

Time-Deniable Signatures

Gabrielle Beck¹, Arka Rai Choudhuri², Matthew Green¹, Abhishek Jain¹, and Pratyush Ranjan Tiwari¹

¹Johns Hopkins University, {becgabri,mgreen,abhishek,pratyush}@cs.jhu.edu
²UC Berkeley, arkarc@berkeley.edu

Abstract

In this work we propose *time-deniable signatures* (TDS), a new primitive that facilitates deniable authentication in protocols such as DKIM-signed email. As with traditional signatures, TDS provide strong authenticity for message content, at least *for a sender-chosen period of time*. Once this time period has elapsed, however, time-deniable signatures can be forged by any party who obtains a signature. This forgery property ensures that signatures serve a useful authentication purpose for a bounded time period, while also allowing signers to plausibly disavow the creation of older signed content. Most critically, and unlike many past proposals for deniable authentication, TDS do not require interaction with the receiver or the deployment of any persistent cryptographic infrastructure or services beyond the signing process (*e.g.*, APIs to publish secrets or author timestamp certificates.)

We first investigate the security definitions for time-deniability, demonstrating that past definitional attempts are insufficient (and indeed, allow for broken signature schemes.) We then propose an efficient construction of TDS based on well-studied assumptions.

1 Introduction

Many communication systems use cryptographic signatures to verify the authenticity of data sent from one party to another over untrusted networks. While cryptographic authentication is standard in end-to-end encrypted messaging systems, it is also increasingly being deployed within traditionally non-encrypted protocols such as SMTP email. Specifically in the email setting, protocols such as DKIM, DMARC and ARC [13] are routinely used to add non-repudiable digital signatures to email in transit between Mail Transfer Agents (MTAs): these signatures allow recipient spam filtering software to verify that it originates from the claimed sender.

While cryptographic authenticity is valuable for preventing spam and spoofing of email traffic, DKIM signatures have been re-purposed for goals that may not have been anticipated by the designers of these protocols.¹ For example, news organizations routinely verify the authenticity of leaked or stolen email collections using DKIM signatures [2, 31, 39]: this is possible because DKIM signing keys are long-lived, and the protocol’s non-repudiable signatures can be verified long after email has been received and processed. To facilitate such verification, organizations such as the Associated Press and Wikileaks even publish detailed instructions and tools for verifying

¹Indeed, many early deployments of DKIM used weak signing keys, and some DKIM standards authors proposed using *e.g.*, 600-bit keys to balance the risks and benefits of DKIM [13].

the authenticity of DKIM signatures in leaked and stolen email corpora. Since email signing is implemented by commercial mail providers rather than end-users, users of popular services cannot opt out. These developments have ignited a technical debate around the desirability of long-term non-repudiability guarantees in widely-used protocols such as email [24], and raised questions around the value of adding cryptographic *deniability* to these systems.

The need for deniability. Cryptographic deniability is a property that allows communication participants to disavow authorship of messages, *e.g.*, in the event that they have been leaked or stolen. This feature has frequently been incorporated in *interactive* messaging protocols [9, 3, 40], which historically realize deniability through the use of interactive key exchange protocols and symmetric authentication primitives such as MACs. Achieving deniable authentication in email authentication protocols such as SMTP/DKIM is more challenging, since these protocols support non-interactive and asynchronous delivery via multiple intermediate recipients. Thus interactive protocols are ruled out, and even designated-verifier solutions can be more challenging due to the presence of intermediaries.

Despite these challenges, the problem of incorporating deniability for the email setting has recently received some attention. For example, in Usenix Security 2021, Specter *et al.* proposed two technical replacements for DKIM signing that are designed to facilitate deniability. Both protocols ensure that messages are digitally signed to enable sender-authenticity verification, but feature a process wherein senders, recipients and even third parties can create deliberate *forgeries* after the necessary anti-spam and spoofing checks have been completed. The two protocols employ different techniques: the first relies on the sender to author forgeries on request and/or publish expired secret keys, while the second employs a trusted *time server* that publishes cryptographic timestamp certificates that allow forgery of signatures after some period of time has elapsed. Others have made even simpler proposals wherein DKIM providers simply rotate and publish existing DKIM signing keys on a periodic basis [12, 24]. Each proposal seeks to build signatures that are unforgeable for a period of time necessary to support short-term transport checks, but become forgeable after this period.

The major limitation of the proposals above is that forgery requires the active cooperation of signers, or else depends on the continuous operation of new trusted infrastructure such as “time servers” that publish keys or timestamp certificates on a periodic basis [38]. The challenge in email systems is that the end-users affected by non-repudiable authentication (*e.g.*, Gmail customers) rely on third-party providers to deploy these infrastructure services and make them available for the often-controversial purpose of forging past email. If this infrastructure is not deployed, then even the Internet-wide adoption of a deniable signature standard will not provide deniability *in practice*. What is needed is a signature scheme that can be used in place of a normal signature scheme within protocols; provides strong authenticity for a period of time; and then subsequently becomes plausibly forgeable *by any party who simply obtains such a signature*, with only the requirement that parties have an (approximately) shared view of time. We refer to such signatures as *time-deniable signatures*.

Properties of time-deniable signatures. Time-deniable signatures operate much like a normal signature scheme, but with some important differences. Like standard digital signatures, time-deniable signatures are designed to be secure and non-repudiable for at least some time period following signing. The duration of this time period is strictly limited, however: any party who obtains a signature on some message M can use it as input to a new forging algorithm called AltSign that, after enforcing some approximate time delay, will output a forgery on a new chosen

message M' . A key requirement of these schemes is that neither signing nor forging should require the cooperation of any other party or infrastructure. This time delay is therefore enforced using a specific computational assumption: the AltSign algorithm requires the forger to perform a pre-specified number of *sequential* operations δ , where the minimum time required for this calculation is roughly as long as the desired length of the unforgeable phase.

Of course, the ability to forge signatures has no bearing on deniability if the resulting forgeries are easily distinguishable from authentic signatures. To achieve plausible deniability, we therefore require that forgeries are indistinguishable from signatures produced using the ordinary signing algorithm, and in fact that even *linking* forgeries to the specific signatures that were used to create them should be challenging. This indistinguishability property is a fundamentally novel property of this work, that is not present in previous attempts to solve this problem [4, 25]. It also has important follow-on implications: since forgeries are indistinguishable from true signatures, this implies that any forgery must be useful to create still further forgeries.

Finally, we wish time deniable signatures to be useful in practice. Given the description above, time-deniable signatures would be of limited usefulness: the revelation of a single signature would allow for an unlimited number of forgeries, rendering the signing key useless for authenticating further messages. To remove this limitation, we slightly relax our forgery and unlinkability requirements. Our constructions allow for renewability via an additional *timestamp* t field that is specified in the signing algorithm and carried with the signature. Forgers can produce a new signature on a message M' provided the new signature carries a timestamp $t' \leq t$. For example, in a practical deployment, the timestamp t can be set to correspond to some real-world time counter, and recipients can choose to accept as authentic any signature with timestamp greater than $t - t_\delta$ where t_δ is the minimal expected time needed to compute a forgery.² This approach requires only that honest senders and receivers possess loosely synchronized clocks.

Our contributions. In this work we investigate the problem of building time-deniable signatures. We first develop formal definitions for this new primitive, and then present a construction that is based on several efficient components. Finally, we implement our approach and show that it is practical enough to deploy today. Concretely, we provide the following contributions:

Defining time-deniable signatures (TDS). We propose new definitions for the concept of time-deniable signatures, and propose strong security definitions for this new primitive. Defining security for time-deniable signatures is surprisingly difficult: in the course of developing our definitions, we found that previous efforts to formalize the security of deniable authentication schemes fall short. For example, we show that the security definitions for some related primitives [25] contain subtle weaknesses that admit practically-insecure constructions. To provide evidence for the robustness of our definitions, we prove that our definitions are strictly stronger than these earlier definitions.

Efficient constructions. To demonstrate that the TDS primitive is practical, we propose an efficient construction of time-deniable signatures based on well-studied cryptographic assumptions. Our constructions improve on previous work [25] in that they do *not* require any *a priori* bound on the number of time epochs that the scheme can handle. We also show that TDS

²This naturally relaxes the unlinkability requirement: given a pair of signatures $\sigma_{t_1}, \sigma_{t_2}$ with timestamps $t_1 < t_2$ it cannot be the case that σ_{t_1} is the original signature and σ_{t_2} is the forgery. However given a sufficiently large collection of signatures containing forgeries and original signatures, this approach still provides a degree of uncertainty for all but the most recent signature.

can be realized using standard assumptions in pairing-based cryptography and sequential puzzles based on repeated-squaring assumption [35], without the need for zkSNARKs or other heavy-weight constructions.

Implementation and performance experiments. To further motivate the usefulness of TDS in systems applications, we implement our TDS constructions and show that the scheme has practical runtime and bandwidth performance for the applications we consider. In particular, we show that our scheme has a fast key setup time, which is particularly important for a scheme with an unbounded number of time epochs.

2 Technical Overview

We now give an overview of the main contributions in this work, starting with formalizing the notion, before moving on to the constructions.

2.1 Defining Time-Deniable Signatures

We study signature schemes where signatures remain valid for a short period of time after creation. Specifically, we consider the notion of an *unforgeability period* that starts when a signer generates a signature for a message using its signing key sk , and the signing algorithm $Sign$. But once the *unforgeability period* elapses, any participant in the system can compute a “fake signature” (aka *forgery*). To allow computation of forgeries, we consider an alternate signing algorithm $AltSign$, that *does not require* the signing key sk to generate signatures. Intuitively, as long as the signatures generated by $Sign$ and $AltSign$ appear indistinguishable, such a notion provides *deniability* after the unforgeability period since a signer can claim that a signature attributed to them could have been generated by anyone.

Key Challenges in the Definition. There are several key considerations towards formalizing the above intuition and defining time-deniable signatures.

Challenge I: Preventing pre-computation of forgeries. Recall that any party can compute a forgery (via the algorithm $AltSign$) after the unforgeability period expires. But how do we ensure that a party cannot execute $AltSign$ *in advance*, thereby having the ability to sign any message *within* the unforgeability period?

One natural approach is to bind signatures to some unpredictable *cryptographic beacon*, perhaps generated at regular intervals by a centralized server or a blockchain [18, 33]. For example, when signing a message m (via $Sign$ or $AltSign$) one might actually sign the pair (m, b) where b is a beacon released at a time known by the receiver. This value b can then be used as the “seed” to allow forgery using $AltSign$, and verifiers can use the known publication time of b to determine whether the signature is still within the unforgeability period. Such models have been considered in prior works, including the TimeForge scheme of Specter *et al.* [38] and a recent proposal by Bonneau *et al.* [4].

In this work, we seek to avoid the use of unpredictable timestamps or centralized servers. In our notion, the $Sign$ and $AltSign$ algorithms do indeed take as input a timestamp t . Assuming that receivers possess loosely synchronized clocks, these timestamps can be used to verify that a received signature was authored within the unforgeability period. However, crucially, these timestamps are simply the output of a predictable clock operated by the signer, which means that we *do not require*

any security properties of this input, nor do we require unpredictable beacons or new infrastructure to produce them. To prevent pre-computation, we instead model AltSign such that it requires a *valid signature* on some pair (m, t) as input. This ensures that forgers do not have the necessary input(s) to pre-compute forgeries until they obtain a signature.³

Challenge II: Selecting forged timestamps. In the proposal above, AltSign requires a valid signature on some time t (and any message) in order to compute a forgery. Naturally the resulting forgery will also need to contain its own timestamp t' . The selection of t' is crucial, however: if this forged timestamp can be chosen arbitrarily by the forger, then an attacker may be able to forge new signatures that appear (to an honest receiver) to be within the unforgeability window, even when the original signature was not. One obvious solution to this problem is to restrict the forged timestamp to $t' = t$. Unfortunately, this restriction weakens the deniability properties of the signature scheme: a signer can deny having signed a particular message at time t , but it *cannot* deny having signed *some* message at time t . To achieve stronger deniability where a signer can also deny having signed *any* message at time t , we further strengthen the AltSign algorithm. Namely, we require that on input a signature on timestamp t , AltSign can compute forgeries for any message m and any time stamp $t' \leq t$.

Challenge III: Avoiding strong clock synchronization. The closely related prior work of *epochal signatures* by Hülsing and Weber [25] considers a security notion that crucially relies on various participants having synchronized clocks. Roughly, in a epochal signature scheme, (real) time is divided into discrete *epochs* where a new key is generated at the start of every epoch. Signatures are associated with the epoch they were generated in, where unforgeability requirements state that no adversary can forge signatures for an epoch during the epoch. As we show in §3, the security definitions for epochal signatures are *fragile*: there exist epochal signature schemes that are secure under the given definitions and yet become *completely insecure when clocks are even slightly out of sync*. This problem stems from the fact that the unforgeability notion proposed for the primitive puts strict time limits on the adversary while it queries a signing oracle. We show that if enforcement of these query restrictions is violated (even slightly) by a real-world signing oracle at epoch e , an epochal signature scheme can become catastrophically insecure *for all future epochs*.

Unfortunately, avoiding such outcomes is not easy, and in this work we seek to strengthen our security definitions to avoid such issues. We do this in two ways: unlike [25], our definitions model the unforgeability period computationally – through the widely-adopted technique of bounding the number of sequential computation steps the adversary may compute [35, 7, 41, 34, 17]. While this still requires conversion when used in the real world, it does not embed the conversion into the security definition. Much more importantly, our definition allows the adversary to participate in a “pre-processing” phase to ensure the robustness of our notion in scenarios where there may be clock synchronization issues. During this phase, the adversary is given free reign (within only a polynomial time bound) to query the signing oracle and forge signatures. This phase significantly loosens the restrictions on the adversary, allowing them to query for signatures and run the AltSign algorithm (or any other process) as many times as they wish. Once the pre-processing phase is complete, the adversary then enters a second *forgery* phase in which their runtime is more strictly bounded. Our sole restriction is that the forgery produced in the second phase must be computed on a timestamp t^* that is greater than any timestamp queried during the pre-processing phase.

³Indeed, we show that the need for AltSign to use an existing signature (or portion thereof) to produce a forgery is seemingly inherent if we do not want to use secure infrastructure. We elaborate on this point in Appendix G.

Our Definition. We are now ready to provide an (informal) definition of time-deniable signatures. We refer the reader to the technical sections for more details.

The protocol is parameterized by Δ , the duration of the *unforgeability period*, and described by the algorithms **KeyGen**, **Sign**, **AltSign** and **Verify**. The **KeyGen** and **Verify** algorithms are the same as standard signature schemes while the **Sign** algorithm, also similar to the standard notion, takes in as input a message m and time stamp t to generate a signature on (m, t) . The main new component is the algorithm **AltSign** which takes as input a message m' , time stamp t' , signature $\sigma_{(m,t)}$ such that $t' \leq t$, and uses the verification key to generate a signature $\sigma_{(m',t')}$.

For correctness of the scheme, we require that **AltSign** generates a verifying signature as long as its given as input the output of the **Sign** algorithm, or (repeated applications) of the **AltSign** algorithm. We now provide an overview of the two key security properties required by our notion.

Unforgeability. This property captures the notion that no adversary capable of computing fewer than Δ sequential steps can generate a forgery. Specifically, we allow an initial *pre-processing stage* for the adversary where it is *not* bounded by the number of sequential steps, gathering as much information as it can. At the end of this stage, say at timestamp t^* , it passes along any information onto the next stage where the adversary that runs in at most Δ sequential steps needs to produce a signature for a message with time stamp $> t^*$.

Deniability. This property asks an adversary to distinguish between a “fresh” signature generated using **Sign**, and a signature generated using **AltSign**. We formalize this by defining two experiments, where the adversary is allowed to specify a tuple $(m_1, t_1, \sigma_1 = \text{Sign}(m_1, t_1), m_2, t_2)$ with $t_2 \leq t_1$. In the first world, the output is simply the signature $\sigma_2 = \text{Sign}(m_2, t_2, \text{sk})$, whereas in the second world the output is $\sigma_2 = \text{AltSign}(m_2, t_2, \sigma_1, \text{vk})$. We say a TDS is deniable if no computationally bounded adversary can distinguish the two with a significant probability.

We refer to the above description of deniability to be “1-hop-deniable”, i.e. a signature generated via **Sign** is indistinguishable from one generated via **AltSign**. In the technical section, we extend this notion to “ k -hop-deniability”, which intuitively corresponds to the indistinguishability between a signature generated via **Sign** and one generated via k applications of **AltSign**.

2.2 Construction

Time-Deniable Signatures from Delegatable Functional Signatures. Our construction centers around the following natural idea: with each signature produced by the signer, we leak a *restricted signing oracle* that can be used to forge later signatures. A signing oracle, as the name suggests, allows a party with access to the aforementioned oracle to sign any message of its choice. For instance, the signing key can be viewed as an oracle since it allows one to sign any message of their choice. A restricted signing oracle limits the messages that can be signed. Thus, continuing with our analogy of signing keys corresponding to an oracle, a restricted signing oracle corresponds to a signing key that is restricted in a fine-grained manner.

When the **Sign** algorithm generates a signature on message m and time stamp t , it also reveals a restricted signing key sk_t that can be used sign any message m' with time stamp $t' \leq t$. Such a key can then be used by the **AltSign** algorithm to create forgeries. Revealing the restricted key with the signature, however, allows anyone in possession of the signature to create forgeries *during* the unforgeability period. To prevent this, we need to *hide* this restricted signing key until after the unforgeability period, and we do so using *time-lock puzzles* [35]. Intuitively, a time-lock puzzle allows one to “lock” a secret s for a predetermined amount of time (i.e., time parameter). Thus,

the output of the Sign algorithm will consist of the signature $\sigma_{m||t}$ along with the time-lock puzzle containing the secret sk_t , computed with time parameter Δ . We note that a similar approach has been considered in constructing notions such as epochal signatures [25], and we refer the reader to Section 3 for a more detailed comparison.

To implement restricted signing keys, we turn to the notion of *functional signatures* (FS) [10, 6, 5]. Functional signatures are equipped with *functional keys* sk_f (instead of “regular” signing keys) such that it allows one to sign $f(m)$ for any message m . We consider the following specific function for our application:

$$f_T(t, m) = \left\{ \begin{array}{ll} t||m & t \leq T \\ \perp & \text{otherwise} \end{array} \right\} \quad (1)$$

We call such functions *prefix functions* (the function prepends the time stamp to the message). It is evident from the above description that with a functional key sk_{f_T} one can generate a signature for any message m and time stamp t as long as $t \leq T$.⁴

For our TDS construction, we leverage specific properties of the functional signature scheme. We provide a more general (and detailed) definition in the technical sections, but for the purposes of the overview, we shall discuss the relevant properties of functional signatures for the specific function f_T described above: (i) *delegatability*: given a key sk_{f_T} for function f_T , using only public parameters, one can derive a key $\text{sk}_{f_{T'}}$ for a function $f_{T'}$ if $T' \leq T$; (ii) *key indistinguishability*: it should be computationally infeasible to differentiate between a fresh key sk_{f_T} and a key derived ; and (iii) *unforgeability*: it should be computationally infeasible to generate signatures $\sigma_{m||t}$ unless one has a key sk_{f_T} where $T \geq t$. While delegatability has previously been studied for functional signatures, the notion of key indistinguishability is new to our work. The latter is crucial towards achieving deniability.

Putting things together, we have:

Sign On input message m and time stamp t , the Sign algorithm generates the key sk_{f_t} (using the *master secret key*, see technical section for details), and uses it to compute the signature $\sigma_{t||m}$. It then encrypts the key sk_{f_t} within a time-lock puzzle with time parameter Δ .

AltSign On input message m' , time stamp t' , and signature $\sigma_{t||m}||\text{TimeLock}(\text{sk}_{f_t})$, the AltSign algorithm first solves the time-lock puzzle to obtain sk_{f_t} . It next uses the delegation functionality to derive a key $\text{sk}_{f_{t'}}$ from sk_{f_t} and then follows the description of the Sign algorithm.

A potentially useful property of the above approach is that the *sequential* part of the computation performed by AltSign, namely, solving the time-lock puzzle, can be *reused for computing many forgeries in parallel*. This is because once the restricted signing key is obtained – a one-time work, it can be used to compute signatures in parallel.

Intuitively, we prove unforgeability by leveraging the unforgeability of the functional signature scheme and the security of time-lock puzzle, while deniability follows from the key indistinguishability property of the functional signature scheme.

Prefix Function FS from Hierarchical Identity Based Encryption. We construct functional signatures for prefix functions from Hierarchical Identity Based Encryption (HIBE). At a high level, HIBE is an encryption scheme that allows one to encrypt to identities, (treated as bit strings in this work) such that only someone in possession of the secret key corresponding to the identity can

⁴Note that while we have thus far described signatures on messages of the form $m||t$, the above description of f_T flips it to be $t||m$. Looking ahead, the change is due to our construction of functional signatures.

decrypt messages. The hierarchical nature of the scheme allows for the delegation of keys, i.e. if one is in possession of a key for identity \mathcal{I} which is a prefix of an identity \mathcal{I}' , one can derive the key for \mathcal{I}' from the key for \mathcal{I} . The identities in our setting will correspond to the nodes of a binary tree with nodes labeled by binary strings corresponding to their path from the root (left is 0, right is 1).

HIBE schemes can be used generically to construct a signature scheme [8] - to sign a message m , use the HIBE scheme to generate a key for the “identity” m with the key corresponding to the signature. The verification of the signature is performed by encrypting a random message to the message (treated as the identity), and using the signature as a key to check whether the decryption is correct.

In our setting, the identities will be the bit strings corresponding to $t||m$. Structuring as above has the following benefit - if one were in possession of a HIBE key for a time stamp t , then one can derive keys for $t||m$ for *any* message m since $t||m$ is “lower” in the hierarchy from t . Therefore to sign a message m at time stamp *exactly* t it suffices to possess the key for t , which serves as the signing key. But recall from the description of f_t in the prior section, the signing key corresponding to f_t should allow one to sign messages for *any* time stamp smaller than t . A naive way would be to generate the signing key for f_t would be to concatenate the HIBE keys for all $t' \leq t$, but this approach is clearly infeasible since the signing key would grow linearly with the total number of possible time stamps.

To overcome this efficiency barrier, we leverage the tree structure of the HIBE scheme with the following insight - it suffices to have a small number of keys as long as we are able to derive keys for any $t' \leq t$. At a high level, the signing key sk_{f_t} will consist of keys for all identities that are the *left siblings* of the nodes along the path from $t + 1$ to the root⁵, resulting in at most $\log(t)$ many keys. A detailed description is provided in the technical sections, but here we provide an illustrative example.




Figure 1: This tree represents the hierarchical nature of the identities in our HIBE scheme. Each node in the tree besides the root represents a HIBE secret key sk_{id} for the identity id . Therefore, by our description $\text{Trace}(\text{root} = \emptyset, 110)$ will constitute the nodes which represent the secret key for f_{110} , i.e. $sk_{f_{110}} = (sk_0, sk_{10}, sk_{110})$. Using this set, all messages with prefixes in the green nodes can be signed.

In the HIBE identity tree of Figure 1, the key corresponding to f_{10} is both sk_0 and sk_{10} . To

⁵One can also view it as the nodes in the stack during the depth first traversal of the (identity) binary tree when node $t + 1$ is visited.

derive a key for f_{00} , one executes the HIBE’s delegate algorithm using sk_0 to create the key sk_{00} . In fact, to derive a key for $sk_{f_{t'}}$ from sk_{f_t} for *any* $t' \leq t$ one can simply use the HIBE delegation algorithm, i.e. there is no need to run the key generation algorithm afresh.

Looking ahead, we want to allow the adversary to choose the message it wants to compute a forgery on *after* it has seen other signatures, we require the HIBE scheme to be *adaptively* secure (i.e. the adversary can choose the identity of the HIBE scheme it wants to break *after* seeing keys for other identities). HIBE schemes satisfying the necessary requirements can be instantiated e.g., assuming the Decisional Linear (DLIN) assumption on Bilinear groups of prime order [27].

3 Related Work

Concurrent work. A concurrent and independent work of Arun et al. [4] also studies a notion similar to time-deniable signatures. Similar to our work, they make use of sequentially-ordered computation as a means to enforce time delay during which signatures are unforgeable, but become forgeable afterwards. However their work considers a different models than ours. Specifically, their system relies on the use of *unpredictable beacons* that are presumably released periodically by some trusted outside source. In contrast, we do not do rely on any randomness beacons or time servers. Unlike our work, they also explore time-based deniability in proof systems.

Epochal Signatures. Our work is closely related to the prior work on epochal signatures [25]. At a very high level, epochal signatures aim to achieve deniability in a manner similar to ours - by leaking a constrained key. In epochal signatures, (real) time is partitioned into discrete *epochs* with a key update mechanism at the start of every epoch. Any signature generated during epoch i additionally include the keys for prior epochs, allowing for forgery of signatures of any epoch $< i$ (but not epoch i).⁶ The constructed epochal signature in [25] leaks only a *single* key with the property that from a key of epoch i , k_i , one can retrieve the key of epoch $i - j$, k_{i-j} with j applications of a “key retrieval” function, but security requires that it is impossible to retrieve keys for epochs $> i$ from k_i .

In the following, we describe some key differences between the two works.

Bounded vs Unbounded Use. Unlike our work, the system proposed in [25] is limited to be *bounded use*. Specifically, the number of epochs that their system can support is bounded ahead of time. This is an outcome of the running time of their system setup, which is linear in the number of epochs.

In practice, the granularity of each epoch and the number of epochs must be fixed *before* their system is initialized, and once the total number of epoch surpass the bound, the entire system needs to be reset from scratch, a scenario that is clearly undesirable. For epochs that are sufficiently small, such a restriction would *limit the total number of signatures* that the system can support. Indeed, for similar reasons, the broad question of *bounded vs unbounded use* has been studied in various contexts in cryptography such as bounded vs unbounded query chosen-ciphertext secure encryption [15], depth-bounded vs depth-unbounded hierarchical identity-based encryption [28] and homomorphic encryption [20], bounded-collusion vs unbounded collusion in functional encryption [36, 23], and more. In all of these cases, there are significant challenges and overheads (in terms

⁶In their work, they consider an additional *deniability* parameter V such that signatures for epoch i include keys for epochs $i - V$ and earlier allowing for V epochs where the signature is valid. But for the purposes of this discussion we describe it in the above simplified manner.

of assumptions, efficiency, etc) in going from bounded system to an unbounded one. As such, we view our improvements over [25] in this regard to be a significant one.

Need for Clock Synchronization. As discussed earlier, the unforgeability notion in [25] requires the participants to have perfectly synchronized clocks. We now demonstrate that if such a requirement is not met, then the consequences can be catastrophic and result in a compromise of security for all future epochs. Specifically, we construct a secure epochal signature scheme where the unforgeability property can be broken when the clocks are slightly out of sync. We also show that the same scheme – translated to the setting of time-deniable signatures – is *not* secure as per our definition, thus demonstrating that the latter is a strictly stronger notion.

In the following, we give an over-simplified presentation of our counter-example to convey the general idea. The full counter-example is more involved (due to technical reasons) and is presented in Appendix H.1.

Intuitively, we exploit the restricted signing oracle in the unforgeability definition of epochal signatures which prevents an adversary from receiving signatures in any epoch e outside of a fixed real time window of size Δt . Our epochal signature scheme makes use of a special *trigger* message m_e^* which differs per epoch. If the adversary queries for a signature on message m_e^* in epoch e , then they receive some “secret information” from the signing oracle which can be used to recover the signing key. If the message space is large enough and m_e^* is chosen uniformly at random, this modification would not make our scheme insecure, as an adversary would only have a negligible chance of guessing m_e^* . We therefore modify the signing oracle so that, in addition to handing out signatures on messages m for epoch e , it time-lock puzzle encrypts m_e^* with difficulty parameter Δ'_t where $\Delta t < \Delta'_t < \Delta t + \varepsilon$. The difficulty parameter of the time-lock puzzle ensures that the puzzle *cannot* be decrypted within the epoch that it is generated, but can be decrypted just after the epoch concludes. Thus, if there is a clock synchronization issue where the challenger’s (the entity generating the signatures) clock is slightly slower, then an adversary can decrypt to obtain m_e^* and query the signing oracle on m_e^* to obtain the “secret information”. In our actual counter-example, this secret information cannot directly be equal to the secret key because the ES scheme must be perfectly deniable even to someone that holds the original signing key. To deal with this, we instead encrypt the signing key with a one time pad that is 2 out of 2 additively secret shared. Querying on different trigger messages reveals different shares of the key. Further details can be found in Appendix H.1.

To argue that this scheme is a secure epochal signature scheme when the clocks are synchronized, we note that in an epochal signature scheme, at the start of an epoch $e + 1$ two things happen: (i) key evolution procedure is applied to the secret signing key to generate the signing key for the next epoch; and (ii) public information pinfo_e is broadcast. Here, pinfo_e allows anyone to produce signatures for epochs $\leq e$ without the signing key such that they are indistinguishable from signatures produced by the real signing key (akin to our definition of deniability). In the above scheme, while secret key material is used to key the signing key, this is not revealed as a part of pinfo_e and does not need to be to create indistinguishable signatures (every field of the signature will be simulatable). Thus, simply having pinfo_e will not allow recovery of sk .

We now argue that the above scheme is *not* a secure time-deniable scheme. Briefly, this is due to the pre-processing phase we allow during the unforgeability definition. In this phase, the adversary can query the same time stamp multiple times (here roughly the time-stamps correspond to an epoch), and therefore can perform the attack described above by decrypting the time-lock puzzles, making the relevant queries, and using the results to obtain the signing key. The key is then passed

	ES [25]	Short-Lived Sigs [4]	KeyForge B [38]	Us
Forging Assumptions	None	RB	Publish Keys	None
Max number epochs	Bounded	Unbounded	Unbounded	Unbounded
Forgers all derive same key	Yes	Yes	Yes	No
Building Blocks	OWF + TLP	tVDF + RB*	BDH	BDH + TLP

Table 1: Comparison of as subset of existing constructions that provide a notion of deniability for signatures. tVDF stands for a trapdoor VDF, while RB is a random beacon.

on to the “online adversary” who uses it to produce a forged signature.

We remark that, again, the above description is oversimplified and the full counter-example is presented in H.1.

We briefly summarize some of the different properties of proposed constructions for expiring signature schemes in Table 1.

4 Preliminaries

Throughout the paper we consider the depth $\text{depth}(C)$ of a circuit C to be defined as the longest path in the circuit from input wires to output wire. The size of a circuit $\text{size}(C)$ corresponds to the number of gates.

Sequential time. In this work sequential time refers to the non-parallelizable time it would take any circuit to compute a particular function. More formally, a function has sequential time d or takes d sequential steps if there exists a circuit C which correctly computes that function and the depth $\text{depth}(C) = d$. This notion attempts to capture inherent limitations in computing a function that cannot be overcome by access to more cores or processors.

Time-lock Puzzles. The concept of a time lock puzzle or time lock encryption was first introduced by Rivest, Shamir, and Wagner [35]. In this definition, a solution is provided as *input* to an algorithm which produces puzzles as output. Security requires that it is hard to figure out what solution was hidden inside the puzzle unless someone is willing to spend sequential time proportionate to a function of t , if t is a time parameter for our scheme. We use the circuit-based adversarial definition which translates to bounding the parallel time of the adversary via bounding the adversary’s depth.

Definition 4.1. A puzzle `TimeLock` is a tuple of algorithms `Gen, Sol` where the signature of the algorithms is defined as below.

`Gen`($1^\lambda, \Delta, s$) $\rightarrow Z$: On input a time/difficulty parameter Δ and a solution $s \in \{0, 1\}^\lambda$, output a puzzle Z

`Sol`(Δ, Z) $\rightarrow s$: This is a deterministic algorithm that when given a puzzle Z and the difficulty parameter Δ produces a solution s .

Correctness. Correctness requires that for all solutions $s \in \{0, 1\}^\lambda$ and difficulty parameters Δ the following holds:

$$\Pr \left[Z \leftarrow \text{Gen}(1^\lambda, \Delta, s) : s \leftarrow \text{Sol}(\Delta, Z) \right] = 1$$

Efficiency. \exists a polynomial p s.t. $\forall \Delta, \lambda \in \mathbb{N}$, $\text{Sol}(\Delta, \cdot)$ runs in time $\Delta \cdot p(\lambda)$

Security. We consider a time lock puzzle to be α -gap secure if \forall functions $T(\lambda) \geq \alpha(\lambda)$ and distinguishers $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ of size $\text{size}(\mathcal{A}_\lambda) \in \text{poly}(\lambda)$ and depth $\text{depth}(\mathcal{A}_\lambda) \leq \frac{T(\lambda)}{\alpha(\lambda)}$, \exists a negligible function μ s.t. $\forall \lambda \in \mathbb{N}$, $\forall s_0, s_1 \in \{0, 1\}^\lambda$,

$$\left| \Pr \left[Z \leftarrow \text{Gen}(1^\lambda, T(\lambda), s_0) : \mathcal{A}_\lambda(Z) = 1 \right] - \Pr \left[Z \leftarrow \text{Gen}(1^\lambda, T(\lambda), s_1) : \mathcal{A}_\lambda(Z) = 1 \right] \right| \leq \mu(\lambda) \quad (2)$$

This is a variation of the time lock puzzle definition of [16], where we define security to hold for adversaries of polynomial size instead of super polynomial.

4.1 Hierarchical Identity Based Encryption

We recall the notion of Hierarchical Identity Based Encryption (HIBE). A HIBE scheme has the following five algorithms:

Setup $(1^\lambda) \rightarrow (msk, pk)$: The setup algorithm generates the master secret key and public parameters.

KeyGen $(msk, I) \rightarrow sk_I$ Generates a key for the identity I using the master secret key msk .

Delegate $(pk, sk_I, I') \rightarrow sk_{I||I'}$: Takes a secret key of some identity I and generates a secret key for the identity $I||I'$.

Encrypt $(pk, m, I) \rightarrow ct$ The encryption algorithm takes the public key, a message and an identity I and outputs the corresponding ciphertext.

Decrypt $(sk_{I'}, ct) \rightarrow m/\perp$: The decryption algorithm takes a secret key and a message and outputs the message if the secret key hierarchy level allows decryption of the ciphertext.

Remark. Throughout this paper, we will make use of HIBE schemes where **Delegate** can take in a child identity I' that is the empty string. In such schemes, $sk'_{I'} \leftarrow \text{Delegate}(pk, sk_I, nil)$ is a re-randomization of the key sk_I for identity I . We note that many HIBE schemes can be modified to have this property [28, 27, 21].

4.1.1 Security and Correctness

We consider HIBE security similarly to the work of Lewko and Waters[28, 29] using the following security game played by a challenger and an adversary.

- **Setup** The challenger runs $(pk, msk) \leftarrow \text{Setup}(1^\lambda)$ and gives the public parameters pk to the adversary. Let set \mathcal{S} be the set of private keys that the challenger creates. At the beginning, $\mathcal{S} = \emptyset$.
- **Phase 1** In this phase, the adversary gets to make three types of queries.
 1. Create queries $\mathcal{QC}(I)$, which are made on some specific identity I . The challenger adds the keys for this identity to the set \mathcal{S} . Note that the adversary does not get these keys.

2. Delegate queries $\mathcal{QD}(I)$, which are made on some identity I such that the corresponding keys are in the set \mathcal{S} . The challenger adds the keys corresponding to the delegated identity I' and adds them to the set \mathcal{S} .
 3. Reveal queries $\mathcal{QR}(I)$, which are also made on some identity I such that the corresponding keys are in the set \mathcal{S} . In response, the challenger gives the corresponding keys to the adversary and removes them from the set \mathcal{S} .
- **Challenge** The adversary gives the challenger messages m_0 and m_1 and a challenge identity I^* . The challenger responds with a random $\beta \in \{0, 1\}$ and encrypts m_β under I^* and sends the ciphertext to the adversary.
 - **Phase 2** The adversary gets to query the challenger similar to Phase 1.
 - **Guess** The adversary outputs its guess β' for β and wins if the following conditions are satisfied:
 1. $\beta' = \beta$.
 2. The challenge identity I^* should satisfy the property that no revealed keys, in either of the query phases, belong to an identity that was a parent of I^* and the I^* 's keys shouldn't have been revealed.

The advantage of the adversary \mathcal{A} is defined as $\mathbf{Adv}_{\mathcal{A}}(1^\lambda) = \Pr[\beta' = \beta] - \frac{1}{2}$.

Definition 4.2 (Adaptive security for HIBE). A HIBE scheme is adaptively-secure if \forall poly size adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ their advantage $\mathbf{Adv}_{\mathcal{A}}(1^\lambda)$ in the HIBE security game defined above is negligible.

4.1.2 Key-indistinguishability

We additionally require that all polynomial-time adversaries have at most a negligible advantage in distinguishing between keys generated via the **KeyGen** algorithm and keys generated via the **Delegate** algorithm even when given access to the master secret key msk . We define this property using the following HIBE key-indistinguishability game $\mathbf{Exp}_{\text{HIBE}}^{\text{IND}}$.

The **Setup** phase is similar to the HIBE security game, except the adversary also gets the master secret key msk . Similarly, the set \mathcal{S} of keys queried and the corresponding query identifier is set to be empty. A bit $\beta \xleftarrow{\$} \{0, 1\}$ is sampled uniformly.

Query phase. In this phase the adversary is allowed to adaptively query a key oracle $\mathcal{QK}(\cdot)$ and a challenge oracle $\mathcal{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$

$\mathcal{QK}(\cdot)$ takes as input an identity I , computes $sk_I \leftarrow \text{KeyGen}(msk, I)$, selects an identifier id and adds (id, I, sk_I) to the set \mathcal{S} and responds with (id, sk_I) .

$\mathcal{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$ takes as input a challenge identifier id , and a pair of identities (I_0, I_1) such that I_0 is a *parent identity* of I_1 . Where parent identity implies that on the hierarchical identity tree (Figure 1) where the root is msk , I_0 is an intermediate node on the shortest path from I_1 to the root. It checks for id in set \mathcal{S} and checks that id corresponds to a key query on I_0 . If no such id is found, the output is \perp . Otherwise, compute $sk_0 \leftarrow \text{KeyGen}(msk, I_0)$, $sk_1 \leftarrow \text{Delegate}(pk, sk_{I_0}, I_1)$ and respond with sk_β .

Guess. The adversary outputs its guess β' for β and wins if $\beta' = \beta$.

The advantage of the adversary \mathcal{A} is defined as $\mathbf{Adv}_{\mathcal{A}}(1^\lambda) = \Pr[\beta' = \beta] - \frac{1}{2}$.

Definition 4.3 (Key-Indistinguishability for HIBE). A HIBE scheme is delegated key indistinguishable if \forall poly size adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ their advantage $\mathbf{Adv}_{\mathcal{A}}(1^\lambda)$ in the $\mathbf{Exp}_{\text{HIBE}}^{\text{IND}}$ game is negligible.

The key indistinguishability property can be easily satisfied by many existing HIBE schemes, provided the sub-key components from earlier levels of the HIBE can be re-randomized. Randomization techniques like these have been used to construct anonymous HIBEs in the past [37]. In Appendix D we show that the prime order variant of the Lewko-Waters HIBE scheme [27] satisfies this property.

5 Time-Deniable Signatures: Definition

A time-deniable signature scheme

$\text{DS} = (\text{KeyGen}, \text{Sign}, \text{Verify}, \text{AltSign})$ is a tuple of possibly probabilistic polynomial-time algorithms:

$\text{KeyGen}(1^\lambda, T = T(\lambda)) \rightarrow (vk, sk)$: On input the security parameter 1^λ and a difficulty parameter for AltSign called T , this randomized algorithm outputs the verification key vk and the signing key sk .

$\text{Sign}(sk, m, t) \rightarrow \sigma$: On input a message m and the signing key sk , this randomized algorithm outputs a signature σ on m for timestamp t .

$\text{Verify}(vk, \sigma, m, t) \rightarrow \{0, 1\}$: On input a signature σ , a message m , verification key vk and timestamp t , this deterministic algorithm outputs a bit.

$\text{AltSign}(vk, (m^*, t^*, \sigma^*), m, t) \rightarrow \sigma$: On input a valid message and signature pair (m^*, σ^*) for timestamp t^* , this randomized algorithm outputs a signature σ on message m for timestamp t .

Definition 5.1 (Efficiency). The algorithms $\text{KeyGen}, \text{Sign}, \text{Verify}$ must run in time poly in the size of the input. For AltSign it is required that there exist a positive polynomial q such that $\forall T = T(\lambda), \forall \lambda \in \mathbb{N}$, AltSign is computable in time $q(\lambda) \cdot T$ where T is the difficulty parameter provided to KeyGen .

Definition 5.2 (Correctness). A time-deniable signature scheme for a message space \mathcal{M} satisfies the correctness property if it satisfies the following two conditions:

1. $\forall m \in \mathcal{M}, (vk, sk) \leftarrow \text{KeyGen}(1^\lambda), \sigma \leftarrow \text{Sign}(sk, m, t)$, it holds that $\text{Verify}(vk, \sigma, m, t) = 1$.
2. Let $\text{AltSign}^k(vk, (m_0, t_0, \sigma_0), \{(m_j, t_j)\}_{j \in [k]})$ be shorthand for the following recursively defined function:

$$\begin{aligned} \text{AltSign}^i(vk, (m_0, t_0, \sigma_0), \{(m_j, t_j)\}_{j \in [i]}) = \\ \text{AltSign}(vk, (m_{i-1}, t_{i-1}, \text{AltSign}^{i-1}(vk, (m_0, t_0, \sigma_0), \{(m_j, t_j)\}_{j \in [i-1]}), m_i, t_i) \end{aligned}$$

where $\text{AltSign}^0(vk, (m_0, t_0), \sigma_0, \{\}) = \sigma_0$. In words, AltSign^k is a signature obtained by applying AltSign k times to a provided signature σ_0 on the message m_0, t_0 . Then we have the following additional correctness property:

$\forall k \in \mathbb{N}$, for all sets of ordered tuples $\{(m_j, t_j)\}_{j \in [k]}$, and $\forall m_0, t_0$ that satisfy $t_{j-1} \geq t_j$ where $j \in [k]$:

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{KeyGen}(1^\lambda, T); \\ \sigma_0 \leftarrow \text{Sign}(sk, m_0, t_0); \\ \sigma \leftarrow \text{AltSign}^k(vk, (m_0, t_0, \sigma_0), \{(m_j, t_j)\}_{j \in [k]}) \end{array} : \text{Verify}(vk, \sigma, m_k, t_k) = 1 \right] = 1 \quad (3)$$

Remark. Property 2 assumes that signatures and “forged” signatures used as input to the AltSign algorithm are computed honestly. One can also consider a stronger notion of correctness, where correctness of AltSign holds even on input signatures (and “forged” signatures) that may not be honestly computed, but nevertheless can be validated by the Verify algorithm. We refer to this as *robust correctness*.

In this work, we focus on the simpler notion and leave the discovery of schemes that satisfy robust correctness to future work.

Definition 5.3 (Robust Correctness). A time-deniable signature scheme satisfies *robust correctness* if:

$\forall (vk, sk) \leftarrow \text{KeyGen}(1^\lambda, T), \forall (m_0, t_0, \sigma_0)$ such that

$\text{Verify}(vk, \sigma_0, m_0, t_0) = 1$, $\forall k \in \mathbb{N}$, for all sets of ordered tuples $\{(m_j, t_j)\}_{j \in [k]}$ where $\forall j \in [k]$, $t_{j-1} \geq t_j$:

$$\Pr \left[\sigma \leftarrow \text{AltSign}^k(vk, (m_0, t_0, \sigma_0), \{(m_j, t_j)\}_{j \in [k]}) : \text{Verify}(vk, \sigma, m_k, t_k) = 1 \right] = 1 \quad (4)$$

5.1 Security Property: (ϵ, T) -Unforgeability

Our unforgeability notion requires that signatures should remain unforgeable within a restricted time window. We capture this via a security game below:

Setup. The challenger generates $(vk, sk) \leftarrow \text{KeyGen}(1^\lambda, T)$ and gives the verification key vk to the adversary.

Phase 1. The adversary is a tuple of two algorithms, \mathcal{A}_0 and \mathcal{A}_1 . In this phase, \mathcal{A}_0 is allowed to adaptively query a signing oracle O_{Sign} which is defined as follows. On input a message m and a timestamp t , the signing oracle $\text{O}_{\text{Sign}}(sk, \cdot, \cdot)$ returns the signature $\sigma \leftarrow \text{Sign}(sk, m, t)$.

Transfer. The adversary \mathcal{A}_0 gives an advice string z to adversary \mathcal{A}_1 .

Phase 2. The adversary \mathcal{A}_1 has to respond to the challenger with a forgery while also being allowed to adaptively query the oracle O_{Sign} .

Forgery. The adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ wins if in the end \mathcal{A}_1 can produce a valid forgery (m^*, t^*, σ^*) under the following constraints:

1. $\forall i, t^* > t_i$ for queries (m_i, t_i) made by \mathcal{A}_0
2. $\forall i, (m_i, t_i) \neq (m^*, t^*)$ where (m_i, t_i) are queries made to O_{Sign} by \mathcal{A}_1

An adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ is considered an ϵ -admissible adversary if it satisfies the above conditions and $\exists T(\lambda)$ where $\forall \lambda \in \mathbb{N}, \epsilon(\lambda) \leq T(\lambda)$, $\text{depth}(\mathcal{A}_1) \leq \frac{T(\lambda)}{\epsilon(\lambda)}$, and $\text{size}(\mathcal{A}) \in \text{poly}(\lambda)$. Note that the depth of \mathcal{A}_0 is allowed to be polynomial in the security parameter whereas the depth of \mathcal{A}_1 is more strictly bounded.

Definition 5.4 ((ϵ, T) -Unforgeability). A time-deniable signature scheme satisfies the ϵ -Unforgeability property if $\forall \epsilon$ -admissible polysize adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}} = \{(\mathcal{A}_{\lambda,0}, \mathcal{A}_{\lambda,1})\}_{\lambda \in \mathbb{N}}$, $\forall T(\lambda)$ satisfying the ϵ -admissibility requirement for \mathcal{A} , there exist a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{KeyGen}(1^\lambda, T(\lambda)), \\ (z) \leftarrow \mathcal{A}_{\lambda,0}^{\text{O}_{\text{Sign}}(sk, \cdot, \cdot)}(vk), \\ (m^*, t^*, \sigma^*) \leftarrow \mathcal{A}_{\lambda,1}^{\text{O}_{\text{Sign}}(sk, \cdot, \cdot)}(z) \end{array} : \text{Verify}(vk, m^*, t^*, \sigma^*) = 1 \right] \leq \mu(\lambda) \quad (5)$$

Remark. The prior definition is more akin to EUF-CMA unforgeability than strong unforgeability since we disallow adversaries from submitting forgeries on message time-stamp pairs they have received in the past. We leave the question of achieving strong unforgeability to future work.

5.2 Deniability

Deniability in our scheme comes from the fact that after T sequential time-steps, anyone can forge a valid signature under the verification key of the original signer. Consequently, a time-deniable signature scheme should ensure the indistinguishability of signatures generated via the `Sign` and the `AltSign` algorithms. Otherwise the original signer could not deny that it signed a message at a particular time. We present below a security game to capture this idea. Our notion would be meaningful even if the adversary did not have access to the signing key, but we give them it as well in order to capture more powerful attacker models.

We now define the security game $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$:

Setup. The challenger generates $(vk, sk) \leftarrow \text{KeyGen}(1^\lambda, T)$ and gives both the verification key vk and the signing key sk to the adversary \mathcal{A} . They also initialize an empty table \mathcal{T} and sample $\beta \xleftarrow{\$} \{0, 1\}$.

Query Phase. In this phase, the adversary is allowed to adaptively query a signing oracle $\text{O}_{\text{Sign}}(sk, \cdot, \cdot)$ and a challenge oracle $\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$.

$\text{O}_{\text{Sign}}(sk, \cdot, \cdot)$ takes as input a message m and a timestamp t to produce $\sigma \leftarrow \text{Sign}(sk, m, t)$.

It randomly chooses a new identifier id (not equal to any previously defined identifiers), records $(\text{id}, m, t, \sigma)$ in \mathcal{T} , and returns (id, σ)

$\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$ takes as input a tuple of identifier, challenge message, and time-stamp (id, m, t) .

It checks \mathcal{T} for id . If it is not present the output is \perp . Let m', t', σ' be the values associated with id . If $t > t'$ the output is also \perp . Compute $\sigma^0 \leftarrow \text{Sign}(sk, m, t)$ and $\sigma^1 \leftarrow \text{AltSign}(vk, (m', t', \sigma'), m, t)$. Finally it responds with σ^β .

Guess. The adversary outputs its guess β' for β . The advantage of the adversary \mathcal{A} is defined as $\mathbf{Adv}_{\mathcal{A}}(1^\lambda) = \Pr[\beta' = \beta] - \frac{1}{2}$. The adversary \mathcal{A} wins if $\beta' = \beta$.

Definition 5.5 (Deniability). A signature scheme is considered to possess the deniability property if \forall poly size adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ their advantage $\mathbf{Adv}_{\mathcal{A}}(1^\lambda)$ in the $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$ game is negligible.

k -hop Deniability: We refer to the definition defined above as 1-hop deniability. It is reasonable to ask if indistinguishability still holds when comparing the output of `Sign` with applying the `AltSign` algorithm k times instead of just once. Intuitively, this notion could be stronger and offer more deniability via a larger pool of indistinguishable forgeries.

A formal definition of the k -hop indistinguishability game $\mathbf{Exp}_{\text{DS}}^{\text{khop}}$ is given below:

Setup. This is the same as the (1-hop) key-indistinguishability game $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$.

Query Phase. The adversary \mathcal{A} has access to two oracles $\text{O}_{\text{Sign}}(\text{sk}, \cdot, \cdot)$ and $\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot, \cdot)$. O_{Sign} is the same oracle as given in $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$.

$\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot, \cdot)$ takes as input a challenge identifier id , one ordered set of messages and time-stamp tuples $\{(m_i, t_i)\}_{i \in [k-1]}$, and a message, time-stamp pair m^*, t^* . It checks that there exists a row in \mathcal{T} with $(\text{id}, \cdot, \cdot, \cdot)$. Let m_0, t_0, σ_0 be the values associated with that row. It ensures that $\forall i \in [k-1], t_{i-1} \geq t_i$, and $t_{k-1} \geq t^*$. If any of these does not hold, the output is \perp . Compute:

$$\begin{aligned} \sigma^0 &\leftarrow \text{Sign}(\text{sk}, m^*, t^*) \\ \sigma^1 &\leftarrow \text{AltSign}^k(vk, (m_0, t_0, \sigma_0), \{(m_1, t_1), \dots, (m_{k-1}, t_{k-1}), (m^*, t^*)\}) \end{aligned}$$

σ^β is returned as the output.

Guess. This is again the same as the $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$ game.

Definition 5.6 (k -hop Deniability). A signature scheme achieves k -hop deniability property if \forall poly size adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ their advantage in the $\mathbf{Exp}_{\text{DS}}^{\text{khop}}$ game is negligible.

Theorem 5.1. Any time-deniable signature scheme satisfying the deniability property as defined in definition 5.5, also satisfies the k -hop deniability property as defined in definition 5.6.

For a proof of Theorem 5.1 see Appendix C.1.

6 Delegatable Functional Signatures

In this section we define and construct delegatable functional signatures and define an additional key indistinguishability property for this primitive.

Functional Signatures. We start by recalling the notion of Functional Signatures as defined by Boyle, Goldwasser, and Ivan[11].

Definition 6.1. A functional signature scheme $\text{FS} = (\text{Setup}, \text{KeyGen}, \text{Sign}, \text{Verify})$ is a tuple of potentially probabilistic, polynomial time algorithms of the following form:

$\text{Setup}(1^\lambda) \rightarrow mvk, msk$: On input the security parameter, this algorithm returns the master verification key mvk and the master signing key msk .

$\text{KeyGen}(msk, f) \rightarrow sk_f$: On input the master signing key msk and a function f , this algorithm outputs a function-specific signing key sk_f .

$\text{Sign}(sk_f, f, m) \rightarrow (f(m), \sigma)$: On input a function-specific signing key, a function f and a message m , this algorithm outputs $f(m)$ and a signature σ .

$\text{Verify}(mvk, f(m), \sigma) \rightarrow \{0, 1\}$: On input a master verification key mvk , a function $f()$ evaluated on message m and a signature, it outputs a bit.

Correctness. Correctness requires that any signature output from the Sign algorithm on a valid functional key and a message verifies correctly. More formally, for all supported functions f , for all messages m ,

$$\Pr \left[\begin{array}{l} mvk, msk \leftarrow \text{Setup}(1^\lambda); \\ sk_f \leftarrow \text{KeyGen}(msk, f); \\ m^*, \sigma \leftarrow \text{Sign}(sk_f, f, m) \end{array} : \text{Verify}(mvk, m^*, \sigma) = 1 \right] = 1$$

Security. For completeness, the unforgeability security game $\text{Exp}_{\text{FS}}^{\text{UNF}}$ between a challenger and adversary \mathcal{A} for functional signatures is provided below.

Setup. The challenger generates $(mvk, msk) \leftarrow \text{Setup}(1^\lambda)$. They also initialize an empty table \mathcal{T} . mvk is given to adversary \mathcal{A} .

Query Phase. In this phase \mathcal{A} gets access to a key oracle $\bar{\text{O}}_{\text{Key}}$ and a signing oracle $\bar{\text{O}}_{\text{Sign}}$.

1. $\bar{\text{O}}_{\text{Key}}(msk, \cdot, \cdot)$ takes as input function description f and an identifier i . The challenger checks if there is a row in \mathcal{T} corresponding to (i, f, \cdot) . If such a row exists then return the corresponding secret key sk_f^i . Otherwise generate $sk_f \leftarrow \text{KeyGen}(msk, f)$, record (i, f, sk_f) in \mathcal{T} and return sk_f .
2. $\bar{\text{O}}_{\text{Sign}}(msk, \cdot, \cdot, \cdot)$ takes as input a function description f , an identifier i , and a message m . If a row in \mathcal{T} corresponds to (i, f, \cdot) then use the secret key sk_f specified in that row. Otherwise, generate $sk_f \leftarrow \text{KeyGen}(msk, f)$ and record (i, f, sk_f) in \mathcal{T} . Let $f(m), \sigma$ be the output of $\text{Sign}(sk_f, f, m)$. Return σ to \mathcal{A} .

Challenge Phase. The adversary \mathcal{A} must return m^*, σ^* to the challenger. An adversary is considered to be *admissible* if that the following conditions are satisfied:

1. \forall values (m_i, σ_i) returned by $\bar{\text{O}}_{\text{Sign}}$, $m_i \neq m^*$
2. there is no key sk_f queried from $\bar{\text{O}}_{\text{Key}}$ such that $\exists m$ where $f(m) = m^*$

A functional signature scheme is said to be secure if for all *admissible* poly-size adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ there exists a negligible function $\mu(\lambda)$ such that

$$\Pr \left[\begin{array}{l} mvk, msk \leftarrow \text{KeyGen}(1^\lambda); \\ m^*, \sigma^* \leftarrow \mathcal{A}_\lambda^{\bar{\text{O}}_{\text{Key}}(msk, \cdot, \cdot), \bar{\text{O}}_{\text{Sign}}(msk, \cdot, \cdot, \cdot)}(mvk) \end{array} : \text{Verify}(mvk, m^*, \sigma^*) = 1 \right] \leq \mu(\lambda)$$

6.1 Key Delegation

In order to create signing keys even without the master signing key, we define an additional probabilistic polynomial time algorithm called **Delegate**. This algorithm takes as input a function f , a corresponding secret key sk_f , and a *restriction* of f , f' . The output is a secret key $sk_{f'}$ or \perp . We say that a function f' is a *restriction* of another function f if the following is true: let $f : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$, then f' has the same domain and codomain as f and $\forall x \in \mathcal{X}$ either $f'(x) = f(x)$ or $f'(x) = \perp$. This captures the ability to create a signing key that can sign some subset of the same messages as the original key. Notice that restriction here is almost a misnomer since technically f can be a restriction of itself.

FS.Delegate(mvk, f, sk_f, f') $\rightarrow sk_{f'}, \perp$: given the verification key mvk , a function f , a signing key sk_f , and another function f' output $sk_{f'}$ if f' is a restriction of f else \perp .

For a delegatable functional signature scheme the following additional correctness property must hold for all functions $f : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ supported by FS, for all restrictions f' of f , and $\forall m \in \mathcal{X}$ where $f'(m) \neq \perp$:

$$\Pr \left[\begin{array}{l} (mvk, msk) \leftarrow \text{FS.Setup}(1^\lambda); \\ sk_f \leftarrow \text{FS.KeyGen}(msk, f); \\ sk_{f'} \leftarrow \text{FS.Delegate}(mvk, f, sk_f, f'); \\ \sigma \leftarrow \text{FS.Sign}(sk_{f'}, f', m) \end{array} : \text{FS.Verify}(mvk, f'(m), \sigma) = 1 \right] = 1 \quad (6)$$

The relevance of the delegation property will be demonstrated in our construction. Furthermore, our construction will require yet another property of these delegatable functional signatures.

Key Indistinguishability We would like it to be the case that keys generated via **KeyGen** and **Delegate** appear the same to any adversary, even if they have access to the master signing key msk and can make adaptive queries. To capture this notion we define the key-indistinguishability game $\text{Exp}_{\text{FS}}^{\text{IND}}$ for delegatable functional signatures.

Setup The challenger runs $(mvk, msk) \leftarrow \text{Setup}(1^\lambda)$ and gives both the master verification key mvk and the master signing key msk to the adversary. Let \mathcal{T} be a table kept by the challenger, initialized to be empty. The challenger also samples $\beta \xleftarrow{\$} \{0, 1\}$ and keeps this value to itself.

Query Phase In this phase, the adversary gets to query two different oracles.

1. Key creation oracle $\text{O}_{\text{Key}}(\cdot)$, which can be queried on some specific function f . On input a function f , the key creation oracle checks \mathcal{T} for keys on function f . Let i be the largest value associated with a row containing f . Run $sk_f \leftarrow \text{FS.KeyGen}(msk, f)$ and record $(i + 1, f, sk_f)$ in \mathcal{T} . Output $(i + 1, sk_f)$.
2. Challenge oracle $\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$ where the first input is an identifier i and the subsequent inputs are functions f_0, f_1 and f_1 is a restriction of f_0 . The challenger checks \mathcal{T} for a row (i, f_0, \cdot) that has secret key sk_{f_0} . If no such key exists, the output is \perp . Otherwise, the oracle computes $sk_0 = \text{FS.Delegate}(mvk, f_0, sk_{f_0}, f_1)$, $sk_1 = \text{FS.KeyGen}(msk, f_1)$ and returns sk_β .

Guess The adversary outputs its guess β' for β and wins if $\beta' = \beta$.

The advantage of the adversary \mathcal{A} is defined as $\mathbf{Adv}_{\mathcal{A}}(1^\lambda) = |\Pr[\beta' = \beta] - \frac{1}{2}|$.

Definition 6.2 (Key-Indistinguishability for Delegatable FS). A delegatable functional signatures scheme is considered key-indistinguishable if \forall poly size adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ their advantage $\mathbf{Adv}_{\mathcal{A}}(1^\lambda)$ in the $\mathbf{Exp}_{\text{FS}}^{\text{IND}}$ game is negligible.

6.2 Construction for Prefix Functions

We now describe how to create delegatable functional signatures for prefixing functions from hierarchical identity-based encryption. We will be concerned with signatures on functions of the form $f_T : \{0, 1\}^\ell \times \{0, 1\}^m \rightarrow \{0, 1\}^{\ell+m}$ that concatenate their arguments. More formally, we consider functions

$$f_T(t, m) = \begin{cases} t || m & t \leq T \\ \perp & \text{otherwise} \end{cases} \quad (7)$$

For the sake of readability, in the following constructions we abuse notation and write T in place of f_T i.e. $\text{FS.Delegate}(mvk, f_y, sk_{f_y}, f_{y'})$ is replaced with $\text{FS.Delegate}(mvk, y, sk_y, y')$.

We also define the notion of *stack trace* which will be useful in the formal description of the protocol.

Definition 6.3. The stack trace of T , $\text{Trace}(r, T)$ is defined as the set of nodes on the stack when executing a depth-first search to find the leaf node $T + 1$ in a binary tree with some root r .

The stack trace can be found efficiently, and as described in the technical overview gives us the set of the $\leq \ell$ identity key nodes required to derive all keys corresponding to timestamps up to T .

The construction is presented in pseudocode in Figure 2.

Theorem 6.1. *If HIBE is adaptively secure then the functional signature scheme for prefix functions constructed in Figure 2 is unforgeable.*

For a proof of Theorem 6.1 see Appendix B.

Theorem 6.2. *If HIBE is key-indistinguishable then the scheme in Figure 2 satisfies the functional signatures key-indistinguishability property.*

For a proof of Theorem 6.2 see Appendix B.2.

7 Construction of Time-Deniable Signatures

This section describes our construction of time-deniable signatures from key indistinguishable, delegatable functional signatures for prefix functions and time lock encryption. To sign a message at time stamp t , we first use the master signing key to construct a signing key for the function f_t . This key is then used to sign the message and is included as a ciphertext, time-lock encrypted with the signature. The alternate signing algorithm decrypts the ciphertext, uses the delegate algorithm to produce an appropriate signing key for $f_{t'}$ with $t' \leq t$, and then signs and time-lock encrypts the key. For security of the scheme to hold, the parameter for the time-lock puzzle Δ cannot be precisely the same as T . The intuitive reason behind the difference is that forging

<p><u>Setup(1^λ) :</u> $(msk', pk) \leftarrow \text{HIBE.Setup}(1^\lambda)$ return $pk, (pk, msk')$</p> <p><u>KeyGen($msk = (pk, msk')$, T) :</u> $list_{sk_T} := []$ $trace \leftarrow \text{Trace}(msk, T)$ for $node \in trace$: $sk_{node} \leftarrow \text{HIBE.KeyGen}(msk', node)$ $add(list_{sk_T}, sk_{node})$ return $pk, list_{sk_T}$</p> <p><u>Sign($sk_T, T, (t, m)$) :</u> if $t > T$: return \perp $t_1 \dots t_\ell \leftarrow \text{parse}(t)$ $pk, list_{sk_T} \leftarrow \text{parse}(sk_T)$ $sk_{t'}, j \leftarrow \text{findPrefix}(list_{sk_T}, t)$ $sk_{t m} \leftarrow \text{HIBE.Delegate}(pk, sk_{t'}, t_{j+1} \dots t_\ell m)$ return $t m, sk_{t m}$</p>	<p><u>Delegate($mvk = pk, T, sk_T = (pk', list_{sk_T}), T'$) :</u> if $T' > T$: return \perp $T'_1 \dots T'_\ell \leftarrow \text{parse}(T')$ $sk, j \leftarrow \text{findPrefix}(list_{sk_T}, T')$ $list_{sk_{T'}} := []$ for $i \in 0, \dots, j-1$: $sk_i \leftarrow \text{HIBE.Delegate}(pk, list_{sk_T}[i], nil)$ $add(list_{sk_{T'}}, sk_i)$ $trace \leftarrow \text{Trace}(T'_1 \dots T'_j, T'_i)$ for $node \in trace$: $sk_{node} \leftarrow \text{HIBE.Delegate}(pk, sk, node)$ $add(list_{sk_{T'}}, sk_{node})$ return $pk, list_{sk_{T'}}$</p> <p><u>Verify($mvk = pk, m^*, \sigma = sk_{t m}$) :</u> $msg \leftarrow \{0, 1\}^n$ $c \leftarrow \text{HIBE.Encrypt}(pk, msg, sk_{t m})$ return $\text{HIBE.Decrypt}(sk_{t m}, c) = msg$ and $t m = m^*$</p>
--	--

Figure 2: The function f_T for each input message m is defined as $f_T(t, m) = \text{parse}(t)||m$ if $t \leq T$ or \perp . $\text{findPrefix}(list, id)$ takes in a list of HIBE secret keys called $list$ and an identity string id . It returns a secret key sk and an index j so that sk is a secret key for a length j prefix of id . Any bit string beginning with t_{j+1} where $j = \ell$ is the empty string. The function $\text{Trace}(root, leaf)$ is specified in definition 6.3.

involves not just breaking the time lock but also executing other algorithms. Let $|A.B|$ denote the depth of the circuit that computes algorithm B of cryptographic primitive A and $z(\lambda) = |\text{FS.Verify}| + 2 + 1 + |\text{FS.Sign}| + |\text{FS.KeyGen}| + |\text{TimeLock.Gen}|$. Our construction is described in Figure 3 and uses $z(\lambda)$ to define Δ .

Theorem 7.1. *Consider the scheme presented in figure 3. Let $\mathcal{A} = \{(\mathcal{A}_{\lambda,0}, \mathcal{A}_{\lambda,1})\}_{\lambda \in \mathbb{N}}$ be a polysize adversary in the (ϵ, T) -unforgeability game, $q(\lambda)$ the max number of queries $\mathcal{A}_{\lambda,1}$ can make, m the number of queries $\mathcal{A}_{\lambda,1}$ actually makes, \mathcal{B} an adversary against the ϵ -gap time lock puzzle scheme, and \mathcal{C} an adversary against the adaptive unforgeability of the functional signature scheme. Then the following holds:*

$$\text{Adv}_A^{\epsilon\text{-UF}}(\lambda) \leq (q(\lambda) + 1) \cdot (m \cdot \text{Adv}_B^{\epsilon\text{-TimeLock}}(\lambda) + \text{Adv}_C^{\text{FS}}(\lambda))$$

For a proof of Theorem 7.1 see Appendix A.

Theorem 7.2. *If the underlying delegatable functional signature scheme is key-indistinguishable then the constructed time-deniable signatures scheme satisfies the deniability property.*

For a proof of Theorem 7.2 see Appendix A.

<u>KeyGen</u> ($1^\lambda, T$) : $(mvk, msk) \leftarrow \text{FS.Setup}(1^\lambda)$ return $((mvk, T, \lambda), (msk, T, \lambda))$	<u>AltSign</u> ($vk, (m^*, t^*, \sigma^*), m, t$) : $mvk, T, \lambda = \text{parse}(vk)$ $c^*, s^* = \text{parse}(\sigma^*)$ $sk_{t^*} = \text{TimeLock.Sol}(c^*)$ sk_t	<u>Verify</u> (vk, σ, m, t) : $c, s = \text{parse}(\sigma)$ return $\text{FS.Verify}(vk, t m, s)$
<u>Sign</u> ($sk = (msk, T, \lambda), m, t$) : $sk_t \leftarrow \text{FS.KeyGen}(msk, f_t)$ $v, s \leftarrow \text{FS.Sign}(sk_t, f_t, (t, m))$ $c \leftarrow \text{TimeLock.Gen}(1^\lambda,$ $T \cdot z(\lambda), sk_t)$ return (c, s)	$sk_t \leftarrow$ $\text{FS.Delegate}(mvk, f_t^*, sk_{t^*}, f_t)$ $v, s \leftarrow \text{FS.Sign}(sk_t, f_t, (t, m))$ $c \leftarrow \text{TimeLock.Gen}(1^\lambda, T \cdot z(\lambda), sk_t)$ return (c, s)	\leftarrow

Figure 3: A construction for a time-deniable signature scheme DS from a key-indistinguishable, delegatable, functional signature scheme FS and a time lock puzzle TimeLock. The function f_t for each input message m and time \bar{t} is defined as $f_t(\bar{t}, m) = \text{parse}(\bar{t}) || m$ if $\bar{t} \leq t$, else \perp . The polynomial $z(\lambda)$ is a multiplicative factor for the difficulty parameter of the time lock puzzle and is described in the text.

8 Implementation and Evaluation

Implementation. To demonstrate the potential efficiency of our scheme, we implemented a variant of it in python. For our time lock puzzle, we modified an existing, open-source implementation of an RSW time-lock puzzle [26]. Our timestamp is capable of supporting 2^{16} different values. This is approximately equivalent to [38] where granularity of the timestamp is 15 minutes and covers a period of 2 years. This is reasonable given at least some motivating applications (i.e. email), where it is assumed that frequent key rotation is being done for domains and coarse granularity may be acceptable.

Construction of FS To instantiate our functional signature scheme, we need a HIBE scheme that is both key-indistinguishable and adaptively secure. The Lewko-Waters scheme mentioned in Section 4.1.2 does satisfy these properties. However, it is not efficient enough for our application. Instead, we implement a version of the more efficient Gentry-Silverberg (GS) [21] scheme using Charm [1]. Modifications are needed since natively GS does not satisfy the key-indistinguishability requirement: when a key is delegated, some of the Q_i values in the secret key are sub-keys given to the user by their parent making distinguishing easy. To fix this, we re-randomize these values by having the parent derive new exponents v_i which are used to update all relevant Q_i values to $v_i P_0 + Q_i$. We also update S_t with $v_{i-1} P_i$ to ensure correctness. These modifications give key-indistinguishability to GS for little overhead. Unfortunately, the Gentry-Silverberg HIBE is only adaptive-secure with exponential loss in the tree depth. While it is true that our depth is constant, the loss is still too great to allow instantiating with reasonable security guarantees. For example, assuming that our tree depth is $t = 5$ with adversaries making billions of key queries, we need $> 156 + k$ bits of security for the BDH problem to instantiate with any security level k . With our own security loss accounted for, the scheme would only be able to achieve at most a 68 bit level of security assuming a pairing-friendly curve with 256 bit security. Due to a library limitation, we instantiate with curve BN-254 (a curve having ~ 100 bits of security), with the understanding that the numbers reported below are merely for approximating benchmarks.

Setting Parameters. Recall that in our construction we make use of a time lock puzzle. There are two main concerns that come with implementing schemes that use time-based crypto assumptions: one is precisely capturing the speed-up (if any) offered by parallelism, the other is accurately estimating the fastest real-world time to do the computational task the assumption is based on.

On the first point, to the best of our knowledge, there are no known improvements from parallelism against the RSW assumption. There are certainly papers that imply solving many instances at once can be done faster than just sequentially solving each instance [32], but it is hard to see how this is helpful in decreasing the time to solve a single instance faster. Therefore, as in prior work, we assume no speed up from parallelism.

For the second, state-of-the-art techniques mostly use an ad-hoc approach, where special hardware or an optimized ASIC implementation is created specifically to determine the best possible times. Of course, this prior approach does not account for the continuous increase in computation power over multiple years or decades and therefore is unreliable for large time windows. A more accurate method would estimate the change in computing power over the time window needed upon instantiation and adjust the time lock puzzle parameter accordingly. Unfortunately, this requires an accurate estimate of the increase in computational power over time which has been historically hard to provide [14]. Therefore, it is advisable to use time-based crypto schemes where the timing parameter Δt is small, perhaps a year or two at most. As we do not have access to an optimized implementation of repeated squaring, we simply benchmark the cost of computing squares modulo a 2048 bit composite on our machine. For us this corresponded to roughly 5,883,206 squares per second.

Experimental Evaluation. Experiments were done on a Intel Xeon E5 with 500GB of memory, running Ubuntu. As previously stated, the implementation itself was done in python and does not use either multi-threading or multiprocessing. Estimates were obtained by running each algorithm 500 times and taking the median. The timestamp value is chosen uniformly at random, as signing time differs significantly depending on the value of t .

We begin analysis by first examining the effect of varying the arity of the tree. This is the number of children per parent node and we denote it throughout this section by N . This is an important parameter for our scheme because the time lock puzzle sent along with the original signature has size that implicitly depends on it. Figure 4 shows the differences in signing time when varying this parameter from 2 to 19. A similar figure in appendix E shows how the signature size varies for the same values. We note that overall signing becomes more efficient as the depth of the tree decreases.

Our microbenchmarks are presented in table 2. While our scheme does incur a slightly large computation cost in some respects to other existing schemes, for our proposed application of DKIM this demand is not unreasonable: email is largely for non-instantaneous transactions and all of our algorithms run in only hundred of milliseconds which should not noticeably effect user experience.

Finally, we note that performance can be expected to improve substantially with proper utilization of parallelism. The signing algorithm requires creating a functional key where each sub-key can be extracted simultaneously. Using back of the envelope calculations, for our 16-ary tree, assuming that the cost of signing is equivalent to KeyForge, computing the time lock encryption takes 50 ms (as it does in our implementation), and we have access to ≥ 5 processors, we could achieve a median of 65.94 ms for signing for our current scheme. Thus, we expect microbenchmarks on optimized implementations to be competitive with existing schemes even if slightly slower on account of working in groups of a larger size for a reasonable security level.

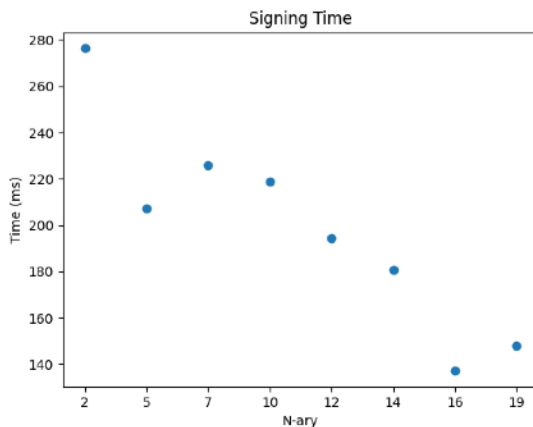


Figure 4: The signing time for a time-deniable signature scheme using the Gentry-Silverberg HIBE with BN-254 and varying values of N

KeyGen (ms)	Sign (ms)	Signature Size (B)	Verify (ms)
109	223	7162	178

Table 2: Relevant operations when instantiating the scheme of Figure 3 using a Gentry-Silverberg HIBE with BN-254 and $N = 16$. All benchmarks were calculated taking the median over 500 attempts while choosing a fresh timestamp value.

9 On the Necessity of Time-Lock Puzzles

Our construction of time-deniable signatures makes use of time-lock puzzles to achieve short-term unforgeability. We show that the use of such a primitive is to an extent inherent. Namely, assuming extractable witness encryption [19, 22], we show that time-deniable signatures imply time-lock puzzles.

We demonstrate this implication in Appendix F. We remark that while extractable witness encryption is a strong tool, it alone is not known to imply time-lock puzzles.⁷

10 Conclusion

In this work we introduced a new notion of deniable signatures that provides strong unforgeability and deniability guarantees without requiring the signer to periodically publish secret key material. We show how to realize our primitive using time lock puzzles and a HIBE scheme that satisfies a special key-indistinguishability property. Important directions for future work include constructing time-deniable signatures from a different set of assumptions (non-HIBE based) and building constructions that satisfy robust correctness. Improvements to the real-world efficiency of HIBE schemes, be it through better implementations or new constructions, will greatly influence the efficiency of our time-deniable signatures construction and allow our scheme to be more competitive with existing proposals.

⁷When supplemented with a computational reference clock, it is known to imply time lock puzzles [30].

References

- [1] Charm: A framework for rapidly prototyping cryptosystems. <https://github.com/JHUISI/charm>.
- [2] emails: Lawyer who met trump jr. tied to russian officials.
- [3] Signal protocol documentation. <https://www.signal.org/docs/>.
- [4] Arasu Arun, Joseph Bonneau, and Jeremy Clark. Short-lived zero-knowledge proofs and signatures.
- [5] Michael Backes, Sebastian Meiser, and Dominique Schröder. Delegatable functional signatures. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part I*, volume 9614 of *LNCS*, pages 357–386. Springer, Heidelberg, March 2016.
- [6] Mihir Bellare and Georg Fuchsbauer. Policy-based signatures. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 520–537. Springer, Heidelberg, March 2014.
- [7] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 757–788. Springer, Heidelberg, August 2018.
- [8] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer, Heidelberg, August 2001.
- [9] Nikita Borisov, Ian Goldberg, and Eric A. Brewer. Off-the-record communication, or, why not to use PGP. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society, WPES 2004, Washington, DC, USA, October 28, 2004*, pages 77–84, 2004.
- [10] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.
- [11] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *International workshop on public key cryptography*, pages 501–519. Springer, 2014.
- [12] Jon Callas. Re: [Ietf-dkim] DKIM key rotation best practice, 2020.
- [13] Jon Callas, Eric Allman, Miles Libbey, Michael Thomas, Mark Delany, and Jim Fenton. Domainkeys identified mail (dkim) signatures.
- [14] Adam Conner-Simons. Programmers solve mit’s 20-year-old cryptographic puzzle. <https://www.csail.mit.edu/news/programmers-solve-mits-20-year-old-cryptographic-puzzle>, 2019.
- [15] Ronald Cramer, Goichiro Hanaoka, Dennis Hofheinz, Hideki Imai, Eike Kiltz, Rafael Pass, abhi shelat, and Vinod Vaikuntanathan. Bounded CCA2-secure encryption. In Kaoru Kurosawa, editor, *ASIACRYPT 2007*, volume 4833 of *LNCS*, pages 502–518. Springer, Heidelberg, December 2007.

- [16] Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. Non-malleable time-lock puzzles and applications. *IACR Cryptol. ePrint Arch.*, 2020:779, 2020.
- [17] Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. SPARKs: Succinct parallelizable arguments of knowledge. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 707–737. Springer, Heidelberg, May 2020.
- [18] Ethereum Foundation . The Ethereum Beacon Chain. Available at <https://ethereum.org/en/eth2/beacon-chain/>.
- [19] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013.
- [20] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [21] Craig Gentry and Alice Silverberg. Hierarchical id-based cryptography. In Yuliang Zheng, editor, *Advances in Cryptology — ASIACRYPT 2002*, pages 548–566, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [22] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2013.
- [23] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 162–179. Springer, Heidelberg, August 2012.
- [24] Matthew D. Green. OK Google: Please publish your DKIM secret keys, Dec 2020.
- [25] Andreas Hülsing and Florian Weber. Epochal signatures for deniable group chats. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, pages 1677–1695, 2021.
- [26] Jonathan Levi. time-lock-puzzle. Available at <https://github.com/drummerjolev/time-lock-puzzle>, 2020.
- [27] Allison B. Lewko. Tools for simulating features of composite order bilinear groups in the prime order setting. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 318–335. Springer, Heidelberg, April 2012.
- [28] Allison B. Lewko and Brent Waters. Unbounded HIBE and attribute-based encryption. In *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, pages 547–567, 2011.

- [29] Allison B. Lewko and Brent Waters. Why proving HIBE systems secure is difficult. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 58–76. Springer, Heidelberg, May 2014.
- [30] Jia Liu, Tibor Jager, Saqib A. Kakvi, and Bogdan Warinschi. How to build time-lock encryption. *Des. Codes Cryptogr.*, 86(11):2549–2586, 2018.
- [31] Jeremy B. Merrill. Authenticating email using dkim and arc, or how we analyzed the kasowitz emails, Jul 2017.
- [32] Andrew Moss, Dan Page, and Nigel P. Smart. Toward acceleration of RSA using 3D graphics hardware. In Steven D. Galbraith, editor, *11th IMA International Conference on Cryptography and Coding*, volume 4887 of *LNCS*, pages 364–383. Springer, Heidelberg, December 2007.
- [33] NIST. NIST Randomness Beacon (prototype implementation). Available at <https://beacon.nist.gov/home>.
- [34] Krzysztof Pietrzak. Proofs of catalytic space. In Avrim Blum, editor, *ITCS 2019*, volume 124, pages 59:1–59:25. LIPIcs, January 2019.
- [35] Ronald L Rivest, Adi Shamir, and David A Wagner. Time-lock puzzles and timed-release crypto. 1996.
- [36] Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
- [37] Elaine Shi and Brent Waters. Delegating capabilities in predicate encryption systems. In *International Colloquium on Automata, Languages, and Programming*, pages 560–578. Springer, 2008.
- [38] Michael A. Specter, Sunoo Park, and Matthew Green. Keyforge: Non-attributable email from forward-forgable signatures. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 1755–1773, 2021.
- [39] Craig Timberg, Matt Viser, and Tom Hamburger. Here’s how The Post analyzed Hunter Biden’s laptop.
- [40] Nik Unger and Ian Goldberg. Improved strongly deniable authenticated key exchanges for secure messaging. *Proc. Priv. Enhancing Technol.*, 2018(1):21–66, 2018.
- [41] Benjamin Wesolowski. Efficient verifiable delay functions. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 379–407. Springer, Heidelberg, May 2019.

A Proofs for Time Deniable Signatures

Theorem A.1. *The time-deniable signature scheme presented in figure 3 is unforgeable.*

In the discussion that follows, let the output of a hybrid game \mathcal{H} be the output of the challenger. We prove the theorem statement using a hybrid argument where \mathcal{H}_0 represents the original (ϵ, T) -unforgeability game. Where details are omitted in the hybrid description of \mathcal{H}_i , it is assumed they are the same as in \mathcal{H}_{i-1} .

$\mathcal{H}_1 =$ Let $q(\lambda) = \text{size}(\mathcal{A}_{\lambda,1})$. Challenger samples $r \xleftarrow{\$} [q(\lambda)] \cup \{0\}$. If the number of queries made to O_{Sign} by $\mathcal{A}_{\lambda,1}$ is not r , output \perp .

Claim 1. $\text{Adv}_{\mathcal{A}}^{\mathcal{H}_1}(\lambda) = \frac{1}{q(\lambda)+1} \text{Adv}_{\mathcal{A}}^{\mathcal{H}_0}(\lambda)$

Note that the win condition is checked whenever the challenger “correctly guesses” how many queries will be made by the adversary in the second phase. Let m be the number of queries made by $\mathcal{A}_{\lambda,1}$, where $0 \leq m \leq q(\lambda)$. This must hold since the adversary cannot make more queries than its size dictates. Therefore,

$$\text{Adv}_{\mathcal{A}}^{\mathcal{H}_1}(\lambda) = \Pr(r = m | r \xleftarrow{\$} \{0, \dots, q(\lambda)\}) \text{Adv}_{\mathcal{A}}^{\mathcal{H}_0}(\lambda)$$

It should be clear the first quantity is $\frac{1}{q(\lambda)+1}$ and the claim is thus true.

Consider the following sequence of hybrids where $2 \leq i \leq m+1$

$\mathcal{H}_i =$ On the $(i-1)^{\text{th}}$ query to O_{Sign} by $\mathcal{A}_{\lambda,1}$, challenger replaces $c = \text{TimeLock.Gen}(T, sk_i)$ with $\text{TimeLock.Gen}(T, 0)$.

Claim 2. $\exists \mathcal{B}$, a $\epsilon - \text{TimeLock}$ adversary, such that $|\text{Adv}_{\mathcal{A}}^{\mathcal{H}_i}(\lambda) - \text{Adv}_{\mathcal{A}}^{\mathcal{H}_{i-1}}(\lambda)| \leq \text{Adv}_{\mathcal{B}}^{\epsilon - \text{TimeLock}}(\lambda)$

WLOG, assume $\text{Adv}_{\mathcal{A}}^{\mathcal{H}_i} > \text{Adv}_{\mathcal{A}}^{\mathcal{H}_{i-1}}$. First consider the following distinguisher \mathcal{D} between \mathcal{H}_{i-1} and \mathcal{H}_i , where output of \mathcal{D} being 0 denotes \mathcal{H}_{i-1} and 1 denotes \mathcal{H}_i .

Description of \mathcal{D} on input b from \mathcal{H}_{i-1} or \mathcal{H}_i :

- If $b = 1$ output $b' = 1$
- If $b = 0$, output $b' = 0$.

Notice that \mathcal{D} 's advantage in distinguishing is dependent on $\text{Adv}_{\mathcal{A}}^{\mathcal{H}_i}(\lambda) - \text{Adv}_{\mathcal{A}}^{\mathcal{H}_{i-1}}(\lambda)$ and that the depth of \mathcal{D} is 2. We will now use \mathcal{D} and \mathcal{A}_{λ} to construct an adversary \mathcal{B}' .

Description of $\epsilon - \text{TimeLock}$ adversary \mathcal{B}' :

- Honestly run the FS.Setup algorithm and answer all queries from $\mathcal{A}_{\lambda,0}$ honestly
- For the j^{th} query from $\mathcal{A}_{\lambda,1}$:
 - if $j < i-1$, $c \leftarrow \text{TimeLock.Gen}(1^\lambda, T, 0)$
 - if $j > i-1$, $c \leftarrow \text{TimeLock.Gen}(1^\lambda, T, sk_{t_j})$
 - if $j = i-1$, compute $sk_{t_{i-1}} \leftarrow \text{FS.KeyGen}(msk, f_{t_{i-1}})$, send the challenger $(sk_{t_{i-1}}, 0)$, and receive z . Set $c = z$.
- From \mathcal{A}_{λ} get output b and give b to \mathcal{D} . Get b' from \mathcal{D} . Output b' to the challenger.

Define $(\mathcal{B}_1, \mathcal{B}_2) = \mathcal{B}'$ where \mathcal{B}_1 represents \mathcal{B}' up until query $i - 1$ is made and \mathcal{B}_2 is all that follows. Let the output of \mathcal{B}_1 be z and hard-code it into \mathcal{B}_2 to get \mathcal{B} .

Analysis of adversary \mathcal{B} :

Let $|\mathbf{A.B}|$ denote the depth bound on the algorithm \mathbf{B} for primitive \mathbf{A} . The depth bound for \mathcal{B} is represented below

$$\begin{aligned} \text{depth}(\mathcal{B}) &\leq \text{depth}(\mathcal{A}_1) + m \cdot \left(|\text{FS.KeyGen}| + |\text{TimeLock.Gen}| + |\text{FS.Sign}| \right) + |\text{FS.Verify}| + 2 \\ &\leq \frac{t(\lambda)}{\epsilon(\lambda)} + m \cdot \left(|\text{FS.KeyGen}| + |\text{TimeLock.Gen}| + |\text{FS.Sign}| \right) + |\text{FS.Verify}| + 2 \\ &\leq \frac{\epsilon(\lambda)(|\text{FS.Verify}| + 2)}{\epsilon(\lambda)} + \frac{t(\lambda)[1 + |\text{FS.KeyGen}| + |\text{TimeLock.Gen}| + |\text{FS.Sign}|]}{\epsilon(\lambda)} \\ &\leq \frac{t(\lambda)z(\lambda)}{\epsilon(\lambda)} \end{aligned}$$

Assuming the ϵ -gap security of the `TimeLock` and because $\Delta = t'(\lambda) = t(\lambda) \cdot z(\lambda)$ in our construction, \mathcal{B} is appropriately bounded.

$$\begin{aligned} \Pr(\mathcal{B} \text{ succeeds}) &= \frac{1}{2} \Pr(\mathcal{B} \text{ succeeds} \mid \beta = 0) + \frac{1}{2} \Pr(\mathcal{B} \text{ succeeds} \mid \beta = 1) \\ &= \frac{1}{2} \Pr(\mathcal{D} \text{ outputs } 0 \text{ when given } \mathcal{H}_{i-1}) + \frac{1}{2} \Pr(\mathcal{D} \text{ outputs } 1 \text{ when given } \mathcal{H}_i) \\ &= \Pr(\mathcal{D} \text{ correctly distinguishes } \mathcal{H}_{i-1} \text{ and } \mathcal{H}_i) \end{aligned}$$

\mathcal{B} 's probability of success entirely depends on \mathcal{D} 's and thus \mathcal{B} 's advantage is the same as \mathcal{D} 's. The claim thus follows.

Claim 3. $\exists \mathcal{C}$, an adversary against the unforgeability of the FS scheme, s.t. $\text{Adv}_{\mathcal{A}}^{\mathcal{H}_{m+1}}(\lambda) \leq \text{Adv}_{\mathcal{C}}^{\text{FS}}(\lambda)$

We now argue that the advantage of any adversary in \mathcal{H}_{m+1} can be translated into equivalent advantage against the unforgeability scheme.

Description of \mathcal{C} :

- Receive mvk from the FS challenger.
- On queries m, t to $\text{O}_{\text{Sign}}(\cdot)$
 - if phase one, query $\bar{\text{O}}_{\text{Key}}(f_t, i)$ and receive sk_{f_t} where $i \in \mathbb{N}$ is next available counter. Compute $\sigma \leftarrow \text{FS.Sign}(sk_{f_t}, f_t, (t, m))$, $c \leftarrow \text{TimeLock.Gen}(1^\lambda, T, sk_{f_t})$. Return (c, σ) to $\mathcal{A}_{\lambda,0}$
 - if phase two, query $\bar{\text{O}}_{\text{Sign}}(f_t, i, (t, m))$ with $i \in \mathbb{N}$ being the next highest counter to get m^*, σ . Compute $c \leftarrow \text{TimeLock.Gen}(1^\lambda, T, 0)$ and return (c, σ) to $\mathcal{A}_{\lambda,1}$
- If \mathcal{A}_λ returns forgery $m, t, \sigma = (c, s)$ return $(t||m, s)$ to the FS challenger.

Analysis. We now show that if \mathcal{A} is successful, then \mathcal{C} must be as well. Say \mathcal{A} returns a forgery $m^*, t^*, \sigma^* = (c, s)$. In order for \mathcal{A} to be admissible, it must be true that \mathcal{A} never received a signature with $t \geq t^*$ during the first phase and during the second phase there was never a query

for (m^*, t^*) specifically. The first point implies \mathcal{C} never queries for a secret key for a function f_t where $t \geq t^*$ so s is a valid signature to give back to the functional challenger. The second point means that \mathcal{A} is not giving \mathcal{C} a signature that \mathcal{C} asked for from the FS challenger with some mangled c' where c' is an incorrectly structured puzzle or does not hide the right secret key. Therefore if \mathcal{A} returns a valid forgery, then \mathcal{C} returns a valid forgery and the claim follows.

Theorem A.2. *If the underlying delegatable functional signature scheme is key-indistinguishable then the constructed time-deniable signatures scheme satisfies the deniability property.*

We prove this by showing how to use an adversary \mathcal{A} who wins the time-deniable signatures $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$ game to construct an adversary \mathcal{B} which wins the delegatable functional signatures key-indistinguishability game $\mathbf{Exp}_{\text{FS}}^{\text{IND}}$.

- \mathcal{B} receives the master verification key and master signing key (mvk, msk) from its challenger and forwards it as is to \mathcal{A} .
- In the query phase, whenever \mathcal{A} makes an $\text{O}_{\text{Sign}}(sk, \cdot, \cdot)$ query on message m and timestamp t , \mathcal{B} performs the following operations:
 - Query the FS key oracle $\bar{\text{O}}_{\text{Key}}(msk, \cdot, \cdot)$ on t to get back (id, sk_t) . Add the response to its internal table of responses. Otherwise, return \perp .
 - Use the received key to create $v, \sigma_{\text{FS}} \leftarrow \text{FS.Sign}(f_t, sk_t, t, m)$ and $c \leftarrow \text{TimeLock.Gen}(\Delta, sk_t)$.
 - Let $\sigma_{\text{DS}} = (c, \sigma_{\text{FS}})$. Send (id, σ_{DS}) to \mathcal{A} . This simulates the functionality of $\text{O}_{\text{Sign}}(sk, t, m)$ which outputs an identifier and $\sigma \leftarrow \text{DS.Sign}(sk, m, t)$.
- In the query phase, whenever \mathcal{A} makes a query to the DS challenge oracle $\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$ on some tuple (id', m', t') , \mathcal{B} performs the following operations:
 - Query the FS challenge oracle on (id', t, t') if id' corresponds to a query on time t in its table and $t' < t$.
 - Using the received key $sk_{t'}$, compute $v', \sigma'_{\text{FS}} \leftarrow \text{FS.Sign}(f_{t'}, sk_{t'}, t, m)$ and $c' \leftarrow \text{TimeLock.Gen}(\Delta, sk_{t'})$.
 - Send $\sigma'_{\text{DS}} = (c', \sigma'_{\text{FS}})$ to \mathcal{A} . This simulates the $\text{DS.Sign}()$ algorithm when the challenger's sampled bit β is 0 as the key used for $\text{FS.Sign}()$ is a freshly generated key. When $\beta = 1$, the key used for $\text{FS.Sign}()$ is a delegated key, which simulates the $\text{DS.AltSign}()$ algorithm.
- At the end, \mathcal{A} outputs its guess β' for β , \mathcal{B} forwards this without change to its challenger. The advantage of \mathcal{B} in winning the $\mathbf{Exp}_{\text{FS}}^{\text{IND}}$ game is same as the advantage of \mathcal{A} in winning the $\mathbf{Exp}_{\text{DS}}^{\text{IND}}$ game as all the responses to \mathcal{A} 's queries are simulated correctly by \mathcal{B} .

B Proofs for Delegatable Functional Signatures

Theorem B.1. *If HIBE is adaptively secure then the functional signature scheme for prefix functions constructed in Figure 2 is unforgeable.*

We prove this by showing how to use an adversary \mathcal{A} who succeeds with non-negligible advantage in $\mathbf{Exp}_{\text{FS}}^{\text{UNF}}$ to construct an adversary \mathcal{B} which succeeds with non-negligible advantage in the HIBE unforgeability game.

Description of \mathcal{B}

- In the setup phase, initialize empty table \mathcal{T} and receive pk from HIBE challenger and forward it to \mathcal{A} .
- When \mathcal{A} queries (f_t, i) to $\bar{\mathcal{O}}_{\text{Key}}$ check if row with i in \mathcal{T} .
 - If it exists, return the list of keys sk_t to \mathcal{A} .
 - Otherwise, use the algorithm FS.KeyGen algorithm in Figure 2 replacing HIBE.KeyGen($msk, node_{id}$) with $\mathcal{QC}(id)$ and $\mathcal{QR}(id)$. Let $list_{sk_t}$ be the resulting list of keys. Record $(i, t, list_{sk_t})$ in \mathcal{T} . Send $pk, list_{sk}$ to \mathcal{A} .
- When \mathcal{A} queries (f_t, i, m) to $\bar{\mathcal{O}}_{\text{Sign}}$ first parse m as $\bar{t}||\bar{m}$. If $\bar{t} > t$ output \perp . Check if there is row in \mathcal{T} with identifier i
 - If there is, let $list_{sk_t}$ be the list of keys associated with that row. Let $sk_{t'}$ be the key in $list_{sk_t}$ associated with an identity that is the prefix of \bar{t} . Compute $sk_{\bar{t}||\bar{m}} \leftarrow \text{HIBE.Delegate}(pk, sk_{t'}, \text{suffix}(t', \bar{t}||\bar{m}))$ where suffix omits the prefix t' from $\bar{t}||\bar{m}$. Return $sk_{\bar{t}||\bar{m}}$
 - Otherwise, use the algorithm FS.KeyGen in Figure 2 replacing HIBE.KeyGen($msk, node_{id}$) with $\mathcal{QC}(id)$. Finally query $\mathcal{QD}(\bar{t}||\bar{m})$ and do a subsequent reveal query $\mathcal{QR}(\bar{t}||\bar{m})$ to get $sk_{\bar{t}||\bar{m}}$. Return $sk_{\bar{t}||\bar{m}}$ to \mathcal{A} .
- When \mathcal{A} outputs its forgery (m^*, σ^*) parses m^* as $t||m$ and σ^* as sk^* .
 - Sample a message msg and check that $\text{Decrypt}(sk^*, \text{HIBE.Encrypt}(pk, t||m, msg)) = msg$. If this does not hold or if \exists a row $(i', t', list'_{sk_{t'}})$ in \mathcal{T} where $t \leq t'$ output $\beta' \xleftarrow{\$} \{0, 1\}$
 - Otherwise randomly sample messages m_0 and m_1 . Let $I^* = t||m$. Send to challenger (m_0, m_1, I^*) and receive ct . Compute $m' \leftarrow \text{HIBE.Decrypt}(sk_{I^*}, ct)$. If $m = m_0$, output 0. If $m = m_1$, output 1. If the response is neither, output $\beta' \xleftarrow{\$} \{0, 1\}$.

Analysis

In order to be an admissible adversary, \mathcal{A} must return a signature σ^* and an m^* that verify where they do not hold a functional key that has m^* in its range. The keys that have m^* in their range are of the form f_T where $T \geq t$. In other words, these functional keys are those that contain some HIBE secret keys that are prefixes of the identity t and no other functional key has such prefix HIBE keys by the design of the construction. Therefore if \mathcal{A} is admissible, $m^* = t||m$ will be a valid identity to challenge on.

If \mathcal{A} is successful, then σ^* passed verification meaning for a random message it acted as a secret key for the identity $t||m$. This implies with high confidence that it is in fact the secret key for this identity. Decrypting with the secret key for identity $t||m$ the ciphertext ct will be successful with overwhelming probability and therefore most of the time when \mathcal{A} succeeds \mathcal{B} succeeds as well.

In any other circumstance, when \mathcal{A} is either not admissible or does not return a valid forgery, \mathcal{B} catches this and responds with a uniform bit. Thus the theorem statement follows.

Theorem B.2. *If HIBE is key-indistinguishable then the scheme in Figure 2 satisfies the functional signatures key-indistinguishability property.*

We prove this by showing how to use an adversary \mathcal{A} who wins the delegatable functional signatures key-indistinguishability game $\mathbf{Exp}_{\text{FS}}^{\text{IND}}$ to construct an adversary \mathcal{B} which wins the HIBE key-indistinguishability game $\mathbf{Exp}_{\text{HIBE}}^{\text{IND}}$.

- \mathcal{B} receives the keys (pk, msk) from its challenger and forwards it as is to \mathcal{A} .
- After this, in the query phase, when \mathcal{A} makes a $\bar{\text{O}}_{\text{Key}}(\cdot)$ query for time t , adversary \mathcal{B} computes $trace \leftarrow \text{Trace}(root, t)$ where $root$ is the position of msk in the HIBE hierarchy tree. This gives \mathcal{B} the list of nodes on which it queries the key oracle $\mathcal{QK}(\cdot)$. Each of the $\mathcal{QK}(\cdot)$ query response has an identifier id along with the key for a node sk_{node} . \mathcal{B} maintains a table with entries of the form $(t, id', \{(id, sk_{node})\}_{node \in trace})$ where id' represents the counter value corresponding to this particular query from \mathcal{A} . \mathcal{B} sends $(id', sk_t = \{sk_{node}\}_{node \in trace})$ to \mathcal{A} . This is the response \mathcal{A} expects as \mathcal{B} simulates the $\text{FS.KeyGen}(msk, t)$ algorithm with its queries to the HIBE key oracle.
- When \mathcal{A} makes a FS challenge oracle $\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$ query with a tuple of the form (i, t_0, t_1) , \mathcal{B} performs the following operations:
 - Check that $t_0 > t_1$ and there is a row starting with (i, t_0) in its table, otherwise return \perp . This guarantees that on the shortest paths from the leaf node t_1 to the root in the HIBE hierarchy tree (Figure 1), there exists an element j such that its corresponding HIBE key is present in the set sk_{t_0} representing the FS key for t_0 .
 - Compute $sk, j \leftarrow \text{findPrefix}(sk_t, t_1)$ and $trace' \leftarrow \text{Trace}(t_0^j, t_1)$ which is the trace of leaf node t_1 in a tree where the root is t_0^j , the first j bits of t_0 . Rerandomize the key sk by computing $sk' \leftarrow \text{HIBE.Delegate}(pk, sk, nil)$, similarly rerandomize all the keys in set sk_{t_0} upto the j 'th position.
 - Query the HIBE challenge oracle $\text{O}_{\text{Ch}}(\cdot, \cdot, \cdot)$ on tuples $(id_j, t_0^j, node)$ for all $node \in trace'$ where id_j corresponds to the $\mathcal{QK}(\cdot)$ response on t_0^j . \mathcal{B} finds id_j in its table, in the row corresponding to t_0 .
 - Compile all the rerandomized keys and the keys received from the key oracle into set sk_{t_1} . Send sk_{t_1} to \mathcal{A} . This is the response \mathcal{A} expect as \mathcal{B} simulates either $\text{FS.KeyGen}()$ or $\text{FS.Delegate}()$ depending on the challenger's sampled bit β .
- At the end, \mathcal{A} outputs its guess β' for β , \mathcal{B} forwards this without change to the challenger. The advantage of \mathcal{B} in winning the $\mathbf{Exp}_{\text{HIBE}}^{\text{IND}}$ game is same as the advantage of \mathcal{A} in winning the $\mathbf{Exp}_{\text{FS}}^{\text{IND}}$ game as all the responses to \mathcal{A} 's queries are simulated correctly by \mathcal{B} .

C Proof of K-hop Deniability

Theorem C.1. *Any time-deniable signature scheme satisfying the deniability property as defined in definition 5.5, also satisfies the k -hop deniability property as defined in definition 5.6.*

We prove this by a hybrid argument. In each subsequent hybrid \mathcal{H} we will slightly modify the security game until we get to one where the advantage in distinguishing is indisputably negligible because both inputs will be drawn from the same distribution.

Let \mathcal{H}_0 be the original k -hop security game and consider the sequence of hybrids $\mathcal{H}_1 \dots \mathcal{H}_k$ that are defined as follows:

$\mathcal{H}_i =$ In each of these hybrids our goal will be to eliminate the dependency of AltSign^k on the message (m_{i-1}, t_{i-1}) . \mathcal{H}_i is identical to \mathcal{H}_{i-1} except when $\beta = 1$, σ is generated as

$$\text{AltSign}^{k-i}(vk, (m_i, t_i, \sigma_i), \{(m_{i+1}, t_{i+1}), \dots, (m^*, t^*)\}) \quad (8)$$

where $\sigma_i \leftarrow \text{Sign}(sk, m_i, t_i)$.

In the discussion that follows let \mathcal{H}_i denote the outcome of \mathcal{H}_i , that is the probability that \mathcal{A} correctly guesses the bit β' chosen by the challenger. We let the notation $\mathcal{H}_{i-1} \approx \mathcal{H}_i$ mean there is negligible difference between the success of \mathcal{A} in \mathcal{H}_{i-1} and \mathcal{H}_i .

Claim 4. $\forall i \in [k], \mathcal{H}_{i-1} \approx \mathcal{H}_i$

For the sake of contradiction, suppose this is not true and \exists a non-negligible function $\rho(\lambda)$ so that $|\Pr(\mathcal{A}_\lambda \text{ outputs } \beta' \mid \mathcal{H}_i) - \Pr(\mathcal{A}_\lambda \text{ outputs } \beta' \mid \mathcal{H}_{i-1})| = \rho(\lambda)$. Without loss of generality, assume $\Pr(\mathcal{A}_\lambda \text{ outputs } \beta' \mid \mathcal{H}_i) > \Pr(\mathcal{A}_\lambda \text{ outputs } \beta' \mid \mathcal{H}_{i-1})$. We now use this difference in success probabilities and the adversary \mathcal{A}_λ to construct an adversary \mathcal{B} that has non-negligible advantage in $\text{Exp}_{\text{DS}}^{\text{IND}}$.

Description of \mathcal{B}

- \mathcal{B} receives vk from its challenger and sends it to \mathcal{A} . It also initializes an empty table \mathcal{T} and uniformly samples $\beta' \xleftarrow{\$} \{0, 1\}$
- In the query phase, \mathcal{B} responds to the sign queries from \mathcal{A} using the O_{Sign} oracle: it forwards (m, t) to O_{Sign} , receives (id, σ) , and records $(\text{id}, m, t, \sigma)$ in \mathcal{T} . It then returns (id, σ) to \mathcal{A}_λ .
- In the query phase, let a challenge query to O_{Ch} from \mathcal{A}_λ be $\text{id}, \{(m_j, t_j)\}_{j \in [k-1]}, m^*, t^*$.
 - Check if \mathcal{T} has an identifier id . If not output \perp .
 - If $\beta' = 0$, query O_{Sign} on m^*, t^* to get (id^*, σ^*) and return σ^* to \mathcal{A}_λ
 - Else if $\beta' = 1$ and $i = 1$, let $\text{id}' = \text{id}$. Otherwise query O_{Sign} on (m_{i-1}, t_{i-1}) to get id', σ . \mathcal{B} makes a query to its challenge oracle with id', m_i, t_i and receives σ_i .
 - Compute $\sigma^* \leftarrow \text{AltSign}^{k-i}(vk, (m_i, t_i, \sigma_i), \{(m_{i+1}, t_{i+1}), \dots, (m^*, t^*)\})$ and send σ^* to \mathcal{A}_λ
- Let b be the output of \mathcal{A}_λ . If $\beta' = b$, return 0. Else return 1.

Analysis

We now analyze \mathcal{B} 's success probability. In the discussion that follows, let β be the bit chosen by the challenger in the $\text{Exp}_{\text{DS}}^{\text{IND}}$ game.

$$\begin{aligned}
\Pr(\mathcal{B} \text{ outputs } \beta) &= \frac{1}{2} \Pr(\mathcal{B} \text{ outputs } \beta \mid \beta = 0) + \\
&\quad \frac{1}{2} \Pr(\mathcal{B} \text{ outputs } \beta \mid \beta = 1) \\
&= \frac{1}{2} \Pr(\mathcal{A}_\lambda \text{ succeeds in } \mathcal{H}_i \mid \beta = 0) + \\
&\quad \frac{1}{2} [1 - \Pr(\mathcal{A}_\lambda \text{ succeeds in } \mathcal{H}_{i-1} \mid \beta = 1)] \\
&= \frac{1}{2} + \frac{1}{2} [\Pr(\mathcal{A}_\lambda \text{ succeeds in } \mathcal{H}_i \mid \beta = 0) - \\
&\quad \Pr(\mathcal{A}_\lambda \text{ succeeds in } \mathcal{H}_{i-1} \mid \beta = 1)] \\
&= \frac{1}{2} + \frac{1}{2} \rho(\lambda) \implies \\
\text{Adv}_{\mathcal{B}}^{\text{Exp}_{\text{DS}}^{\text{IND}}} &= \frac{1}{2} \rho(\lambda)
\end{aligned}$$

This advantage is non-negligible because ρ is non-negligible. And this contradicts our assumption that our scheme is secure in the sense of definition 5.5. Our claim follows.

Claim 5. \forall distinguishers \mathcal{D} , the advantage of \mathcal{D} in \mathcal{H}_k is 0

Recall that in \mathcal{H}_k we have replaced AltSign^k with $\text{AltSign}^{k-k}(vk, (m^*, t^*, \sigma), \{\}) = \text{AltSign}^0(vk, (m^*, t^*, \sigma), \{\}) = \sigma$ where $\sigma \leftarrow \text{Sign}(sk, m^*, t^*)$. Therefore in \mathcal{H}_k when computing any challenge queries the resulting computed signatures are

$$\sigma^0 \leftarrow \text{Sign}(sk, m^*, t^*) \quad (9)$$

$$\sigma^1 \leftarrow \text{Sign}(sk, m^*, t^*) \quad (10)$$

The distributions of σ^0 and σ^1 are thus identical and trivially no distinguisher can have an advantage in distinguishing between them.

D Lewko's Prime-Order HIBE Scheme

This is a description of Lewko's[27] prime order translation for an unbounded HIBE scheme. This scheme performs some operations over vectors of n -dimensional space, similar to Lewko's work we describe the scheme for $n = 10$.

Setup(1^λ) \rightarrow (pk, msk): The setup algorithm takes as input the security parameter 1^λ . A bilinear group G of sufficiently large prime order p is selected, where the bilinear map is denoted by $e : G \times G \rightarrow G_T$. The random dual orthonormal bases required for the scheme are also sampled as part of this algorithm $(\mathbb{D}, \mathbb{D}^*) \xleftarrow{R} \text{Dual}(\mathbb{Z}_p^n)$. Let $\mathbb{D} = \{\vec{d}_1, \dots, \vec{d}_n\}$ and $\mathbb{D}^* = \{\vec{d}_1^*, \dots, \vec{d}_n^*\}$. It also chooses random exponents $\alpha_1, \alpha_2, \theta, \sigma, \gamma, \xi \in \mathbb{Z}_p$. The public parameters, which we describe as part of the public key are

$$pk = \left\{ G, p, e(g, g)^{\alpha_1 \vec{d}_1 \cdot \vec{d}_1^*}, e(g, g)^{\alpha_2 \vec{d}_2 \cdot \vec{d}_2^*}, g^{\vec{d}_1}, \dots, g^{\vec{d}_6} \right\}$$

and the master secret key is

$$msk = \left\{ G, p, \alpha_1, \alpha_2, g^{\vec{d}_1^*}, g^{\vec{d}_2^*}, g^{\gamma \vec{d}_1^*}, g^{\xi \vec{d}_2^*}, g^{\theta \vec{d}_3^*}, g^{\theta \vec{d}_4^*}, g^{\sigma \vec{d}_5^*}, g^{\sigma \vec{d}_6^*} \right\}$$

KeyGen($msk, (ID_1, \dots, ID_j)$) \rightarrow sk_{ID} : The key generation algorithm samples random values $r_1^i, r_2^i \xleftarrow{\$} \mathbb{Z}_p$ for each $i \in [j]$. It also samples random values $y_i \xleftarrow{\$} \mathbb{Z}_p$ and $w_i \xleftarrow{\$} \mathbb{Z}_p$ for $i \in [j]$ under the constraints that $y_1 + y_2 + \dots + y_j = \alpha_1$ and $w_1 + w_2 + \dots + w_j = \alpha_2$. For each $i \in [j]$, it computes:

$$K_i := g^{y_i \vec{d}_1^* + w_i \vec{d}_2^* + r_1^i ID_i \vec{d}_3^* - r_1^i \theta \vec{d}_4^* + r_2^i ID_i \sigma \vec{d}_5^* - r_2^i \sigma \vec{d}_6^*}$$

The secret key is output as:

$$sk_{ID} = \left\{ g^{\gamma \vec{d}_1^*}, g^{\xi \vec{d}_2^*}, g^{\theta \vec{d}_3^*}, g^{\theta \vec{d}_4^*}, g^{\sigma \vec{d}_5^*}, g^{\sigma \vec{d}_6^*}, K_1, \dots, K_j \right\}$$

$\text{Delegate}(sk_{ID}, ID_{j+1}) \rightarrow sk_{ID|ID_{j+1}}$: The delegation algorithm samples random values $\omega_1^i, \omega_2^i \xleftarrow{\$} \mathbb{Z}_p$ for each $i \in [j+1]$. It also samples random values $y'_i, w'_i \xleftarrow{\$} \mathbb{Z}_p$ for $i \in [j+1]$ subject to the constraint that $y'_1 + \dots + y'_{j+1} = 0 = w'_1 + \dots + w'_{j+1}$. Let $g^{\gamma \vec{d}_1^*}, g^{\xi \vec{d}_2^*}, g^{\theta \vec{d}_3^*}, g^{\theta \vec{d}_4^*}, g^{\sigma \vec{d}_5^*}, g^{\sigma \vec{d}_6^*}, K_1, \dots, K_j$ denote the elements of sk_{ID} . The delegated key $sk_{ID|ID_{j+1}}$ is formed as follows:

$$\begin{aligned} sk_{ID|ID_{j+1}} &:= \left\{ g^{\gamma \vec{d}_1^*}, g^{\xi \vec{d}_2^*}, g^{\theta \vec{d}_3^*}, g^{\theta \vec{d}_4^*}, g^{\sigma \vec{d}_5^*}, g^{\sigma \vec{d}_6^*} \right. \\ &K_1 \cdot g^{y'_1 \gamma \vec{d}_1^* + w'_1 \xi \vec{d}_2^* + \omega_1^1 ID_1 \theta \vec{d}_3^* - \omega_1^1 \theta \vec{d}_4^* + \omega_2^1 ID_1 \sigma \vec{d}_5^* - \omega_2^1 \sigma \vec{d}_6^*} \\ &\dots, K_j \cdot g^{y'_j \gamma \vec{d}_1^* + w'_j \xi \vec{d}_2^* + \omega_1^j ID_1 \theta \vec{d}_3^* - \omega_1^j \theta \vec{d}_4^* + \omega_2^j ID_1 \sigma \vec{d}_5^* - \omega_2^j \sigma \vec{d}_6^*} \\ &g^{y'_{j+1} \gamma \vec{d}_1^* + w'_{j+1} \xi \vec{d}_2^* + \omega_1^{j+1} ID_{j+1} \theta \vec{d}_3^*} \\ &\left. \cdot g^{-\omega_1^{j+1} \theta \vec{d}_4^* + \omega_2^{j+1} ID_{j+1} \sigma \vec{d}_5^* - \omega_2^{j+1} \sigma \vec{d}_6^*} \right\} \end{aligned}$$

$\text{Encrypt}(pk, M, ID) \rightarrow ct$: The encryption algorithm samples random values t_1^i, t_2^i for each $i \in [j]$, as well as random values $s_1, s_2 \xleftarrow{\$} \mathbb{Z}_p$. It computes

$$C_0 := Me(g, g)^{\alpha_1 s_1 \vec{d}_1 \cdot \vec{d}_1^*} e(g, g)^{\alpha_2 s_2 \vec{d}_2 \cdot \vec{d}_2^*}$$

and

$$C_i := g^{s_1 \vec{d}_1 + s_2 \vec{d}_2 + t_1^i \vec{d}_3 + ID_i t_1^i \vec{d}_4 + t_2^i \vec{d}_5 + ID_i t_2^i \vec{d}_6}$$

for each $i \in [j]$. The ciphertext is $ct := \{C_0, C_1, \dots, C_j\}$

$\text{Decrypt}(ct, sk_{ID}) \rightarrow M$: The decryption algorithm computes

$$B := \prod_{i=1}^j e(C_i, K_i)$$

and computes the message as:

$$M = C_0/B$$

D.1 Proof of Key-Indistinguishability

Theorem D.1. *The prime order variant of Lewko's scheme from [27] satisfies the key-indistinguishability property.*

The delegation algorithm (in Appendix D) in Lewko's prime order scheme [27] re-randomizes each exponent in a secret key. Each group element (the ones which are unique for an identity) in the key generated by $\text{KeyGen}()$, is of the form:

$$K_i := g^{y_i \vec{d}_1^* + w_i \vec{d}_2^* + r_1^i ID_i \theta \vec{d}_3^* - r_1^i \theta \vec{d}_4^* + r_2^i ID_i \sigma \vec{d}_5^* - r_2^i \sigma \vec{d}_6^*}$$

Where the values $y_i \xleftarrow{\$} \mathbb{Z}_p$ and $w_i \xleftarrow{\$} \mathbb{Z}_p$ for $i \in [j+1]$ are under the constraints that $y_1 + y_2 + \dots + y_{j+1} = \alpha_1$ and $w_1 + w_2 + \dots + w_{j+1} = \alpha_2$.

Whereas, in the key generated by $\text{Delegate}()$, each group element (the ones which are unique for an identity) is of the form:

$$K'_i = K_i^* \cdot g^{y'_i \gamma \vec{d}_1^* + w'_i \xi \vec{d}_2^* + \omega_1^i ID_1 \theta \vec{d}_3^* - \omega_1^i \theta \vec{d}_4^* + \omega_2^i ID_1 \sigma \vec{d}_5^* - \omega_2^i \sigma \vec{d}_6^*}$$

Where K_i^* is the i 'th group element of the key of the identity which was used as an input for `Delegate()`. The following variables are sampled uniformly at random, $\omega_1^i, \omega_2^i \xleftarrow{\$} \mathbb{Z}_p$ for each $i \in [j+1]$. It also samples random values $y'_i, w'_i \xleftarrow{\$} \mathbb{Z}_p$ for $i \in [j+1]$ subject to the constraint that $y'_1 + \dots + y'_{j+1} = 0 = w'_1 + \dots + w'_{j+1}$.

The key fact to note is that the exponent of g in K'_i the variables $y_1^* + \gamma y'_1, \dots, y_j^* + \gamma y'_j, \gamma y'_{j+1}$ are *randomly distributed* up to the constraint that their sum is α_1 , and similarly $w_1 + \xi w'_1, \dots, w_j + \xi w'_j, \xi w'_{j+1}$ are *randomly distributed* up to the constraint that their sum is α_2 . Also, $r_1^i + \omega_1^i$ and $r_2^i + \omega_2^i$ are uniformly random for each i . The keys generated via `KeyGen()` are also sampled from the same distributions with the exact same constraints. This gives us the fact that the distribution of a secret key obtained through any sequence of delegations is the same as the distribution of a secret key for the same identity generated via `KeyGen()` making them statistically indistinguishable. In fact, this is noted by Lewko in the description of the scheme as well. Moreover, the fact that the adversary has the master secret key msk , doesn't give it any advantage because the two keys generated only differ in the randomness used to generate them, having the msk doesn't give the adversary any way to distinguish between these two because they are statistically indistinguishable. Therefore, the challenge key pairs $(sk_\beta, sk_{1-\beta})$ for the HIBE key-indistinguishability game $\mathbf{Exp}_{\text{HIBE}}^{\text{IND}}$ are indistinguishable to any PPT adversary.

E Analysis of Signature Size for various values of N

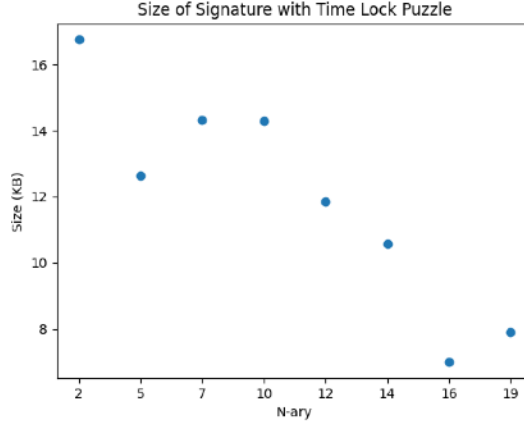


Figure 5: The signature size - including the encrypted functional key in the time lock puzzle - for a time-deniable signature scheme using the Gentry-Silverberg HIBE with BN-254 and varying values of N

F On the Necessity of Time-Lock Encryption

Definition F.1 (Time-Lock Encryption). A time-lock encryption scheme for computational reference clock $\mathcal{C}(1^\lambda)$, consists of the following two PPT algorithms:

- $\text{TL.Encrypt}(1^\lambda, m, \tau) \rightarrow ct$: The encryption algorithm takes as input the message, a security parameter, and an integer to represent time.
- $\text{TL.Decrypt}(ct, w) \rightarrow m / \perp$: The decryption algorithm takes as input a ciphertext, and a witness and outputs either a message m or the symbol \perp .

It has the following **correctness** property:

$$\Pr \left[\begin{array}{l} c \stackrel{\$}{\leftarrow} \text{TL.Encrypt}(1^\lambda, \tau_{\text{dec}}, m); \\ w_\tau \stackrel{\$}{\leftarrow} \mathcal{C}(1^\lambda, \tau); \\ m' := \text{TL.Decrypt}(w_\tau, c); \\ m = m' \end{array} \right] = 1$$

Security. We say an *admissible* polysize adversaries $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}} = \{(\mathcal{A}_{\lambda,0}, \mathcal{A}_{\lambda,1})\}_{\lambda \in \mathbb{N}}$, $\mathcal{A}_\lambda = (\mathcal{A}_0, \mathcal{A}_1)$, (t, τ, ε) -breaks time-lock encryption scheme $(\text{TL.Encrypt}, \text{TL.Decrypt})$ for a clock \mathcal{C} , if $\text{size}(\mathcal{A}) \in \text{poly}(\lambda)$ and $\tau \in \mathbb{N}$ such that

$$\Pr \left[\begin{array}{l} (m_0, m_1, st) \stackrel{\$}{\leftarrow} \mathcal{A}_0^{\mathcal{C}}(1^\lambda, \tau); \quad b \stackrel{\$}{\leftarrow} \{0, 1\}; \\ c \leftarrow \text{TL.Encrypt}(1^\lambda, \tau, m_b); \quad b' \leftarrow \mathcal{A}_1^{\mathcal{C}}(1^\lambda, c, st); \\ b = b' \end{array} \right] - 1/2 > \varepsilon$$

Where admissibility of an adversary is defined similar to subsection 5.1.

Witness encryption schemes allow encrypting a message m to an instance x of an NP language and allows decryption using a valid witness w such that $(x, w) \in R$.

Definition F.2 (Extractable Witness Encryption). A witness encryption scheme for an NP relation R consists of the following polynomial time algorithms:

- $\text{WE.Encrypt}(1^\lambda, x, m) \rightarrow ct$: This algorithm takes as input the security parameter 1^λ , an unbounded string x which represents an instance of R and a message m . It outputs a ciphertext ct which encrypts message m .
- $\text{WE.Decrypt}(ct, w) \rightarrow m$: This algorithm takes as input a ciphertext ct and a witness w and outputs a message m or the symbol \perp .

An extractable WE scheme has the following properties:

- **Correctness.** For all $(x, w) \in R$ and every message m :

$$\Pr \left[\text{WE.Decrypt}(\text{WE.Encrypt}(1^\lambda, x, m), w) = m \right] = 1 - \text{negl}(\lambda)$$

- **Extractable Security.** Let $(\text{WE.Encrypt}, \text{WE.Decrypt})$ be a witness encryption scheme for an NP relation R , such a scheme is secure if for all PPT adversaries \mathcal{A} and all polynomials q , there exists a PPT extractor E and a polynomial p , such that for all auxiliary inputs z and for all $x \in \{0, 1\}^*$:

$$\Pr \left[\begin{array}{l} b \xleftarrow{\$} \{0, 1\}; ct \leftarrow \text{WE.Encrypt}(1^\lambda, x, m) : \\ \mathcal{A}(x, ct, z) = m \end{array} \right] \geq 1/2 + 1/q(|x|)$$

$$\implies \Pr [E(x, z) = w : (x, w) \in R] \geq 1/p(|x|)$$

Time-Lock Encryption from Time-Deniable Signatures. To construct a time-lock encryption scheme using deniable signatures and extractable witness encryption, consider $(vk, sk) \leftarrow \text{DS.KeyGen}(1^\lambda, T(\lambda))$ and $\sigma \leftarrow \text{DS.Sign}(sk, m, t)$. Now consider a witness encryption scheme which encrypts to statements of the form $x = (m, t, \sigma, vk)$ for a relation R where for witnesses of the forms $w = (m^*, t^*, \sigma^*)$, $(x, w) \in R$ if $\text{DS.Verify}(vk, \sigma^*, m^*, t^*) = 1$. We also provide the intuition behind why this scheme is secure. The time-lock encryption algorithms proceed as follows

1. $\text{TL.Encrypt}(1^\lambda, m, t + T)$: It outputs $ct \leftarrow \text{WE.Encrypt}(1^\lambda, x, m)$, where $x = (m, t, \sigma, vk)$.
2. $\text{TL.Decrypt}(ct, w)$: Since it takes time T to create σ^* from σ , after time T a valid witness is available to run the decryption algorithm for WE with witness (m^*, t^*, σ^*) . Output $m' \leftarrow \text{WE.Decrypt}(ct, w)$.

The intuition behind the security argument is essentially that no admissible adversary should be able to distinguish an encryption of m_0 from an encryption of m_1 as this adversary is depth bounded. Otherwise, such an adversary computes $w \in R$, i.e., a different signature σ^* on some message, timestamp pair (m^*, t^*) by performing significantly less operations than the number of operations required. This adversary is solving the time-lock puzzle in sequential time less than T . Given such a distinguishing adversary we can leverage the extractor for witness encryption to break the unforgeability property for deniable signatures.

G On the Necessity of Secure Timestamps

Recall that in our definition, the AltSign algorithm takes as input a previously computed valid signature (or forgery). In particular, our notion does not rely upon the use of cryptographic timestamps.

An alternative notion discussed in Section 2 is one where AltSign does not require a previously computed signature as input; instead it only uses a timestamp issued by an external server to create a forgery. We argue that in the latter case, the timestamps issued by the server must be cryptographic (and in particular, unpredictable or unforgeable, depending on the implementation).

Suppose this is not the case. Then we can devise a simple attack using the AltSign algorithm to break the unforgeability of the signature scheme. Consider a (non-uniform) adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ that wants to generate a forged signature for *any* message m^* , and *any* time-stamp t^* . Since we allow for arbitrary polynomial time pre-processing in the unforgeability game, \mathcal{A}_0 runs AltSign on input m^* and $g(t^*)$ to compute a forged signature, where $g(\cdot)$ computes the output of the time server for time t^* (this also captures the scenario that the time stamp is entirely ignored by $\text{Sign}/\text{AltSign}$). Since there is no security property associated with the timestamps issued by the server, $g(\cdot)$ is a function that can be computed efficiently, so \mathcal{A}_0 is polynomial time.

Let σ^* be the forged signature computed by \mathcal{A}_0 , who passes it along to \mathcal{A}_1 to output as its forgery. Since the above strategy works for *any* (m^*, t^*) , and \mathcal{A}_1 needs only a single computational

step (to output the forged signature received from \mathcal{A}_0), this attack constitutes a valid forgery of the time-deniable signature scheme.

H Epochal Signatures

We recall the unforgeability game and definitions from the work [25] by Hülsing and Weber. An epochal signature scheme ES has the following four algorithms:

ES.KeyGen($1^\lambda, \Delta t, E, V$) $\rightarrow (pk, sk)$: Takes as input a security-parameter 1^λ , an epoch-length Δt , the maximum number of epochs $E \in \text{poly}(\lambda)$ and the number of epochs $V < E$ for which the signatures are valid and generates the long-term key pair (pk, sk) .

ES.Evolve(sk) $\rightarrow (pinfo_e, sk' / \perp)$: Takes as input the long-term secret key sk and returns the public epoch information $pinfo_e$ and an updated secret key sk' or \perp if sk has already been evolved E times.

ES.Sign(sk_e, m) $\rightarrow \sigma$: Takes as input a secret key sk_e and a message m and outputs a signature σ for the corresponding epoch e .

ES.Verify(pk, e, σ, m) $\rightarrow b$: Takes as input the public key pk , an epoch e , a signature σ and a message m and returns a bit b .

The unforgeability definition is defined with respect to the unforgeability game $\mathbf{Exp}_{\text{ES}}^{\text{UNF}}$ defined as follows:

Setup. The challenger runs $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda, \Delta t, E, V)$ and gives the public key pk to the adversary \mathcal{A} . It sets t_0 to the current time, sets $e = 0$ and initializes the set of queries $q = [\emptyset, \dots, \emptyset]$, where the k 'th entry in the set corresponds to the set of queries asked in epoch k .

Query Phase. In this phase, the adversary gets to query two different oracles.

1. The key evolution oracle $\text{O}_{\text{Evolve}}(\cdot)$ takes as input some wall clock time t , checks that $t \geq t_0 + e \cdot \Delta t$ which indicates that the current time is the e 'th epoch. It computes $(pinfo_e, sk_e) \leftarrow \text{Evolve}(sk)$ if t satisfies the above properties. At the end of an $\text{O}_{\text{Evolve}}(sk, \cdot)$ query, the adversary also gets access to the corresponding sign oracle for epoch e , $\text{O}_{\text{Sign}}(sk_e, \cdot)$.
2. The sign oracle $\text{O}_{\text{Sign}}(sk_e, \cdot, \cdot)$ takes as input a secret key sk_e , wall-clock time t and a message m . If $t < t_0 + e \cdot \Delta t$, it outputs a signature $\sigma \leftarrow \text{Sign}(sk_e, m)$ on the message and updates the corresponding epoch in the query set $q[e] \cup \{m\}$.

Forge. The adversary outputs its forgery (σ', m') and wins (game outputs 1) if:

1. For the corresponding epoch e' , there is no query corresponding to this message $(m', e') \notin q[e']$.
2. $\text{Verify}(pk, e', \sigma', m')$ outputs 1.

The advantage of the adversary \mathcal{A} is defined as $\mathbf{Adv}_{\mathcal{A}}(1^\lambda, \Delta t) = \Pr[\mathbf{Exp}_{\text{ES}}^{\text{UNF}}(1^\lambda, \Delta t, E, V) = 1]$.

Remark: To the best of our knowledge the authors do not define the function $\text{now}()$. We assume that this is the current time value and hence implies the existence of a wall clock.

Definition H.1 (Unforgeability of Epochal Signatures). An epochal signature scheme Σ is unforgeable if there is no efficient adversary that has a non-negligible chance of winning the unforgeability game $\mathbf{Exp}_{\text{ES}}^{\text{UNF}}$:

$$\forall \mathcal{A} \in \text{PPT}, \lambda \in \mathbb{N}, E \in \text{poly}(\lambda), V \in \{1, \dots, E - 1\} : \\ \mathbf{Adv}_{\mathcal{A}}(1^\lambda, \Delta t) \leq \text{negl}(\lambda)$$

Because the time-lock puzzle security definition in their work is different from ours, we give the indistinguishability game below and denote it as $\mathbf{Exp}_{\text{TL}}^{\text{IND}}$. It is assumed in this game that the challenger has access to a wall-clock.

Setup. The challenger sends the security parameter 1^λ and the difficulty parameter Δt to the adversary.

Challenge. The adversary \mathcal{A} picks the challenge messages (m_0, m_1) and sends them to the challenger. The challenger performs the following operations:

- Picks a random bit $b \xleftarrow{\$} \{0, 1\}$.
- Computes $c \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t, m_b)$, sends c to \mathcal{A} and sets current time as t_0 .

Response. \mathcal{A} sends its guess b' for b to the challenger and wins (game outputs 1) if the following conditions are satisfied:

- $b = b'$.
- Let the wall-clock time for when the challenger receives \mathcal{A} 's guess be t_1 , then it should be true that $t_1 - t_0 < \Delta t$.

The advantage of the adversary \mathcal{A} is defined as $\mathbf{Adv}_{\mathcal{A}}(1^\lambda, \Delta t) = \Pr[\mathbf{Exp}_{\text{TL}}^{\text{IND}}(1^\lambda, \Delta t) = 1]$.

Definition H.2 (TimeLock security according to [25]). A time-lock puzzle is secure if $\forall \mathcal{A} \in \text{poly}(\lambda)$, $\mathbf{Adv}_{\mathcal{A}}(1^\lambda, \Delta t) \leq \text{negl}(\lambda)$.

H.1 Faulty Epochal Signature Construction

Given a secure epochal signature construction Σ , we use it to construct another secure epochal signature scheme Σ' which has undesirable properties as discussed in Section 3. However, due to certain properties implicit in the definition of an epochal signature scheme, the scheme we present as the counter example is fairly intricate. We begin with an informal description of the counter example, which suffices for a relaxed notion of epochal signatures. We build upon this to present our final scheme Σ' - we formally argue that Σ' (i) is a secure epochal signature scheme; and (ii) is not a time-deniable signature scheme.

We first consider a counter example that satisfies a weaker but still reasonable notion of deniability where the judge never gets access to any secret key material. In this setting, our counter

<p><u>KeyGen</u>($1^\lambda, \Delta t, E, V$) :</p> <p>$pk, sk \leftarrow \Sigma.\text{KeyGen}(1^\lambda, \Delta t, E, V)$ return (pk, sk)</p> <p><u>Evolve</u>(sk'_e) :</p> <p>if $e == 0$: $sk_e \leftarrow sk'_e$ return $\Sigma.\text{Evolve}(sk_e)$</p> <p>$\begin{pmatrix} sk_e, r_0, \\ r_1, r_2, r_3 \end{pmatrix} \leftarrow \text{parse}(sk'_e)$</p> <p>$\begin{pmatrix} pinfo_{e+1}, \\ sk_{e+1} \end{pmatrix} \leftarrow \Sigma.\text{Evolve}(sk_e)$</p> <p>$r_0, r_1, r_2, r_3 \leftarrow \mathcal{R}$ return $pinfo_{e+1}, (sk_{e+1}, r_0, r_1, r_2, r_3)$</p> <p><u>Verify</u>($pk, e, \sigma', m$) :</p> <p>if $e == 0$: $\sigma \leftarrow \sigma'$ else: $\sigma, -, -, - \leftarrow \sigma'$ return $\Sigma.\text{Verify}(pk, e, \sigma, m)$</p>	<p><u>Sign</u>(sk'_e, m) :</p> <p>if $e == 0$: $sk_e \leftarrow \text{parse}(sk'_e)$ return $\Sigma.\text{Sign}(sk_e, m)$</p> <p>$\begin{pmatrix} sk_e, r_0, \\ r_1, r_2, r_3 \end{pmatrix} \leftarrow \text{parse}(sk'_e)$</p> <p>$\sigma \leftarrow \Sigma.\text{Sign}(sk_e, m)$ $c_{sk} \leftarrow sk_e \oplus (r_2 \oplus r_3)$ $c_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r_0)$ $c'_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r_1)$</p> <p>if $m == r_0$: $r' \leftarrow r_2$ else if $m == r_1$: $r' \leftarrow r_3$ else: $r' \leftarrow \mathcal{R}$ return $\sigma, c_{tlp}, c'_{tlp}, c_{sk}, r'$</p>
---	---

Figure 6: Our counterexample ES construction Σ' built from another secure epochal signature scheme Σ . The domain of \mathcal{R} is $\{0, 1\}^\lambda$ and $\forall i \in \{0, 1, 2, 3\}, r_i \in \mathcal{R}$. ϵ can be any value that is not 0, but we are especially interested in *small* ϵ .

example is fairly simple: each epoch e has a special trigger message m_e^* associated with it. If a signature ever needs to be made on m_e^* , in epoch e , then the signature contains the master secret key msk of the scheme. Included with every signature is a time lock puzzle that holds the special trigger message m_e^* , where the difficulty parameter on the puzzle is slightly more than Δt . It is straightforward to see that this scheme is still secure under the ES unforgeability game: the epoch e will always expire by the time the adversary could attempt to use m_e^* by solving the time lock puzzle. However, in the unforgeability game of time-deniable signatures, such a scheme is trivially defeated by an adversary in the pre-processing phase since the time lock puzzle can be solved during this phase.

The main problem with this counter-example is that the construction does not satisfy *perfect deniability*. Perfect deniability requires that one can simulate a *single* signature perfectly without revealing whether or not the signature was a forgery. Specifically in our counter example, we must ensure that one reply can never give away msk . To accomplish this, we encrypt msk under a key that is (2,2) secret shared. In order to recover the key, the adversary must make *two* queries which is explicitly disallowed in the deniability definition of epochal signatures. This ensures that only one share of the key is ever recovered, and thus we can simulate this share correctly without knowledge of msk . The final construction is presented in Figure 6. We would like to emphasize that this counter example is meant to show weaknesses in the unforgeability definition and that almost all of the complexity is added because of the deniability definition.

Theorem H.1. *The epochal signature construction in figure 6 is unforgeable under the security game of definition H.1*

Proof. We prove this by a standard hybrid argument starting with the real unforgeability game of \mathcal{H}_0 . For any hybrid \mathcal{H}_i , the output of \mathcal{H}_i is considered to be the challenger's output bit i.e. the adversary's success probability. This is related but not necessarily equal to $\mathbf{Adv}_{\mathcal{A}}^{\mathcal{H}_i}$. Therefore, the statement $\mathcal{H}_i \approx \mathcal{H}_{i+1}$ means that the probability of \mathcal{A} succeeding between two different hybrid games is negligible. For security, we want $\mathbf{Adv}_{\mathcal{A}}^{\mathcal{H}_0}(\lambda) \leq \text{negl}(\lambda)$ where $\text{negl}(\lambda)$ is a negligible function in the security parameter λ . Where details are omitted from the description of hybrid \mathcal{H}_i , it is assumed that they are the same as in \mathcal{H}_{i-1} .

\mathcal{H}_1 : Guess the challenge epoch e^* of \mathcal{A} by uniformly choosing an epoch in $\{1, \dots, E\}$ where E is the maximum number of epochs. If the guess is incorrect, the challenger's output is 0.

Analysis:

$$\mathbf{Adv}_{\mathcal{A}}^{\mathcal{H}_1}(\lambda) = \frac{1}{E} \cdot \mathbf{Adv}_{\mathcal{A}}^{\mathcal{H}_0}(\lambda)$$

□

Let r_0, r_1, r_2, r_3 be the ephemeral randomness used in an epoch e . Consider the following hybrid,

\mathcal{H}_3 : If $\exists e \in \{1, \dots, e^*\}$ such that \mathcal{A} queries the sign oracle on messages r_0 or r_1 , the challenger outputs 0.

Let F be the event that on at least one epoch $e \in \{1, \dots, e^*\}$, \mathcal{A} queries on one or both of r_0 and r_1 . If $\Pr[F] \leq \text{negl}(\lambda)$, then $\mathcal{H}_1 \approx \mathcal{H}_3$.

Lemma H.1. $\Pr[F] \leq \text{negl}(\lambda)$

We first provide a sequence of hybrids $\mathcal{H}_{i,j,k}^2$ where $i \in \{1, \dots, e^*\}, j \in \mathbb{N}, k \in \{0, 1\}$.

$\mathcal{H}_{i,j,k}^2$: Let r_0, r_1, r_2, r_3 be the ephemeral randomness used in epoch i . If j is not 0, on the j^{th} sign query of \mathcal{A} the time lock puzzle that normally holds r_k instead encrypts another random value \hat{r}_k where $\hat{r}_k \neq r_k$.

Note that $\mathcal{H}_1 = \mathcal{H}_{1,0,k}^2$ for any k . For a given epoch i , we will be concerned with the possibly infinite sequence of hybrids $(\mathcal{H}_{i,0,0}^2, \mathcal{H}_{i,0,1}^2, \mathcal{H}_{i,1,0}^2, \mathcal{H}_{i,1,1}^2, \dots)$. If this sequence is actually *finite* with the last query of \mathcal{A} in i being $r(\lambda)$, then $\mathcal{H}_{i,r(\lambda),1}^2 = \mathcal{H}_{i+1,0,0}^2$.

Claim 6. $\forall i \in \{1, \dots, e^*\}, \forall j \in \{1, \dots, \}, \mathcal{H}_{i,j-1,1}^2 \approx \mathcal{H}_{i,j,0}^2$

Suppose this is not true and there exists a distinguisher D with non-negligible advantage that distinguishes between \mathcal{A} 's success probability in $\mathcal{H}_{i,j-1,1}^2$ and $\mathcal{H}_{i,j,0}^2$ that outputs 0 when it thinks \mathcal{A} 's success is from $\mathcal{H}_{i,j-1,1}^2$ and 1 otherwise. Consider the following TimeLock adversary \mathcal{B} which uses \mathcal{A} and D to break the security of the TimeLock.

Description of TimeLock adversary \mathcal{B} :

- Execute $\Sigma.\text{KeyGen}(1^\lambda, \Delta t, E, V)$ to get pk, sk . Give pk to \mathcal{A}
- On queries in epochs $e > i$, respond as normal.
- For the j^{th} query in i , sample m uniformly from \mathcal{R} subject to the constraint that m is not equal to r_0 . Send to the TimeLock challenger (r_0, m) . Receive s . Construct $\sigma, c'_{tlp}, c_{sk}, r'$ as normal and send to \mathcal{A} the signature $\sigma, s, c'_{tlp}, c_{sk}, r'$.
- Let b' be the output of \mathcal{A} . Send $b = b'$ to D . If D outputs \hat{b} respond with \hat{b} .

Analysis:

The probability that \mathcal{B} wins is equivalent to the probability of D distinguishing correctly. Since D by assumption has non-negligible advantage so does \mathcal{B} , in contradiction to the security of the TimeLock.

Claim 7. $\forall i \in \{1, \dots, e^*\}, \forall j \in \mathbb{N}, \mathcal{H}_{i,j,0}^2 \approx \mathcal{H}_{i,j,1}^2$

Proof sketch. The argument here is equivalent to the previous one, except instead of \mathcal{B} challenging on (r_0, m) the challenge is (r_1, m) . \square

In order to be a successful adversary, \mathcal{A} must run in polynomial time, which means that the number of queries that \mathcal{A} can make in any particular epoch is also bound by some polynomial. Therefore, $\forall i$ the sequences $(\mathcal{H}_{i,0,0}^2, \mathcal{H}_{i,0,1}^2, \dots)$ must also be bounded by $\text{poly}(\lambda)$. Since $e^* \leq E \in \text{poly}(\lambda)$ is also polynomially bounded by the security parameter and the last hybrid in sequence $i - 1$ is actually the first hybrid in i , there is a negligible difference between $\mathcal{H}_{1,0,0}^2$ and the last hybrid in the sequence beginning with $\mathcal{H}_{e^*,0,1}^2$.

The probability that \mathcal{A} asks for either appropriate trigger in any of the epochs $e = 1 \dots e^*$ can now be calculated by multiple union bounds. For an epoch i the probability that \mathcal{A} guesses either r_0 or r_1 is $\frac{2}{|\mathcal{M}|}$ where \mathcal{M} is the space of all possible messages that can be signed. If $\mathcal{M} = \{0, 1\}^\lambda$ the probability this happens in any epoch before or at the challenge epoch is upper bounded by

$\frac{2 \cdot e^*}{2^\lambda}$ which is negligible. Let $e^* = r(\lambda)$ and $q(\lambda)$ be the maximum number of queries \mathcal{A} asks for in any epoch. Then we have

$$\begin{aligned} \Pr[F] &= q(\lambda) \cdot e^* \cdot \mathbf{Adv}_{\mathcal{B}}^{\text{Exp}_{\text{TL}}^{\text{IND}}}(\lambda) + \frac{2e^*}{2^\lambda} \\ &= q(\lambda)r(\lambda)\mathbf{Adv}_{\mathcal{B}}^{\text{Exp}_{\text{TL}}^{\text{IND}}}(\lambda) + \frac{2r(\lambda)}{2^\lambda} \end{aligned}$$

By assumption $\mathbf{Adv}_{\mathcal{B}}^{\text{Exp}_{\text{TL}}^{\text{IND}}}(\lambda)$ is negligible and $\frac{1}{2^\lambda}$ is clearly negligible therefore $\Pr[F]$ is as well. This completes our proof of the lemma.

Recall that in \mathcal{H}_3 we know that \mathcal{A} will not ask for a query on the trigger message r^* during the challenge epoch e^* or any earlier epoch. We can now argue security by reducing to the security of the original epochal signature scheme ES' . Let \mathcal{B} be an adversary for the ES unforgeability game that is constructed from \mathcal{A} in \mathcal{H}_3 as described below.

Description of ES adversary \mathcal{B} :

- Receive pk from the challenger and pass pk on to \mathcal{A}
- For any evolve query asked by \mathcal{A} before or during e^* , make the same query to the challenger. Receive $pinfo_e$. Re-sample new randomness r_0, r_1, r_2, r_3 . Send $pinfo_e$ to \mathcal{A}
- For any sign query m for epoch $e \leq e^*$, make the same query to the challenger and receive σ . Forward σ to \mathcal{A} along with correctly structured values $c_{tlp}, c'_{tlp}, c_{sk}, r'$
- When \mathcal{A} gives forgery (m^*, σ^*) , parse out the first component $\bar{\sigma}$ and give the forgery $(m^*, \bar{\sigma})$ to the challenger

Analysis: It should be clear that in this case if \mathcal{A} succeeds \mathcal{B} must as well since \mathcal{A} never asks for a query on either r_0 or r_1 in the relevant epochs before the forgery and \mathcal{B} is merely making the same queries to its challenger that \mathcal{A} is. Therefore,

$$\mathbf{Adv}_{\mathcal{A}}^{\mathcal{H}_3}(\lambda) = \mathbf{Adv}_{\mathcal{B}}^{\text{Exp}_{\text{ES}}^{\text{UNF}}}(\lambda) \leq \text{negl}(\lambda).$$

Since the advantage of \mathcal{A} in \mathcal{H}_3 is negligible, the construction is unforgeable.

Theorem H.2. *The epochal signature construction in figure 6 is deniable according to the definition of [25].*

We first give a description of the simulator Sim for the deniability game. It makes use of $\Sigma.\text{Sim}$ which is the simulator for the deniability of Σ .

```

Sim( $m, e_0, \text{pinf}o_{e_0+e_1}$ ) :
   $r \leftarrow \mathcal{R}$ 
   $c_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r)$ 
   $c_{sk} \leftarrow \mathcal{R}$ 
   $r' \leftarrow \mathcal{R}$ 
   $\sigma \leftarrow \Sigma.\text{Sim}(m, e_0, \text{pinf}o_{e_0+e_1})$ 
  if  $e_0 == 0$  or  $e_0 == 1$  :
    return  $\sigma$ 
  return  $\sigma, c_{tlp}, c_{sk}, r'$ 

```

In order to prove that our scheme is deniable we need the distributions of $\sigma, c_{tlp}, c_{sk}, r'$ to be indistinguishable when generated via `Sign` or `Sim` when the distinguisher has access to $pk, sk_{e_0+e_1}, \text{pinf}o_{e_0+e_1}$.

We can safely ignore σ because σ is conditionally independent from the tuple (c_{tlp}, c_{sk}, r') given sk and it does not reveal information on what it was derived from because of the deniability of Σ . We can also restrict our analysis to $e > 1$, since when $e = 1$, σ is the only component of the signature.

This can be simplified to us needing the following to hold for all messages m , for all valid epochs $e_0 \neq 0$ and $e_0 \neq 1$, for all valid $e_1, \forall V$, and all key generation outputs $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda, \Delta t, E, V)$:

$$\left\{ \begin{array}{l} r_0, r_1, r_2, r_3 \leftarrow \mathcal{R}, \\ (-, sk_{e_0}) \leftarrow \text{Evolve}^{e_0}(sk), \\ (pi_{e_0+e_1}, sk_{e_0+e_1}) \leftarrow \text{Evolve}^{e_1}(sk_{e_0}), \quad c_{tlp}, c'_{tlp}, c_{sk}, r', \\ c_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r_0), \quad pk, sk_{e_0+e_1}, pi_{e_0+e_1} \\ c'_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r_1) \\ c_{sk} \leftarrow (r_2 \oplus r_3) \oplus sk_{e_0}, r' \leftarrow \mathcal{R}^{***} \end{array} \right\} \approx$$

$$\left\{ \begin{array}{l} r_0, r_1 \leftarrow \mathcal{R}, \\ (pi_{e_0+e_1}, sk_{e_0+e_1}) \leftarrow \text{Evolve}^{e_0+e_1}(sk) \\ c_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r_0), \quad c_{tlp}, c'_{tlp}, c_{sk}, r', \\ c'_{tlp} \leftarrow \text{TimeLock.Gen}(1^\lambda, \Delta t + \epsilon, r_1) \quad pk, sk_{e_0+e_1}, pi_{e_0+e_1} \\ c_{sk} \leftarrow \mathcal{R}, r' \leftarrow \mathcal{R} \end{array} \right\}$$

The asterisks on the left distribution for r' denote that the value of r' is mostly random, except when $m = r_0$ or $m = r_1$. In that case, $r' = r_2$ or $r' = r_3$ respectively.

Since $V \geq 1$, the updated key $sk_{e_0+e_1}$ does not contain the randomness used in epoch e_0 . However, both “trigger” messages, are made available to the judge by breaking the time lock puzzles c_{tlp} and c'_{tlp} . Luckily this, at most, gives the judge access to one of r_2 or r_3 : the judge only

<u>KeyGen($1^\lambda, \Delta t, E, V$) : Sign($sk_e, m$) :</u>	
$vk, sk \leftarrow$	$sk, e \leftarrow \text{parse}(sk_e)$
$\text{DS.KeyGen}(1^\lambda, \epsilon(\lambda) \cdot \sigma$	\leftarrow
$V \cdot \Delta t \cdot d^*)$	$\text{DS.Sign}(sk, m, e)$
return ($vk, (sk, 0)$)	return σ
<u>Evolve(sk') :</u>	<u>Verify(pk, e, σ, m) :</u>
$sk, e \leftarrow \text{parse}(sk')$	return
$e = e + 1$	$\text{DS.Verify}(pk, \sigma, m, e)$
$\sigma \leftarrow \text{DS.Sign}(sk, 0, e)$	
return ($\sigma, (sk, e)$)	

Figure 7: An ES construction from a time-deniable signature scheme DS. $\epsilon(\lambda)$ is the admissibility parameter for the time-deniable signature scheme. The sentinel message for *pinfo* is $m = 0$.

gets to see the output of one signing query so if the judge uses the message $m = r_0$ or $m = r_1$ it receives either r_2 or r_3 as r' . Recall that $c_{sk} = (r_2 \oplus r_3) \oplus sk_{e_0}$ whenever Sign is used. Because the judge can only get access to at most one share, c_{sk} is indistinguishable from a random element of \mathcal{R} . This also means that the random pad r' is independent of c_{sk} to the perspective of the judge, regardless of if it is supposed to be a share of the one time pad or not. Therefore the joint distribution of (c_{sk}, c_{tlp}, r') in the left-hand side of the equation is one where each value is independent from the others, c_{sk} and r' are uniform, and c_{tlp} locks a uniformly random value. This is precisely what the right hand side is and thus completes our proof.

H.2 Time-Deniable Signatures as Epochal Signatures

In this section we show any secure DS scheme can be generically transformed into a secure ES scheme. At a high level, our construction is a simple transformation where verification and signing uses the time-deniable signature scheme and the evolve algorithm keeps track of the current epoch. *pinfo_e* contains a signature on a dummy, sentinel message at the timestamp corresponding to epoch e . Because of the different models of time considered by the two primitives, we do a translation between wall-clock time and circuit depth. We make the following assumption: for any circuit C that terminates in wall-clock time t , the depth of C when it terminates is $d_C \cdot t$ where d_C is a constant that depends on C . Let \mathbf{C} be the set of circuits for which some input x causes C to terminate before or at wall-clock time t and let $d^* = \max_{C \in \mathbf{C}} d_C$. d^* is needed to correctly set the time parameter given to the time-deniable signature scheme's KeyGen algorithm. The construction appears in Figure 7.

Theorem H.3. *The ES scheme presented in Figure 7 is unforgeable if the time-deniable signature scheme DS is unforgeable.*

We prove this by contradiction, supposing there exists an adversary \mathcal{B} who succeeds with non-negligible advantage in the ES unforgeability game and then using that adversary to construct an adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ for the DS unforgeability game. For the ES unforgeability definition we assume that before the counter e is given to the Verify algorithm it is appropriately advanced.

Description of \mathcal{A} :

- Receive vk from the challenger. Give vk to \mathcal{B} . Initialize a counter $ctr = 0$ and the wall clock time $t_{init} = \text{now}()$.
- On queries O_{Evolve} check if $t' \geq t_{init} + (e + 1)\Delta t$ where t' is the current wall clock time. If yes, set $ctr = ctr + 1$ and query O_{Sign} on message 0 to receive σ . Return σ as *pinfo*. Else do not advance the counter and output \perp .
- On queries O_{Sign} from \mathcal{B} with message m , query the challenger's O_{Sign} oracle with m and timestamp ctr . Return σ to \mathcal{B} .
- If \mathcal{B} returns signature σ^* in epoch e^* on message m^* , then \mathcal{A} returns (m^*, e^*, σ^*) as its forgery.

First, we argue that \mathcal{A} is an admissible adversary in the time-deniable signature unforgeability game. Let \mathcal{A}_0 denote the interactions of \mathcal{A} with \mathcal{B} before epoch e^* begins. Because for an ES scheme, $\text{size}(\mathcal{B}) \in \text{poly}(\lambda)$ and since \mathcal{A} is also doing $\text{poly}(\lambda)$ work while interacting with \mathcal{B} , we have that $\text{size}(\mathcal{A}_0) \in \text{poly}(\lambda)$. If we let the output of \mathcal{A} and \mathcal{B} after this interaction be an advice string z , we can then split off the rest of \mathcal{A} and \mathcal{B} 's interaction as \mathcal{A}_1 . For \mathcal{B} to be a successful adversary, they must be able to produce m^*, σ^* before wall clock time $V\Delta t$ has past since the start of epoch e^* . Then we have an upper bound on the circuit depth of \mathcal{B} from the start of e^* until termination as $d^*V\Delta t$. Since \mathcal{A} just forwards queries between the challenger and \mathcal{B} the overhead it adds is minimal (on the order of the number of queries made by \mathcal{B}) and can be ignored for the sake of this proof sketch. $\text{depth}(\mathcal{A}_1)$ is therefore appropriately bounded as $d^*\Delta t \cdot V \leq \frac{d^*\Delta t \cdot V \cdot \epsilon(\lambda)}{\epsilon(\lambda)}$.

We now argue that if \mathcal{B} is successful in its forgery so is \mathcal{A} . As said earlier, in order for \mathcal{B} to succeed it must produce a valid pair (m^*, σ^*) before the wall clock time bound where validity means that DS.Verify succeeds given the current timestamp is e^* and that \mathcal{B} has never asked for a signature on m^* at time e^* . The tuple (m^*, e^*, σ^*) is thus a valid forgery for \mathcal{A} as well.

Theorem H.4. *The ES scheme presented in Figure 7 is deniable if the time-deniable signature scheme DS is deniable.*

Suppose this is not true and there exists a judge \mathcal{J} that succeeds with non-negligible advantage in the ES deniability game. Then we will construct an adversary \mathcal{A} that succeeds with non-negligible advantage in the DS deniability game.

Description of \mathcal{A} :

- Receive vk, sk from the challenger. Forward sk to \mathcal{J} .
- Uniformly sample a random message m . When \mathcal{J} specifies its challenge (m^*, e_0, e_1) , query O_{Sign} with $(m, e_0 + e_1)$ and receive (id, σ) .
- Query O_{Ch} with id, m^*, e_0 to get σ^* . Send σ^* to \mathcal{J} . If \mathcal{J} responds with b send b to the challenger.

\mathcal{J} expects to see one of two signatures. One creates the signature by evolving sk e_0 times while the other evolves the key $e_0 + e_1$ times and uses $\text{pinfo}_{e_0+e_1}$. When the challenger's bit $b = 0$, \mathcal{A} produces the output of ES.Sign in Figure 7 which is simply DS.Sign . When $b = 1$, the output is produced by the simulator \mathcal{S} in the ES deniability game which is the equivalent to DS.AltSign in our construction. Therefore, the distributions \mathcal{J} sees are correct and if \mathcal{J} is successful in distinguishing then so is \mathcal{A} .