SEEM*less*: Secure End-to-End Encrypted Messaging with *less* Trust

Melissa Chase
Microsoft Research
melissac@microsoft.com

Apoorvaa Deshpande
Brown University
apoorvaa_deshpande@brown.edu

Esha Ghosh
Microsoft Research
esha.ghosh@microsoft.com

Harjasleen Malvai
Cornell University
hm553@cornell.edu

ABSTRACT

End-to-end encrypted messaging (E2E) is only secure if participants have a way to retrieve the correct public key for the desired recipient. However, to make these systems usable, users must be able to replace their keys (e.g. when they lose or reset their devices, or reinstall their app), and we cannot assume any cryptographic means of authenticating the new keys. In the current E2E systems, the service provider manages the directory of public keys of its registered users; this allows a compromised or coerced service provider to introduce their own keys and execute a man in the middle attack.

Building on the approach of CONIKS (Melara et al, USENIX Security '15), we formalize the notion of a *Privacy-Preserving Verifiable Key Directory* (VKD): a system which allows users to monitor the keys that the service is distributing on their behalf. We then propose a new VKD scheme which we call SEEMless, which improves on prior work in terms of privacy and scalability. In particular, our new approach allows key changes to take effect almost immediately; we show experimentally that our scheme easily supports delays less than a minute, in contrast to previous work which proposes a delay of one hour.

CCS CONCEPTS

• **Security and privacy** → **Key management**; **Privacy-preserving protocols**; **Privacy protections**; *Public key encryption*; *Pseudonymity, anonymity and untraceability*; Social network security and privacy;

KEYWORDS

Privacy-preserving verifiable directory service; Zero knowledge sets; Accumulators; Persistent Patricia Trie; History Tree; PKI; Transparency; Security definitions

1 INTRODUCTION

A number of popular messaging apps such as iMessage, WhatsApp and Signal have recently deployed end-to-end encryption (E2EE) in an attempt to mitigate some of the serious privacy concerns that arise in these services. E2EE is a system of communication where only encrypted messages leave the sender's device. These messages are then downloaded and decrypted on the recipient's device, ensuring that only the communicating users can read the messages. But E2EE relies on a Public Key Infrastructure (PKI); this in practice requires the provider of the messaging service (such as Apple, Facebook, Microsoft etc.) to maintain a centralized directory of the public keys of its registered users. To be accessible to average users, these systems assume the user will store her secret key on her personal device, and do not assume she has any other way to store long term secrets. When a user loses her device (and thus her secret key), she will need to generate a new (secret key, public key) pair and replace her old public key stored in the PKI with the newly generated public key.

Such a system naturally places a lot of trust in the service provider – a malicious service provider (or one who is compelled to act maliciously, possibly because of a compromise) can arbitrarily set and reset users' public keys. It might, for example replace an honest user's public key with one whose secret key it knows, and thus implement a man-in-the-middle attack without the communicating users ever noticing. Ironically, this defeats the purpose of E2EE. Without some way of verifying that the service provider is indeed returning the correct keys, E2E encryption does not provide any protection against malicious (or coerced) service providers. This problem has been well recognized as an important and challenging open problem [14, 20, 25].

Some service providers provide a security setting option to notify a sender when a recipient's public key changes. In WhatsApp, the sender can scan a QR code on the recipient's device to verify the authenticity of the new public key. Skype encrypted messaging provides a similar interface for checking fingerprints. This option, however, is turned off by default to provide a seamless user experience. Moreover, the communicating users will be able to verify each other's QR codes only if their devices are physically close to each other, which is most often not the case. And these features are something few users use, as is evidently uncovered in this Man-in-the-Middle attack [1].

To enable E2EE with real security, we need to keep the inherent constraints of the system in mind. To begin with, the primary objective of any E2EE messaging system is to provide a *secure and seamless* communication service between remote users. This problem is made even more challenging by the fact that we must assume that a user can lose her device and along with it all of her secrets. Moreover, in our attempt to reduce trust on the service provider, we must not introduce any new attack surface. For example, if we introduce a mechanism that will enable a user to verify that she is receiving the correct key of the intended recipient, this mechanism should not leak any additional information about the other users and the public keys they have registered with the service provider. Privacy may not be very important in a traditional PKI, where all the players are usually public entities like businesses, but in the context of private messaging, privacy is very important. Hiding usernames may help prevent spam messaging. And the user's update pattern may itself be sensitive information. Users change their keys primarily when they change devices, or detect that their devices/keys have been compromised, either of which may be sensitive. Moreover, if a user rarely changes her key, then compromising her secret key gives the attacker the ability to decrypt large volumes of messages at once, making her a more vulnerable and attractive target. Or if a device is compromised but the user does not update her key, then the attacker knows the compromise has gone undetected.

Keeping these inherent constraints in mind, we design and provide a prototype implementation of SEEMless, *a verifiable key directory service* for end-user key verification and initiate the study of the privacy and security guarantees of such a key directory service formally. In our design, we follow the basic approach introduced by by Melara et al. in [20], in which the service provider maintains the public key directory as described above, but each user is in charge of verifying that the public key that the service provider is presenting on their behalf is correct. (We call this key monitoring.) To facilitate this, the service provider will at regular intervals (called epochs) publish commitment of the latest directory. When responding to queries (either from Bob looking up Alice's key or from Alice monitoring her own key) it will then prove that the response is correct with respect to the current commitment.

Our major contributions are the following.

Contributions: Here we highlight the significant contributions of this work and contrast them with the existing privacy-preserving key directory services for encrypted messaging [3, 20, 26].

(A) Formalizing key directory

(1) We formalize the security and privacy requirements of a verifiable key directory service, (such CONIKS [20, 26], EthIKS [3]), in terms of a new primitive that we call Verifiable Key Directories (VKD) (Section 2). To the best of our knowledge, this is the first formal treatment of VKD, which is absent in all the prior works [3, 20, 26].

(2) Proving security of any VKD system (including our instantiation, SEEMless) in our framework requires specifying a precise leakage function. This leads to a better understanding of the privacy guarantee of any VKD system. As a concrete example, we were able to identify a *tracing attack* in existing VKD systems [3, 20, 26].

(B) New primitives, modular and principled design

(1) We take a modular and principled approach in designing our VKD system, SEEMless (Section 4). We first define a new primitive, *Append-Only Zero Knowledge Set* (aZKS) and construct SEEMless using aZKS in a blackbox way. Then we construct aZKS modularly from a *strong accumulator* [4] and a *verifiable random function* [6]. None of the prior works [3, 20, 26] has a modular approach in designing their systems which makes it hard to reason about those constructions and improve on them. With our modular design, we significantly simplify the presentation and analysis and allow for independent study of each of the building blocks: an improvement in any of them would directly translate to improvement in SEEMless.

(2) Our new primitive aZKS (that acts as the building block for SEEMless) generalizes the definition of a traditional static zero-knowledge set [5, 21] by accounting for verifiable updates and could be of independent interest.

(C) Persistent Patricia Trie

(1) SEEMless requires the server to maintain history of all the key updates (along with proofs) from its registered users. CONIKS [20, 26] has an identical requirement in order to support users who are not always online. A naive way of supporting this is to maintain one authentication data-structure (or aZKS) per server epoch. This is very space-consuming and wasteful since there might be a significant overlap between the aZKS of consecutive epochs. Motivated by this constraint, we develop a *Persistent Patricia Trie* (PPTr) data structure and use it to build a space-efficient persistent aZKS (that can also answer queries relative to older versions the aZKS efficiently). CONIKS took the naive approach of building a new authentication data-structure at every epoch, which resulted in serious scalability problem, as we uncovered in the experiments.

(2) Our PPTr construction extends history tree [23] and presents a way of maintaining any temporal, append-only, authenticated log of data. We believe PPTr will find applications in several contexts, (e.g. tamper-evident logging, distributed identity providers, distributed certificate authorities) and may be of independent interest.

(D) Privacy and Efficiency Improvements Compared to the prior work, we significantly improve on the following parameters. For experimental comparison, we implemented a prototype of our system, SEEMless in Java and compared it with CONIKS, which is also implemented in Java.

(1) Server storage: CONIKS [20, 26], uses a space-inefficient data-structure in order to retain the history of all user key updates, where the size of the data-structure is proportional to the number of server epochs. In contrast, thanks to our Persistent Patricia Trie construction, the size of our data-structure (aZKS) is proportional to the total number of key updates in the system as opposed to the number of epochs. The space overhead in EthIKS [3] is even more since the data-structure at the server is posted on the Ethereum blockchain, thereby replicating it at each network node.

(2) Key Monitoring cost: The cost that a user has to pay to monitor her key revision history at the server is an important metric in a VKD system, especially because the users are often on a weak device, like a mobile phone. We significantly improve on

the monitoring cost compared to [3, 20, 26]. CONIKS [20, 26] required that each user queries the server *every epoch* to ensure that her key was correct at that time; a user who was offline for a period would have to perform all the checks for that time on coming back online. EthIKS [3] lets a user audit only for the *latest version* of her key. In SEEMless, a user can monitor her key revisions *any time* with cost that depends on the *number of times her key has been updated* rather than the number of epochs that have passed. For example, she can be offline for several server epochs and can verify her entire key history by downloading around 2MB data (even if she is updating her key often) when she comes back online; in contrast, a CONIKS user will need to download about 577MB.

(3) Frequency of server updates: In a VKD system, when a user updates her key, the change is not reflected immediately; it is reflected in the server's commitment of the subsequent epoch. Therefore, until the next epoch, either the server can continue to respond with the stale user key, or the server can respond with a key which is not consistent with the published commitment and update it in the next epoch. CONIKS proposes this latter approach; it is unclear how the user can verify that this new key is in fact correct. In our solution, we propose making the epochs short enough that we can take the first approach without significantly affecting security or usability of the system. Thanks to our space-efficient design, we can afford to have short epochs. This makes SEEMless scale well with frequent server updates, as opposed to CONIKS. For the same update frequency and experimental setup, CONIKS cannot handle more than 4.5M users, owing to the inefficiency of their underlying data structure and hangs with more frequent epochs whereas our system could handle 10M users seamlessly. In EthIKS, the frequency of server updates is at least as slow as the block frequency of Ethereum.

(4) Privacy: SEEMless provides provably stronger privacy guarantees as compared to [3, 20, 26]. In fact, we identify some significant privacy leakage in [3, 20, 26] through which it might be possible to trace the entire update history of a particular user. This means once Alice queried for Bob's key, she might be able to completely trace when Bob's key changed without ever querying for his key again. Even if Bob has deleted Alice from his contact list (and the server does not give out Alice's key to Bob after he has removed Alice from his contact list), Alice will still be able to trace when Bob's key changed just by looking at the proof of her own key. We call this *tracing attack* and explain its technical details in Appendix C.

(5) User cost: In our experiments, the cost of lookup query in SEEMless is slightly more expensive than in CONIKS (at most 3×) due to our stronger privacy guarantees, but it is still about 10*ms* from request to complete verification, which is reasonably fast. In EthIKS each user lookup requires interaction with the Ethereum blockchain, which which introduces significant overhead compared to CONIKS and SEEMless.

(6) Auditors: We introduce the concept of auditors in a VKD — potentially untrusted parties who verify that the server's commitments are well formed and updated in a valid way. This is crucial to our more efficient key monitoring: essentially where in CONIKS the users do a lot of redundant work to verify that

each server update is correct with respect to their keys, we only need one honest party to check each server update. By design, this verification is not privacy-sensitive in SEEMless, so we can open this verification up to anyone, including privacy concerned users or privacy advocacy groups. See Section 2.1 for more discussion. In EthIKS the auditing is implicitly performed by Etherum smart contracts adding a huge overhead.

Organization of the paper In Section 2, we define the primitive of Verifiable Key Directories (VKD) along with the security properties. In Section 3, we define append-only Zero Knowledge Sets (aZKS) which is a building block for VKD construction. In Section 4, we describe our SEEMless construction starting with an overview of the same. In Section 5, we give concrete instantiations of aZKS. In Section 6 we describe our space efficient Persistent Patricia Trie construction. We present our prototype implementation and performance results in Section 7. In Section 8 we discuss related work.

2 VERIFIABLE KEY DIRECTORY (VKD)

In this section, we will define the primitive of a *Verifiable Key Directory* (VKD) and formalize its properties. The goal of a VKD is to capture the functionality and security of a privacy-preserving verifiable key directory system.

A VKD consists of three types of parties: an identity provider or server, clients or users and external auditors. The server stores a directory Dir with the names of the users (which we call labels) and their corresponding public keys (the values corresponding to the labels). For the ease of exposition, let Alice and Bob be two users. VKD provides the following query interface to the users. 1) Alice can add her (username, key), i.e., (label=username, val=key), to Dir. 2) Alice can update her key and request that Dir be updated with the new key value. 3) She can query the server periodically to obtain history of her key updates over time (VKD.KeyHistory). 4) Bob can also query for the key corresponding to username Alice (VKD.Query) at the current time.

The functionality VKD.KeyHistory warrants further discussion since this is not a functionality that one usually expects from a key directory service. In a privacy-preserving verifiable key directory service, intuitively, we expect the server to be able to prove to Bob that he is seeing Alice's latest key without leaking any additional information about the directory. This is trivial to achieve if we assume that Alice can always sign her new public key with her old secret key. But this is a completely unreasonable assumption from an average user who may lose her device or re-install the software, thereby losing her previous secret key; the user will only have access to her latest secret key which is stored on her latest device. It is crucial to *not assume* that Alice or Bob can remember any cryptographic secret. Under this constraint, we need Alice to monitor her key sufficiently often to make sure her latest key is in the server directory. Only then, we can talk about Bob getting Alice's latest key in a meaningful way. Alice could of course check every epoch to make sure that her key is being correctly reported, but this becomes costly, particularly when epochs are short. Instead, we allow Alice to query periodically and retrieve a list of all the times her key has changed and the resulting values using the VKD.KeyHistory interface.

The server applies updates from its users (of type 1 and 2 described above) in batches, and publishes a commitment to the current state of the database com and proof Π^{Upd} that a valid update has been performed (VKD.Publish) periodically. The server also publishes a public datastructure which maintains information about all the commitments so far. The auditors in the system keep checking the published update proofs and the public datastructure in order to ensure global consistency (VKD.Audit) of Dir.

The updates should be sufficiently frequent, so that the user keys are not out-of-date for long. The exact interval between these server updates, or epochs has to be chosen as a system parameter. We use time and epoch interchangeably in our descriptions.

The server also produces proofs for VKD.Query and VKD.KeyHistory. At a very high level, the users verify the proofs (VKD.QueryVer, VKD.HistoryVer) to ensure that the server is not returning an incorrect key corresponding to a username or an inconsistent key history for the keys corresponding to a username. VKD also requires the proofs to be privacy-preserving, i.e., the proofs should not leak information about any other key (that has not been queried) in Dir. The auditors are not trusted for privacy and hence, the proofs that the server produces as part of VKD.Publish need to be privacy-preserving as well.

DEFINITION 1. A Verifiable Key Directory is comprised of the algorithms (VKD.Publish, VKD.Query, VKD.QueryVer, VKD.KeyHistory, VKD.HistoryVer, VKD.Audit) and all the algorithms have access to the system parameters. We do not make the system parameters explicit. The algorithms are:

Periodic Publish:

▷$(\mathsf{com}_t, \Pi_t^{\mathsf{Upd}}, \mathsf{st}_t, \mathsf{Dir}_t) \leftarrow \mathsf{VKD.Publish}(\mathsf{Dir}_{t-1}, \mathsf{st}_{t-1}, S_t)$:

This algorithm takes the previous state of the server and the key directory at previous epoch $t-1$ and also a set S_t of elements to be updated. Whenever a client submits a request to add a new label or update an existing label from epochs $t-1$ to t, the corresponding (label, val) pair is added to S_t to be added in the VKD at epoch t. The algorithm produces a commitment to the current state of the directory com_t and a proof of valid update Π^{Upd} all of which it broadcasts at epoch t. It also outputs the updated directory Dir_t and an updated internal state st_t. If this is the first epoch, i.e., no previous epoch exists, then the server initializes an empty directory and its internal state first, then updates them as described above.

Querying for a Label:

▷ $(\mathsf{val}, \pi) \leftarrow \mathsf{VKD.Query}(\mathsf{st}_t, \mathsf{Dir}_t, \mathsf{label})$: This algorithm takes the current state of the server for epoch t, the directory Dir_t at that epoch and a query label label and returns the corresponding value if it is present in the current directory, $\perp$ if it is not present, a proof of membership or non-membership respectively.

▷ $1/0 \leftarrow \mathsf{VKD.QueryVer}(\mathsf{com}, \mathsf{label}, (\mathsf{val}, \pi))$:
This algorithm takes a commitment with respect to some epoch, a label, value pair and verifies the above proof.

Checking Consistency of Key Updates:

▷ $(\{(\mathsf{val}_i, t_i)\}_{i=1}^n, \Pi^{\mathsf{Ver}}) \leftarrow \mathsf{VKD.KeyHistory}(\mathsf{st}_t, \mathsf{Dir}_t, \mathsf{label})$: This algorithm takes in the server state, the directory at current time t and a label. It outputs $\{(\mathsf{val}_i, t_i)\}_{i=1}^n$ which are all the times at which the value corresponding to label was updated so far, the resulting val's, along with a proof Π^{Ver}.

▷ $1/0 \leftarrow \mathsf{VKD.HistoryVer}(\mathsf{com}_t, \mathsf{label}, \{(\mathsf{val}_i, t_i)\}_{i=1}^n, \Pi^{\mathsf{Ver}})$: This algorithm takes the commitment published by the server for the current time t, a label, and $\{(\mathsf{val}_i, t_i)\}_{i=1}^n$, and verifies Π^{Ver}.

Auditing the VKD:

▷ $1/0 \leftarrow \mathsf{VKD.Audit}(t_1, t_n, \{(\mathsf{com}_t, \Pi_t^{\mathsf{Upd}})\}_{t=t_1}^{t_n})$: This algorithm takes the epochs t_1 and t_n between which audit is being done, the server's published pub for all the epochs from times t_1 to t_n. It outputs a boolean indicating whether the audit is successful.

Now we discuss the security properties we require from a VKD. We give the informal descriptions here and defer the formal definitions to Appendix B. The security properties are the following.

• **Completeness:** We want to say that if a VKD is set up properly and if the server behaves honestly at all epochs, then all the following things should happen for any label updated at $t_1, \ldots, t_n$ with $\mathsf{val}_1, \ldots, \mathsf{val}_n$: 1) their history proof with $\{(\mathsf{val}_i, t_i)\}_{i=1}^n$ and Π^{Ver} should verify at t_n 2) the query proof for the label at any $t_j \leq t^* < t_{j+1}$ should verify with respect to the value consistent with the versions proof at t_j which is val_j and 3) the audit from epochs t_1 to t_n should verify.
Note that for KeyHistory and HistoryVer, we consider epochs $t_1, t_2, \ldots, t_n$ when the updates have happened for a label. These will be epochs distributed in the range $[t_1, t_n]$. However for Audit, we consider all possible pairwise epochs between t_1 and t_n. For example, for $t_1 = 3$ to $t_n = 10$, there might be updates at $3, 5, 8, 10$ but for audit we need to consider all of the epochs $3, 4, 5, 6, 7, 8, 9, 10$.

• **Soundness:** VKD soundness guarantees that if Alice has verified the update history of her key till time t_n and if for each t_i from the beginning of time till t_n, there exists at least one honest auditor whose audits have been successful then, whenever Bob queried before t_n, he would have received Alice's key value that is consistent with the key value reported in Alice's history query. Thus soundness is derived from all of VKD.Publish, VKD.QueryVer, VKD.HistoryVer and VKD.Audit.
Note that the onus is on the user, Alice, to make sure that the server is giving out the most recent and *correct* value for her key. Soundness guarantees that under the circumstances described above, Bob will always see a key consistent with what Alice has audited. But Alice needs to verify that her key as reported in the history query is consistent with the actual key that she chose.

• **Privacy:** The privacy guarantee of a VKD is that the outputs of Query, HistoryVer or Audit should not reveal anything beyond the answer and a well defined leakage function on the directory's state. So, the proofs for each of these queries should be simulatable given the output of the leakage function and the response.

2.1 Model Assumptions

Here we discuss the assumptions we have in our VKD model and discuss why we think these are justified.

What we assume from users. Our goal is to make VKDs usable by average users who cannot be assumed to remember or correctly store long term cryptographic secrets. We do assume that the user's device can store a secret (the user's current key), although that device might be lost/re-imaged, etc in which case the secret would be lost. We assume that the user has some way of authenticating to the service provider (this could be through a second factor text

message, or a phone call to answer security questions, etc); this is already a requirement in existing end-to-end encryption systems or other messaging services so we are not increasing the attack surface. We also assume that the server enforces some kind of access control on the public key directory, for example only responding to Bob's query for Alice's key if he is on her contact list. Finally, we assume that the user can recognize the approximate times when she updated her key (enough to be able to identify if extra updates have been included in the list). To help with this we could also have the server store a note from the user about each update, e.g. "bought new iphone" or "re-installed app". VKDs could of course support options for power users to make KeyHistory queries whenever they want, to explicitly compare public keys with their friends (as in WhatsApp), or to sign their key updates.

Assumptions on system. We do assume that Alice, Bob, and the server have clocks that are approximately in sync. We also assume some way of ensuring that all parties have consistent views of the current root commitment or at least that they periodically compare these values to make sure that their views have not forked. A simple way of doing this would be for the server to publish the head periodically to a blockchain (as in Catena [26]), but we could also implement it with a gossip protocol as in CONIKS [20]. If we implement this by having the server post the commitment on the blockchain, then this means we assume the client software periodically interact with the blockchain. (Because of the hash chain, even if clients do not check the root on the blockchain every epoch, if two different client's views of the root diverge, they will eventually detect that when they do contact the blockchain.) Similarly, our epoch length need not be limited by the blockchain block time – it is enough if the server periodically posts the current commitment on the blockchain; a misbehaving server will be caught as soon as the next blockchain post happens. This vulnerability window is a tunable parameter. The bandwidth requirements (for interacting with the blockchain) of log auditors of the system and of thin clients (e.g. cell phones) can be significantly reduced by using an efficiently-verifiable blockchain witnessing service, like Catena [26].

Cost of auditors. In a VKD system, the audit functionality ensures that the server updates are consistent. The audit proofs are not privacy sensitive, so we do not have to place any restrictions on who can perform audits. Audit can be performed by many auditors in parallel and it is sufficient to have *at least one honest auditor* perform audits over each adjacent server epoch. This means, auditor A could audit for server update between epoch $t, t+1$, while a different auditor B could audit between epoch $t+1, t+2$. It is reasonable to assume that *some privacy conscious* users of a VKD system or *privacy-advocacy groups* or even altruistic nodes[7], would be willing to periodically audit server updates. Furthermore, in the prototype implementation of our VKD system, SEEMless, we observed that the bandwidth cost for an audit to check server update between two consecutive epochs is less than 5MB and takes less than 0.3s to verify. For scale, 5MB is smaller than the size of an average image taken from a 12MP smartphone camera. Hence, auditing a single server update is not a significant burden. We discuss the details of the audit experiments in Section 7.

User Experience. Having an end-user seamlessly interact with the functionality of any VKD is crucial for successful deployment.

Designing a user interface that exposes the VKD functionality to end users without drastically changing their experience of using existing messaging software requires usability studies, which we leave for future work. However, here we discuss a blueprint of what we envision the user interface of a VKD system to be.

The high level goal is to have an end-user's software run the queries and verification the background *in a timely manner* and alert an end-user only when some verification fails. When Alice first registers for the service, the client software installed on her device generates keys in the background, and makes a KeyHistory query to the server to verify that this is the first key that was issued for her, and a Query to verify that the key is stored correctly. This happens in the background and is invisible to Alice unless verification of one of these fails, in which case she will be alerted of possible misbehavior. Similarly, when Alice requests Bob's key, the client software will run verification of the proof received from the server in the background and will only notify Alice if the verification fails. In addition to that, we want Alice's client to check her key history *sufficiently often.* This entails running KeyHistory query in the background periodically and notifying Alice if it is not consistent with the updates she has made.*$keyver$ must also be run when Alice changes her device or reinstalls the client software (thereby forcing a key update) in which cases the software would display to Alice a list of the times when her key was updated.

Multiple devices. We have described the system as storing only a single key per user, but in reality Alice might use her account on multiple devices, and the system would be used to store a list of public keys, one for each of her devices. Bob would then encrypt under all of these keys, and the service would determine which to deliver. This works without any modification of the VKD system. The only change in user experience is that when Alice's device runs periodic updates, it might find updates made by other devices: the times for these updates should be displayed to Alice for verification.

Distributing the service. We have described the service in terms of a single server, but in practice, the "server" can be implemented using a distributed network of servers, for reliability and redundancy. Our model captures the server as a *single logical entity*, so it can accommodate a distributed implementation fairly easily. Constructing proofs just requires reading the shared data structures for previous time epochs, so that can easily be done by many processes in parallel. Once all of the updates for a given epoch have been collected, then the authentication data structure needs to be updated. In our VKD system, SEEMless, this can also be parallelized fairly easily because of the tree structure (details in Section 5). We can support queries even during the data structure update by keeping a snapshot of the last authentication data structure until the update epoch completes. Once the epoch completes, the snapshot can be discarded (so this does not blow up memory).

3 APPEND-ONLY ZERO KNOWLEDGE SET (aZKS)

In this section we introduce a new primitive, *Append-Only Zero Knowledge Set* (aZKS) which we will use to build SEEMless. Zero Knowledge Set [5, 21] (ZKS) is a primitive that lets a (potentially

---

*Determining how often the client software should run KeyHistory depends on how quickly users want to detect misbehavior and requires user studies.

malicious) prover commit to a static collection of (label,value) pairs (where the labels form a set) such that: 1) the commitment is succinct and does not leak any information about the committed collection 2) the prover can prove statements about membership/non-membership of labels (from the domain of the labels) in the committed collection with respect to the succinct commitment 3) the proofs are efficient and do not leak any information about the rest of the committed collection. Our primitive, *Append-Only Zero Knowledge Set* (aZKS) generalizes the traditional zero-knowledge set primitive by accounting for append-only updates and characterizing the collection with a leakage function.

Here it is worth pointing out that the notion of soundness one would expect from updates in a ZKS is not obvious. For example, if the expectation is that updates leak absolutely no information about the underlying sets or type of updates (inserts/deletes), then there is no reasonable definition of soundness of updates: any set the prover chooses will be the result of some valid set of updates. In [19], Liskov did not define any soundness notion for updates. In our context, we want to be able to define an append-only ZKS, which makes the expectation of update soundness clear: it should ensure for any label, its value never gets modified and in particular, it never gets deleted.

Here we describe the primitive and informally define its security properties. The formal security definition is in Appendix D.

DEFINITION 2. Append-Only Zero Knowledge Set is comprised of the algorithms (ZKS.CommitDS, ZKS.Query, ZKS.Verify, ZKS.UpdateDS, ZKS.VerifyUpd)[†] described as follows:
▷ $(\text{com}, \text{st}_{\text{com}}) \leftarrow \text{ZKS.CommitDS}(1^\lambda, D)$: This algorithm takes the security parameter and the datastore to commit to as input, and produces a commitment to the data store and an internal state to pass on to the Query algorithm. Datastore D will be a collection of (label, val) pairs.
▷ $(\pi, \text{val}) \leftarrow \text{ZKS.Query}(\text{st}_{\text{com}}, D, \text{label})$: This algorithm takes the state output by ZKS.CommitDS, the datastore and a query label and returns its value ($\bot$ if not present) and a proof of (non-)membership.
▷ $1/0 \leftarrow \text{ZKS.Verify}(\text{com}, \text{label}, \text{val}, \pi)$: This algorithm takes a (label, value) pair, its proof and a commitment by ZKS.CommitDS and verifies the above proof.
▷ $(\text{com}', \text{st}'_{\text{com}}, D', \pi_S) \leftarrow \text{ZKS.UpdateDS}(\text{st}_{\text{com}}, D, S)$: This algorithm takes in the current server state st_{com}, the current state of the datastore and a set $S = \{(\text{label}_1, \text{val}_1), \dots, (\text{label}_k, \text{val}_k)\}$ of new (label, value) pairs for update. It outputs an updated commitment to the datastore, an updated internal state and an updated version of the datastore and proof π_S that the update has been done correctly.
▷ $0/1 \leftarrow \text{ZKS.VerifyUpd}(\text{com}, \text{com}', \pi_S)$: This algorithm takes in two commitments to the datastore before and after an update and verifies the above proof.

We require the following security properties of an *append-only* ZKS:
Soundness: For soundness we want to capture two things: First, a malicious prover $\mathcal{A}^*$ algorithm should not be able to produce two verifying proofs for two different values for the same label

---

†The original ZKS definition also included a setup algorithm run by a trusted party to generate public parameters used in all the algorithms. Our construction does not need such a set up (we show security in the random oracle model), so we omit it here.

with respect to a com. Second, since the aZKS is append-only, a malicious server should be unable to modify an existing label.
Zero-Knowledge with Leakage: We generalize the definition ZKS [5, 21] by introducing leakage functions in the classical definition. The goal of our privacy definition is to capture the following: we want the query proofs and update proofs to leak no information beyond the query answer (which is a val/$\bot$ in case of query and a bit indicating validity of the update operation). But often, it is reasonable to tolerate a small leakage to gain more efficiency. To capture this sort of leakage formally, we parameterize our definition with a leakage function. If the leakage function is set to null, then our definition reduces to the classical ZKS definition.
Append-Only Strong Accumulator: Finally, we remark that the primitive *strong accumulator* (SA) [4] can be extended to a new primitive of *append-only strong accumulator* (aSA) trivially from the definition of aZKS. An aSA is essentially a aZKS with completeness and soundness and *without* the privacy requirement. In Section 5, we will first construct an efficient aSA and then construct aZKS using aSA in a blackbox way.

4 SEEMLESS CONSTRUCTION

In this section, we will describe SEEMless: our construction of VKD from *append-only* zero-knowledge sets (aZKS). We first review the construction for CONIKS, then give an informal overview of our SEEMless construction and then describe the construction more formally. In our construction, we will be using a *hash chain*. Hash Chain is a classical authentication data structure that chains multiple data elements by successively applying cryptographic hashes, e.g. a hash chain that hashes elements a, b, c is $H(c, H(b, H(a)))$.

4.1 An overview of the CONIKS construction

In a verifiable directory, as described in the introduction, every epoch the server will publish a commitment to the current state of the directory; this will allow the server to respond to username queries with the corresponding public key and a proof that that is correct according to the committed database. CONIKS is implicitly based on the zero knowledge set we describe in Section 5, although they do not explictly describe it as such and they do not use the append-only property. Recall that a zero knowledge set, as described in Section 5, allows the server to give a concise commitment to a set of (label, val) pairs, and to give proofs for each query that reveal nothing beyond the query, the response, and the number of pairs in the directory. Each epoch the server will modify the zks with any updates to correspond to the latest directory and publish the resulting zks commitment; it will then respond to each Query with a zks proof. The system should also ensure that if two users at some point compare their root commitments and find that they match, they will be assured that all previous responses were consistent. To allow this, the server will actually publish the head of a hash chain containing all of the zks commitments.

This construction has two issues. First, in terms of privacy, because of the way the the zks is modified, it reveals additional timing information about when Alice's key is updated, leading to the tracing attack described in Appendix C. On the efficiency side, the issue with this construction is that in order to audit her key history Alice

will have to obtain a zks proof for every epoch, even if she has only made a few updates.

Key History The KeyHistory functionality is not provided in CONIKS, as such. To ensure that the key served by a CONIKS provider for her was always correct, a user `alice` would either need to always remain online, verifying the responses to the Query requests for her key at each epoch, or, retroactively check them.

Auditing CONIKS has no explicit auditors. Each user is assumed to verify the server's update each epoch to ensure that each commitment is correct w.r.t. her own key.

4.2 Intuition behind our SEEMless construction

To explain the intuition behind our construction, we take the CONIKS construction as a starting point. However, here we want two additional properties. First, we want to allow a user to check the history of when her key has been updated in time proportional to the number of updates rather than the number of edits. That means we need to do something more sophisticated than having the user query her key in every epoch. The second is that we want that all of the proofs leak as little information as possible about users' update patterns. We build up to our construction gradually. (We will describe our construction generally in terms of (label, val) pairs, but in our application label will be a username, and val will be an associated public key.)

Attempt 1: To address the efficiency issue, we need to ensure user can see changes to his key even if he doesn't check the update when they occur. For instance, suppose that in the above construction Alice performs a KeyHistory query by only verifying ZKS proofs for the epochs in which she changed her key. An adversarial server could change Alice's key for a few epochs, and then change it back to the correct value, and if Alice does not request ZKS proofs for those epochs (because she doesn't expect her key to have changed during that time), she will never detect it. We prevent this by using an *append only* ZKS as described in Section 5, which guarantees that all changes will be retained. We use the aZKS to store ((label||version), val) pairs, and when Bob queries he will get Alice's latest version number and an aZKS proof for this (label||version). Alice on KeyHistory query will get proofs for the values associated with (label||1) . . . (label||n) where n is the number of updates she's made so far.

We note that this also addresses the privacy issue discussed above. Instead of modifying the value for an existing label in the ZKS we add a new label; from the privacy of the aZKS, the update proof and commitment reveal nothing about which label was added.

The issue with this construction is that we have no way to ensure that the server shows Bob the correct latest version. For example a corrupt server could show Bob an out of date key. Or if Alice has performed n updates, it could correctly show Alice the values for label||1, . . . label||n, but then show Bob a much higher version.

Attempt 2: To prevent the server from showing Bob an out of date key, we modify the construction so that at every epoch we store 2 aZKS, one aZKS_{all} as described above, and a second one aZKS_{old} that stores all of the out of date versions. When Bob queries for Alice's key he will receive a proof for the current version with respect to aZKS_{all} as above, but he will now also recive a proof for the previous version with respect to aZKS_{old}. When Alice performs a KeyHistory check, she now additionally checks that the checks that at each update the previous version is added to aZKS_{old}.

This does not prevent the server from showing Bob a version where version is much higher than Alice's real latest update. One approach to prevent this would be to have the server provide a that |version is not in the aZKS_{all} for any higher version than Alice's current version. However, this potentially has a very high cost for Alice, proportional to the total possible number of updates, i.e. the number of epochs. Instead, we want to achieve the same guarantee but reduce Alice's work. To do this, we have the server add a "marker" node, on every 2^ith update Alice makes. When Bob queries for Alice's key, he will also check that the previous marker node is in the aZKS_{all}. When Alice performs a KeyHistory check, she will also check that no higher marker nodes are included in the aZKS_{all}. Because we mark only the 2^ith updates, this cost is now only logarithmic in the number of epochs.

Final construction The previous attempt captures the core of our construction. The one additional aspect is that the append only ZKS requires that someone verify the update proofs to ensure that each update is indeed only adding values to the set. We could have the users all perform this check, but that would again require them to do work linear in the number of epochs. Instead, we observe that the privacy guarantees of the aZKS mean that anyone can be given access to these proofs, so we can open verification up to anyone willing to perform this task; we call these parties auditors. See Section 2.1 for more discussion.

4.3 SEEMless construction:

In the construction we assume that the server's identity and public key is known to each user and auditor and all the messages from the server are signed under the server's key, so that the server cannot be impersonated.

Along with the steps of our construction, we will provide a running example for expositional clarity.

> Consider a VKD with 2 chat client users, in which the labels are usernames and the values are the corresponding public keys. Suppose at some point in time between server epochs $t-1$ and t, `alice` requested registration with her first ever public key $\text{PK}_{a,1}$ and an existing user `bob` requested to update his public key for the 2nd time to $\text{PK}_{b,3}$. These values will reflect in the VKD at epoch t. Previously, `bob` registered his first key $\text{PK}_{\text{bob},1}$ at server epoch 10 and updated it to $\text{PK}_{\text{bob},2}$ at server epoch 100.

▷ VKD.Publish($\text{Dir}_{t-1}, \text{st}_{t-1}, S_t$) : At every epoch, the server gets a set S_t of (label, value) pairs that have to be added to the VKD. The server first checks if the label already exists for some version $\alpha - 1$, else sets $\alpha = 1$. It adds a new entry (label | α, val) to the "all" aZKS and also adds (label | $\alpha - 1$,) to the "old" aZKS if $\alpha > 1$. If the new version $\alpha = 2^i$ for some i, then the server adds a marker entry (label | mark | i, "marker") to the "all" aZKS. The server computes commitments to both the aZKS, and adds them to the hash chain to obtain a new head com_t. It also produces a proof Π^{Upd}

aZKS$_{\text{all},t-1}$		aZKS$_{\text{old},t-1}$
(bob|1, PK$_{b,1}$)	(bob|mark|0, . . .)	(bob|1, null)
(bob|2, PK$_{b,2}$)	(bob|mark|1, . . .)	

Entries at epoch t − 1 (before optimization in Remark 3)

aZKS$_{\text{all},t}$		aZKS$_{\text{old},t}$
(bob|1, PK$_{b,1}$)	(bob|mark|0,. . .)	(bob|1, null)
(bob|2, PK$_{b,2}$)	(bob|mark|1,. . .)	(bob|2, null)
(bob|3, PK$_{b,3}$)		
(alice|1, PK$_{a,1}$)	(alice|mark|0, . . .)	

Entries at epoch t (before optimization in Remark 3)

Table 1: An example with two users, with views of data structures at epochs $t − 1$ and t.

consisting of the previous and new pair of aZKS commitments com$_{\text{all},t-1}$, com$_{\text{all},t}$ and com$_{\text{old},t-1}$, com$_{\text{old},t}$ and the corresponding aZKS update proofs.

REMARK 1. *Note that the directory of (username, public-key) pairs is never published. The published commitments are privacy-preserving.*

▷ VKD.Query(st$_t$, Dir$_t$, label) : When a client Bob queries for Alice's label, he should get the val corresponding to the latest version α for Alice's label and a proof of correctness. Bob gets three proofs in total: First is the membership proof of (label $\mid \alpha$, val) in the "all" aZKS. Second is the membership proof of the most recent marker entry (label $\mid$mark $\mid a$) for $\alpha \geq 2^a$. And third is non membership proof of label $\mid \alpha$ in the "old" aZKS. Proof 2 ensures that Bob is not getting a value higher than Alice's current version and proof 3 ensures that Bob is not getting an old version for Alice's label.

> In our example, if alice requested to see bob's public key at epoch t, alice would receive proofs for bob$\mid 3 \in$ aZKS$_{\text{all},t}$ with value PK$_{b,3}$ and bob$\mid$mark$\mid 1 \in$ aZKS$_{\text{all},t}$ and lastly, bob$\mid 3 \notin$ aZKS$_{\text{old},t}$. Additionally, alice will receive com$_{\text{all},t}$, com$_{\text{old},t}$, com$_{t-1}$.

▷ VKD.QueryVer(com$_t$, label, val$_t$, π_t, α) : The client checks each membership or non-membership proof, and the hash chain. Also check that version α as part of proof is less than current epoch t.

▷ VKD.KeyHistory(st$_t$, Dir$_t$, label): The server first retrieves all the update epochs $t_1, \ldots, t_\alpha$ for label versions $1, \ldots, \alpha$ from T, the corresponding com$_{\text{all},t_1-1}$, com$_{\text{all},t_1}, \ldots,$ com$_{\text{all},t_\alpha-1}$, com$_{\text{all},t_\alpha}$ and com$_{\text{old},t_1}, \ldots,$ com$_{\text{old},t_\alpha}$ and the hashes necessary to verify the hash chain: $H(\text{com}_{\text{all},0}, \text{com}_{\text{old},0}), \ldots, H(\text{com}_{\text{all},t}, \text{com}_{\text{old},t})$. For versions $i = 1$ to n, the server retrieves the val$_i$ for t_i and version i of label from Dir$_{t_i}$. Let $2^a \leq \alpha < 2^{a+1}$ for some a where α tis the current version of the label. The server generates the following proofs (together called as Π):

(1) **Correctness of** com$_{t_i}$ **and** com$_{t_i-1}$: For each i, output com$_{t_i}$ com$_{t_i-1}$. Also output the values necessary to verify the hash chain: $H(\text{com}_{\text{all},0}, \text{com}_{\text{old},0}), \ldots, H(\text{com}_{\text{all},t}, \text{com}_{\text{old},t})$.

(2) **Correct version i is set at epoch t_i:** For each i: Membership proof for (label$\mid i$) with value val$_i$ in the "all" aZKS with respect to com$_{t_i}$.

(3) **Server couldn't have shown version $i − 1$ at or after t_i :** For each i: Membership proof in "old" aZKS with respect to com$_{t_i}$ for (label$\mid i − 1$).

(4) **Server couldn't have shown version i before epoch t_i:** For each i: Non membership proof for (label$\mid i$) in "all" aZKS with respect to com$_{t_i-1}$

(5) **Server can't show any version from $\alpha + 1$ to 2^{a+1} at epoch t or any earlier epoch:** Non membership proofs in the "all" aZKS with respect to com$_t$ for (label$\mid i + 1$), (label$\mid i + 2$), $\ldots,$ (label$\mid 2^{a+1} − 1$).

(6) **Server can't show any version higher than 2^{a+1} at epoch t or any earlier epoch:** Non membership proofs in "all" aZKS with respect to com$_t$ for marker nodes (label$\mid$ mark$\mid a + 1$) up to (label$\mid$ mark$\mid \log t$).

▷ VKD.HistoryVer(com$_t$, label, $\{(\text{val}_i, t_i)\}_{i=1}^n$, Π^{Ver}): Verify each of the above proofs.

> In our example, if bob queried for his key history at epoch t, he would check the following,
> (1) com$_{10}$, com$_9$, com$_{100}$, com$_{99}$, com$_t$ and com$_{t-1}$ and the hashes necessary to verify the the hashchain $H(\text{com}_{\text{all},0}, \text{com}_{\text{old},0}), \ldots, H(\text{com}_{\text{all},t}, \text{com}_{\text{old},t})$.
> (2) bob$\mid 1$ exists and has value PK$_{b,1}$ in aZKS$_{\text{all},10}$, bob$\mid 2$ exists and has value PK$_{b,2}$ in aZKS$_{\text{all},100}$ and bob$\mid 3$ exists and has value PK$_{b,3}$ in aZKS$_{\text{all},t}$.
> (3) bob$\mid 1 \in$ aZKS$_{\text{old},100}$ and bob$\mid 2 \in$ aZKS$_{\text{old},t}$.
> (4) bob$\mid 1 \notin$ aZKS$_{\text{all},9}$, bob$\mid 2 \notin$ aZKS$_{\text{all},99}$ and bob$\mid 3 \notin$ aZKS$_{\text{all},t-1}$.
> (5) Since bob's version is $3 < 2^2 = 4$, nothing to check here.
> (6) bob$\mid$mark$\mid 2 \ldots$, bob$\mid$mark$\mid \log t \notin$ aZKS$_{\text{all},t}$

REMARK 2. *The client software runs* VKD.HistoryVer *to monitor the history of the user keys. This software either downloads the entire proof from the server each epoch it runs* VKD.HistoryVer *(when the software is re-installed or the user installs it on a new device) or caches parts of the proof from the first run of* VKD.HistoryVer *to use in the subsequent verifications. We experimentally evaluate the performance for* VKD.HistoryVer *with and without caching in Section 7.*

▷ VKD.Audit(t_1, t_n, $\{(\text{com}_t, \Pi_t^{\text{Upd}})\}_{t=t_1}^{t_n}$): Auditors will audit the commitments and proofs to make sure that no entries ever get deleted in either aZKS. They do so by verifying the update proofs Π^{Upd} output by the server. They also check that at each epoch both aZKS commitments are added to the hash chain. Note that, while the Audit interface gives a monolithic audit algorithm, our audit is just checking the updates between each adjacent pair of aZKS commitments, so it can be performed by many auditors in parallel. For security, it is sufficient to have at least one honest auditor perform audits over each adjacent pair.

REMARK 3. *In our implementation, the marker entry doubles up as the normal key entry for that version, where the version number is $\log(k)$, k being the count of number of updates for a specific user. We use a special symbol* mark *for the marker entries. For example, the 3rd key for* bob *will be saved as* bob$\mid 3$ *and the 1st key for* alice *will be saved as* alice$\mid$mark$\mid 0$, *(since $2^0 = 1$), hence only two new entries are made for any user update request.*

REMARK 4. *We describe our construction with a hash chain to for simplicity of exposition. However, we note that this could instead be replaced with a Merkle tree (we describe this data structure in*

Section 5) built over the list of all commitments to date. This would result in slightly higher update and audit costs (adding a new entry to the end of this list would require up to a logarithmic number of hashes), but would significantly reduce the cost of history queries (from linear in the number of epochs to logarithmic). We discuss this in Section 7 (Update, Gethistory and Audit experiments).

4.4 Privacy of SEEMless

We prove in Appendix B that the following leakage for SEEMless:
Publish: For each label that was updated, if this is the first update since the adversary queried for it via Query then add it to set Q_{Query}, and if it was previously queried to KeyHistory then add it to set $Q_{\text{KeyHistory}}$. The leakage from this query is the number of new registrations and of key updates in the latest epoch, and the sets $Q_{\text{Query}}, Q_{\text{KeyHistory}}$.
Query: The leakage from this query is the version number of the queried key and the epoch at which it was last updated.
KeyHistory: There is no additional leakage from this query.

Interpreting this leakage: A party who only acts as an auditor learns only the numbers of keys added and keys updated each epoch. If that party additionally acts as a user (Alice) performing KeyHistory queries, the combined leakage may reveal when her keys are updated (even if she does not perform more KeyHistory queries), but that is expected to be something Alice knows since she is the one requesting the updates. If Alice additionally queries for Bob's key, the leakage reveals the version number of Bob's current key and the epoch when it was last updated, and may reveal when that key is no longer valid (because Bob performed an update), but will not reveal anything about subsequent or previous updates.

Remark 5. *Note that, while the functionality inherently allows Bob to learn Alice's entire update history if he queries for her key every epoch, we assume as discussed in Section 2.1 that the server limits who is allowed to query for Alice's key to e.g. her contacts. Our goal here then is to guarantee that Bob cannot learn about Alice's update history without performing these queries.*

5 aZKS INSTANTIATIONS

In this section we will give a concrete instantiation for the aZKS used for SEEMless. We refer the reader to Appendix A for definitions of the standard cryptographic primitives used in the construction of the aZKS.

5.1 Append-Only Strong Accumulator (aSA) Construction

Here, we give a construction of an *append-only* strong accumulator over a data collection of (label,value) pairs. The high level idea is to build a Patricia Trie (PTr) [17] over the labels. PTr is a succinct representation of the labels such that each child has a unique suffix string associated with it and the leaf nodes constitute the actual label values. See Fig 1 for an illustrative example. Our aSA construction is built on a PTr. We use a collision-resistant hash function $H : \{0, 1\}^* \mapsto \{0, 1\}^m$ in our construction.

▷ SA.CommitDS(1^λ, D) : Datastore D $= \{(l_1, v_1), \ldots, (l_n, v_n)\}$, a collection of label-value pairs. Choose a constant $k_D \xleftarrow{\$} \{0, 1\}^\lambda$. Let $\{y_1, \ldots, y_n\}$ be a lexicographic ordering corresponding to

$\{l_1, \ldots, l_n\}$. Build a Patricia trie on $\{y_i\}$ and output com $=$ (h_{root}, k_D). For the nodes in the tree, hashes are computed as follows.
Leaf nodes: For a node y_i, compute: $h_{y_i} = H(k_D|y_i|v_i)$. For example, in the tree in Figure 1, $h_2 = H(k_D|0100|v_2)$.
Interior nodes: For an interior node x, let $x.s_0$ and $x.s_1$ be the labels of its children. Compute: $h_x = H(k_D|x|h_{x.s_0}|h_{x.s_1}|s_0|s_1)$. For example in Figure 1, $h_6 = H(k_D|0|h_1|h_7|010|1)$

▷ SA.Query(D, l) : If $l \in$ D, output value v associated to it. Let h_l be the hash value of the node. Give the sibling path for h_l in the Patricia trie along with the common prefix x at each sibling node and the suffixes s_0, s_1 which form the nodes $x.s_0$ and $x.s_1$. For example, proof for 0100 will be its value and $[h_2, (h_3, 01, 00, 11), (h_1, 0, 010, 1), (h_8, \epsilon, 0, 1)]$
If $l \notin$ D, let z be the longest prefix of l such that z is a node in the Patricia tree. Let zu_0, zu_1 be its children. Output $z, h_z, u_0, u_1, h_{zu_0}, h_{zu_1}$ along with sibling path of z. For example, proof for 1010 will be $\perp$ and $[1, 1000, 1100, h_8, h_4, h_5, (h_6, \epsilon, 0, 1)]$.

▷ SA.Verify(com, l, v, π) : Parse π as the hash values and the auxiliary prefix information at each node. Compute h_l according to leaf node calculation and verify the given value. Compute the hash values upto the root with help of the proof Let this hash value be h. Verify that $h = h_{\text{root}}$. In case of a non-membership proof, additional verification is required. Let z_l and z_r be the labels of the left and right child (respectively) of the returned node z. Verify that 1) z is the longest common prefix of z_l and z_r 2) z_l and z_r are distinct.

▷ SA.UpdateDS(D, S): First we check that S is a valid set of updates which means that for all (label$_i$, val$_i$) $\in S$, label$_i \notin$ D. Initialize sets $Z_{\text{new}}, Z_{\text{old}}, Z_{\text{const}}$ to empty. For all label$_j \in S$, compute: $h_{\text{label}_j} = H(k_D|l_j|\text{val}_j)$ and add h_{label_j} to the appropriate position in the trie and change the appropriate hash values. Add the old hash values to Z_{old} and the updated to Z_{new} and those that remain unchanged after all updates to Z_{const}. Output the final updated root hash h'_{root} as com$'$, output the final st$'$ and D$' =$ D $\cup \{(\text{label}_j, \text{val}_j)\}$ and output $\pi_S = (Z_{\text{old}}, Z_{\text{new}}, Z_{\text{const}})$ and set S.

▷ SA.VerifyUpd(com, com$'$, S, π_S): Parse $\pi_S = (Z_{\text{old}}, Z_{\text{new}}, Z_{\text{const}})$. Compute the root hash from all the values in Z_{old} and Z_{const} and check that it equals com. Similarly, compute the root hash from all the values in Z_{new} and Z_{const} and check that it equals com$'$. Let Z_S be the set of roots of the subtrees formed by labels in S. Check that $Z_{\text{old}}, Z_{\text{new}}$ are exactly the hash values of all nodes in Z_S. Output 1 if all checks go through.

Security This aSA construction is secure if H is a collision-resistant hash function. The security of this primitive is proven in [23].
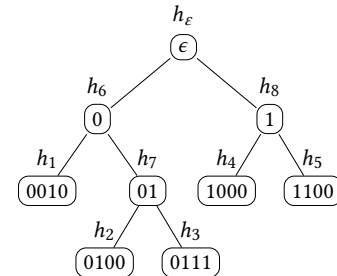
Figure 1: For the universe of be all 4 bit binary strings, a Patricia trie built on subset $P = \{0010, 0100, 0111, 1000, 1100\}$

Merke Hash Tree We use the Merkle Hash Tree variant used in Certificate Transparency [18]. This is essentially the append-only strong accumulator construction we described, where, the labels start from 0 and gets appended incrementally, i.e., the first label that gets added is 0, next one is 1 and so on.

5.2 aZKS **Instantiations**

Our construction builds upon append-only strong accumulators which directly give us completeness and soundness. To get zero-knowledge we will use a *simulatable Verifiable Random Function* (hitherto denoted as an sVRF) and a *simulatable Commitment Scheme* (hitherto denoted as an sCS). For definitions of these primitives, see Appendix A. At a high level, the sVRF helps map an element in the datastructure to a random position and the commitment helps in hiding the exact value that is being accumulated.

Let (sVRF.KeyGen, sVRF.Eval, sVRF.Prove, sVRF.Verify) be a simulatable VRF and let (CS.Commit, CS.Open, CS.VerifyOpen) be a simulatable commitment scheme as described before. Given $D = \{(label_1, val_1), \ldots, (label_n, val_n)\}$, ZKS.Commit will first generate an sVRF key pair $SK, PK \leftarrow sVRF.KeyGen(1^\lambda)$, build a new $D' = \{(l_1, v_1), \ldots, (l_n, v_n)\}$ where $l_i = sVRF.Eval(SK, label_i)$, $v_i = CS.Commit((label_i, val_i) ; r_i)$ for random r_i. It will then build an append-only SA on D' and output that together with PK. ZKS.Query will return the appropriate l_i along with the sVRF proof, the opening to the commitment, and the strong accumulator membership/non-membership proof for (l_i, v_i). ZKS.UpdateDS will compute the l_i, v_i pairs for the new datastore entries and then run SA.UpdateDS. The formal proof of security is in Appendix D.

Leakage In our aZKS construction, ZKS.Query leaks the size of the datastore, and when the queried element was added (assuming it is a member). ZKS.UpdateDS leaks the size of the datastore before and after an update. For each element that was added, it also leaks whether and when the adversary previously queried for it.

6 PERSISTENT PATRICIA TRIE

In our construction, SEEMless, the server maintains two aZKS at every epoch. In Section 4, we described how to construct a aZKS using a aSA, which in turn, we implemented using Patricia Trie. Recall that, in SEEMless, the server needs to store the entire history of the key directory and the two corresponding aSA (used in the aZKS construction) to be able to answer key history queries. Naively, the server can maintain all the aSA from the beginning of time, but this will blow up the server's storage significantly and hinder scalability of the system. This is particularly wasteful, when in fact, there might be a significant overlap between the aSA of consecutive epochs. To address this problem, we build a persistent data structure that retains information about every epoch, while being space efficient. We call this persistent data structure *Persistent Patricia Trie* (PPTr). We believe that PPTrs will find a wide number of applications (such as the tamper-evident logging in [23]).

Challenges: The idea is similar to that of history tree (HT) [23], built on Sparse Merkle Tree. However, we can not naively use the same technique to build a Persistent Patricia Trie (PPTr). Recall that Patricia Trie (PTr) is compressed for efficiency — there is no empty node in the tree. Compression introduces several subtle challenges in building its persistent version, i.e., PPTr. For example, unlike

in HT [23], nodes in a Patricia Trie do not have a fixed position. A node with a certain label may be at lower depth at an earlier epoch (root being at depth 0) and fall to a higher depth at a later epoch. A node with a given label may change depth several times from its birth epoch until the latest epoch. Please see Fig 2 for an illustrative example. Therefore, a parent to child pointer does not remain fixed throughout the life time of a PTr. In the example, the node 00 pointed to 0000 as its left child at time t_2, and to node 000 at time t_4. The HT construction in [23] crucially relies on parent to children pointers being fixed throughout the lifetime of the tree. We will build our final construction of PPTr gradually.

Attempt 1: First, let us observe some invariants in an append-only PTr (aPTr): (1) The depth of a PTr is defined with respect to the root node of the tree, therefore the root node is always at depth 0 (2) Any node in a append-only PTr can only fall to higher depths over epochs, it never goes to lower depths. (3) Any node with k bit label can be at depth at most k in a PTr.

Now let us attempt to build a PPTr. Every node in a PTr is identified by its label with a binary string. For every node u, we store some information corresponding to every epoch when the hash value at node u changed. Let us denote this information as nodestate. These epochs and the corresponding nodestates are stored in a hashtable. Corresponding to every epoch t when the hash value at node u changed, nodestate$_t$ stores the following.

(1) h_t: hash value of node u at epoch t
(2) leftskip: a binary string indicating how many levels have to be skipped to get the left child of u at epoch t. Let the left child label at epoch t is l
(3) leftepoch: the key to lookup l's hashtable and get the corresponding nodestate
(4) rightskip: a binary string indicating how many levels have to be skipped to get the left child of u at epoch t. Let the right child label at epoch t is r
(5) rightepoch: the key to lookup r's hashtable and get the corresponding nodestate

Since we are constructing a aPTr, we need to support only two operations efficiently on the data structure:

Generate proof for a historical query: Recall (non-) membership proof for certain node u at epoch t consists of the siblings of the nodes on the path from the leaf (or an internal node with both children's labels non-prefixes of label u) to the root. For any previous epoch t, this information can be easily extracted traversing the tree from the root and following at each node u the (leftskip, leftepoch) and (rightskip, rightepoch) pointers from hashtable entry of epoch t. Thus, the cost of generating membership and non-membership proofs for any epoch in the past is proportional to the tree height, and is independent of the number of epochs.

Insert new node in aPTr: When a new node gets inserted to a aPT, all nodes on its path from the root gets updated. But note that, each node on this path has a hash table (of (epoch, nodestate)) associated with it. For each of these entries, the skip entries might need to get updated while updating the path. Hence, the cost of updating each node on the path is not constant, it is proportional to the size of the hashtable at that node. Consequently, the cost of inserting a new node is not proportional to the tree height.

Attempt 2: Our next attempt is to retain the cost for generating historical proofs while bringing down the cost of insertion. Notice
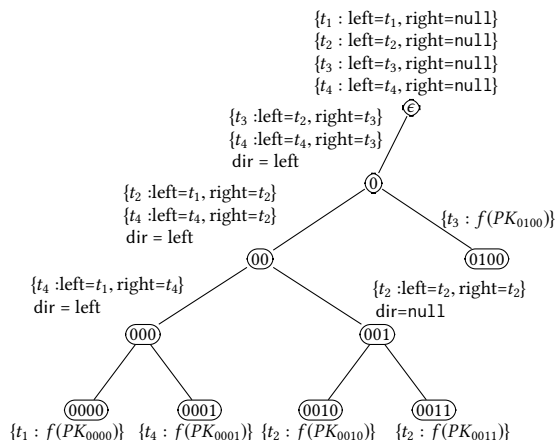
Figure 2: Append only Patricia Trie: final view at epoch t_4

that the skip entry at the parent of node u gets updated if and only if a new node with label v is inserted to the tree where v is a proper prefix of u. v becomes the new parent of u. We keep a direction flag with every node v that indicates whether u became the left or right child of the newly inserted node v. Let us denote this flag as dir. Along with this, at every node we store birthepoch: the epoch at which the label u was inserted.

Now we do not need to keep the skip fields any more, because dir will help us get to the correct node. More specifically, for a historical query (u, t), we will go down from the root following as long as birthepoch $\geq t$ guided by the dir bit. For example, in Figure 2, at t_1, 0000 was the leftchild of the root but at time t_2 00 replaced 0000 and 0000 became the left child of 00. So the dir flag at t_2 is set to left. The birthepoch of 00 is t_2 which equals the queried epoch t_2. Now, if we query $(0000, t_2)$, we will get to node 00 whose node state entry t_2 will have leftepoch = t_1 and dir is set to Left. So we go down the leftchild of 00 and find the entry 0000 at epoch t_2.

Here insertion cost is proportional to the height of the tree since cost of insertion on every node on the path is constant. But we have introduced yet another problem in the historical query. The proof needs the sibling path and we may need to traverse several levels from a parent (following dir) to get the correct sibling.

Final construction: To overcome this problem, we keep at the nodestate of every node, the label and the hashes of its current children. This way, we get the required sibling directly from the parent node. Nodestate now contains the following.

(1) h_t: hash value of node u at epoch t
(2) leftlabel: the label l of the left child of node u at time t.
(3) lefthash: the hash value of node l at epoch t
(4) leftepoch: the key to lookup l's hashtable and get the corresponding nodestate
(5) rightlabel: the label of the right child of node u at time t. Let r denote the label
(6) righthash: the hash value of node r at epoch t
(7) rightepoch: the key to lookup r's hashtable and get the corresponding nodestate

With this new nodestate, combined with the additional information of dir, birthepoch, we have a persistent Patricia Trie data structure where both historical query and insertion works in time proportional to the height of the Trie.

7 EXPERIMENTS

In this section, we provide an experimental evaluation of SEEMless implemented in Java. As discussed in Section 8, EthIKS [3] and Catena [26] are essentially wrappers for CONIKS. So, we only provide explicit comparison with the Java implementation of CONIKS.

7.1 Experimental Setup

Our experiments were conducted on a Linux VM running on a 2.30GHz Intel Xeon E5-2673 CPU with 125GB of memory and a 64GB heap allocated to the JVM for the server-side experiments for both SEEMless and CONIKS. All client experiments are done on a 2.8 GHz Intel Core i7 laptop. We used Google's protocol buffers [27] to implement communication between the client and the server, due to their efficiency and compression. All sizes below are computed using the getSerializedSize() API call for protocol buffer messages, which returns the number of bytes used to encode a message.

We implemented the VRF using the *Secp384r1* curve using the Bouncy Castle library, Icart's function [12] for hashing to the curve and SHA384. We use the technique of [11] of hashing the input twice (using SHA384($0||input$) and SHA384($1||input$)) and applying Icart on both these hashes, since applying Icart once does not provide a distribution indistinguishable from random.

7.2 Performance Evaluation

Server updates: The graph in Fig. 5 shows the time taken for server update in the following experiment: we increased the total users by 100k at each epoch and then took measurements for 1k new user registrations and 1k key updates by existing users. For a server with 10M registered users, it takes $0.28s$ on average to update the authentication structure (the two PPTrs) and its directory. As expected, the update time is proportional to the log of total number of registrations and updates in the system. The server adds the aZKS commitments in a Merkle Tree (instead of a hashchain) as discussed in Section 4. We estimate this cost using the Patricia trie used in CONIKS (an over-estimate). The time to insert a aZKS commitment in the Mekle Tree of aZKS commitments is $61\mu s$ after 10M server epochs have passed.

Both CONIKS and SEEMless compute the VRFs for updates and registrations between epochs, when a client message arrives. The online update phase only needs to update the authentication data structure and does not incur the VRF cost. For SEEMless, registering a new user requires a single VRF computation (for inserting into the "all" aZKS) and an update requires 2 VRF computations – one for the new key entering the "all" aZKS and the other for the old key entering the "old" aZKS. Therefore the offline VRF computation cost at the server (per request) is at least $1.3ms$ and at most $2.5ms$, which is relatively low computation between epochs. For CONIKS, the cost is a single VRF computation (for registration, no VRF is computed for update), hence the cost is at most $1.3ms$.[‡]

Scalability comparison We tried to run the same update experiment with CONIKS, on the same machine and saw that the heap filled up by the time we had around 4.5M users. Hence, we decided to run a limited experiment with CONIKS which only measured the time

---

[‡]In the current implementation, CONIKS doesn't cache the VRF, re-computing it for every update and Query. Keeping a hash map for the VRFs would be a simple optimization. Note that CONIKS did not implement the VRF for their evaluation.

	Prove Query	Verify Query
SEEMless	6.03 *ms*	10.51 *ms*
CONIKS	2.01 *ms*	3.47 *ms*

Table 2: Simulated times for proof generation at the server, verification at the client in Query **experiments at** 10**M users.**

for 1k new users and 1k updates at selected epochs. This illustrates the ability of SEEMless to scale in contrast to the limitation of a CONIKS server to less frequent epochs. CONIKS performs even worse (and hangs) with more frequent epochs. This issue is an outcome of the repeated copying of the Patricia trie in CONIKS and is language agnostic, since the RAM overflows with frequent epochs. Even if older states were delegated to storage to prevent overflowing the RAM, any KeyHistory requests in CONIKS would require expensive accesses. The PPTr construction from Section 6 used in SEEMless helps achieve scalability by solving this problem and allowing us to run our experiments only using RAM.

Query: To evaluate the cost of Query, we measure the computation and verification times at the server and the client respectively, and the proof-size. The number of users registered with the server vary up to 10M. All costs are averaged over 100 users with 10 trials each.

Time The time spent in computing and verifying a Query proof has two components: the VRF proof and the authentication path in the Patricia trie. Table 2 shows the time taken to compute and verify Query proofs including the VRF cost, which is the most expensive operation and dominates the Query cost. We computed the VRF costs as follows. The compute, prove and verify functions for our VRF implementation, take, on average, 1.3, 1.9 and 3.4 *ms* respectively over 5k trials. Since the time for each of these computations is independent of the input, we used them as constants to simulate the VRF costs for the Query experiments (recall that each Query response has 1 VRF proof in CONIKS and 3 in SEEMless).

In our experiments, we saw the cost of authentication path generation and verification is negligible compared to the VRF – $0.2-0.3ms$ for SEEMless and $0.05 - 0.06ms$ for CONIKS[§]. Since SEEMless requires 3 VRF proofs, its VRF cost is triple that of CONIKS. However, performance of both systems remains comparable: on the order of 10 milliseconds, i.e. not a noticeable delay for a user. We can gain further speedup by parallelizing the authentication paths and VRF computations in SEEMless proof generation and verification.

Size of Proofs In our implementation, the proof contains the username and public key (which are sent in addition to the proofs as part of the Querys) along with a commitment to the public key, its opening, the VRF proof and the authentication path in the Patricia trie. The average size of the proof is about 8600B when the number of users is 10M. For scale, the length of a Twitter post is famously 140 characters (1,120B)[8], so, the bandwidth consumed by the proof is less that of reading 8 Twitter posts in plaintext.

The proof size is proportional to the logarithm of the total number of users registered, as expected (Fig 3). We also measure the size of the CONIKS proofs. Since we have 3 authentication paths per Query proof in our system, whereas CONIKS has one, the size is at most 3× the proof of CONIKS, as expected.

KeyHistory: Here we measure the computation and verification times of KeyHistory query and the size of the proofs. We note that

[§]Note that the CONIKS Java implementation stores public keys in plaintext and not as commitments (a crucial privacy omission). This slightly reduces their cost.
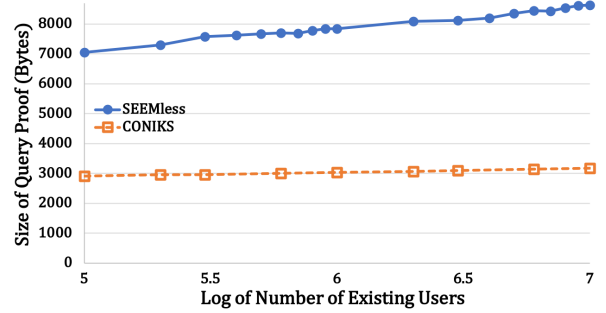
Figure 3: Mean size of Query **proofs for** 100 **nodes as the number of users in the VKD varies. The** x**-axis is logarithmic in base** 10 **and each data point is the mean of** 10 **trials.**
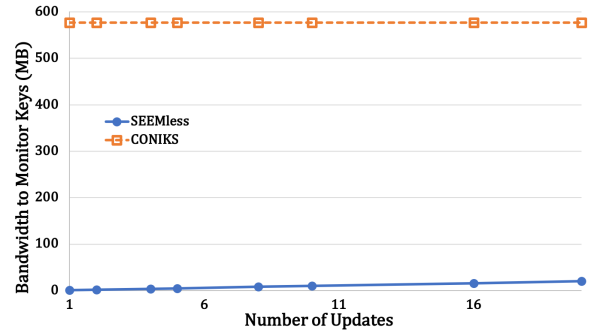
Figure 4: Mean bandwidth to verify the history proofs over 10 **users (with caching) in** S **and 1M users in** R **and** 200**k epochs as the number of key updates by users in** S **varies and each user checks her key history after each update.**

our implementation of KeyHistory functionality currently implements some redundant checks (for verification without caching) which we optimized in the protocol description in Section 4.

The cost of KeyHistory is determined by 3 parameters: 1) the number of server side epochs , 2) the latest key version or number of key updates made by the user in question and 3) the total number of updates on the server side (i.e. total number of entries in the history trees). We have already seen how the Query proofs grow with respect to parameter 3, instead, we measured how the first two parameters affect the cost of the history proofs.

Both our experiments include a set S of 10 'special' users, each of whom not only register, but also update their keys as per the experiments. The set of other users, let us call it the set R of 'regular' users, serves as a control: these users register and then update their keys only once. In both experiments, we try to simulate *special* users updating their keys at regular intervals. This means, if there are 200k epochs and the *special* users update their keys 10 times: they register initially and then update every 20k epochs. As a control, the users in R register at a fixed rate. For instance, if $|R| = 1$M, with a total of 200k epochs, 5 users in R register and 5 update their keys at every epoch. We first consider the experiments sans VRF. We do not include the cost of verifying the aZKS commitments in the Merkle Hash Tree (as described in Section 4) in these experiments. We simulate the cost of the Merkle Tree by using the Patricia trie of CONIKS (an over-estimate). The verification time for one leaf is about $49\mu s$ even for 10M server epochs in this Patricia trie, which would be close to running the server for several years with half
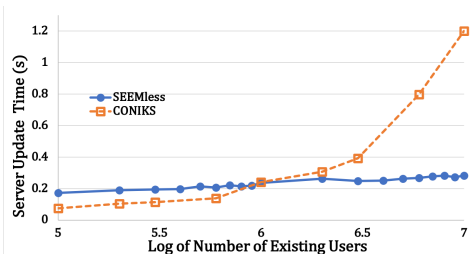
Figure 5: Mean times to update SEEMless and CONIKS for a new epoch with 1k new users and 1k updates averaged over 10 trials. The x-axis is logarithmic in base 10.
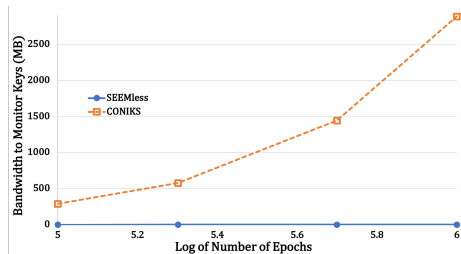

Figure 6: Mean bandwidth for the history proofs over 10 users (with caching) in S with 10 updates each and 1M users in R as the number of epochs varies. The growth for SEEMless bandwidth is slow: it is ≈ 1.03MB with 10^5 epochs and ≈ 1.09MB with 10^7 epochs.

minute epochs. Hence, if a user updated her key on 10 occasions, she would have less than an additional $0.5ms$ of verification time.

Dependence on epochs To measure dependence of KeyHistory response sizes and verification times on epochs, we fixed the size of R at 1M, the number of updates by the users in S at 10 and varied the number of epochs (adding an equal number of elements in R at each epoch). Note that the number of minutes in a year is 525, 600, so 1M is close to the number of epochs which would pass if the length of 1 epoch were half a minute.

In our experiments, the average bandwidth consumed by a user who updates 10 times and checks her history after every update, with a server which has 1M epochs is ≈ 1.09MB – not much more than at 100k epochs which is ≈ 1.03MB (Fig 6). The slight growth in size is due to the increasing number of marker entries to be checked as the number of epochs increases. The verification of the history proofs took the client a total computation time of 0.02s for 1M epochs excluding the cost for the VRF. In contrast, CONIKS requires a user to check her key at each epoch – the total cost of monitoring her key grows rapidly (Fig 6). These experiments show that SEEMless scales with short epochs as opposed to CONIKS.

Dependence on key updates The graph in Fig. 4 shows the total bandwidth consumed by the proofs for various numbers of updates by users (running verification with caching) in S, where the epochs are fixed at $200k$ and $|R| = 1$M. In the case of our system, even a user who updates her keys almost every week for about five months (20 updates), has to download only an average of 2.01MB to verify her entire key history and the downloaded proof can be verified in an average of less than 0.05s excluding the VRF cost. This includes sending the keys and usernames themselves (at least 1000 bytes per key update), since the user may want to verify the used values. If a user changes her key somewhat less frequently, say, every two weeks, then she can monitor the entire history of her keys by downloading about 1.05MB. This does not change significantly as the number of epochs grows. On the other hand, a CONIKS user would still need to download a proof every epoch to monitor her key binding – amounting to almost 576.8MB over as many epochs.

Caching VRF labels at the client Each time a user runs KeyHistory, she can cache the new VRFs returned from the server, amortizing verification cost. Recall that, to verify a key version, i, such that $2^a \leq i < 2^{a+1}$ for some non-negative integer a, a user:

- $i - 1$: needs to check for any version that the previous entry is in the "old" aZKS. One additional version is moved to the "old" aZKS for each new update.

- $i + 1, ..., 2^{a+1} - 1$: will have the VRF values for versions in the set $i + 1, ..., 2^{a+1} - 1$ already verified and cached. So, they do not need to be verified, just checked for equality, unless she updates to a new version i such that $i = 2^k$ for some $k \in \mathbb{Z}$. If indeed $i = 2^k$ for some $k \in \mathbb{Z}$, she will have to download the VRF values $i + 1, ..., 2^{k+1}$ (2^k in number). This amortizes her cost to 1 VRF verification per update.

- $2^{a+1}, ..., 2^{\lfloor \log t \rfloor}$: will need to verify an additional VRF only when verifying after some new 2^kth epoch. For 10 updates at regular intervals, once 1M epochs pass, she will have to verify a total of 20 VRFs, getting an additional amortized cost of 2 VRF verifications, per update.

Once verified, she can just save the 32B Patricia trie labels for the "all" aZKS and the "old" aZKS which results in the following formula

$$\text{client cache} \approx [2i + 2^{\lfloor \log i \rfloor + 1} + \log t - (\lfloor \log i \rfloor + 1)] \times 32\text{B}$$

where i is the number of updates by the user and t is the number of server epochs which have passed. In practice, this means that even a user who has 20 updates over 1M epochs only needs to cache ≈ 2.8KB. Note that this is much smaller than even a low resolution image taken on a flip phone. Figures 7 and 8 show the total computation times for verifying the history proofs including simulated costs for the VRF verifications with or without caching as well as amortized costs. Compare the numbers in Fig. 7 with a total of over 349s in CONIKS for a user needing to verifying her key every epoch regardless of number of updates when she does not save the VRF locally at just 100k epochs. Alternatively, if she were to store the tree label corresponding to her name, it would require constant monitoring and take a total of over 7s to ensure the server never shows an incorrect key for her. The numbers for CONIKS would increase significantly when the number of epochs is higher to 3495s with 1M epochs, if she doesn't save the VRF, and 71s when she does. In the experiment in Fig. 8, the epochs are fixed at 200k, so even as the number of updates in CONIKS increases, its computation time would remain fixed at over 699s to verify sans caching and over 14s to verify with caching.

Caching VRF proofs at the server If the server were to cache the VRF proofs for each of 10M users, each with 20 updates, it would correspondingly have to store ≈ 28GB of data. However, updates and KeyHistorys are relatively infrequent, (such as when a user reinstalls or updates her app), it is reasonable to assume that even about a second of startup time for an app to verify key history is not a barrier to usability. To respond to Query without the need
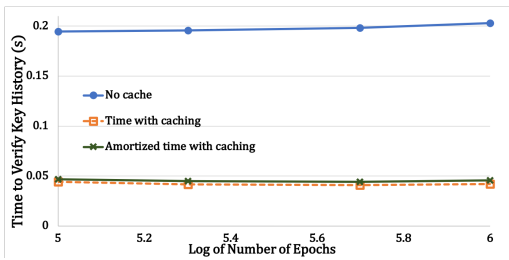
Figure 7: Mean time to verify SEEMless history proofs over 10 users in S and 1M users in R and 10 key updates as the number of epochs varies including the time VRF verifications. In CONIKS, the total time for a non-caching user depends on the number of epochs and is over 349s, even at 10^5 epochs.
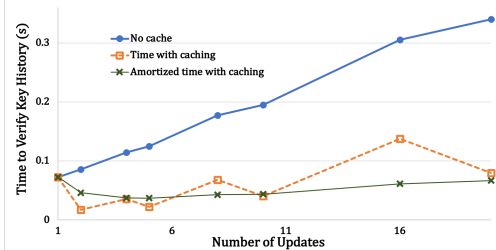
Figure 8: Mean computation time to verify SEEMless history proofs over 10 users in S and 1M users in R and 200k epochs as the number of key updates by users in S varies including the time for VRF verifications. In CONIKS, the time would only depend on the number of epochs and equal over 699s for a non-caching user and over 14s with caching at 200k epochs.

to compute a VRF will require caching the VRF value and proof for the current key of each user. This amounts to $3 \times 48B$ (the VRF value is 48B) for each registered user, which is 1.44GB over 10M users. Further optimizations can be made, such as saving values for inactive users to disk. We leave such optimizations as future work.

Auditing: Recall that the fresh and stale keys are stored in the 'all' and 'old' aZKS respectively. Commitments to the roots of these aZKSs are stored in a Merkle tree. The auditors of SEEMless verify that (1) the commitments to the aZKSs are correctly added to the Merkle Tree. Since they use the same data structure, the cost for verifying entries in this tree is the same as verifying a membership proof for a label in CONIKS: it takes about $\approx 49\mu s$ and about 3KB of bandwidth for a tree on 10M epochs (as already discussed in KeyHistory experiments). (2) the aZKS are growing in an append-only manner. For (2), the auditors need to verify append-only growth of the Patricia tries which instantiate the aZKS. To monitor that the Patricia tries are growing in an append-only manner, its auditors receive parts of the tree which are unchanged and the changed leaves, from which they reconstruct the new version of the tree. We simulate this cost by considering the number of new hashes our aZKS trees have to perform, since our implementation only computes the hashes that changed. For the "all" aZKS, less than 47k new hashes had to be computed and this took 0.17s, and the "old" aZKS computed less than 17k new hashes and it took 0.05s. So the total cost of hashing is 0.22s.

The auditors need to download the respective new leaves and the nodes whose subtrees did not change. For downloading the new leaves, the auditors need to download 1k leaves for the "old" aZKS,

2k leaves for the "all" aZKS, which are 2×256 bits each (leaf label and commitment), totaling 192KB for the two aZKS. To estimate an upper-bound on the number of roots of the unchanged aZKS subtrees, consider the following. Any node which is the root of an unchanged subtree must have a sibling which has changed. If not, then the node as well as its sibling are unchanged, the parent of this node must be unchanged and thus, the parent or some other ancestor must be the root of the unchanged subtree. Therefore, contenders for roots of unchanged subtrees in an aZKS tree are only siblings of nodes along the path from the root to new leaves. Recall that the average depth of a node in a compressed Patricia Trie is $\log n$ (we experimentally confirmed this) where n is the total number of leaves. Hence, the number of unchanged roots of unchanged subtrees is at most $k \times \log n$ where k is the number of new leaves. Corresponding to our experiment in Section 7, this amounts to about 1000 log 100k unchanged nodes for the "old" aZKS and 2000 log 10.1M for the "all" aZKS, totaling an upper bound of 1.06MB+2.98MB= 4.04MB (including hash values and labels). Thus, the total proof size that the auditors need to download is less than 4.24MB. Note that, sizes of the labels becomes smaller at each level, but this estimate counts them all as 256 bits. This estimate also double counts nodes at the intersection of paths to multiple new leaves. So, in practice, this proof size will be much lower.

8 RELATED WORK

Our work broadly falls in the category of building provably secure and efficient privacy-preserving key directory service, particularly relevant in the context of end-to-end secure messaging services. In the recent past, this problem has received significant attention in both the academic community and industry [2, 3, 10, 13, 20, 22, 24, 26]. This line of work is related to transparency logs [9, 18], but here we focus on the works that are most relevant to us.

The work closest to us in the literature is CONIKS [20], a directory service that lets users of an end-to-end encrypted communication system verify that their keys are being correctly reported to all other users. We have already discussed this construction in Section 4. EthIKS [3] implements CONIKS using the Ethereum blockchain for storing the server's authentication data-structures in an Ethereum contract. Not only does this weaken the privacy guarantee of CONIKS, invoking the blockchain for every operation leads to a significantly high overhead and makes this approach quickly infeasible. Catena [26] provides a more general infrastructure for managing application-specific logs of *append only* statements using the `OP_RETURN` function of the Bitcoin blockchain and implements CONIKS using Catena. Catena can be used for storing SEEMless digests with the same overhead as for storing CONIKS digests. We discuss how we can integrate SEEMless with Catena in Section 2.1.

Recently, Keybase has rolled out an auditable key directory service [15, 16] but without any privacy; all the key changes by end-users in Keybase are publicly available. Keybase also differs from SEEMless, CONIKS and EthIKS in the user assumptions, they assume that users have multiple trusted devices and access to long-term cryptographic signature keys with which they sign all their encryption key updates. For an end user, it is incredibly difficult to manage keys, so we believe this assumption isn't very realistic.

REFERENCES

[1] 2017. WhatsApp Security Vulnerability. https://www.schneier.com/blog/archives/2017/01/whatsapp_securi.html. Accessed: 2019-01-25.

[2] Muneeb Ali, Jude Nelson, Ryan Shea, and Michael J. Freedman. 2016. Blockstack: A Global Naming and Storage System Secured by Blockchains. In *2016 USENIX Annual Technical Conference (USENIX ATC 16)*. USENIX Association, Denver, CO, 181–194. https://www.usenix.org/conference/atc16/technical-sessions/presentation/ali

[3] Joseph Bonneau. 2016. EthIKS: Using Ethereum to audit a CONIKS key transparency log. In *International Conference on Financial Cryptography and Data Security*. Springer, 95–105.

[4] Philippe Camacho, Alejandro Hevia, Marcos Kiwi, and Roberto Opazo. 2008. Strong accumulators from collision-resistant hashing. In *International Conference on Information Security*. Springer, 471–486.

[5] Melissa Chase, Alexander Healy, Anna Lysyanskaya, Tal Malkin, and Leonid Reyzin. 2005. Mercurial commitments with applications to zero-knowledge sets. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 422–439.

[6] Melissa Chase and Anna Lysyanskaya. 2007. Simulatable VRFs with applications to multi-theorem NIZK. In *Annual International Cryptology Conference*. Springer, 303–322.

[7] Ruben Cuevas, Michal Kryczka, Angel Cuevas, Sebastian Kaune, Carmen Guerrero, and Reza Rejaie. 2010. Is content publishing in BitTorrent altruistic or profit-driven?. In *Proceedings of the 6th International COnference*. ACM, 11.

[8] developer.twitter.com. 2010. Counting Characters. https://developer.twitter.com/en/docs/basics/counting-characters.html. Accessed: 2018-12-25.

[9] Saba Eskandarian, Eran Messeri, Joe Bonneau, and Dan Boneh. 2017. Certificate Transparency with Privacy. *arXiv preprint arXiv:1703.02209* (2017).

[10] Mohammad Etemad and Alptekin Kupcu. 2015. Efficient Key Authentication Service for Secure End-to-end Communications. Cryptology ePrint Archive, Report 2015/833. https://eprint.iacr.org/2015/833.

[11] Reza R Farashahi, Pierre-Alain Fouque, Igor Shparlinski, Mehdi Tibouchi, and J Voloch. 2013. Indifferentiable deterministic hashing to elliptic and hyperelliptic curves. *Math. Comp.* 82, 281 (2013), 491–512.

[12] Thomas Icart. 2009. How to hash into elliptic curves. In *Advances in Cryptology-CRYPTO 2009*. Springer, 303–316.

[13] Keybase.io. 2014. Keybase is now writing to the Bitcoin blockchain. https://keybase.io/docs/server_security/merkle_root_in_bitcoin_blockchain. Accessed: 2018-10-05.

[14] Keybase.io. 2019. Keybase is not softer than TOFU. https://keybase.io/blog/chat-apps-softer-than-tofu. Accessed: 2019-05-05.

[15] Keybase.io. 2019. Managing Teams and Keys with Keybase. https://keybase.io/docs-assets/blog/NCC_Group_Keybase_KB2018_Public_Report_2019-02-27_v1.3.pdf. Accessed: 2019-05-05.

[16] Keybase.io. 2019. Protocol Security Review. https://rwc.iacr.org/2019/slides/keybase-rwc2019.pdf. Accessed: 2019-05-05.

[17] Donald Ervin Knuth. 1998. *The art of computer programming: sorting and searching*. Vol. 3. Pearson Education.

[18] Ben Laurie, Adam Langley, and Emilia Kasper. 2013. *Certificate transparency*. Technical Report.

[19] Moses Liskov. 2005. Updatable zero-knowledge databases. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 174–198.

[20] Marcela S Melara, Aaron Blankstein, Joseph Bonneau, Edward W Felten, and Michael J Freedman. 2015. CONIKS: Bringing Key Transparency to End Users.. In *Usenix Security*. 383–398.

[21] Silvio Micali, Michael Rabin, and Joe Kilian. 2003. Zero-knowledge sets. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium*. IEEE, 80–91.

[22] Namecoin. 2014. https://namecoin.org. Accessed: 2018-10-05.

[23] Alina Oprea and Kevin D Bowers. 2009. Authentic time-stamps for archival storage. In *European Symposium on Research in Computer Security*. Springer, 136–151.

[24] LEAP Encryption Access Project. 2012. Nicknym. https://leap.se/en/docs/design/nicknym. Accessed: 2018-10-05.

[25] signal.org. 2016. Identity binding. https://www.signal.org/docs/specifications/x3dh. Accessed: 2019-05-05.

[26] Alin Tomescu and Srinivas Devadas. 2017. Catena: Efficient non-equivocation via Bitcoin. In *IEEE Symp. on Security and Privacy*.

[27] Kenton Varda. 2008. Protocol buffers: Google's data interchange format. *Google Open Source Blog, Available at least as early as Jul* 72 (2008).

A CRYPTOGRAPHIC PRIMITIVES FOR aZKS

Collision Resistant Hash Function (CRHF): A hash function H is collision resistant if it is hard to find two inputs that hash to the same output; that is, two inputs x and y such that $H(x) = H(y)$, and $x \neq y$. In some of our building blocks, we treat a hash function as a random oracle. This roughly provides security against attackers who use the hash function as a black box.

INSTANTIATION. For our experiments we use SHA384 in our VRF construction and SHA256 everywhere else.

Simulatable Commitment Scheme (sCS) : A simulatable commitment scheme consists of algorithms (CS.Commit, CS.Open, CS.VerifyOpen): (1) $\text{com}_\sigma \leftarrow \text{CS.Commit}(1^\lambda, m; r)$: takes security parameter λ and produces commitment com_σ to message m using randomness r. (2) $\tau \leftarrow \text{CS.Open}(\sigma, m, r, \text{com}_\sigma)$: outputs a decommitment value corresponding to commitment com_σ for message m and randomness r. (3) $1/0 \leftarrow \text{CS.VerifyOpen}(\sigma, \text{com}_\sigma, m, \tau)$: accepts or rejects the decommitment of com_σ to message m in terms of the decommitment value τ.

An sCS satisfies the standard requirements of commitment schemes with respect to hiding and binding: commitments hide the committed message, and it is hard to open a commitment to two different messages. In addition, sCS requires that there exist a hypothetical "simulator" with special powers (either a trapdoor or in our case control of the random oracle) who could form commitments that can later be opened to any value m. This then shows that the commitment and opening give the adversary no additional information.

INSTANTIATION. We can construct this commitment scheme from a hash function which we model as a random oracle: $\text{Commit}(\sigma, m; r)$ samples random string $r \leftarrow \{0, 1\}^\lambda$ and outputs $H(m, r)$.

Simulatable Verifiable Random Function (sVRF) [6]: A Verifiable Random Function (VRF) is similar to a pseudorandom function, with the additional property of verifiability: corresponding to each secret key SK, there is a public key PK, such that, for any $y = \text{VRF.Eval}(\text{SK}, x)$, it is possible to verify that y is indeed the value of the VRF seeded by SK evaluated on input x. A simulatable VRF (sVRF) is a VRF for which this proof can be simulated, so a hypothetical simulator with special powers (e.g. controlling the random oracle) can fake a proof that the value of $\text{VRF.Gen}(\text{SK}, x)$ is any y. This guarantees that the proof does not compromise the pseudorandomness of this or any other output. An sVRF is comprised of the following algorithms: 1) $(\text{PK}, \text{SK}) \leftarrow \text{sVRF.KeyGen}(1^\lambda)$: takes security parameter λ and outputs the public key PK and secret key SK. 2) $y \leftarrow \text{sVRF.Eval}(\text{SK}, x)$: takes the secret key SK and x and outputs the sVRF evaluation of x as y 3) $\pi \leftarrow \text{sVRF.Prove}(\text{SK}, x)$: takes the secret key SK and x and outputs a proof for the sVRF evaluation of x 4) $1/0 \leftarrow \text{sVRF.Verify}(\text{PK}, x, y, \pi)$: verifies the proof.

INSTANTIATION. We describe an efficient constructions of sVRFs based on the DDH assumption as proposed in [20]. The proof for pseudorandomness follows from DDH as observed in [20]. Let $\mathbb{G}$ be a DDH group of order q and let g be a generator. Let $H_1 : \{0, 1\}^* \mapsto \mathbb{G}$ and $H_2 : \{0, 1\}^* \mapsto \mathbb{Z}_q^*$ be two hash functions. The sVRF construction is the following. (1) sVRF.KeyGen: Choose $k \xleftarrow{\$} \mathbb{Z}_q^*$ and output $\text{SK} = k$ and $\text{PK} = g^k$ (2) sVRF.Eval: Output $y = (H_1(x))^{\text{SK}}$ (4) sVRF.Prove: The proof is the Fiat-Shamir transformation of Schnorr protocol for proving common exponent; proving that prover knows $\text{SK} = k$ such that $\text{PK} = g^k$ and $y = h^k$ for $h = H_1(x)$. Prover chooses $r \xleftarrow{\$} \mathbb{Z}_q^*$ and outputs proof as $\pi = (s, t)$

for $s = H_2(x, g^r, h^r)$ and $t = r - sk$. (5) sVRF.Verify: Parse π as (s, t). Check that $s = H_2(x, g^t(\text{PK})^s, h^t y^s)$ for $h = H_1(x)$

B VERIFIABLE KEY DIRECTORY

Notation: A function $v : \mathbb{Z}^+ \mapsto \mathbb{R}^+$ is a negligible function if for all $c \in \mathbb{Z}^+$, there exists k_0 such that for all $k \geq k_0$, $v(k) < k^{-c}$. An algorithm $\mathcal{A}$ is said to have oracle access to machine O if $\mathcal{A}$ can write an input for O on a special tape, and tell the oracle to execute on that input and then write its output to the tape. We denote this oracle access by $\mathcal{A}^O$. For the rest of the sections, we will use (label, val) and (username, public key) interchangeably.

Security Properties: Here we give the formal definitions that we informally described in Section 2. In the security definitions, t_1 denotes the first server epoch at which the label in the definitions gets registered for the first time with the directory Dir.

- **Completeness:** Let $\text{Dir}_0 = \{\}$, $\text{st}_0 = \perp$.
  For all possible labels, for all t_{current}, n, for all sets $\{\text{val}_i\}_{i=1}^n$ for all update sets $S_1, \ldots, S_{t_{\text{current}}}$ such that $(\text{label}, \text{val}_i) \in \{S_{t_i}\}_{i=1}^n$ is the set of all occurrences of label in $S_1, \ldots, S_{t_{\text{current}}}$, and $\forall t^* < t_{\text{current}}$:

$$\Pr[(\{((\text{com}_t, \Pi_t^{\text{Upd}}), \text{st}_t, \text{Dir}_t) \leftarrow \text{VKD.Publish}(\text{Dir}_{t-1}, \text{st}_{t-1}, S_t)\}_{t=1}^{t_{\text{current}}} \wedge$$

$$(\{(\text{val}_i', t_i')\}_{i=1}^{n'}, \Pi^{\text{Ver}}) \leftarrow \text{VKD.KeyHistory}(\text{st}_{t_{\text{current}}}, \text{Dir}_{t_{\text{current}}}, \text{label}) \wedge$$

$$n = n' \wedge \forall i \in [1, n](t_i' = t_i \wedge \text{val}_i' = \text{val}_i) \wedge$$

$$\text{VKD.HistoryVer}(\text{com}_{t_{\text{current}}}, \text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}) \wedge$$

$$\text{VKD.Audit}(t_1, t_{\text{current}}, \{\text{com}_k, \Pi_k^{\text{Upd}}\}_{k=t_1}^{t_{\text{current}}}) \wedge$$

$$(\pi, \text{val}) \leftarrow \text{VKD.Query}(\text{st}_{t^*}, \text{Dir}_{t^*}, \text{label}) \wedge$$

$$\text{VKD.QueryVer}(\text{com}_{t^*}, \text{label}, \text{val}, \pi) \wedge$$

$$\exists j \in [1, n] s.t. (t_j \leq t^* < t_{j+1} \wedge (\text{val} = \text{val}_j))] = 1$$

- **Soundness:** We want to capture that for any label label if versions proofs verifies with respect to $\{(\text{val}_i, t_i)\}_{i=1}^n$ and if the audit verifies from t_1 to t_n then at any time t^* between an interval $[t_j, t_{j+1}]$ for some $j \in [n]$, a malicious server cannot give out a proof for label with a value which is inconsistent with the corresponding versions proof at t_j that is, $\text{val} \neq \text{val}_j$ for t_j. Checking this is enough because if the server has given an incorrect key at time t^*, he will have to introduce an additional update sometime later to potentially fix it and will hence be caught with high probability. Hence a malicious server S^* should not be able to come up with a label, $\{(\text{val}_i, t_i)\}_{i=1}^n$ with versions proof Π^{Ver}, commitments and update proofs Π^{Upd} for all times between t_1 to t_n and query proof (π, val) for some t^* for $\text{val} \neq \text{val}_j$.
  For all PPT S^*, there exists a negligible function $v()$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr[(\text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}, \{\text{com}_k, \Pi_k^{\text{Upd}}\}_{k=t_1}^{t_{\text{current}}},$$

$$t^*, j, (\pi, \text{val})) \leftarrow S^*(1^\lambda):$$

$$\text{VKD.QueryVer}(\text{com}_{t^*}, \text{label}, \text{val}, \pi)$$

$$\wedge \text{VKD.Audit}(t_1, t_{\text{current}}, \{\text{com}_k, \Pi_k^{\text{Upd}}\}_{k=t_1}^{t_{\text{current}}})$$

$$\wedge \text{VKD.HistoryVer}(\text{com}_{t_{\text{current}}}, \text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}})$$

$$\wedge (\text{val} \neq \text{val}_j) \wedge (t_j \leq t^* < t_{j+1})$$

$$\wedge j \in [1, n] \wedge t_1 \leq \ldots \leq t_n \leq t_{\text{current}}] \leq v(\lambda)$$

THEOREM B.1. *The construction in Section 4 satisfies VKD soundness as defined above.*

PROOF. First, we claim that every pair of proofs $\Pi_k^{\text{Upd}}, \Pi_{k+1}^{\text{Upd}}$ must contain consistent values for $\text{com}_{\text{all}, k}, \text{com}_{\text{old}, k}$. If not, we could directly build a reduction breaking collision resistance of H. Now, we observe that we have a chain of aZKS commitments and associated aZKS update proofs, all of which verify. We claim that the aZKS commitments $\text{com}_{\text{all}, t^*}, \text{com}_{\text{old}, t^*}$ and $\text{com}_{\text{all}, t_i}, \text{com}_{\text{old}, t_i}$ contained in π and Π^{Ver} must also be consistent with those given in the Π^{Upd} proofs. If not, again we can break collision resistance of H.

Next, say that π and Π^{Ver} are aZKS-inconsistent if they contain a pair of proofs w.r.t. either the "old" or "all" aZKS such that one is a membership proof at time x and the other is a nonmembership proof for the same label at some time $y \geq x$, or a pair of proofs with contradictory vals for the same label (even at different times). Note that an adversary who with non-negligible probability wins the above soundness game with π, Π^{Ver} that are aZKS-inconsistent can be directly used to build an adversary attacking the aZKS soundness property. Thus, we have only to show that any adversary who successfully breaks VKD soundness must produce aZKS-inconsistent proofs. We argue that as follows. Let α be the version number contained in proof π produced by the adversary. Consider the following cases:

$\boldsymbol{\alpha < j}$: In this case Π^{Ver} (step (3) of KeyHistory) contains a membership proof for $(\text{label}||\alpha)$ w.r.t. $\text{com}_{\text{old}, t_{\alpha+1}}$ and π contains a non-membership proof for $(\text{label}||\alpha)$ w.r.t. $\text{com}_{\text{old}, t^*}$, where $t^* > t_j \geq t_{\alpha+1}$. (The first inequality follows from the definition, the second inequality follows from $\alpha < j$.)

$\boldsymbol{\alpha = j}$: In this case Π^{Ver} (step (2) of KeyHistory) contains a membership proof for $(\text{label}||\alpha)$ with value val_j in $\text{com}_{\text{all}, t_j}$, while π contains a membership proof for $(\text{label}||\alpha)$ with value $\text{val} \neq \text{val}_j$ in $\text{com}_{\text{all}, t^*}$.

$\boldsymbol{j < \alpha \leq n}$: In this case Π^{Ver} (step (4) of KeyHistory) contains a nonmembership proof for $(\text{label}||\alpha)$ w.r.t. $\text{com}_{\text{all}, t_\alpha - 1}$, while π contains a membership proof for $(\text{label}||\alpha)$ w.r.t. $\text{com}_{\text{all}, t^*}$. Note that $t_\alpha - 1 \geq t_{j+1} - 1 \geq t^*$, where the first inequality follow from $j < \alpha$ and the second follows from the definition.

$\boldsymbol{n < \alpha < 2^{a+1}}$ where a is the largest integer s.t. $2^a \leq n < 2^{a+1}$: In this case Π^{Ver} (step (5) of KeyHistory) contains a nonmembership proof for $(\text{label}||\alpha)$ w.r.t. $\text{com}_{\text{all}, t_{\text{current}}}$, while π contains a membership proof for $(\text{label}||\alpha)$ w.r.t. $\text{com}_{\text{all}, t^*}$. Note that $t_{\text{current}} \geq t^*$, from the definition.

$\boldsymbol{\alpha \geq 2^{a+1}}$ where $a \in \mathbb{Z}$ s.t. $2^a \leq n < 2^{a+1}$: Let b be the largest integer s.t. $2^b \leq \alpha$. Π^{Ver} (step (6) of KeyHistory) contains a nonmembership proof for $(\text{label}|\text{mark}|b)$ w.r.t. $\text{com}_{\text{all}, t_{\text{current}}}$, while π contains a membership proof for $(\text{label}|\text{mark}|b)$ w.r.t. $\text{com}_{\text{all}, t^*}$. Note, $t_{\text{current}} \geq t^*$, from the definition. □

- **$\mathcal{L}$-Privacy:** We will say that a VKD is private for leakage function $\mathcal{L} = (L_{\text{Publish}}, L_{\text{Query}}, L_{\text{KeyHistory}})$ if there exists a simulator $\mathcal{S} = (\mathcal{S}_{\text{Publish}}, \mathcal{S}_{\text{Query}}, \mathcal{S}_{\text{KeyHistory}})$ such that for any PPT client C^*, the outputs of the following two experiments are computationally indistinguishable:
  In the real game, C^* is given access to $O_P, O_\pi, O_{\Pi^{\text{Ver}}}$, three stateful oracles that share state. O_π is the proofs oracle

which on query a label label, will output π generated by $(\pi, \text{val}) \leftarrow \text{VKD.Query}(\text{st}_t, \text{Dir}_t, \text{label})$. $O_{\Pi^{\text{Ver}}}$ is the key history oracle, which on query, label outputs Π^{Ver} generated by $(\{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}) \leftarrow \text{VKD.KeyHistory}(\text{st}_t, \text{Dir}_t, \text{label})$. O_P is the publish oracle which on input update set S, updates directory Dir_t and outputs $(\text{com}_t, \Pi_t^{\text{Upd}})$ as computed by VKD.Publish().

In the simulated game C^* is given access to an oracle which maintains Dir_t, but calls the simulators to produce commitments and proofs. On a Publish query, the oracle $\mathcal{S}_{\text{Publish}}$ is given leakage on the update set S given by $L_{\text{Publish}}(S)$, emulates the publish oracle O_P. On a Query query, the oracle looks up val for label in the current directory Dir_t, and calls $\mathcal{S}_{\text{Query}}(L_{\text{Query}}(\text{label}, \text{val}), \text{label}, \text{val})$¶ to emulate the proof oracle O_π. On a KeyHistory query, the oracle looks up the history for val_i in its directories $\text{Dir}_1, \ldots, \text{Dir}_t$, calls $\mathcal{S}_{\text{KeyHistory}}(L_{\text{KeyHistory}}(\text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n), \text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n)$‖ to emulate the key history oracle $O_{\Pi^{\text{Ver}}}$.

Theorem B.2. *The construction in Section 4 when implemented with an aZKS with the leakage profile described in Section 4 satisfies VKD privacy as defined above, with the above leakage functions.*

Proof. We first describe the necessary simulators. Recall that the simulators share state. In this case that state will include three tables: T_{current}, which stores the most recent version for labels which are in $Q_{\text{KeyHistory}}$, T_{Query}, which stores the version for labels which are in Q_{Query} at the time of their latest Query query, and $T_{\text{KeyHistory}}$, which stores the version and epoch for the latest KeyHistory query. Our simulators will run two instances of the aZKS simulator, which we will denote aZKS.Sim$^{\text{all}}$ and aZKS.Sim$^{\text{old}}$.

Sim$_{\text{Publish}}$ is given as input the number of new registration and number of key updates, and the sets $Q_{\text{Query}}, Q_{\text{KeyHistory}}$. On the first query, Sim$_{\text{Publish}}$ calls aZKS.Sim$^{\text{all}}_{\text{CommitDS}}$ and aZKS.Sim$^{\text{old}}_{\text{CommitDS}}$ on an empty directory D_0 to generate $\text{com}_{\text{all},0}, \text{com}_{\text{old},0}$. On subsequent queries it behaves as follows:

First, note that the number of additions to ZKS.$_{\text{all}}$ will be the total number of new registrations and updates, while the number of additions to ZKS.$_{\text{old}}$ will be the number of updates. Next the VKD simulator must construct the sets of labels Q_{all} and Q_{old} in the current ZKS update sets for which there has been a previous non-membership query. We do this as follows: for each label $\in Q_{\text{Query}}$, look up the version number α in T_{Query} and add label$|\alpha$ to Q_{old} (unless $\alpha = 2^a$ in which case we add label$|$mark$|a$). For each query label $\in Q_{\text{KeyHistory}}$, lookup the current version α_{current} in T_{current} and the version $\alpha_{\text{KeyHistory}}$ and query epoch $t_{\text{KeyHistory}}$ from the most recent key history query in $T_{\text{KeyHistory}}$. Let $\alpha_{\text{current}} = \alpha_{\text{current}} + 1$, and update the value in T_{current}. Let a be the maximum value such that $2^a \leq \alpha_{\text{KeyHistory}}$. If $\alpha_{\text{current}} < 2^{a+1}$ add (label$|\alpha_{\text{current}}$) to Q_{all}. If $\alpha_{\text{current}} = 2^b$ for $b \geq a+1$ and $\alpha_{\text{current}} \leq t_{\text{KeyHistory}}$, add (label$|$mark$|b$) to Q_{all}. Finally, it calls aZKS.Sim$^{\text{all}}_{\text{CommitDS}}$ and aZKS.Sim$^{\text{old}}_{\text{CommitDS}}$ with these leakage values to obtain updated commitments and proofs.

Sim$_{\text{Query}}$ is given as input the (label, val), and the leakage L_{Query}, which is the version number α of label in the current directory, and the epoch t_{prev} at which it was last updated.

- It will call the simulator aZKS.Sim$^{\text{all}}_{\text{Query}}(t_{\text{current}}, (\text{label}|\alpha), \text{val}, t_{\text{prev}})$ where t_{current} is the current epoch to generate the membership proof in ZKS.$_{\text{all}}$ (or if $\alpha = 2^a$, it will query with aZKS.Sim$^{\text{all}}_{\text{Query}}(t_{\text{current}}, (\text{label}|\text{mark}|a), \text{val}, t_{\text{prev}})$).
- It will call the simulator aZKS.Sim$^{\text{old}}_{\text{Query}}(t_{\text{current}}, (\text{label}|\alpha), \perp, \perp)$ to generate the non-membership proof in ZKS.$_{\text{old}}$.

Finally, it updates the entry for label in T_{Query} to store version α (or adds it if it does not exist).

Sim$_{\text{KeyHistory}}$ is given label, $\{(\text{val}_i, t_i)\}_{i=1}^{\alpha_{\text{current}}}$ and generates the simulated proof as follows:

(1) For each i it outputs the commitments $\text{com}_{t_i}, \text{com}_{t_{i-1}}$ produced by Sim$_{\text{Publish}}$, and similarly for the hash values.
(2) For each i it will call the simulator aZKS.Sim$^{\text{all}}_{\text{Query}}(t_i, (\text{label}|i), \text{val}_i, t_i)$ to generate the membership proof in ZKS.$_{\text{all}}$ (or if $i = 2^a$, it will query with aZKS.Sim$^{\text{all}}_{\text{Query}}(t_i, (\text{label}|\text{mark}|a), \text{val}_i, t_i)$).
(3) For each i it will call the simulator aZKS.Sim$^{\text{old}}_{\text{Query}}(t_i, (\text{label}|i - 1), \text{null}, \perp)$ to generate the membership proof in ZKS.$_{\text{old}}$.
(4) For each i it will call the simulator aZKS.Sim$^{\text{all}}_{\text{Query}}(t_i - 1, (\text{label}|i), \text{null}, \perp)$ to generate the nonmembership proof in ZKS.$_{\text{all}}$.
(5) For each j from $\alpha_{\text{current}} + 1$ to $2^{a+1} - 1$ it will run aZKS.Sim$^{\text{all}}_{\text{Query}}(t_{\text{current}}, (\text{label}|j), \text{null}, \perp)$ to generate the non-membership proof in ZKS.$_{\text{all}}$.
(6) For each j from 2^{a+1} to $\log(t_{\text{current}})$ it will run aZKS.Sim$^{\text{all}}_{\text{Query}}(t_{\text{current}}, (\text{label}|j), \text{null}, \perp)$ to generate the nonmembership proof in ZKS.$_{\text{all}}$.

Finally, it updates the entry for label in $T_{\text{KeyHistory}}$ to store $(\alpha_{\text{current}}, t_{\text{current}})$ (or adds it if it does not exist). If there is no entry in T_{current} for label, it adds it.

Now that we have defined our simulator in terms of the aZKS simulators, the proof is very straightforward. We introduce one hybrid game which proceeds as in the simulated game except that the aZKS simulator for old is replaced by the real aZKS algorithms. Then we can argue that the real game is indistinguishable from the hybrid game by the aZKS privacy property, and similarly that the hybrid game is indistinguishable from the simulated game by the aZKS privacy property. □

C TRACING VULNERABILITY IN CONIKS

In [3, 20, 26], when a user queries for the same label several times, she gets values in the proof which depend on the position of her label in the authentication tree. These values can be positions of other labels and give information about a label that was not queried. Hence if Alice gets the proofs for the same label over time, she can infer about other labels, whether they were updated or deleted. For example, consider a system with 4 users: Alice, Bob, Charlie, Mary with P(Alice) = 010, P(Bob) = 011 P(Charlie) = 101, P(Mary) = 110, P() denotes the position of the label in the tree. The proof for Alice's key will contain 011, being its sibling, which is P(Bob). Since the position is fixed for the entire lifetime of the directory, now Alice

---

¶Recall that the leakage functions share state, so this also implicitly gets access to the leakage from all previous Publish, Query, or KeyHistory queries.

‖see ¶

can trace when Bob's key changes just by querying for her own key and observing when sibling node changes. While the username is not directly leaked, once Alice queried for Bob's key, she will be able to completely trace when the key changed without ever querying for his key again. This means, even if Bob has deleted Alice from his contact list, Alice will still be able to trace when Bob's key changed just my looking at the proof of her own key. We call this attack *tracing vulnerability*.

D AZKS DEFINITIONS

Soundness. We allow $\mathcal{A}^*$ to win the soundness game if it is able to do either of the following: Output com, label, $\mathrm{val}_1, \mathrm{val}_2, \pi_1, \pi_2$ such that both proofs verify for $\mathrm{val}_1 \neq \mathrm{val}_2$. Or output $\mathrm{com}_1, \ldots \mathrm{com}_n$, label, $\mathrm{val}_1, \mathrm{val}_2, \pi_1, \pi_2, S, \pi_S$ such that π_1 verifies for $\mathrm{com}_1, \mathrm{val}_1$ for $\mathrm{val}_1 \neq \bot$ and π_2 verifies for $\mathrm{com}_n, \mathrm{val}_2$ for $\mathrm{val}_2 \neq \mathrm{val}_1$ and the update verifies for $\mathrm{com}_1, \ldots \ldots, \mathrm{com}_n,$.
In general, we want that for all PPT $\mathcal{A}^*$ algorithm there exists a negligible function $v()$ such that for all n, λ:

$$\mathrm{Pr}[(\mathrm{com}_1, \{(\mathrm{com}_i, \pi^i)\}_{i=2}^n, \mathrm{label}, \mathrm{val}_1, \mathrm{val}_2, \pi_1, \pi_2) \leftarrow \mathcal{A}^*(1^\lambda, \mathrm{pp}) :$$
$$\{\mathrm{ZKS.VerifyUpd}(\mathrm{com}_{i-1}, \mathrm{com}_i, \pi^i) = 1\}_{i=2}^n \wedge (\mathrm{val}_1 \neq \bot) \wedge (\mathrm{val}_1 \neq \mathrm{val}_2)$$
$$\wedge \mathrm{ZKS.Verify}(\mathrm{com}_1, \mathrm{label}, \mathrm{val}_1, \pi_1) = 1$$
$$\wedge \mathrm{ZKS.Verify}(\mathrm{com}_n, \mathrm{label}, \mathrm{val}_2, \pi_2) = 1] \leq v(\lambda)$$

THEOREM D.1. *The construction in Section 5.2 satisfies* aZKS *soundness as defined above.*

PROOF. Here we consider three cases
- The sVRF values presented for label in π_1 and π_2 are different. In this case we can directly reduce to the verifiability property of the sVRF.
- $\mathrm{val}_2 \neq \bot$ and the leaf commitments presented for label in π_1, π_2 are the same. In this case we can directly reduce to the binding property of the commitment scheme.
- The sVRF values are the same, and either the commitments are different, or π_2 is a non-membership proof. In this case we can directly reduce to soundness of the strong accumulator. □

Privacy. In our definition, we have an initial commitment leakage function L_{CommitDS}, a query leakage function L_{Query}, and a leakage function on the updates L_{UpdateDS}. L_{CommitDS} captures what we leak about the collection/datastore when initializing the system, L_{Query} captures what is leaked by the proofs for each query, and L_{UpdateDS} captures what we leak during an update. Note that all of these are stateful functions that share state, so e.g. L_{Query} may depend on previous CommitDS and UpdateDS queries.

We will say that an updatable ZKS is zero knowledge for leakage function $L = (L_{\mathrm{CommitDS}}, L_{\mathrm{Query}}, L_{\mathrm{UpdateDS}})$ if there exists a simulator $\mathrm{Sim} = (\mathrm{Sim}_{\mathrm{CommitDS}}, \mathrm{Sim}_{\mathrm{Query}}, \mathrm{Sim}_{\mathrm{UpdateDS}})$ such that for any PPT malicious client algorithms C^*, the outputs of the following two experiments are computationally indistinguishable:

In the real game, the adversary C^* produces an initial directory D_0, and receives the output $(\mathrm{com}, \mathrm{st}_{\mathrm{com}}) \leftarrow \mathrm{ZKS.CommitDS}(1^\lambda, D_0)$. It then gets oracle access to two oracles, O_Q, and O_U. O_U is the update oracle which on input a set of updates $S = \{(\mathrm{label}_i, \mathrm{val}_i)\}$ will output $(\mathrm{com}', \mathrm{st}_{\mathrm{com}, t+1}, \pi_S) \leftarrow$

ZKS.UpdateDS($\mathrm{st}_{\mathrm{com}}, t, D_t, S$). The game keeps track of the state of the datastore D_t after every update. The second oracle is O_Q, the query oracle which on query datastore version t and label_i, will output $(\pi_i, \mathrm{val}_i) \leftarrow \mathrm{ZKS.Query}(\mathrm{st}_{\mathrm{com}}, t, D_t, \mathrm{label}_i)$.

In the simulated game, the C^* produces an initial datastore D_0, and receives the output $\mathrm{com} \leftarrow \mathrm{Sim}_{\mathrm{CommitDS}}(1^\lambda, L_{\mathrm{CommitDS}}(D_0))$. C^* then gets access to simulated versions of the two oracles. On UpdateDS queries, $\mathrm{Sim}_{\mathrm{UpdateDS}}$ emulates the update oracle O_U: it gets a leakage on the set to be updated outputs $(\mathrm{com}', \pi_S) \leftarrow \mathrm{Sim}_{\mathrm{UpdateDS}}(\mathrm{Ł}_{\mathrm{UpdateDS}}(S))$. On Query queries, the oracle outputs val_i and $\pi_i \leftarrow \mathrm{Sim}_{\mathrm{Query}}(t, \mathrm{label}_i, \mathrm{val}_i, L_{\mathrm{Query}}(t, \mathrm{label}_i, \mathrm{val}_i))$.

Leakage. The concrete leakage for our ZKS is: $L_{\mathbf{CommitDS}}$ reveals the size of the datastore. $L_{\mathbf{Query}}$ reveals when each queried item was added to the datastore. $L_{\mathbf{UpdateDS}}$ reveals the number of items added. It also reveals the set Q of labels of items in this update for which there had been a previous non-membership query.

THEOREM D.2. *The construction in Section 5.2 satisfies* aZKS *privacy as defined above with the specified leakage functions.*

PROOF. We first define the simulator:
$\mathbf{Sim}_{\mathbf{CommitDS}}$ takes as input the size N of the data store. It uses the sVRF simulator to generate an sVRF public key. It chooses N random strings as the output of the sVRF (we will refer to them below as leaf strings), and uses the sCS simulator to form the commitments. Then it builds a tree as in the real protocol.
$\mathbf{Sim}_{\mathbf{Query}}$ takes as input the zks version t, the label label, the corresponding value val (or $\bot$ if label is not in D_t), and the update t_{prev} when label was added to the datastore. If $\mathrm{val} \neq \bot$ it chooses an unused leaf string from the t_{prev}th update, simulates the sVRF proof to show that that is the correct value for label, and simulates an opening of sCS to val. If $\mathrm{val} = \bot$, it chooses a random leaf string, simulates the sVRF proof to show that that is the correct value for label, and constructs the rest of the proof as in the real protocol. Finally, it records that that leaf string has been assigned to label. (Future Query's for label will use the same leaf string.)
$\mathbf{Sim}_{\mathbf{UpdateDS}}$ takes as input the number of items added and the set Q of new items for which there had been a previous non-membership query. For each label in Q, it looks up the leaf string that was assigned to label by $\mathrm{Sim}_{\mathrm{Query}}$, for the remaining number of items it chooses random leaf strings. For each item it also uses the sCS simulator to form the corresponding commitment. Then it performs the rest of the update algorithm as in the real protocol. The proof that this simulator satisfies the privacy definition follows from a fairly straightforward series of games:
Game 1: Real game.
Game 2: As in game 1, but commitments and openings are simulated. *Indistinguishable by the hiding property of the sCS.*
Game 3: As in game 2, but the sVRF public key and sVRF proofs are generated by the sVRF simulator, and the leaf strings are chosen at random. *Indistinguishable by simulatability of sVRF.*
Game 4: Simulated game. *Identical to game 3.*

□