i\}_{i \in \mathcal{C}})$ 6: $\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, m_b, \text{id}^*)$ 7: $b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Keylss}}}(\text{ct}^*)$ 8: if $ \mathcal{C} \leq t \wedge \text{id}^* \notin \mathcal{Q}_{\text{Keylss}}$: 9: return b' 10: return 0 | 1: $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ 2: $\mathcal{C} \leftarrow \mathcal{A}(\text{pp})$ 3: $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$ 4: $(m^*, \text{id}^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Keylss}}}(\text{mpk}, \{\text{sk}_i\}_{i \in [n]})$ 5: $\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, m^*, \text{id}^*)$ 6: $k_{\text{id}^*} \leftarrow \mathcal{O}_{\text{Keylss}}(\text{id}^*)$ 7: if $ \mathcal{C} \leq t \wedge \text{Dec}(\text{ct}^*, k_{\text{id}^*}) \neq m^*$: 8: return 1 9: return 0 |
| Oracle $\mathcal{O}_{\text{Keylss}}(\text{id})$: | |
| 1: $\mathcal{Q}_{\text{Keylss}} \leftarrow \mathcal{Q}_{\text{Keylss}} \cup \{\text{id}\}$ <i>// Π-Keylss is run among all n parties where \mathcal{A} <i>controls the parties in \mathcal{C}</i> 2: $k_{\text{id}} \leftarrow \Pi\text{-Keylss}(\text{id})$ 3: return k_{id} </i> | |

Fig. 1: Security game $\text{Expt}_{\mathcal{A},b}^{\text{tIBE-cpa}}$ and $\text{Expt}_{\mathcal{A}}^{\text{tIBE-rob}}$ and for a tIBE scheme

Definition 2 (Correctness of tIBE). For all $n, t < n$, all identity id , all message $m \in \mathcal{M}$, here \mathcal{M} is the message space, the following holds.

$$\Pr \left[\begin{array}{l} \text{Dec}(k_\tau, \text{ct}) = m : \\ \text{mpk}, \{\text{sk}_i\}_{i \in [n]} \leftarrow \text{KeyGen}(n, t) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, m, \text{id}) \\ k_\tau \leftarrow \Pi\text{-Keylss}(\tau) \end{array} \right] = 1 - \text{negl}(\lambda)$$

Definition 3 (IND-CPA security of tIBE). Let $\text{Expt}_{\mathcal{A},b}^{\text{tIBE-cpa}}$ be the game as defined in Fig. 1. A tIBE scheme is IND-CPA secure if for all $n, t < n$, all PPT adversaries \mathcal{A} , we have

$$\left| \Pr \left[\text{Expt}_{\mathcal{A},0}^{\text{tIBE-cpa}}(n, t, \lambda) \Rightarrow 1 \right] - \Pr \left[\text{Expt}_{\mathcal{A},1}^{\text{tIBE-cpa}}(n, t, \lambda) \Rightarrow 1 \right] \right| = \text{negl}(\lambda)$$

Definition 4 (Robustness of tIBE). Let $\text{Expt}_{\mathcal{A}}^{\text{tIBE-rob}}$ be the game as defined in Fig. 1. A tIBE scheme is robust if for all $n \in \mathbb{N}, t < n/2$, and all PPT adversaries \mathcal{A} , we have

$$\Pr \left[\text{Expt}_{\mathcal{A}}^{\text{tIBE-rob}}(n, t, \lambda) \Rightarrow 1 \right] = \text{negl}(\lambda)$$

CPA vs chosen-ciphertext attack (CCA) secure tIBE. In our tTLES construction, we issue identity keys only publicly. More specifically, we will associate each identity id with a time (measured in block height) and publicly issue k_{id} , once the block at height id has been finalized. Therefore, all these keys are also available to the adversary after issuance. Hence, the IND-CPA security of tIBE suffices.

Concretely, we will instantiate the tIBE scheme using a threshold version of the IBE scheme from [2].

2.2 Batched Threshold Encryption

Our tTLES protocol also requires a chosen-ciphertext attack (IND-CCA) secure batched threshold encryption (BTE) primitive with *associated data* [6,9,1]. We will summarize the interface of BTE and the properties

| Game $\text{Expt}_{\mathcal{A},b}^{\text{BTE-cca}}(n, t, \lambda, \mathbf{B}_{\max})$: | Game $\text{Expt}_{\mathcal{A}}^{\text{BTE-rob}}(n, t, \lambda, \mathbf{B}_{\max})$: |
|---|---|
| 1: $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ 2: $\mathcal{Q}_{\text{BDec}} \leftarrow \emptyset$ 3: $\mathcal{C} \leftarrow \mathcal{A}(\text{pp})$ 4: $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$ 5: $(m_0, m_1, \text{ad}^*) \leftarrow \mathcal{A}^{\text{OBDec}}(\text{mpk}, \{\text{sk}_i\}_{i \in \mathcal{C}})$ 6: $\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, m_b, \text{ad}^*)$ 7: $b' \leftarrow \mathcal{A}^{\text{OBDec}}(\text{ct}^*)$ 8: if $ \mathcal{C} \leq t \wedge \text{ct}^* \notin \mathcal{Q}_{\text{BDec}}$: 9: return b' 10: return 0 | 1: $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ 2: $\mathcal{C} \leftarrow \mathcal{A}(\text{pp})$ 3: $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$ 4: Let $\text{inp} := (\text{mpk}, \{\text{sk}_i\}_{i \in [n]})$ 5: $(B, S \subseteq [B]) \leftarrow \mathcal{A}^{\text{OBDec}}(\text{inp})$ 6: if $B > \mathbf{B}_{\max}$: return 0 7: $\{(m_i, \text{ad}_i)\}_{i \in [B] \setminus S} \leftarrow \mathcal{A}^{\text{OBDec}}(\text{inp})$ 8: for all $i \in [B] \setminus S$: 9: $\text{ct}_i \leftarrow \text{Enc}(\text{mpk}, m_i, \text{ad}_i)$ 10: $\{\text{ct}_i\}_{i \in S} \leftarrow \mathcal{A}^{\text{OBDec}}(\{\text{ct}_i\}_{i \in [B] \setminus S})$ 11: $\{\tilde{m}_i\}_{i \in [B]} \leftarrow \mathcal{O}_{\text{BDec}}(\{\text{ct}_i\}_{i \in [B]})$ 12: if $ \mathcal{C} \leq t \wedge (\exists i \in [B] \setminus S \text{ s.t. } \tilde{m}_i \neq m_i)$: 13: return 1 14: return 0 |
| Oracle $\mathcal{O}_{\text{BDec}}(\{\text{ct}_i\}_{i \in [B]})$: <i>// Π-BDec is run among all n parties where \mathcal{A} controls the parties in \mathcal{C}</i> 12: Let $\{\tilde{m}_i\}_{i \in [B]} \leftarrow \Pi\text{-BDec}(\text{mpk}, \{\text{ct}_i\}_{i \in [B]})$ 13: $\mathcal{Q}_{\text{BDec}} \leftarrow \mathcal{Q}_{\text{BDec}} \cup \{\text{ct}_i\}_{i \in [B]}$ 14: return $\{\tilde{m}_i\}_{i \in [B]}$ | |

Fig. 2: Security and robustness game $\text{Expt}_{\mathcal{A},b}^{\text{BTE-cca}}$ and $\text{Expt}_{\mathcal{A}}^{\text{BTE-rob}}$, respectively, for a BTE scheme

it provides. We note that existing BTE protocols do not guarantee IND-CCA security with associated data. However, using techniques from [4], we can generically transform those schemes to guarantee CCA security even with associated data.

Definition 5 (Batched Threshold Encryption with Associated Data). *A batched threshold encryption scheme BTE with associated data consists of the following PPT algorithms and protocols:*

- $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^{\mathbf{B}_{\max}})$: *The randomized setup algorithm takes as input the security parameter λ and the maximum batch size \mathbf{B}_{\max} and outputs public parameters pp which is an implicit input to all subsequent interfaces.*
- $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$: *The randomized key-generation algorithm takes as input the total number of parties n , the corruption threshold t , and outputs a master public key mpk to all, and the i -th secret key sk_i to party i .*
- $\text{ct} \leftarrow \text{Enc}(\text{mpk}, \text{msg}, \text{ad})$: *The randomized encryption algorithm takes as input a public key mpk , and a message msg , and associated data ad , and outputs a ciphertext ct .*
- $\{\text{msg}_\ell\}_{\ell \in [B]} / \perp \leftarrow \Pi\text{-BDec}(\text{mpk}, \{\text{ct}_\ell\}_{\ell \in [B]})$: *This is an interactive protocol among all parties for batch decryption, where each party P_i has a private input sk_i , and all parties have common input mpk and a batch of ciphertexts $\{\text{ct}_\ell\}_{\ell \in [B]}$ where $B \leq \mathbf{B}_{\max}$. The protocol outputs the original decrypted messages $\{\text{msg}_\ell\}_{\ell \in [B]}$, otherwise \perp .*

We require the BTE scheme to satisfy the *correctness*, IND-CCA security with associated data, and *robustness* properties. We formalize these properties in Definitions 6 to 8 and describe them next.

Naturally, the *correctness* property ensures that correctly encrypted ciphertexts successfully decrypt to the original encrypted messages. Formally,

Definition 6 (Correctness of BTE). *A batched threshold encryption scheme BTE is correct for all $n, t < n$, all \mathbf{B}_{\max} , all $B < \mathbf{B}_{\max}$, all messages $(m_1, m_2, \dots, m_B) \in \mathcal{M}^B$, and all associated data $(\text{ad}_1, \text{ad}_2, \dots, \text{ad}_B)$, here \mathcal{M} is the message space, the following holds.*

$$\Pr \left[\begin{array}{l} \forall i \in [B] \\ \tilde{m}_i = m_i \end{array} : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, 1^{\mathbf{B}_{\max}}) \\ (\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t) \\ \{\text{ct}_i \leftarrow \text{Enc}(\text{mpk}, m_i, \text{ad}_i)\}_{i \in [B]} \\ \{\tilde{m}_j\}_{j \in [B]} \leftarrow \Pi\text{-BDec}(\text{mpk}, \{\text{ct}_i\}_{i \in [B]}) \end{array} \right] = 1 - \text{negl}(\lambda)$$

The IND-CCA security with associated data for BTE ensures that an adversary cannot tell apart between encryption of two equal-length messages, even if it can invoke the batch decryption protocol (Π -BDec) on arbitrary batches of its choice not containing the challenge ciphertext. We include the *associated data* property in our IND-CCA definition to prevent any adversary from copying honestly generated ciphertexts for any time τ and submitting them to the batch decryption protocol by claiming them to be for a different time $\tau' \neq \tau$. As we discuss later, preventing such malleability attacks is crucial to guarantee the non-inclusion privacy guarantee we discuss in Section 1. We note that this should hold, even when the adversary statically corrupts up to t parties during the protocol. More formally,

Definition 7 (IND-CCA security of BTE with associated data). Let $\text{Expt}_{\mathcal{A},b}^{\text{BTE-cca}}$ be the game as defined in Fig. 2. A batched threshold encryption scheme BTE is IND-CCA secure with associated data if for all $n, t < n$, and all PPT adversaries \mathcal{A} , we have

$$\left| \Pr \left[\text{Expt}_{\mathcal{A},0}^{\text{BTE-cca}}(n, t, \lambda, \mathbf{B}_{\max}) \Rightarrow 1 \right] - \Pr \left[\text{Expt}_{\mathcal{A},1}^{\text{BTE-cca}}(n, t, \lambda, \mathbf{B}_{\max}) \Rightarrow 1 \right] \right| = \text{negl}(\lambda)$$

Finally, the *robustness* property prevents an adversary \mathcal{A} , that can corrupt up to t parties, from invoking the batch decryption protocol (Π -BDec) on malformed ciphertexts to disrupt decryption of correctly generated ciphertexts. We note that this is a generalization of the rouge-ciphertext security property defined in [6], where \mathcal{A} can not corrupt any of the decrypting parties. Moreover, we want the robustness property to hold even when the adversary learns the secret keys of all n parties. More formally,

Definition 8 (Robustness of BTE). Let $\text{Expt}_{\mathcal{A}}^{\text{BTE-rob}}$ be the game as defined in Fig. 2. A batched threshold encryption scheme BTE is robust if for all $n \in \mathbb{N}, t < n/2, \mathbf{B}_{\max} \in \mathbb{N}$, we have the following:

$$\Pr[\text{Expt}_{\mathcal{A}}^{\text{BTE-rob}}(n, t, \lambda) \Rightarrow 1] = \text{negl}(\lambda)$$

2.3 Pseudorandom Generator

Finally, our tTLES protocol relies on a pseudorandom generator (PRG) scheme $\text{PRG} : \{0, 1\} \rightarrow \{0, 1\}^{\ell(\lambda)}$ for some function $\ell(\lambda) > \lambda$. A PRG is secure, if the following holds:

Definition 9 (Pseudorandom Generator Security). Consider the games Expt_{PRG} and $\text{Expt}_{\text{Rand}}$ defined in Fig. 3. A pseudorandom generator scheme PRG is secure, if for all PPT adversary \mathcal{A} , we have the following:

$$\left| \Pr \left[\text{Expt}_{\text{PRG}}(1^\lambda) \Rightarrow 1 \right] - \Pr \left[\text{Expt}_{\text{Rand}}(1^\lambda) \Rightarrow 1 \right] \right| = \text{negl}(\lambda).$$

| | |
|---|--|
| Game $\text{Expt}_{\text{PRG}}(1^\lambda)$: | Game $\text{Expt}_{\text{Rand}}(1^\lambda)$: |
| 1: Sample $x \leftarrow \{0, 1\}^\lambda$ | 1: Sample $y \leftarrow \{0, 1\}^{\ell(\lambda)}$ |
| 2: return $\mathcal{A}(\text{PRG}(x))$ | 2: return $\mathcal{A}(y)$ |

Fig. 3: Security games for pseudorandom generator

3 Defining Threshold Time-lock Encrypted Storage

A threshold time-lock encrypted storage (tTLES) is a storage service provided by an entity \mathcal{E} (which can be a blockchain) and a set of n servers, among which up to t servers can be malicious. tTLES has clients, who can store encrypted data in tTLES and specify a future decryption time $\tau \in \mathbb{N}$ (time is just a monotonic counter here). We assume that \mathcal{E} provides the following guarantees, which can, for instance, be easily achieved using a blockchain.

1. \mathcal{E} maintains a monotonic current time τ_{curr} as time.
2. At any time τ_{curr} , \mathcal{E} receives a batch of encrypted storage requests $C_{\tau_{\text{curr}}} := \{\text{ct}_i\}_{i \in [B_{\tau_{\text{curr}}}]}$ where $B_{\tau_{\text{curr}}}$ is an integer upper bounded by B_{max} . Each encrypted value $\text{ct}_i \in C_{\tau_{\text{curr}}}$ is encrypted towards some future time $\tau_i > \tau_{\text{curr}}$. We note that each time τ_i can be distinct.
3. At any time τ_{curr} , upon receiving the batch of encrypted values $C_{\tau_{\text{curr}}}$, \mathcal{E} can run arbitrary public validation checks on $C_{\tau_{\text{curr}}}$. Looking ahead, \mathcal{E} can enforce inclusion policies based on the associated data in clients' requests.
4. At any time τ_{curr} , \mathcal{E} maintains the encrypted storage values in storage $\text{St}_{\tau'}$, for all future time $\tau' > \tau_{\text{curr}}$ for which any client has submitted an encrypted storage request.

With access to an entity \mathcal{E} that provides the guarantees we describe above, we define a time-lock encrypted storage protocol **tTLES** as follows.

Definition 10. A threshold time-lock encrypted storage (tTLES) protocol **tTLES** consists of the following PPT algorithms:

- $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^{B_{\text{max}}})$. The setup algorithm takes as input the security parameter λ , an upper bound B_{max} on the batch size, and outputs the public parameters for the scheme. We assume that all the interfaces below take pp as input.
- $(\text{pk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$. The key generation algorithm takes as input the total number of parties n , a threshold $t < n$, and outputs a public key mpk and n secret keys $\{\text{sk}_i\}_{i \in [n]}$, where the i -th party receives sk_i .
- $k_\tau \leftarrow \text{II-Keylss}(\tau)$. This is an interactive protocol that is executed among all parties, where each party P_i has a private input sk_i , and all parties have a common input τ . The protocol outputs the decryption key k_τ for time τ .
- $\text{ct} \leftarrow \text{Enc}(\text{mpk}, m, \tau)$. A client uses the randomized encryption algorithm Enc to encrypt a message m , under the public key mpk , and a future time τ .
- $\{\tilde{m}_\ell\}_{\ell \in [B]} / \perp \leftarrow \text{II-BDec}(\text{mpk}, \{\text{ct}_j\}_{j \in [B]})$. This is an interactive protocol among all parties for batch decryption, where each party P_i has a private input sk_i , and all parties have common input mpk and a batch of ciphertexts $\{\text{ct}_\ell\}_{\ell \in [B]}$, possibly encrypted towards different times, where $B \leq B_{\text{max}}$. The protocol outputs partially decrypted messages $\{\tilde{m}_\ell\}_{\ell \in [B]}$, otherwise \perp .
- $m := \text{Dec}(k_\tau, \tilde{m})$. The deterministic decryption algorithm takes a decryption key k_τ and a partially decrypted message \tilde{m} , and outputs the message m .

Intuition behind tTLES formulation. We formalize the tTLES primitive envisioning the following usage scenarios. In a blockchain, consider a future target time (or block height) τ_{tgt} , that a client wants to encrypt its data towards. Then, any time before τ_{tgt} , the client can encrypt a value towards time τ_{tgt} using the Enc algorithm and submit it to get included in the blockchain (modeled by \mathcal{E} in our definition) before the block at height τ_{tgt} gets finalized.

For decrypting a ciphertext ct (created w.r.t time τ_{tgt}), instead of defining a single decryption protocol, say $\text{II-SDec}(\tau_{\text{tgt}}, \text{ct})$ which outputs the message m iff ct has been included at time $\tau_{\text{incl}} \leq \tau_{\text{tgt}}$, we split the decryption process across three protocols and algorithms (II-Keylss , II-BDec , Dec). This is motivated by the following efficiency reason: In the former approach, the II-SDec protocol will need to be executed on the entire set of encrypted storage values in $\text{St}_{\tau_{\text{tgt}}}$ which can be of unbounded size, and hence will be infeasible to execute in a timely manner in practice. Our approach avoids this pitfall by splitting decryption across three protocols, where the resources needed for II-Keylss and Dec are independent of the storage size $\text{St}_{\tau_{\text{tgt}}}$. Although the cost of running II-BDec depends on the size of $\text{St}_{\tau_{\text{tgt}}}$, in our design, parties run II-BDec in batches, once for every time leading up to the target time τ_{tgt} . As a result, the cost of running II-BDec gets amortized over a long time interval, and does not result in a bottleneck in practice.

Properties of tTLES. Naturally, we want any tTLES protocol to satisfy correctness, which ensures that honestly generated ciphertexts get decrypted correctly.

Definition 11 (Correctness). For all $n, t < n$, all λ, B_{\max} , all $B < B_{\max}$, all messages $(m_1, m_2, \dots, m_B) \in \mathcal{M}^\ell$, and all time $(\tau_1, \tau_2, \dots, \tau_B)$, here \mathcal{M} is the message space, the following holds.

$$\Pr \left[\begin{array}{l} \forall i \in [B] \\ \text{Dec}(k_{\tau_i}, \tilde{m}_i) = m_i \end{array} : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda, 1^{B_{\max}}) \\ (\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t) \\ \{\text{ct}_i \leftarrow \text{Enc}(\text{mpk}, m_i, \tau_i)\}_{i \in [B]} \\ \{k_{\tau_i} \leftarrow \Pi\text{-KeyIss}(\tau_i)\}_{i \in [B]} \\ \{\tilde{m}_j\}_{j \in [B]} \leftarrow \Pi\text{-BDec}(\text{mpk}, \{\text{ct}_i\}_{i \in [B]}) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

In terms of security, we intuitively require a tTLES to reveal the encrypted storage values when both: (i) the encrypted values are included in the blockchain, as per some inclusion condition c (say before the target decryption time), and (ii) the decryption time for the encrypted value has passed. Informally, we refer to these properties as *non-inclusion privacy* and *time-lock indistinguishability* property. The non-inclusion privacy ensures privacy of all the encrypted storage values that are not included in the chain, and the time-lock indistinguishability ensures that the standard timed delay privacy property. We formalize these requirements using the security game $\text{Expt}_{\mathcal{A}, b}^{\text{tTLES-sec}}$ in Figure 4.

In the security game $\text{Expt}_{\mathcal{A}, b}^{\text{tTLES-sec}}$, the game maintains the current time τ_{curr} and a set $\mathcal{Q}_{\text{BDec}}$ of partially decrypted values. Additionally, a PPT adversary \mathcal{A} gets access to the time increment oracle $\mathcal{O}_{\text{TInc}}$, which upon each invocation with a batch of ciphertexts, (i) increments the current time (line 12), (ii) runs the partial batch-decryption protocol on the input batch of ciphertexts after removing the invalid ciphertexts based on the inclusion condition (line 13-16), and (iii) issues a key for the current time τ_{curr} (line 17). The winning condition of \mathcal{A} is similar to the IND-CCA security game of public key encryption schemes. Specifically, for the challenge ciphertext \mathcal{A} chooses two equal length messages (m_0, m_1) along with a decryption time τ^* (see line 5), and receives a challenge ciphertext ct^* . To win the game, \mathcal{A} can then either request a key issuance for time τ^* or the batch decryption of a batch containing ct^* , but not both. We check this constraint in line 9. Intuitively, here the check $\tau_{\text{curr}} \geq \tau^*$ ensures the time-lock indistinguishability, and the check $\text{ct}^* \in \mathcal{Q}_{\text{BDec}}$ ensures the non-inclusion privacy. \mathcal{A} wins the game $\text{Expt}_{\mathcal{A}, b}^{\text{tTLES-sec}}$, if it can successfully guess the challenge bit b , without violating this constraint. We formalize this in Definition 12.

Definition 12 (TLES security). Let $\text{Expt}_{\mathcal{A}, b}^{\text{tTLES-sec}}$ be the game as defined in Fig. 4. A tTLES scheme is secure if for all $n, t < n$, all security parameter λ , all $B_{\max} \in \mathbb{N}$ and all PPT adversaries \mathcal{A} , we have

$$\left| \Pr \left[\text{Expt}_{\mathcal{A}, 0}^{\text{tTLES-sec}}(n, t, \lambda, B_{\max}) \Rightarrow 1 \right] - \Pr \left[\text{Expt}_{\mathcal{A}, 1}^{\text{tTLES-sec}}(n, t, \lambda, B_{\max}) \Rightarrow 1 \right] \right| = \text{negl}(\lambda)$$

Finally, from the usability perspective, we also require the tTLES protocol tTLES to be *robust* in the following sense. Under the honest-majority assumption (i.e., with $t < n/2$), an adversary cannot prevent honest parties from successfully decrypting the encrypted values, even when the adversary includes these honestly encrypted ciphertexts in batches containing arbitrarily crafted adversarial ciphertexts. We formalize this property using the game $\text{Expt}_{\mathcal{A}}^{\text{tTLES-rob}}$ in Fig. 4. Note that we want robustness to hold, even when the adversary learns the secret keys of all honest parties.

Definition 13 (TLES robustness). Let $\text{Expt}_{\mathcal{A}}^{\text{tTLES-rob}}$ be the game as defined in Fig. 4. A tTLES scheme is secure if for all $n, t < n/2$, all security parameter λ , all $B_{\max} \in \mathbb{N}$ and all PPT adversaries \mathcal{A} , we have

$$\Pr \left[\text{Expt}_{\mathcal{A}}^{\text{tTLES-rob}}(n, t, \lambda, B_{\max}) \Rightarrow 1 \right] = \text{negl}(\lambda)$$

Efficiency requirements. To eliminate expensive solutions, we require a tTLES protocol to satisfy the following efficiency properties:

(E1) Let $\text{ct} \leftarrow \text{Enc}(\text{mpk}, m, \tau)$, then we want that size of the ciphertext is proportional to the length of the underlying message, i.e., $|\text{ct}| = O(|m|)$.

| Game $\text{Expt}_{\mathcal{A},b}^{\text{tTLES-sec}}(n, t, \lambda, \mathbf{B}_{\max})$: | Game $\text{Expt}_{\mathcal{A}}^{\text{tTLES-rob}}(n, t, \lambda, \mathbf{B}_{\max})$: |
|--|--|
| 1: $\tau_{\text{curr}} := 0$ 2: $\mathcal{Q}_{\text{BDec}} \leftarrow \emptyset$ 3: $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^{\mathbf{B}_{\max}})$ 4: $\mathcal{C} \leftarrow \mathcal{A}(\text{pp})$ // we assume \mathcal{A} to be stateful 5: $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$ 6: $(m_0, m_1, \tau^*) \leftarrow \mathcal{A}^{\text{OTinc}}(\text{mpk}, \{\text{sk}_i\}_{i \in \mathcal{C}})$ 7: $\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, m_b, \tau^*)$ 8: $b' \leftarrow \mathcal{A}^{\text{OTinc}}(\text{ct}^*)$ 9: if $\tau_{\text{curr}} \geq \tau^* \wedge \text{ct}^* \in \mathcal{Q}_{\text{BDec}}$: 10: return 0 11: return b' | 1: $\tau_{\text{curr}} := 0$ 2: $\mathcal{Q}_{\text{BDec}} \leftarrow \emptyset$ 3: $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^{\mathbf{B}_{\max}})$ 4: $\mathcal{C} \leftarrow \mathcal{A}(\text{pp})$ // we assume \mathcal{A} to be stateful 5: $(\text{mpk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}(n, t)$ 6: $(B, U \subseteq [B]) \leftarrow \mathcal{A}^{\text{OTinc}}(\text{mpk}, \{\text{sk}_i\}_{i \in [n]})$ 7: if $B > \mathbf{B}_{\max}$: return 0 8: $\{(m_i, \tau_i)\}_{i \in [B] \setminus U} \leftarrow \mathcal{A}^{\text{OTinc}}(\text{mpk}, \{\text{sk}_i\}_{i \in \mathcal{C}})$ // Π -BDec and Π -Keylss is run among all n parties where \mathcal{A} controls the parties in \mathcal{C} 9: for all $i \in [B] \setminus U$: 10: $\text{ct}_i \leftarrow \text{Enc}(\text{mpk}, m_i, \tau_i)$ 11: $k_{\tau_i} \leftarrow \Pi\text{-Keylss}(\tau_i)$ 12: $\{\text{ct}_i\}_{i \in U} \leftarrow \mathcal{A}^{\text{OTinc}}(\{\text{ct}_i\}_{i \in [B] \setminus U})$ 13: $\{\tilde{m}_i\}_{i \in [B]} \leftarrow \Pi\text{-BDec}(\text{mpk}, \{\text{ct}_i\}_{i \in [B]})$ 14: if $ \mathcal{C} \leq t \wedge (\exists i \in [B] \setminus U \text{ s.t. } \text{Dec}(k_{\tau_i}, \tilde{m}_i) \neq m_i)$: 15: return 1 16: return 0 |
| Oracle $\mathcal{O}_{\text{Tinc}}(\{\text{ct}_j\}_{j \in [B]})$: 12: $\tau_{\text{curr}} := \tau_{\text{curr}} + 1$ 13: parse each ct_j as (τ_j, ct'_j) // checking inclusion condition 14: $S := \{\text{ct}_j \mid j \in [B] \wedge \tau_j > \tau_{\text{curr}}\}$ 15: $\mathcal{Q}_{\text{BDec}} \leftarrow \mathcal{Q}_{\text{BDec}} \cup S$ // Π -BDec and Π -Keylss is run among all n parties where \mathcal{A} controls the parties in \mathcal{C} 16: $\{\tilde{m}_i\}_{i \in [S]} \leftarrow \Pi\text{-BDec}(\text{mpk}, S)$ 17: $k_{\text{curr}} \leftarrow \Pi\text{-Keylss}(\tau_{\text{curr}})$ 18: return $(k_{\text{curr}}, \{\tilde{m}_i\}_{i \in [S]})$ | |

Fig. 4: Security and robustness game $\text{Expt}_{\mathcal{A},b}^{\text{tTLES-sec}}$ and $\text{Expt}_{\mathcal{A}}^{\text{tTLES-rob}}$, respectively, for a tTLES protocol

- (E2) For any time τ , let St_τ be the encrypted storage maintained by \mathcal{E} , then we want $|\text{St}_\tau|$ to be unbounded, i.e., \mathcal{E} should be able to store arbitrary number of values in the encrypted storage. Note that this also implies that $|\text{St}_\tau| \gg \mathbf{B}_{\max}$.
- (E3) Let $m := \text{Dec}(k_\tau, \tilde{m})$ for some message \tilde{m} , then the running time of Dec is $O(|\tilde{m}|, \lambda)$.
- (E4) For any invocation of $\Pi\text{-BDec}$, we require the per party communication cost to be $O(\lambda)$.
- (E5) For any time τ , the running time of the entity \mathcal{E} is bounded by $\tilde{O}(\mathbf{B}_{\max})$.

4 Design and Analysis of Time-lock Encrypted Storage

In this section, we present our tTLES protocol and its security analysis. Our protocol is extremely simple and relies on threshold identity-based encryption (tIBE) and batched threshold encryption (BTE) in a black-box manner. Before we describe our approach, we want to note issues with designing tTLES using just tIBE or BTE. As discussed in Section 1, protocols such as tlock [13] that rely only on tIBE do not provide the non-inclusion privacy property. Below, we describe an approach that only relies on BTE, and discuss its limitations.

An inefficient solution using only BTE. Using just a BTE scheme, we can design a tTLES protocol as follows. Clients can submit their encrypted storage values towards a target time τ_{tgt} , anytime before τ_{tgt} . Let $\text{St}_{\tau_{\text{tgt}}}$ be the set of encrypted values to be decrypted at time τ_{tgt} . To decrypt these values, at time τ_{tgt} , parties run the $\text{BTE}.\Pi\text{-BDec}$ protocol to decrypt all the values in $\text{St}_{\tau_{\text{tgt}}}$ simultaneously.

This approach is correct and ensures both time-lock indistinguishability and non-inclusion privacy. However, it is not very efficient, as parties need to run the $\text{BTE}.\Pi\text{-BDec}$ on a large batch size $|\text{St}_{\tau_{\text{tgt}}}|$ at once, which can be very expensive [1,6,9]. Moreover, if $|\text{St}_{\tau_{\text{tgt}}}|$ can be unbounded, as required by the efficiency property (E2), this approach would be infeasible.

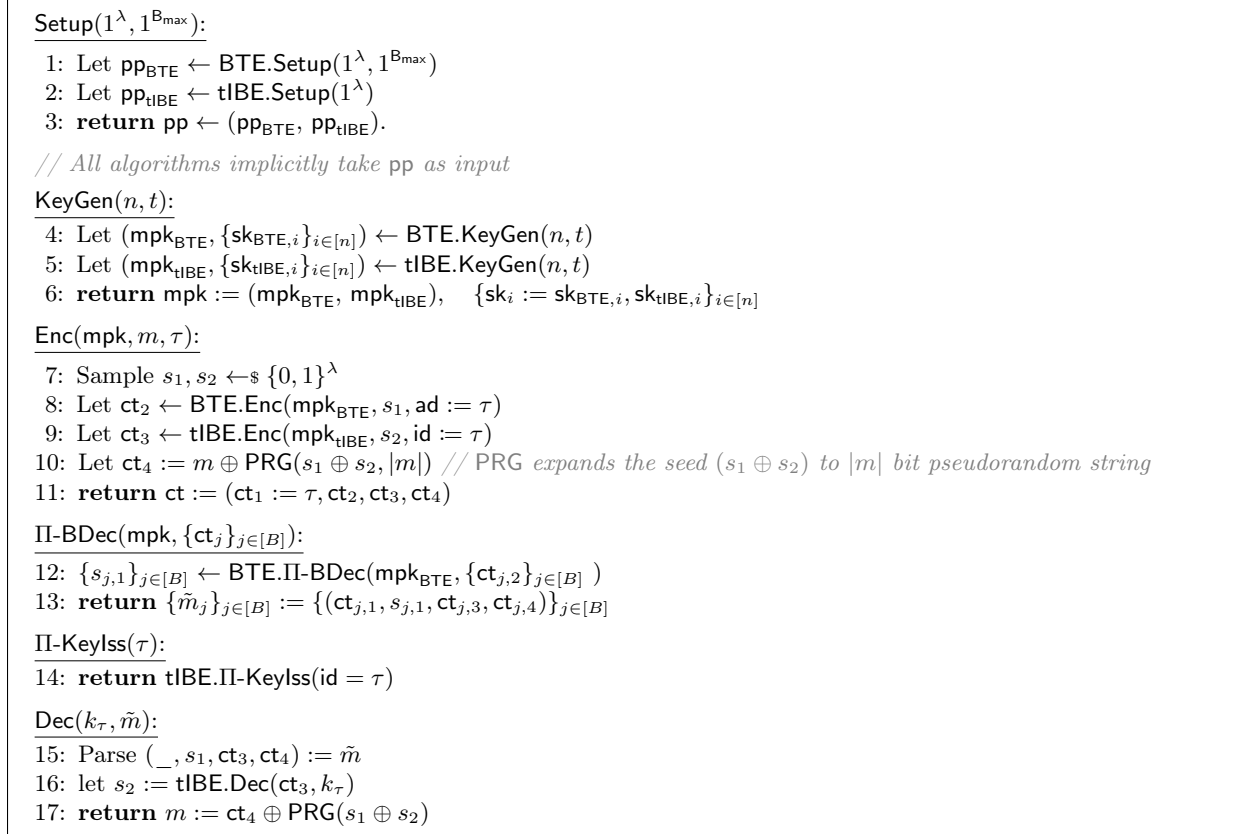


Fig. 5: Description of our tTLES protocol.

4.1 Our Approach

We summarize our tTLES protocol tTLES in Fig. 5, and describe it next. tTLES is simple, efficient, and relies on threshold identity based encryption and batched threshold encryption in a black-box manner.

Setup (line 1-3). The setup algorithm runs the setup algorithm of BTE and tIBE, and outputs $(\text{pp}_{\text{BTE}}, \text{pp}_{\text{tIBE}})$ as its public parameters.

Key generation (line 4-6). Similar to the setup algorithm, our key generation algorithm runs the key generation algorithm of both BTE and tIBE. It then outputs the combined public key $(\text{mpk}_{\text{BTE}}, \text{mpk}_{\text{tIBE}})$ as its public key. Similarly, the secret key of the i -th party is the combined secret key $(\text{sk}_{\text{BTE},i}, \text{sk}_{\text{tIBE},i})$.

Encryption (line 7-11). To encrypt a message msg towards a target time τ_{tgt} , a client samples two secrets $s_1, s_2 \in \{0, 1\}^\lambda$. The client then encrypts s_1 using the BTE encryption algorithm with τ_{tgt} as the associated data, and encrypts s_2 using tIBE with τ_{tgt} as the identity. The client then expands $s_1 \oplus s_2$ to $|m|$ bit pseudorandom string using a pseudorandom generator, and uses it as a one-time pad to encrypt the message m . Here, the secrets (s_1, s_2) can be seen as 2-out-of-2 secret shares of some secret $s := s_1 \oplus s_2$ [17]. Note that by using τ_{tgt} as the associated data in the BTE encryption, we ensure that the generated ciphertext ct is non-malleable even with respect to time τ_{tgt} .

Batch decryption at inclusion time (line 12-13). Let us first describe how the inputs to the tTLES.II-BDec protocol are created. At any given time τ_{curr} , the bulletin-board entity \mathcal{E} receives batch $\{\text{ct}_j\}_{j \in [B]}$ of B ciphertexts. \mathcal{E} parses each ciphertext ct_j as $\text{ct}_j = (\tau_j, \text{ct}_{j,2}, \star)$, where τ_j is the target decryption time of the ciphertext ct_j . Next, \mathcal{E} checks that $\tau_j > \tau_{\text{curr}}$, and creates a batch containing only such valid ciphertexts.

Without loss of generality, let us assume that all the ciphertexts in $\{\text{ct}_j\}_{j \in [B]}$ satisfy the abovementioned inclusion condition. Then, our batch decryption protocol, invokes the BTE.II-BDec protocol on input

$(\text{mpk}_{\text{BTE}}, \{\text{ct}_{j,2}\}_{j \in [B]})$. Let $\{s_{j,1}\}$ be the output of the BTE.II-BDec protocol. Our protocol, then outputs $\{\tilde{m}_j\}_{j \in [B]} := \{(\text{ct}_{j,1}, s_{j,1}, \text{ct}_{j,3}, \text{ct}_{j,4})\}_{j \in [B]}$ as the partially decrypted values.

Key issuance (line 14). To issue a key for any time τ , our protocol runs the tIBE key issuance algorithm with $\text{id} = \tau$, and outputs its output.

Decryption (line 15-17). Given the tIBE key k_τ for any time τ , and a partially decrypted ciphertext $\tilde{m} = (\tau, s_1, \text{ct}_3, \text{ct}_4)$, the decryption algorithm first recovers the share s_2 by running the tIBE.Dec with k_τ as the key. The algorithm outputs as $\text{ct}_4 \oplus \text{PRG}(s_1 \oplus s_2)$ as the message.

4.2 Analysis

The correctness of our protocol is clear. Next, we will argue security, assuming that the underlying BTE and the tIBE scheme are secure as per Definitions 3 and 7, respectively. Formally, we prove the following:

Theorem 1 (tTLES Security). *Assuming the existence of an IND-CPA secure threshold Identity Based Encryption (IBE) scheme (Definition 1) and IND-CCA secure Batched Threshold Encryption (BTE) scheme with associated data (Definition 5), and a pseudorandom generator (PRG), our tTLES protocol tTLES (Fig. 5) is secure as per Definition 12.*

Proof. The security of our scheme follows from the IND-CCA security of the BTE scheme and IND-CPA security of the tIBE scheme. To see this, consider the following set of experiments:

- Expt_0 : This is the experiment $\text{Expt}_{\mathcal{A},b}^{\text{tTLES-sec}}(n, t, \lambda, \mathbf{B}_{\max})$ as defined in Fig. 4 with $b = 0$ and TLES scheme instantiated with our construction in Fig. 5. Let $\ell = |m_0| = |m_1|$, be the length of the challenge messages.
- Hyb_0 : This is the same as Expt_0 except that when forming the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \text{ct}_3^*, \text{ct}_4^*)$, instead of setting ct_4^* as $m_0 \oplus \text{PRG}(s_1 \oplus s_2, \ell)$, we sample a random string $r \leftarrow_{\$} \{0, 1\}^\lambda$ and set ct_4^* as $m_0 \oplus \text{PRG}(r, \ell)$
- Hyb' : This is the same as Hyb_0 except that when forming the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \text{ct}_3^*, \text{ct}_4^*)$, we set ct_4^* as $r \leftarrow_{\$} \{0, 1\}^\ell$.
- Hyb_1 : This is the same as Hyb' except that when forming the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \text{ct}_3^*, \text{ct}_4^*)$, we sample a random string $r \leftarrow_{\$} \{0, 1\}^\lambda$ and set ct_4^* as $m_1 \oplus \text{PRG}(r, \ell)$
- Expt_1 : This is the experiment $\text{Expt}_{\mathcal{A},b}^{\text{tTLES-sec}}(n, t, \lambda, \mathbf{B}_{\max})$ as defined in Fig. 4 with $b = 1$ and TLES scheme instantiated with our construction in Fig. 5.

We now prove that $\text{Expt}_0 \approx_c \text{Expt}_1$. To do this, we argue that each intermediate pairs of hybrids are computationally indistinguishable. At a high-level, the proof structure looks as follows:

- $\text{Expt}_0 \approx_c \text{Hyb}_0$: This follows from the IND-CCA security of the BTE scheme and IND-CPA security of the threshold IBE scheme.
- $\text{Hyb}_0 \approx_c \text{Hyb}'$: This follows from the security of the pseudorandom generator.
- $\text{Hyb}' \approx_c \text{Hyb}_1$: This again follows from the security of the pseudorandom generator.
- $\text{Hyb}_1 \approx_c \text{Expt}_1$: This follows from the IND-CCA security of the BTE scheme and IND-CPA security of the threshold IBE scheme in a manner identical to the indistinguishability argument between Expt_0 and Hyb_0 .

We will now rigorously prove each of the four aforementioned indistinguishability claims.

Lemma 1. $\text{Expt}_0 \approx_c \text{Hyb}_0$.

Proof. Consider an interaction of \mathcal{A} with a game (either Expt_0 or Hyb_0). In this interaction, let E be the event that \mathcal{A} makes oracle query on the ciphertext ct_2^* with associated data τ^* . Similarly, let F be the event that \mathcal{A} makes $\geq \tau^*$ oracle queries.

Intuitively, we define these events to separately analyze the cases where the indistinguishability of Expt_0 and Hyb_0 is reliant on the BTE scheme and tIBE scheme. Specifically, we will now establish the following three lemmas, and then combine them to prove that $\text{Expt}_0 \approx_c \text{Hyb}_0$.

Reduction algorithm \mathcal{A}_{BTE} :

1. Initialize $\tau_{\text{curr}} := 0$ and $\mathcal{Q}_{\text{BDec}} \leftarrow \emptyset$
2. Receive pp_{BTE} from \mathcal{C}_{BTE} , and sample $\text{pp}_{\text{tIBE}} \leftarrow \text{tIBE.Setup}(1^\lambda)$
3. Set $\text{pp} := (\text{pp}_{\text{BTE}}, \text{pp}_{\text{tIBE}})$ and send pp to \mathcal{A}
4. Receive the corruption set \mathcal{C} from \mathcal{A} and forward it to the challenger \mathcal{C}_{BTE}
5. Receive $(\text{mpk}_{\text{BTE}}, \{\text{sk}_{\text{BTE},i}\}_{i \in \mathcal{C}})$ from \mathcal{C}_{BTE} , and sample $(\text{mpk}_{\text{tIBE}}, \{\text{sk}_{\text{tIBE},i}\}_{i \in [n]}) \leftarrow \text{tIBE.KeyGen}(n, t)$
6. Send $\text{mpk} := (\text{mpk}_{\text{BTE}}, \text{mpk}_{\text{tIBE}})$, $\{\text{sk}_i := \text{sk}_{\text{BTE},i}, \text{sk}_{\text{tIBE},i}\}_{i \in \mathcal{C}}$
7. Simulate the oracle $\mathcal{O}_{\text{TInc}}$ for adversary \mathcal{A} as follows:
 - (a) Update $\tau_{\text{curr}} := \tau_{\text{curr}} + 1$
 - (b) For every query $Q = \{\text{ct}_j\}_{j \in [B]}$ from \mathcal{A} , compute S as in Line 14 in Figure 4, and forward S to \mathcal{C}_{BTE} .
 - (c) Once \mathcal{C}_{BTE} initiates an execution of the protocol $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ on behalf of the set of honest parties $[n] \setminus \mathcal{C}$, forward the protocol messages back and forth between \mathcal{C}_{BTE} and \mathcal{A} (which is executing $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ on behalf of the set of corrupt parties \mathcal{C}). Let $\{\tilde{m}_i\}_{i \in [|\mathcal{S}|]}$ be the output of the execution of $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ computed by the \mathcal{C}_{BTE} on behalf of honest parties. Receive $\{\tilde{m}_i\}_{i \in [|\mathcal{S}|]}$ from \mathcal{C}_{BTE} .
 - (d) Update $\mathcal{Q}_{\text{BDec}} \leftarrow \mathcal{Q}_{\text{BDec}} \cup S$
 - (e) Initiate an execution of the protocol $\Pi\text{-KeyIsc}(\tau_{\text{curr}})$ by simulating the set of honest parties $[n] \setminus \mathcal{C}$ using the secret keys $\{\text{sk}_{\text{tIBE},i}\}_{i \in [n] \setminus \mathcal{C}}$, forward the protocol messages back and forth between the simulated set of honest parties and \mathcal{A} (which is executing $\Pi\text{-KeyIsc}(\tau_{\text{curr}})$ on behalf of the set of corrupt parties \mathcal{C}). Let k_{curr} be the output of the execution of $\Pi\text{-KeyIsc}(\tau_{\text{curr}})$ computed by the simulated honest parties.
 - (f) Send $(k_{\text{curr}}, \{\tilde{m}_i\}_{i \in [|\mathcal{S}|]})$ to \mathcal{A} .
8. Simulate the challenge phase for adversary \mathcal{A} in the following way:
 - (a) Receive the challenge (m_0, m_1, τ^*) from \mathcal{A}
 - (b) Sample $s_{1,0}, s_{1,1} \leftarrow_s \{0, 1\}^\lambda$ and forward $(s_{1,0}, s_{1,1}, \text{ad} := \tau^*)$ to \mathcal{C}_{BTE} .
 - (c) Receive the response ciphertext ct_{BTE}^* from \mathcal{C}_{BTE} where $\text{ct}_{\text{BTE}}^* \leftarrow \text{BTE.Enc}(\text{mpk}_{\text{BTE}}, s_{1,0}, \text{ad} := \tau^*)$ or $\text{ct}_{\text{BTE}}^* \leftarrow \text{BTE.Enc}(\text{mpk}_{\text{BTE}}, s_{1,1}, \text{ad} := \tau^*)$
 - (d) Sample $s_2 \leftarrow_s \{0, 1\}^\lambda$ and compute $\text{ct}_{\text{tIBE}}^* \leftarrow \text{tIBE.Enc}(\text{mpk}_{\text{tIBE}}, s_2, \text{id} := \tau^*)$
 - (e) Compute $\text{ct}_{\text{PRG}}^* := m_0 \oplus \text{PRG}(s_{1,0} \oplus s_2, \ell)$
 - (f) Send $\text{ct}^* := (\tau^*, \text{ct}_{\text{BTE}}^*, \text{ct}_{\text{tIBE}}^*, \text{ct}_{\text{PRG}}^*)$ to \mathcal{A}
9. Let b' be the bit output by \mathcal{A} . Define, \tilde{b} as:

$$\tilde{b} = \begin{cases} 0 & : \text{if } \tau_{\text{curr}} \geq \tau^* \wedge \text{ct}^* \in \mathcal{Q}_{\text{BDec}} \\ b' & : \text{otherwise} \end{cases}$$

10. If $\tau_{\text{curr}} \geq \tau^* \wedge \text{ct}^* \notin \mathcal{Q}_{\text{BDec}}$, output \tilde{b} , otherwise output 0.

Fig. 6: Reduction algorithm \mathcal{A}_{BTE} to break the IND-CCA security of BTE.

Claim 2 $\left| \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg E \wedge F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg E \wedge F] \right| = \text{negl}(\lambda)$.

Proof. Suppose the lemma is false w.r.t. an adversary \mathcal{A} that participates in either Expt_0 and Hyb_0 . We can use \mathcal{A} to construct an adversary \mathcal{A}_{BTE} that contradicts the IND-CCA security of the BTE scheme. Let \mathcal{C}_{BTE} be the challenger for BTE scheme which executes the steps outlined in BTE security game $\text{Expt}_{\mathcal{A}}^{\text{BTE-cca}}$ (see Fig. 2) while interacting with the reduction adversary \mathcal{B} and providing it access to the oracle $\mathcal{O}_{\text{BDec}}$.

Let $\text{Expt}_{\mathcal{A},0}^{\text{BTE-cca}}$ and $\text{Expt}_{\mathcal{A},1}^{\text{BTE-cca}}$ denote the experiment where \mathcal{C}_{BTE} encrypts $s_{1,0}$ and $s_{1,1}$, respectively. The output of these experiments, as defined in BTE security game $\text{Expt}_{\mathcal{A},b}^{\text{BTE-cca}}$ (see Fig. 2), is either 0 if the ciphertext ct_{BTE}^* with associated data $\text{ad} := \tau^*$ is queried to \mathcal{C}_{BTE} or is equal to the bit received from the reduction adversary \mathcal{A}_{BTE} , which we describe in Fig. 6.

By the construction of \mathcal{A}_{BTE} , we note the following:

- When \mathcal{C}_{BTE} encrypts $s_{1,0}$, the view of \mathcal{A} is identical to its view in Expt_0 and the bit \tilde{b} (defined in Line 13 of the adversary \mathcal{A}_{BTE}) is identical to the output of Expt_0 (as defined in Fig. 4).

- Similarly, when \mathcal{C}_{BTE} encrypts $s_{1,1}$, the view of \mathcal{A} is identical to its view in Hyb_0 and the bit \tilde{b} is identical to the output of Hyb_0 .

Moreover, the adversary \mathcal{A}_{BTE} outputs 1 when all the following conditions are true:

- $\tilde{b} = 1$ (equivalent to Expt_0 or Hyb_0 outputting 1)
- $\text{ct}^* \notin \mathcal{Q}_{\text{BDec}}$ (equivalent to the event $\neg E$ occurring)
- $\tau_{\text{curr}} \geq \tau^*$ (equivalent to the event F occurring)

Hence, it follows that,

$$\Pr[\text{Expt}_{\mathcal{A},0}^{\text{BTE-cca}} \Rightarrow 1] = \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg E \wedge F] \quad (1)$$

$$\Pr[\text{Expt}_{\mathcal{A},1}^{\text{BTE-cca}} \Rightarrow 1] = \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg E \wedge F] \quad (2)$$

Subtracting Eq. (2) from Eq. (1), we get

$$\left| \Pr[\text{Expt}_{\mathcal{A},0}^{\text{BTE-cca}} \Rightarrow 1] - \Pr[\text{Expt}_{\mathcal{A},1}^{\text{BTE-cca}} \Rightarrow 1] \right| = \left| \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg E \wedge F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg E \wedge F] \right| \neq \text{negl}(\lambda)$$

This contradicts the IND-CCA security of the BTE scheme. \square

Claim 3 $\left| \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg F] \right| = \text{negl}(\lambda).$

Proof. Suppose the lemma is false w.r.t. an adversary \mathcal{A} that participates in Expt_0 and Hyb_0 . We can use \mathcal{A} to construct an adversary $\mathcal{A}_{\text{tIBE}}$ that contradicts the IND-CPA security of the tIBE scheme. Let $\mathcal{C}_{\text{tIBE}}$ be the challenger for tIBE scheme which executes the steps outlined in tIBE security game (Fig. 1) while interacting with the reduction adversary $\mathcal{A}_{\text{tIBE}}$ and providing it access to the oracle $\mathcal{O}_{\text{KeyIss}}$.

Let $\text{Expt}_{\mathcal{A},0}^{\text{tIBE-cpa}}$ and $\text{Expt}_{\mathcal{A},1}^{\text{tIBE-cpa}}$ denote the experiment where $\mathcal{C}_{\text{tIBE}}$ encrypts $s_{2,0}$ and $s_{2,1}$, respectively. The output of these experiments, as defined in tIBE security game $\text{Expt}_{\mathcal{A},b}^{\text{tIBE-cpa}}$ (see Fig. 1), is 0 if $\text{id} := \tau^*$ is queried to $\mathcal{C}_{\text{tIBE}}$, else the output is equal to the bit received from the reduction adversary $\mathcal{A}_{\text{tIBE}}$, which we describe in Fig. 7.

By the construction of $\mathcal{A}_{\text{tIBE}}$, we note the following:

- When $\mathcal{C}_{\text{tIBE}}$ encrypts $s_{2,0}$, the view of \mathcal{A} is identical to its view in Expt_0 and the bit \tilde{b} (defined in Line 13 of the adversary $\mathcal{A}_{\text{tIBE}}$) is identical to the output of Expt_0 (as defined in Fig. 4).
- Similarly, when $\mathcal{C}_{\text{tIBE}}$ encrypts $s_{2,1}$, the view of \mathcal{A} is identical to its view in Hyb_0 and the bit \tilde{b} is identical to the output of Hyb_0 .

Note that $\mathcal{A}_{\text{tIBE}}$ outputs 1 when both of the following conditions are true:

- $\tilde{b} = 1$ (equivalent to Expt_0 or Hyb_0 outputting 1); and
- $\tau_{\text{curr}} < \tau^*$ (equivalent to the event $\neg F$ occurring)

Hence, it follows that,

$$\Pr[\text{Expt}_{\mathcal{A},0}^{\text{tIBE-cpa}} \Rightarrow 1] = \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg F] \quad (3)$$

$$\Pr[\text{Expt}_{\mathcal{A},1}^{\text{tIBE-cpa}} \Rightarrow 1] = \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg F] \quad (4)$$

Subtracting Eq. (4) from Eq. (3), we get

$$\left| \Pr[\text{Expt}_{\mathcal{A},0}^{\text{tIBE-cpa}} \Rightarrow 1] - \Pr[\text{Expt}_{\mathcal{A},1}^{\text{tIBE-cpa}} \Rightarrow 1] \right| = \left| \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg F] \right| \neq \text{negl}(\lambda)$$

This contradicts the IND-CPA security of the tIBE scheme. \square

Reduction algorithm $\mathcal{A}_{\text{tIBE}}$:

1. Initialize $\tau_{\text{curr}} := 0$ $\mathcal{Q}_{\text{BDec}} \leftarrow \emptyset$
2. Receive pp_{tIBE} from $\mathcal{C}_{\text{tIBE}}$ and sample $\text{pp}_{\text{BTE}} \leftarrow \text{BTE.Setup}(1^\lambda)$
3. Set $\text{pp} := (\text{pp}_{\text{BTE}}, \text{pp}_{\text{tIBE}})$ and send pp to \mathcal{A}
4. Receive the corruption set \mathcal{C} from \mathcal{A} and forward it to the challenger $\mathcal{C}_{\text{tIBE}}$
5. Receive $(\text{mpk}_{\text{tIBE}}, \{\text{sk}_{\text{tIBE},i}\}_{i \in \mathcal{C}})$ from $\mathcal{C}_{\text{tIBE}}$ and sample $(\text{mpk}_{\text{BTE}}, \{\text{sk}_{\text{BTE},i}\}_{i \in [n]}) \leftarrow \text{BTE.KeyGen}(n, t)$
6. Send $\text{mpk} := (\text{mpk}_{\text{BTE}}, \text{mpk}_{\text{tIBE}})$, $\{\text{sk}_i := \text{sk}_{\text{BTE},i}, \text{sk}_{\text{tIBE},i}\}_{i \in [\mathcal{C}]}$ to \mathcal{A}
7. Simulate the oracle $\mathcal{O}_{\text{TInc}}$ for adversary \mathcal{A} in the following way:
 - (a) Update $\tau_{\text{curr}} := \tau_{\text{curr}} + 1$
 - (b) For every query $Q = \{\text{ct}_j\}_{j \in [B]}$ from \mathcal{A} , compute the set S as defined in Line 14 in Figure 4. Initiate an execution of the protocol $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ by simulating the set of honest parties $[n] \setminus \mathcal{C}$ using the secret keys $\{\text{sk}_{\text{BTE},i}\}_{i \in [n] \setminus \mathcal{C}}$, forward the protocol messages back and forth between the simulated set of honest parties and \mathcal{A} (which is executing $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ on behalf of the set of corrupt parties \mathcal{C}). Let $\{\tilde{m}_i\}_{i \in [|\mathcal{S}|]}$ be the output of the execution of $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ computed by the simulated honest parties.
 - (c) Update $\mathcal{Q}_{\text{BDec}} \leftarrow \mathcal{Q}_{\text{BDec}} \cup S$
 - (d) Send $\text{id} := \tau_{\text{curr}}$ to $\mathcal{C}_{\text{tIBE}}$.
 - (e) Once $\mathcal{C}_{\text{tIBE}}$ initiates an execution of the protocol $\Pi\text{-KeyIss}(\text{id} := \tau_{\text{curr}})$ on behalf of the set of honest parties $[n] \setminus \mathcal{C}$, forward the protocol messages back and forth between $\mathcal{C}_{\text{tIBE}}$ and \mathcal{A} (which is executing $\Pi\text{-KeyIss}(\text{id} := \tau_{\text{curr}})$ on behalf of the set of corrupt parties \mathcal{C}). Let k_{curr} be the output of the execution of $\Pi\text{-KeyIss}(\text{id} := \tau_{\text{curr}})$ computed by the $\mathcal{C}_{\text{tIBE}}$ on behalf of honest parties. Receive k_{curr} from $\mathcal{C}_{\text{tIBE}}$.
 - (f) Send $(k_{\text{curr}}, \{\tilde{m}_i\}_{i \in [|\mathcal{S}|]})$ to \mathcal{A} .
8. Simulate the challenge phase for adversary \mathcal{A} in the following way:
 - (a) Receive the challenge (m_0, m_1, τ^*) from \mathcal{A}
 - (b) Sample $s_{2,0}, s_{2,1} \leftarrow_{\$} \{0, 1\}^\lambda$ and forward $(s_{2,0}, s_{2,1}, \text{id} := \tau^*)$ to $\mathcal{C}_{\text{tIBE}}$.
 - (c) Receive the response ciphertext $\text{ct}_{\text{tIBE}}^*$ from $\mathcal{C}_{\text{tIBE}}$ where $\text{ct}_{\text{tIBE}}^* \leftarrow \text{tIBE.Enc}(\text{mpk}_{\text{tIBE}}, s_{2,0}, \text{id} := \tau^*)$ or $\text{ct}_{\text{tIBE}}^* \leftarrow \text{tIBE.Enc}(\text{mpk}_{\text{tIBE}}, s_{2,1}, \text{id} := \tau^*)$
 - (d) Sample $s_1 \leftarrow_{\$} \{0, 1\}^\lambda$ and compute $\text{ct}_{\text{BTE}}^* \leftarrow \text{BTE.Enc}(\text{mpk}_{\text{BTE}}, s_1, \text{ad} := \tau^*)$
 - (e) Compute $\text{ct}_{\text{PRG}}^* := m_0 \oplus \text{PRG}(s_1 \oplus s_{2,0}, \ell)$
 - (f) Set $\text{ct}^* := (\tau^*, \text{ct}_{\text{BTE}}^*, \text{ct}_{\text{tIBE}}^*, \text{ct}_{\text{PRG}}^*)$ and send ct^* to \mathcal{A}
9. Let b' be the bit output by \mathcal{A} . Define, \tilde{b} as:

$$\tilde{b} = \begin{cases} 0 & \text{if } \tau_{\text{curr}} \geq \tau^* \wedge \text{ct}^* \in \mathcal{Q}_{\text{BDec}} \\ b' & \text{otherwise} \end{cases}$$

10. If $\tau_{\text{curr}} < \tau^*$, output \tilde{b} , otherwise output 0.

Fig. 7: Reduction algorithm $\mathcal{A}_{\text{tIBE}}$ to break the IND-CPA security of tIBE.

Claim 4 $\left| \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \text{E} \wedge \text{F}] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \text{E} \wedge \text{F}] \right| = 0$

Proof. The proof is trivial. When events E and F both occur, then Expt_0 always outputs 0 (see line 8 in Fig. 4), and so does Hyb_0 . Therefore, $\Pr[\text{Expt}_0 \Rightarrow 1 \wedge \text{E} \wedge \text{F}] = \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \text{E} \wedge \text{F}] = 0$. \square

Now that we have proved all the three intermediate lemmas, we move on to proving the claim. Note that the following holds.

$$\Pr[\text{Expt}_0 \Rightarrow 1] = \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg \text{F}] + \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg \text{E} \wedge \text{F}] + \Pr[\text{Expt}_0 \Rightarrow 1 \wedge \text{E} \wedge \text{F}] \quad (5)$$

$$\Pr[\text{Hyb}_0 \Rightarrow 1] = \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg \text{F}] + \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg \text{E} \wedge \text{F}] + \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \text{E} \wedge \text{F}] \quad (6)$$

Subtracting Eq. (6) from Eq. (5), we get

Reduction algorithm \mathcal{A}_{PRG} :

1. Initialize $\tau_{\text{curr}} := 0$ and $\mathcal{Q}_{\text{BDec}} \leftarrow \emptyset$
2. Sample $\text{pp}_{\text{BTE}} \leftarrow \text{BTE.Setup}(1^\lambda, 1^{\text{Bmax}})$ and $\text{pp}_{\text{tIBE}} \leftarrow \text{tIBE.Setup}(1^\lambda)$
3. Set $\text{pp} := (\text{pp}_{\text{BTE}}, \text{pp}_{\text{tIBE}})$ and send pp to \mathcal{A} .
4. Receive the corruption set \mathcal{C} from \mathcal{A} .
5. Sample $(\text{mpk}_{\text{BTE}}, \{\text{sk}_{\text{BTE},i}\}_{i \in [n]}) \leftarrow \text{BTE.KeyGen}(n, t)$ and $(\text{mpk}_{\text{tIBE}}, \{\text{sk}_{\text{tIBE},i}\}_{i \in [n]}) \leftarrow \text{tIBE.KeyGen}(n, t)$
6. Send $\text{mpk}, \{\text{sk}_i\}_{i \in \mathcal{C}}$ to \mathcal{A} .
7. Simulate the oracle $\mathcal{O}_{\text{Tinc}}$ for adversary \mathcal{A} in the following way:
 - (a) Update $\tau_{\text{curr}} := \tau_{\text{curr}} + 1$
 - (b) For every query $Q = \{\text{ct}_j\}_{j \in [B]}$ from \mathcal{A} , compute the set S as defined in Line 14 in Figure 4. Initiate an execution of the protocol $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ by simulating the set of honest parties $[n] \setminus \mathcal{C}$ using the secret keys $\{\text{sk}_{\text{BTE},i}\}_{i \in [n] \setminus \mathcal{C}}$, forward the protocol messages back and forth between the simulated set of honest parties and \mathcal{A} (which is executing $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ on behalf of the set of corrupt parties \mathcal{C}). Let $\{\tilde{m}_i\}_{i \in [|S|]}$ be the output of the execution of $\Pi\text{-BDec}(\text{mpk}_{\text{BTE}}, S)$ computed by the simulated honest parties.
 - (c) Update $\mathcal{Q}_{\text{BDec}} \leftarrow \mathcal{Q}_{\text{BDec}} \cup S$
 - (d) Initiate an execution of the protocol $\Pi\text{-Keylss}(\tau_{\text{curr}})$ by simulating the set of honest parties $[n] \setminus \mathcal{C}$ using the secret keys $\{\text{sk}_{\text{tIBE},i}\}_{i \in [n] \setminus \mathcal{C}}$, forward the protocol messages back and forth between the simulated set of honest parties and \mathcal{A} (which is executing $\Pi\text{-Keylss}(\tau_{\text{curr}})$ on behalf of the set of corrupt parties \mathcal{C}). Let k_{curr} be the output of the execution of $\Pi\text{-Keylss}(\tau_{\text{curr}})$ computed by the simulated honest parties.
 - (e) Send $(k_{\text{curr}}, \{\tilde{m}_i\}_{i \in [|S|]})$ to \mathcal{A} .
8. Simulate the challenge phase for adversary \mathcal{A} in the following way:
 - (a) Receive the challenge (m_0, m_1, τ^*) from \mathcal{A}
 - (b) Sample $s_1, s_2 \leftarrow_{\$} \{0, 1\}^\lambda$
 - (c) Compute $\text{ct}_{\text{BTE}}^* \leftarrow \text{BTE.Enc}(\text{mpk}_{\text{BTE}}, s_1, \text{ad} := \tau^*)$ and $\text{ct}_{\text{tIBE}}^* \leftarrow \text{tIBE.Enc}(\text{mpk}_{\text{tIBE}}, s_2, \text{id} := \tau^*)$
 - (d) Receive a challenge string u of length ℓ from \mathcal{C}_{PRG} where $u := \text{PRG}(s)$ for $s \leftarrow_{\$} \{0, 1\}^\lambda$ or $u \leftarrow_{\$} \{0, 1\}^\ell$
 - (e) Send $\text{ct}^* := (\tau^*, \text{ct}_{\text{BTE}}^*, \text{ct}_{\text{tIBE}}^*, \text{ct}_{\text{PRG}}^* := m_0 \oplus u)$ to \mathcal{A}
9. Let b' be the bit output by \mathcal{A} . If $\tau_{\text{curr}} \geq \tau^* \wedge \text{ct}^* \in \mathcal{Q}_{\text{BDec}}$, output 0. Otherwise, output b' .

Fig. 8: Reduction algorithm \mathcal{A}_{PRG} to break the pseudorandom property of PRG.

$$\begin{aligned}
\Pr[\text{Expt}_0 \Rightarrow 1] - \Pr[\text{Hyb}_0 \Rightarrow 1] &= \left(\Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg F] \right) \\
&\quad + \left(\Pr[\text{Expt}_0 \Rightarrow 1 \wedge \neg E \wedge F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge \neg E \wedge F] \right) \\
&\quad + \left(\Pr[\text{Expt}_0 \Rightarrow 1 \wedge E \wedge F] - \Pr[\text{Hyb}_0 \Rightarrow 1 \wedge E \wedge F] \right) \tag{7}
\end{aligned}$$

Now, by invoking Theorem 2, Theorem 3 and Theorem 4 in Eq. (7), it follows that:

$$\left| \Pr[\text{Expt}_0 \Rightarrow 1] - \Pr[\text{Hyb}_0 \Rightarrow 1] \right| = \text{negl}(\lambda)$$

This completes the proof of our claim. \square

Lemma 2. $\text{Hyb}_0 \approx_c \text{Hyb}'$.

Proof. Suppose that the claim is false w.r.t an adversary \mathcal{A} that participates in Hyb_0 and Hyb' . We can use \mathcal{A} to construct an adversary \mathcal{A}_{PRG} that contradicts the security of the PRG scheme. Let \mathcal{C}_{PRG} be the challenger for PRG scheme.

Let Expt_{PRG} and $\text{Expt}_{\text{Rand}}$ denote the experiment where \mathcal{C}_{PRG} sets $u := \text{PRG}(s)$ and $u \leftarrow_{\$} \{0, 1\}^\ell$ respectively. The output of these experiments, as defined in PRG security game, is equal to the bit received by the \mathcal{C}_{PRG} from the reduction adversary \mathcal{B} . By the construction of \mathcal{B} , we note the following:

- When \mathcal{C}_{PRG} sets $u := \text{PRG}(s)$ for $s \leftarrow_{\$} \{0, 1\}^\lambda$, the view of \mathcal{A} is identical to its view in Hyb_0 and the output of \mathcal{A}_{PRG} (see Line 12 of Fig. 8) is identical to the output of Hyb_0 .
- Similarly, when \mathcal{C}_{PRG} samples $u \leftarrow_{\$} \{0, 1\}^\ell$, the view of \mathcal{A} is identical to its view in Hyb' and the output of \mathcal{A}_{PRG} is identical to the output of Hyb' .

Therefore, we get:

$$\left| \Pr[\text{Expt}_{\text{PRG}} \Rightarrow 1] - \Pr[\text{Expt}_{\text{Rand}} \Rightarrow 1] \right| = \left| \Pr[\text{Hyb}_0 \Rightarrow 1] - \Pr[\text{Hyb}' \Rightarrow 1] \right| \neq \text{negl}(\lambda).$$

This contradicts the security of the PRG scheme and concludes the proof of our claim. \square

Lemma 3. $\text{Hyb}' \approx_c \text{Hyb}_1$.

Proof. The proof for this claim follows from the security of the PRG scheme. Formally, the proof is identical to the proof of Lemma 2 ($\text{Hyb}_0 \approx_c \text{Hyb}'$) with the exception that m_0 is replaced with m_1 when the reduction adversary \mathcal{B} creates the ciphertext ct_{PRG}^* . \square

Lemma 4. $\text{Hyb}_1 \approx_c \text{Expt}_1$.

Proof. The proof for this claim follows from the IND-CCA security of the BTE scheme and IND-CPA security of the threshold IBE scheme. Formally, the proof is identical to the proof of Lemma 1 ($\text{Expt}_0 \approx_c \text{Hyb}_0$) with the exception that m_0 is replaced with m_1 when the reduction adversary \mathcal{B} creates the ciphertext ct_{PRG}^* . \square

Next, we will proof robustness of our scheme, assuming that the underlying schemes BTE and tIBE is robust as per Definitions 4 and 8, respectively. Formally,

Theorem 5 (tTLES Robustness). *Assuming the existence of an robust threshold Identity Based Encryption (IBE) scheme (Definition 4) and a robust Batched Threshold Encryption (BTE) scheme (Definition 8), our tTLES protocol tTLES (Fig. 5) is robust as per Definition 13.*

Proof. Any ciphertext ct in our tTLES is a tuple of four elements, i.e., $\text{ct} = (\text{ct}_1, \text{ct}_2, \text{ct}_3, \text{ct}_4)$ where ct_2 is a BTE ciphertext of random $s_1 \in \{0, 1\}^\lambda$ with associated data τ , ct_3 is a tIBE ciphertext encrypting another random value s_2 under identity τ , and $\text{ct}_4 = m \oplus \text{PRG}(s_1 \oplus s_2, |m|)$.

Consequently, the tTLES decryption has two main steps: (i) decrypting ct_2 to recover s_1 , and (ii), decrypting ct_3 to recover s_2 . Clearly, if s_1 and s_2 are correct, since ct_4 is a deterministic function of s_1, s_2 and the underlying message m , we will correctly recover original m .

Next, we will show that except with negligible probability, both s_1 and s_2 are correctly recovered for all honestly encrypted messages in a batch. Let F_{BTE} be the event that BTE decryption of any honestly encrypted message in the batch output by the adversary fails, and F_{tIBE} be the event that tIBE decryption of any honestly encrypted message in the batch fails. Then, from Definitions 4 and 8, we have that there exist negligible functions $\text{negl}_1(\cdot), \text{negl}_2(\cdot)$ such that:

$$\Pr[F_{\text{BTE}}] = \Pr[\text{Expt}_{\mathcal{A}}^{\text{BTE-rob}}(1^\lambda) \Rightarrow 1] \leq \text{negl}_1(\lambda), \quad \text{and} \quad \Pr[F_{\text{tIBE}}] = \Pr[\text{Expt}_{\mathcal{A}}^{\text{tIBE-rob}}(1^\lambda) \Rightarrow 1] \leq \text{negl}_2(\lambda)$$

Now, using union bound, we have that:

$$\Pr[F_{\text{BTE}} \vee F_{\text{tIBE}}] \leq \text{negl}_1(\lambda) + \text{negl}_2(\lambda)$$

Finally, if neither F_{BTE} nor F_{tIBE} occurs, then all honestly encrypted messages in the batch are correctly decrypted, i.e.,

$$\Pr[\text{Expt}_{\mathcal{A}}^{\text{tTLES-rob}} \Rightarrow 1] \leq \Pr[F_{\text{BTE}} \vee F_{\text{tIBE}}] \leq \text{negl}(\lambda)$$

for some negligible function $\text{negl}(\cdot)$. This completes the proof that our tTLES protocol is robust. \square

5 Discussion and Conclusion

In this paper, we presented time-lock encrypted storage (tTLES), a storage service provided by blockchains where clients can submit (using transactions) encrypted values to be decrypted at a future time τ_{tgt} (measured in block height). The security guarantees of tTLES require that the submitted values can only be decrypted after the time τ_{tgt} has elapsed (the time-lock indistinguishability property) and the transactions submitting the encrypted values have been included in the blockchain (the non-inclusion privacy property). We have provided a formal definition of tTLES and proposed an efficient protocol that relies on (in a black-box manner) a threshold identity-based encryption (tIBE) scheme and a recent batch threshold decryption (BTE) scheme. Our proposed solution allows for an unbounded number of ciphertexts to be encrypted towards any decryption time τ_{tgt} while keeping the decryption cost (both communication and computation) at time τ_{tgt} independent of the number of such ciphertexts.

We suggest two natural extensions of our tTLES design.

Generalizing to other decryption policies. Our tTLES protocols guarantee time-lock indistinguishability and non-inclusion privacy by secret sharing the ephemeral secret key ($s_1 \oplus s_2$) with 2-out-of-2 secret sharing, encrypting the share s_1 with tIBE and the other share s_2 with BTE. This ensures that the decryption requires both shares s_1 and s_2 . A natural extension of this approach is to support more expressive decryption policies, such as t -out-of- n decryption policies by secret sharing the ephemeral key using threshold secret sharing [17]. We note that these decryption conditions can be either internal to the blockchain protocols or can be fed in externally using external oracles.

Relaxing the encrypted storage inclusion condition. Our tTLES construction works with the inclusion condition that each encrypted storage encrypted towards a target decryption time τ_{tgt} gets included in a block before the time τ_{tgt} has elapsed. Although this inclusion condition is sufficient for the applications we discussed, some applications can benefit from a more flexible inclusion condition. For example, we may allow the client to specify two timestamps: the target decryption time τ_{tgt} , and an expiry time τ_{exp} , with possibly $\tau_{\text{exp}} \geq \tau_{\text{tgt}}$. Under this relaxed condition, encrypted values may still be included after τ_{tgt} , enabling use cases where users want ciphertexts to remain valid candidates for inclusion even after τ_{tgt} has passed.

Acknowledgements

We thank John Bergschneider, Dan Boneh, Mike Setrin and Ayush Yadav for their feedback and helpful discussions related to this paper. We specially thank Michael Setrin for discussions that motivated this work.

References

1. Agarwal, A., Fernando, R., Pinkas, B.: Efficiently-thresholdizable batched identity based encryption, with applications. In: Annual International Cryptology Conference. pp. 69–100. Springer (2025)
2. Boneh, D., Franklin, M.: Identity-based encryption from the weil pairing. In: Annual international cryptology conference. pp. 213–229. Springer (2001)
3. Boneh, D., Laufer, E., Tas, E.N.: Batch decryption without epochs and its application to encrypted mempools. Cryptology ePrint Archive (2025)
4. Boneh, D., Shoup, V.: A graduate course in applied cryptography. Draft 0.6 (2023)
5. Bormet, J., Choudhuri, A.R., Faust, S., Garg, S., Othman, H., Policharla, G.V., Qu, Z., Wang, M.: Beast-mev: Batched threshold encryption with silent setup for mev prevention. Cryptology ePrint Archive (2025)
6. Bormet, J., Faust, S., Othman, H., Qu, Z.: {BEAT-MEV}: Epochless approach to batched threshold encryption for {MEV} prevention. In: 34th USENIX Security Symposium (USENIX Security 25). pp. 3457–3476 (2025)
7. Cerulli, A., Connolly, A., Neven, G., Preiss, F.S., Shoup, V.: vetkeys: How a blockchain can keep many secrets. Cryptology ePrint Archive (2023)
8. Choudhuri, A.R., Garg, S., Piet, J., Policharla, G.V.: Mempool privacy via batched threshold encryption: Attacks and defenses. In: 33rd USENIX Security Symposium (USENIX Security 24). pp. 3513–3529 (2024)

9. Choudhuri, A.R., Garg, S., Policharla, G.V., Wang, M.: Practical mempool privacy via one-time setup batched threshold encryption. In: 34th USENIX Security Symposium (USENIX Security 25). pp. 3477–3495 (2025)
10. Ding, Q., Liebau, D., Wang, Z., Xu, W.: A survey on decentralized autonomous organizations (daos) and their governance. *World Scientific Annual Review of Fintech* **1**, 2350001 (2023)
11. Dujmovic, J., Garg, R., Malavolta, G.: Time-lock puzzles with efficient batch solving. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 311–341. Springer (2024)
12. Dziembowski, S., Faust, S., Luhn, J.: Shutter network: Private transactions from threshold cryptography. *Cryptology ePrint Archive* (2024)
13. Gailly, N., Melissaris, K., Romailier, Y.: tlock: Practical timelock encryption from threshold bls. *Cryptology ePrint Archive* (2023)
14. Galal, H.S., Youssef, A.M.: Succinctly verifiable sealed-bid auction smart contract. In: International Workshop on Data Privacy Management. pp. 3–19. Springer (2018)
15. Mahmoody, M., Moran, T., Vadhan, S.: Time-lock puzzles in the random oracle model. In: Annual Cryptology Conference. pp. 39–50. Springer (2011)
16. Rivest, R.L., Shamir, A., Wagner, D.A.: Time-lock puzzles and timed-release crypto (1996)
17. Shamir, A.: How to share a secret. *Commun. ACM* **22**(11), 612–613 (Nov 1979). <https://doi.org/10.1145/359168.359176>, <https://doi.org/10.1145/359168.359176>
18. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Workshop on the theory and application of cryptographic techniques. pp. 47–53. Springer (1984)
19. Xiong, J., Wang, Q.: Anonymous auction protocol based on time-released encryption atop consortium blockchain. *arXiv preprint arXiv:1903.03285* (2019)