Transparency Overlays and Applications

Melissa Chase
Microsoft Research Redmond
melissac@microsoft.com

Sarah Meiklejohn
University College London
s.meiklejohn@ucl.ac.uk

ABSTRACT

In this paper, we initiate a formal study of transparency, which in recent years has become an increasingly critical requirement for the systems in which people place trust. We present the abstract concept of a transparency overlay, which can be used in conjunction with any system to give it provable transparency guarantees, and then apply the overlay to two settings: Certificate Transparency and Bitcoin. In the latter setting, we show that the usage of our transparency overlay eliminates the need to engage in mining and allows users to store a single small value rather than the entire blockchain. Our transparency overlay is generically constructed from a signature scheme and a new primitive we call a dynamic list commitment, which in practice can be instantiated using a collision-resistant hash function.

1. INTRODUCTION

In the past decade, the trust that society places in centralized mechanisms run by government, network operators, and financial institutions has been eroding, with various incidents demonstrating that high integrity cannot be achieved solely through trust in one or a handful of parties. As a reaction to this erosion in trust, two alternative architectures have emerged: users have either taken matters into their own hands and flocked to systems that have no central point of trust, or they have increased pressure on central entities to provide more openness and accountability.

A prominent example of a system with no central point of trust is Bitcoin [28], which was deployed in January 2009. Bitcoin is a monetary system that is not backed by any government and is managed through a consensus mechanism over a peer-to-peer network; there is thus no single entity that issues bitcoins or validates individual transactions, and users of Bitcoin operate using pseudonyms that are not inherently tied to their real-world identity. Bitcoin has achieved staggering success: as of this writing, its market capitalization is over 8 billion USD and its underlying structure has inspired hundreds of alternative cryptocurrencies; payment gateways such as Bitpay and Coinbase allow thousands of vendors to accept it; a number of governments have taken steps to legitimize Bitcoin via interfaces with traditional financial and regulatory infrastructures; and major financial institutions such as JPMorgan Chase [30] and Nasdaq [29] have announced plans to develop Bitcoin-based technologies.

Bitcoin and its variants have achieved a large degree of success, but denying all forms of central authority arguably limits their ability to achieve widespread adoption. Thus, technological solutions have emerged that instead seek to provide more visibility into currently centralized systems. One key example of such a system is Certificate Transparency (CT) [21], which addresses shortcomings with SSL certificate authorities (CAs) — which have ranged from failing to verify the identity of even major website owners such as Google before issuing a cryptographic certificate [10, 19] to suffering major hacks [25] that result in hundreds of forged certificates being issued [22] — and empowers users to verify for themselves the correct functioning of a system with which they interact many times a day (e.g., any time they log in to a secure website, such as their email provider). Unlike Bitcoin's approach, CT does not substantially alter the underlying infrastructure (i.e., the issuance of a certificate is largely unchanged), but instead provides a way for anyone to monitor and audit the activities of CAs to ensure that bad certificates can be detected quickly, and misbehaving authorities identified and excluded.

While Bitcoin and Certificate Transparency provide solutions in different settings, they in fact share some common features; most notably, they rely on *transparency* as a means to achieve integrity. In Bitcoin, the ledger of transactions — called the blockchain — is completely transparent, meaning all Bitcoin transactions are globally visible. A similar property is provided in Certificate Transparency, in which a distributed set of servers each maintain a globally visible log of all the issued certificates of which they are aware.

Furthermore, both Bitcoin and CT adopt a *distributed* solution, which is essential to avoid placing trust in any single entity. Indeed, relying on one party creates (at worst) a system in which this central party has unilateral control over the information that is released, or (at best) a system with one central point of failure on which attackers could target their efforts. By using a solution that is both transparent and distributed, these systems intuitively provide some notion of *public auditability*: individual users can check for themselves that only "good" events have taken place (e.g., in the case of Bitcoin, that all bitcoins have been spent at

CCS '16, October 24–28, 2016, Vienna, Austria.

ISBN 978-1-4503-4139-4/16/10…$15.00

DOI: http://dx.doi.org/10.1145/2976749.2978404

most once) and detect misbehavior on the part of all actors within the system. Understanding the link between transparency and the types of misbehavior that can be detected across a variety of settings is one of the main motivations behind this work.

1.1 Our contributions.

Systems such as Bitcoin and CT seem to provide important transparency benefits (namely, the public auditability mentioned above), but the similarities and differences between their benefits are not well understood, and no formal analysis has demonstrated either the level of transparency that they provide or how this transparency provides the intended benefits. In this paper, we initiate such a formal study. In doing so, we seek to not only compare the different guarantees provided by these systems (although our analysis does accomplish this), but more importantly to create an abstract *transparency overlay* that may be used to provide these guarantees in a variety of applications beyond financial transactions and certificate issuance.

Before we can analyze these protocols or construct a transparency overlay, we must first consider the crucial components that make up these systems. Our first step is thus to formalize — in Section 3.2 — a primitive that we call a *dynamic list commitment* (DLC); a DLC can be thought of as a generalization of a rolling hash chain or hash tree, and serves as the foundation for our construction of a transparency overlay. After defining this underlying primitive, we then go on to present transparency overlays in Section 4; here our design is heavily inspired by the design of CT. We begin with a formal model for transparency overlays, and then go on to present an abstract transparency overlay and prove its security.

Armed with this abstract secure transparency overlay, we go on in Section 5 to demonstrate that CT is a secure transparency overlay. We also demonstrate that our formal notion of security implies more intuitive notions of security in this setting (i.e., that users should accept only "good" certificates) and discuss some practical considerations.

In Section 6, we continue by turning our attention to the Bitcoin protocol. Here, we do not use the protocol directly (as we argue that it clearly cannot satisfy our notions of security), but rather plug crucial components of the protocol into our abstract transparency overlay. While this allows us to achieve a provably secure transparency overlay for Bitcoin, it more importantly also implies that "regular" Bitcoin users (i.e., users interested only in transacting in bitcoin, rather than engaging in the mining process) can operate significantly more efficiently, provided they are willing to outsource some trust to a distributed set of parties. This result demonstrates that, in any setting in which users are willing to trust any distributed set of parties, the full decentralization of Bitcoin is not needed, and the same goals can in fact be accomplished by a CT-like structure, in which regular users store significantly less information about the transaction ledger and the mining process is superfluous; i.e., the quadrillion hashes per second expended on Bitcoin mining (as of March 2016) can be eliminated without sacrificing security. Our formal analysis thus reveals the fine line separating fully decentralized (and expensive) solutions like Bitcoin from distributed (and relatively cheap) solutions like CT, and we hope that our results can help to inform future decisions about which protocol to adopt.

1.2 Related work.

We consider research that is related both in terms of the applications of Bitcoin and Certificate Transparency, and in terms of the underlying primitives used to construct our transparency overlay.

An emerging line of work has both formalized some of the properties provided by the Bitcoin network and bootstrapped Bitcoin to obtain provably secure guarantees in other settings. Garay et al. [17] analyzed the so-called "backbone" protocol of Bitcoin and prove that it satisfies two important properties as long as the adversary controls some non-majority percentage of the hashing power. Similarly, Bentov and Kumaresan [8] provided a two-party computation built on top of (an abstracted version of) Bitcoin that provably achieves a notion of fairness, and Andrychowicz et al. [3] used Bitcoin to build a provably fair system for multi-party computation. Andrychowicz and Dziembowski [2] further formalized some of the fairness properties they require from Bitcoin (and more generally from systems based on proof-of-work) and used them to construct a broadcast protocol. Finally, on the privacy side, the Zerocash project [6] provides a cryptocurrency that has provable anonymity guarantees, and Garman et al. [18] showed how to adapt the decentralized approach of Bitcoin to achieve anonymous credentials. To the best of our knowledge, ours is the first paper to focus on the transparency property of Bitcoin, rather than its privacy or fairness guarantees.

Aside from CT, a number of other solutions exist for changing the way we interact with certificate authorities; many of these solutions require a ground-up redesign of the CA ecosystem, which is why we chose to examine CT instead and use it as our inspiration for an overlay system. Fromknecht et al. [16] propose a decentralized PKI, based on Bitcoin and Namecoin, that eliminates the trust in centralized authorities altogether. CONIKS [24] provides an approach to logging certificates that differs from CT in two key ways: it focuses on user rather than website certificates, and largely because of this it provides a privacy-preserving solution, in which certain aspects of the stored certificates (e.g., usernames) are kept hidden. The Accountable Key Infrastructure [20] and the related ARPKI [4] both require a distributed infrastructure for not only the storage of issued certificates (as CT does), but also for their issuance, thus focusing on the prevention rather than just detection of misbehavior. Ryan [33] demonstrated how to extend CT to handle revocation of certificates. In a concurrent work, Dowling et al. [15] provided a different security model for CT and demonstrated that if various properties of the underlying Merkle trees are satisfied then CT is provably secure in their model. Although somewhat overlapping with our own work, their paper is focused firmly on CT and not on the abstract properties of transparency overlays and how they can be applied across a variety of settings.

Finally, the main primitive underlying our transparency overlay (a dynamic list commitment) is primarily a generalization of a Merkle tree [26], and is similar to the definition of a tamper-evident log given by Crosby and Wallach [13]. It is also related to the notion of an authenticated data structure (ADS) [1, 32, 31] and the notion of a cryptographic accumulator [7, 12, 11, 23]; indeed the application of ADSs to Bitcoin has already been touched on in previous work [27] (but from the perspective of programming languages, and thus without any consideration of security). Dynamic list

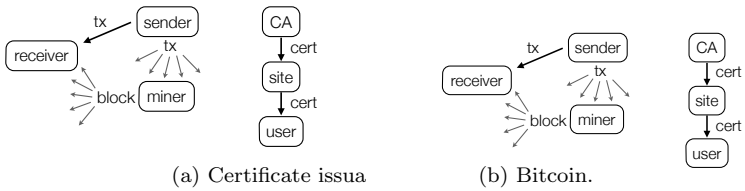

(a) Certificate issua (b) Bitcoin.

Figure 1: The basic structure for each of the settings in which we apply transparency.

commitments differ from these related primitives in terms of the security model, however, and as a result we can provide more efficient constructions while still satisfying a notion of provable security.

2. BACKGROUND

2.1 Certificate Transparency

Certificate Transparency (CT) was proposed in 2011 by Ben Laurie and Adam Langley as a way to increase transparency in the process of issuing certificates, so that certificate authorities (CAs) can be held responsible for their actions and bad certificates can be caught and revoked early on. The basic process of issuing certificates operates as depicted in Figure 1a: a CA issues a certificate to a website operator, who then publishes this certificate so that users can check it. Certificate Transparency then provides an extra layer on top of this basic interaction to provide transparency; in fact, as we will see in Section 4, our design of a transparency overlay is heavily inspired by the CT design.

Briefly, CT introduces three additional actors: a log server, who is responsible for keeping track of issued certificates, an auditor, who is responsible (on behalf of the client) for keeping track of whether given certificates are in the log or not, and a monitor, who is responsible for checking the quality of the certificates in the log. As we use these additional actors in our general transparency overlay, we defer further discussion of their roles and actions to Section 4.1. In Section 5, we prove that CT provides a provably secure transparency overlay, thus (provably) providing the intuitive security properties that one would hope to achieve.

2.2 Bitcoin

Bitcoin is a decentralized cryptocurrency that was introduced in 2008 [28] and deployed on January 3 2009. We briefly sketch the main properties of Bitcoin and its underlying *blockchain* technology here, and refer the reader to Bonneau et al. [9] for a more comprehensive overview.

Briefly, Bitcoin operates as depicted in Figure 1b. A sender, identified using a pseudonym or *address*, has some number of bitcoins stored with this address; i.e., within the Bitcoin network, this address is acknowledged as the owner of these bitcoins. To transfer ownership of these bitcoins to some receiver, the sender first creates a transaction to send them to the receiver, as identified by whichever address she has given to the sender. The transaction is signed to ensure that only the sender can give away his own bitcoins.

After forming this transaction, the sender broadcasts it to the Bitcoin network, where it eventually reaches a miner, who acts to seal the transaction into a block. The miner broadcasts this block, containing the transaction, to the network, where it eventually reaches the receiver, who can confirm the transaction and its position within the Bitcoin ledger (i.e., the blockchain) to satisfy herself that she is now the owner of the bitcoins.

Because the Bitcoin blockchain is globally visible, it already provides a degree of transparency that is higher than that of traditional financial transactions. In Section 6, we apply a transparency overlay on top of Bitcoin and demonstrate that it provides a significantly more efficient way for Bitcoin users to participate in transactions and allows hashing to be eliminated from the system.

3. DEFINITIONS AND NOTATION

In this section, we define various notions that will be used throughout the rest of the paper. In particular, we formalize *dynamic list commitments* in Section 3.2, which can be thought of as a generalization of Merkle trees and allow us to construct high-integrity logs.

3.1 Preliminaries

If x is a binary string then $|x|$ denotes its bit length. If S is a finite set then $|S|$ denotes its size and $x \xleftarrow{r} S$ denotes sampling a member uniformly from S and assigning it to x. $\lambda \in \mathbb{N}$ denotes the security parameter and 1^λ denotes its unary representation. ε denotes the null value.

Algorithms are randomized unless explicitly noted otherwise. "PT" stands for "polynomial-time." By $y \leftarrow A(x_1, \ldots, x_n; R)$ we denote running algorithm A on inputs $x_1, \ldots, x_n$ and random coins R and assigning its output to y. By $y \xleftarrow{r} A(x_1, \ldots, x_n)$ we denote $y \leftarrow A(x_1, \ldots, x_n; R)$ for coins R sampled uniformly at random. By $[A(x_1, \ldots, x_n)]$ we denote the set of values that have positive probability of being output by A on inputs $x_1, \ldots, x_n$. Adversaries are algorithms.

For interactive protocols, we use the notation of Bellare and Keelveedhi [5]. For completeness, we include the formal notion of defining and executing interactive protocols in the full version of the paper. Briefly, the behavior of a stateful participant party that is given m during the i-th round of the j-th execution of a protocol Prot can be defined as $(\mathsf{state}_{\mathsf{party}}, m', p, \mathsf{out}) \xleftarrow{r} \mathsf{Prot}[\mathsf{party}, i, j](1^\lambda, \mathsf{state}_{\mathsf{party}}, m)$, where p indicates the party to which it is sending m'; the execution of the entire interactive protocol can be defined by $\mathsf{outputs} \xleftarrow{r} \mathsf{Run}(1^\lambda, \mathsf{Prot}, \mathsf{Parties}, \mathsf{inputs})$; and the message sent during the protocol (i.e., the transcript) can be defined by $M \xleftarrow{r} \mathsf{Msgs}(1^\lambda, \mathsf{Prot}, \mathsf{Parties}, \mathsf{inputs})$.

We use games in security definitions and proofs. A game G has a MAIN procedure whose output is the output of the game. $\Pr[\mathsf{G}]$ denotes the probability that this output is true.

3.2 Dynamic list commitments

We define a *dynamic list commitment* (DLC), which allows one to commit to a list of elements in such a way that (1) the list represented by the commitment can be updated only by having new elements appended to the end, and (2) given just the list commitment, one can efficiently prove both the append-only property of the list and that a given element is in the list.

One common example of a DLC is a hash tree, and in particular a Merkle tree, in which the root hash acts as the commitment and one can use the hashes of intermediate nodes to prove the above properties. (Indeed, this is what CT uses.) Our basic formalization is similar to the definition

of a tamper-evident history system [13], but we also include an augmented version that considers additional properties one can use when operating on *ordered* lists.

3.2.1 A basic formalization for general lists.

We define a dynamic list commitment DLC as a collection of the following algorithms:

- $c \leftarrow \mathbf{Com(list)}$ creates the commitment c and $\mathbf{0/1} \leftarrow \mathbf{CheckCom}(c, \mathbf{list})$ checks that c is a commitment to list;
- $c_{\mathsf{new}} \leftarrow \mathbf{Append}(\mathbf{list}_{\Delta}, c_{\mathsf{old}})$ updates the commitment to take into account the new elements in list_{Δ};
- $\pi \leftarrow \mathbf{ProveAppend}(c_{\mathsf{old}}, c_{\mathsf{new}}, \mathbf{list})$ proves that c_{new} was obtained from c_{old} solely by appending elements to an earlier version of list and $\mathbf{0/1} \leftarrow \mathbf{CheckAppend}(c_{\mathsf{old}}, c_{\mathsf{new}}, \pi)$ checks this proof;
- $\pi \leftarrow \mathbf{ProveIncl}(c, \mathbf{elmt}, \mathbf{list})$ proves that elmt is in list (as represented by c); and $\mathbf{0/1} \leftarrow \mathbf{CheckIncl}(c, \mathbf{elmt}, \pi)$ checks this proof.

We say that a DLC is *compact* if $\mathsf{Com}(\mathsf{list})| \ll |\mathsf{list}|$ for all sufficiently long lists list. Formal definitions of the basic security properties of a DLC can be found in the full version of the paper. Informally, a DLC should be

1. *binding*, which means that a commitment cannot represent two different lists, so an adversary should be unable to output a commitment c and two lists list_1 and list_2 such that c represents both lists (i.e., $\mathsf{CheckCom}(c, \mathsf{list}_1) = \mathsf{CheckCom}(c, \mathsf{list}_2) = 1$) but they are not equal;

2. *sound*, which means that it should be hard to produce a proof of inclusion for an element not in the list, so an adversary should be unable to output a commitment c, list list, element elmt, and proof π such that c represents list, $\mathsf{CheckIncl}(c, \mathsf{elmt}, \pi) = 1$, but $\mathsf{elmt} \notin \mathsf{list}$; and

3. *append-only*, which means that it should be hard to produce a proof that a list has been only appended to, so an adversary should be unable to produce a list list_2, two commitments c_1 and c_2, and a proof π such that c_2 represents list_2, $\mathsf{CheckAppend}(c_1, c_2, \pi) = 1$, but c_1 is not a commitment to any prefix of list_2.

3.2.2 An augmented formalization for ordered lists.

It will also be useful for us to consider a special type of DLC, in which the elements in the list have some kind of order imposed on them. In particular, this allows us to more efficiently perform two additional operations: demonstrate that two DLCs are inconsistent (i.e., that they are commitments to strictly distinct or forking lists), and demonstrate that a given element is not in the list represented by a given commitment. As we will see in our applications later on, these operations are crucial for providing evidence that certain types of misbehavior have taken place.

In addition to the algorithms required for a basic DLC, we now require a notion of timing (which may not be the actual time, but rather any representation that allows us to impose an ordering): for every element elmt in a list, we assume there exists a function $\mathsf{time}(\cdot)$ that returns a value t, and that a global ordering exists for this function, so that we can also define a Boolean function $0/1 \leftarrow \mathsf{isOrdered}(\mathsf{list})$. Using this, we define a notion of *consistency* for DLCs as follows:

DEFINITION 3.1. *A tuple* (c, t, list) *is* consistent *if* c *is a commitment to the state of* list *at time* t*. Formally, we consider a function* isConsistent *such that* $\mathsf{isConsistent}(c, t, \mathsf{list}) = 1$ *if and only if there exists a* j*,* $1 \le j \le \mathsf{len}(\mathsf{list})$*, such that (1)* $\mathsf{CheckCom}(c, \mathsf{list}[1 : j]) = 1$*, (2)* $\mathsf{time}(\mathsf{list}[j]) \le t$*, (3)* $j = \mathsf{len}(\mathsf{list})$ *or* $\mathsf{time}(\mathsf{list}[j+1]) \ge t$*), and (4)* $\mathsf{isOrdered}(\mathsf{list})$*.*

We can now define four additional algorithms as follows:

- $\pi \leftarrow \mathbf{DemoInconsistent}(\mathbf{list}, c', t')$ proves that list is inconsistent with c' at time t' and
- $\mathbf{0/1} \leftarrow \mathbf{CheckInconsistent}(c', t', c, \pi)$ checks this proof;
- $\pi \leftarrow \mathbf{DemoNotIncl}(\mathbf{list}, \mathbf{elmt})$ proves that elmt is not in the ordered list list; and
- $\mathbf{0/1} \leftarrow \mathbf{CheckNotIncl}(c, \mathbf{elmt}, \pi)$ checks this proof.

Formal definitions of the augmented security properties can be found in the full version of the paper. Informally, in the augmented setting a DLC should satisfy

1. *provable inconsistency*, which means any inconsistent tuple should be demonstrably inconsistent, so an adversary should be unable to produce a tuple (c, t, list) such that the tuple is inconsistent but an honestly generated proof of inconsistency fails verification;

2. *provable non-inclusion*, which means it should be possible to demonstrate that an element is not in a list, so an adversary should be unable to produce a list list and an element elmt such that $\mathsf{elmt} \notin \mathsf{list}$ but the honestly generated proof of non-inclusion fails verification;

3. *unforgeable inconsistency*, which means it should be impossible to demonstrate an inconsistency that does not exist, so an adversary should be unable to produce $(c_1, t, c_2, \mathsf{list}_2, \pi)$ such that c_2 represents list_2, c_1 is consistent with list_2 at time t, and $\mathsf{CheckInconsistent}(c_1, t, c_2, \pi) = 1$; and

4. *unforgeable non-inclusion*, which means it should be impossible to prove non-inclusion of an element that is in a list, so an adversary should be unable to produce $(c, \mathsf{list}, \mathsf{elmt}, \pi)$ such that c represents list, $\mathsf{elmt} \in \mathsf{list}$, and $\mathsf{CheckNotIncl}(c, \mathsf{elmt}, \pi) = 1$.

3.2.3 Two instantiations of augmented DLCs.

To demonstrate that dynamic list commitments exist, we provide two instantiations; both can be found in the full version of the paper and derive their security from the collision resistance of a hash function. Briefly, our first instantiation is essentially a rolling hash chain: new elements appended to the list are folded into the hash (i.e., $c_{\mathsf{new}} \leftarrow H(c_{\mathsf{old}} \| \mathsf{elmt}_{\mathsf{new}})$), and proofs about (in)consistency and (non-)inclusion reveal selective parts of the list. This first instantiation thus demonstrates the feasibility of dynamic list commitments (and is conceptually quite simple), but the proofs are linear in the size of the list, which is not particularly efficient. Thus, our second instantiation is essentially a Merkle tree, which allows us to achieve proofs that are logarithmic in the size of the list.

4. TRANSPARENCY OVERLAYS

In this section, we present our main contributions. First, in Sections 4.1 and 4.2, we introduce both basic and augmented formal models for reasoning about transparency.

Then, in Sections 4.3 and 4.4 we present a generic transparency overlay and prove its security. To instantiate this securely (as we do in Sections 5 and 6), one then need only provide a simple interface between the underlying system and the overlay.

4.1 Basic overlays

In order for a system to be made transparent, we must provide an efficient mechanism for checking that the system is running correctly. Our setting overlays three additional parties on top of an existing system Sys: a *log server* LS, an *auditor* Auditor, and a *monitor* Monitor. The role of the log server is to take certain events in the system's operation and enter them into a publicly available log. The role of the auditor is to check — crucially, without having to keep the entire contents of the log — that specific events are in the log. Finally, the role of the monitor is to flag any problematic entries within the log. Collectively then, the auditor and monitor act to hold actors within the system responsible for the creation of (potentially conflicting) events.

We assume that each of these parties is stateful: the log server maintains the log as state, so $\mathsf{state}_{\mathsf{LS}} = \mathsf{log}$; the auditor maintains a *snapshot* (i.e., some succinct representation of the current log) as state, so $\mathsf{state}_{\mathsf{Au}} = \mathsf{snap}$; and the monitor maintains a snapshot, a list of bad events, and a list of all events, so $\mathsf{state}_{\mathsf{Mo}} = (\mathsf{snap}, \mathsf{events}_{\mathsf{bad}}, \mathsf{events})$.

A transparency overlay then requires five interactive protocols; these are defined abstractly as follows:[1]

`GenEventSet` is an interaction between the actor(s) in the system that produces the events to be logged. The protocol is such that $\mathsf{eventset} \xleftarrow{r} \mathsf{Run}(1^\lambda, \texttt{GenEventSet}, \mathsf{Sys}, \mathsf{aux})$.

`Log` is an interaction between one or more of the actors in the system and LS that is used to enter events into the log. The protocol is such that $(b, \varepsilon) \xleftarrow{r} \mathsf{Run}(1^\lambda, \texttt{Log}, (\mathsf{Sys}, \mathsf{LS}), (\mathsf{eventset}, \varepsilon))$, where b indicates whether or not the system actor(s) believes the log server behaved honestly.

`CheckEntry` is an interaction between one or more of the actors in the system, Auditor, and LS that is used to check whether or not an event is in the log. The protocol is such that $(b, b', \varepsilon) \xleftarrow{r} \mathsf{Run}(1^\lambda, \texttt{CheckEntry}, (\mathsf{Sys}, \mathsf{Auditor}, \mathsf{LS}), (\mathsf{event}, \varepsilon, \varepsilon))$, where b indicates whether or not the system actor(s) believes the event to be in the log and b' indicates whether or not the auditor believes the log server behaved honestly in the interaction.

`Inspect` is an interaction between Monitor and LS that is used to allow the monitor to inspect the contents of the log and flag any suspicious entries. The protocol is such that $(b, \varepsilon) \xleftarrow{r} \mathsf{Run}(1^\lambda, \texttt{Inspect}, (\mathsf{LS}, \mathsf{Monitor}), (\varepsilon, \varepsilon))$, where b indicates whether or not the monitor believes the log server behaved honestly in the interaction.

`Gossip` is an interaction between Auditor and Monitor that is used to compare versions of the log and detect any inconsistencies. If any misbehavior on behalf of the log server is found, then both parties are able to output *evidence* that this has taken place. The protocol is such that $(\mathsf{evidence}, \mathsf{evidence}) \xleftarrow{r} \mathsf{Run}(1^\lambda, \texttt{Gossip}, (\mathsf{Auditor}, \mathsf{Monitor}), (\varepsilon, \varepsilon))$.

We also require the following (non-interactive) algorithms:

$(pk_{\mathsf{LS}}, sk_{\mathsf{LS}}) \xleftarrow{r} \mathsf{GenLogID}(1^\lambda)$ is used to generate a public and secret identifier for the log server; and

$0/1 \leftarrow \mathsf{CheckEvidence}(pk_{\mathsf{LS}}, \mathsf{evidence})$ is used to check if the evidence against the log server identified by pk_{LS} is valid.

From a functionality standpoint, we would like the protocols to be *correct*, meaning all parties should be satisfied by honest interactions, and *compactly auditable*, meaning the size of a snapshot is much smaller than the size of the log.

We define security for a basic transparency overlay in terms of two properties: *consistency*, which says that a potentially dishonest log server cannot get away with presenting inconsistent versions of the log to the auditor and monitor, and *non-frameability*, which says that potentially dishonest auditors and monitors (and even actors in the original system) cannot blame the log server for misbehavior if it has behaved honestly. Participants can thus be satisfied that they are seeing the same view of the log as all other participants, and that the interactions they have really are with the log server.

To formalize consistency, we consider a game in which the adversary takes on the role of the log server and is allowed to interact (via the MsgAu and MsgMo oracles, respectively) with the auditor and monitor. The adversary wins if there is an event that is not in the list maintained by the monitor but that the auditor nevertheless perceives as being in the log (the third winning condition of Definition 4.1), yet the auditor and monitor are unable to produce valid evidence of this inconsistency (the first two winning conditions). For ease of formal exposition, we (1) assume that in the `CheckEntry` protocol the first message sent to the auditor is the event to be checked and the last message sent by the auditor is a bit indicating whether the event is in the log, and (2) require that the monitor must have a newer snapshot than the auditor, but can naturally extend our definition to cover other configurations as well.

Definition 4.1. *Define* $\mathbf{Adv}^{cons}_{trans,\mathcal{A}}(\lambda) = Pr[\mathsf{G}^{cons}_{\mathcal{A}}(\lambda)]$, *where* $\mathsf{G}^{cons}_{\mathcal{A}}(\lambda)$ *is defined as follows:*

$\underline{\text{MAIN } \mathsf{G}^{cons}_{\mathcal{A}}(\lambda)}$
$\mathsf{events} \leftarrow \emptyset$; $\mathsf{events}_{\mathsf{pass}} \leftarrow \emptyset$
$pk_{\mathsf{LS}} \xleftarrow{r} \mathcal{A}^{\text{MsgAu},\text{MsgMo}}(1^\lambda)$
$\mathsf{evidence} \xleftarrow{r} \mathsf{Run}(1^\lambda, \texttt{Gossip}, (\mathsf{Auditor}, \mathsf{Monitor}), (\varepsilon, \varepsilon))$
return $((\mathsf{CheckEvidence}(pk_{\mathsf{LS}}, \mathsf{evidence}) = 0) \quad \wedge$
$(\mathsf{time}(\mathsf{state}_{\mathsf{Mo}}[\mathsf{snap}]) \geq \mathsf{time}(\mathsf{state}_{\mathsf{Au}}[\mathsf{snap}])) \quad \wedge$
$(\mathsf{events}_{\mathsf{pass}} \setminus \mathsf{state}_{\mathsf{Mo}}[\mathsf{events}] \neq \emptyset))$

$\underline{\text{MsgAu}(i, j, m)}$
$(\mathsf{state}_{\mathsf{Au}}, m', p, \mathsf{out}) \xleftarrow{r} \texttt{CheckEntry}[\mathsf{Auditor}, i, j](1^\lambda, \mathsf{state}_{\mathsf{Au}}, m)$
if $(i = 1)$ $\mathsf{events}[j] \leftarrow m$
if $(\mathsf{out} \neq \perp) \wedge (m' = 1)$ $\mathsf{events}_{\mathsf{pass}} \leftarrow \mathsf{events}_{\mathsf{pass}} \cup \{\mathsf{events}[j]\}$
return m'

$\underline{\text{MsgMo}(i, j, m)}$
$(\mathsf{state}_{\mathsf{Mo}}, m', p, \mathsf{out}) \xleftarrow{r} \texttt{Inspect}[\mathsf{Monitor}, i, j](1^\lambda, \mathsf{state}_{\mathsf{Mo}}, m)$
return m'

[1] In each protocol, we also allow the participants to output fail, which indicates that they believe they were given improperly formatted inputs.

Then the transparency overlay satisfies consistency *if for all PT adversaries $\mathcal{A}$ there exists a negligible function $\nu(\cdot)$ such that $\mathbf{Adv}^{cons}_{trans,\mathcal{A}}(\lambda) < \nu(\lambda)$.*

Next, to formalize non-frameability, we consider an adversary that wants to frame an honest log server; i.e., to produce evidence of its "misbehavior." In this case, we consider a game in which the adversary takes on the role of the auditor, monitor, and any actors in the system, and is allowed to interact (via the MSG oracle) with the honest log server. The adversary wins if it is able to produce evidence that passes verification.

DEFINITION 4.2. *Define* $\mathbf{Adv}^{frame}_{trans,\mathcal{A}}(\lambda) = Pr[\mathsf{G}^{frame}_{\mathcal{A}}(\lambda)]$, *where* $\mathsf{G}^{frame}_{\mathcal{A}}(\lambda)$ *is defined as follows:*

$\underline{\text{MAIN } \mathsf{G}^{frame}_{\mathcal{A}}(\lambda)}$
$(pk_{\mathsf{LS}}, sk_{\mathsf{LS}}) \xleftarrow{r} \mathsf{GenLogID}(1^\lambda)$
$\mathsf{evidence} \xleftarrow{r} \mathcal{A}^{\text{MSG}}(1^\lambda, pk_{\mathsf{LS}})$
return $\mathsf{CheckEvidence}(pk_{\mathsf{LS}}, \mathsf{evidence})$

$\underline{\text{MSG}(\mathsf{Prot}, i, j, m)}$
if $(\mathsf{Prot} \notin \{\mathsf{Log}, \texttt{CheckEntry}, \texttt{Inspect}\})$ *return* $\perp$
$(\mathsf{state}_{\mathsf{LS}}, m', p, \mathsf{out}) \xleftarrow{r} \mathsf{Prot}[\mathsf{LS}, i, j](1^\lambda, \mathsf{state}_{\mathsf{LS}}, m)$
return m'

Then the transparency overlay satisfies non-frameability *if for all PT adversaries $\mathcal{A}$ there exists a negligible function $\nu(\cdot)$ such that $\mathbf{Adv}^{frame}_{trans,\mathcal{A}}(\lambda) < \nu(\lambda)$.*

We then say that a basic transparency overlay is *secure* if it satisfies consistency and non-frameability.

Comparison with concurrent work.

With respect to the security model of Dowling et al. [15], their model requires only that the monitor and auditor produce evidence of misbehavior in the case where the log fails to include an event for which it has issued a receipt (which we consider in the next section). Our model, on the other hand, also produces evidence in the case where the log has given inconsistent views to the two parties; this type of evidence seems particularly valuable since this type of misbehavior is only detected after the fact. This difference allows them to present a simpler definition of non-frameability, as they do not have to worry about malicious monitors and auditors forging this type of evidence.

4.2 Pledged overlays

In the basic setting described, log servers can be held responsible if they attempt to present different views of the log to the auditor and monitor. If log servers simply fail to include events in the log in the first place, however, then there is currently no way to capture this type of misbehavior. While in certain settings the log server could plausibly claim that it never received an event rather than ignoring it, if the log server issues promises or *receipts* to include events in the log then we can in fact enforce inclusion, or at least blame the log server if it fails to do so.

Formally, we capture this as an additional security property, *accountability*, which says that evidence can also be used to implicate log servers that promised to include events but then did not. In the game (which we defer to the full version of the paper — included as supplemental material — due to the formal notational overhead), the adversary then takes on the role of the log server and is allowed to interact arbitrarily with the actor(s) in the system, auditor, and monitor. It wins if there is an event that it has pledged to include but that the auditor and monitor do not believe to be in the log, yet the auditor and monitor are unable to produce evidence of this omission.

We then say that a pledged transparency overlay is *secure* if it satisfies consistency, non-frameability, and accountability.

4.3 A generic pledged transparency overlay

We now present a generic version of a pledged transparency overlay. We begin by introducing algorithms for performing various operations in the overlay, and then describe the interactive protocols from Section 4.1 in terms of these algorithms and the algorithms for a dynamic list commitment (DLC) and a signature scheme ($\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify}$). For ease of exposition we assume that various objects (snapshots, receipts, etc.) contain only the fields necessary to make the protocol work, but could naturally extend our algorithms to cover more general configurations as well.

To start, an event set $\mathsf{eventset}$ contain (at least) a list of events events; a snapshot $\mathsf{snap} = (c, t, \sigma)$ contains a DLC, timing information, and an unforgeable signature; a receipt $\mathsf{rcpt} = (pk, t, \sigma)$ contains a public key, timing information, and an unforgeable signature; and a log $\mathsf{log} = (\mathsf{snap}, \mathsf{events})$ contains a snapshot and a list of events. We denote these subcomponents using bracket notation; e.g., we use $\mathsf{snap}[c]$, or — where subscripts make it appropriately clear — use c_i to denote $\mathsf{snap}_i[c]$.

To perform basic operations on these objects, we also introduce algorithms for forming and verifying snapshots and receipts, and for updating the log. These are defined — with respect to a notion of timing t and a keypair $(pk_{\mathsf{LS}}, sk_{\mathsf{LS}})$ — as follows:

$\underline{\mathsf{FormSnap}(c, t)}$
return $(c, t, \mathsf{Sign}(sk_{\mathsf{LS}}, (c, t)))$

$\underline{\mathsf{CheckSnap}(\mathsf{snap})}$
return $\mathsf{Verify}(pk_{\mathsf{LS}}, (\mathsf{snap}[c], \mathsf{snap}[t]), \mathsf{snap}[\sigma])$

$\underline{\mathsf{FormRcpt}(\mathsf{log}, \mathsf{event})}$
return $(pk_{\mathsf{LS}}, t, \mathsf{Sign}(sk_{\mathsf{LS}}, (t, \mathsf{event})))$

$\underline{\mathsf{CheckRcpt}(\mathsf{event}, \mathsf{rcpt})}$
return $\mathsf{Verify}(pk_{\mathsf{LS}}, (\mathsf{rcpt}[t], \mathsf{event}), \mathsf{rcpt}[\sigma])$

$\underline{\mathsf{UpdateLog}(\mathsf{log}, \mathsf{events})}$
$\mathsf{events}' \leftarrow \mathsf{log}[\mathsf{events}] \| \mathsf{events}$
$c' \leftarrow \mathsf{Append}(\mathsf{events}, \mathsf{log}[\mathsf{snap}][c])$
$\mathsf{snap}' \leftarrow \mathsf{FormSnap}(c', t)$
return $(\mathsf{snap}', \mathsf{events}')$

Briefly, in the `Log` protocol (Figure 2), an event set is given as input to the actor(s) in the system; this is created by `GenEventSet`, which we describe for our individual applications in Sections 5 and 6 but leave here as an abstract interaction. This event set is sent to the log server, who first checks if it is well formed. The log server then provides a receipt for every event in the set, and sends the receipts back to the system actor(s). If any of the receipts are invalid, the system rejects the interaction, and otherwise it

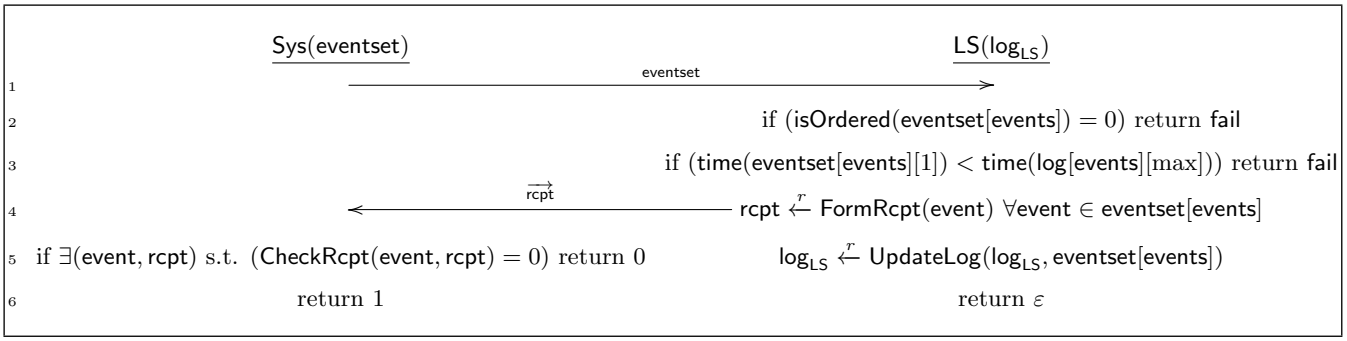

Figure 2: The `Log` protocol for pledged transparency overlays.

accepts. Either way, the log server updates the log; in our protocol specification here, the log server updates the log immediately, but we discuss in Section 5.3 how this process can be batched and the promises of the log server altered accordingly.

Next, in the `CheckEntry` protocol (Figure 3), some actor in the system sends an event and a receipt to the auditor, who first checks that the receipt is valid. If it is, then the auditor checks if the event already falls within its current purview; i.e., if it falls within the log that the auditor already knows about (according to its snapshot). If it does, then the auditor skips to asking the log server for a proof of inclusion of this event; if not, the auditor must update its snapshot and check that the new snapshot is consistent with the old one. Once the auditor has the proof of inclusion (either after updating or not), it returns to the client whether or not the proof verifies; the client returns whatever it receives from the auditor, and the auditor returns $b = 0$ if the protocol has failed in some way (i.e., the updated snapshot was inconsistent with the old one) and $b = 1$ otherwise.

Next, in the `Inspect` protocol (Figure 4), the monitor sends its current snapshot to the log server, and the log server responds with all events that have been logged since then, along with an updated snapshot. If this list of appended events is valid (i.e., ordered and consistent with the new snapshot), the monitor can update its records and look for any bad events in this new list. It returns $b = 1$ if the protocol has gone smoothly; i.e., if the new list and snapshot seem to have been formed appropriately.

Finally, in the `Gossip` protocol (Figure 5), the auditor and monitor begin by exchanging snapshots, and by ensuring that each snapshot is valid. The monitor then attempts to demonstrate any inconsistencies between the two snapshots (i.e., demonstrate that they represent forking or distinct logs) and — if any inconsistencies do exist — this is returned as evidence of the log server's misbehavior.

To augment the protocol for pledged overlays, we include in Figure 5 a further optional interaction in which the auditor sends to the monitor all events for which the `CheckEntry` protocol failed, to see if they are being monitored; these are stored in a list $\mathsf{events}_{\mathsf{bad}}$ that is now part of the auditor's state and updated in the `CheckEntry` protocol (line 15 of Figure 3). This allows the auditor and monitor to detect and provide evidence for the additional type of misbehavior in which the log server simply drops events from the log. This means that the auditor and monitor can provide two types of evidence: evidence that the log server presented them with forked or distinct views of the log, or evidence that the log server reneged on the promise it gave in a receipt. We thus instantiate the algorithm CheckEvidence as follows:

$\underline{\mathsf{CheckEvidence}(pk_{\mathsf{LS}}, \mathsf{evidence})}$
if $(\mathsf{evidence} = \bot)$ return 0
$(\mathsf{snap}_1, \mathsf{snap}_2, (\mathsf{event}, \mathsf{rcpt}), \pi) \leftarrow \mathsf{evidence}$
if $(\mathsf{CheckSnap}(\mathsf{snap}_1) = 0)$ return 0
if $(\mathsf{CheckSnap}(\mathsf{snap}_2) = 0)$ return 0
if $((\mathsf{event}, \mathsf{rcpt}) = (\bot, \bot))$ return
  $(\mathsf{CheckInconsistent}(c_1, t_1, c_2, \pi) \wedge (t_1 \leq t_2))$
return $(\mathsf{CheckRcpt}(\mathsf{event}, \mathsf{rcpt}) \wedge (\mathsf{rcpt}[t] \leq t_2) \wedge$
  $\mathsf{CheckNotIncl}(c_2, \mathsf{event}, \pi)$

Finally, our gossip protocol assumes the monitor has a more up-to-date snapshot than the auditor, which protects against an adversarial log server trivially winning the consistency game (Definition 4.1) by ignoring the monitor. One could also imagine a protocol in which the monitor pauses, updates (using the `Inspect` protocol), and then resumes its interaction with the auditor, in which case the extra winning condition in Definition 4.1 could be dropped.

THEOREM 4.3. *If the DLC is secure in the augmented setting and the signature scheme is unforgeable (i.e., EUF-CMA secure), then the protocols presented in Figures 2-5 and the algorithms presented above comprise a secure pledged transparency overlay, as defined in Section 4.2.*

A proof of this theorem can be found in the full version of the paper. Briefly, consistency follows from three properties of the dynamic list commitment: provable inconsistency, append-only, and soundness. Together, these ensure that if the log server presents an inconsistent view of the log to the auditor and monitor, then the commitment seen by the auditor in its snapshot — which, crucially, was updated using ProveAppend and used to demonstrate the inclusion of events — is inconsistent with the list seen by the monitor. By provable inconsistency, the monitor can thus provide a proof of inconsistency that comprises valid evidence of the log server's misbehavior. Non-frameability, on the other hand, follows from the unforgeability of the signature scheme and from the difficulty of forging either a proof of inconsistency or a proof of non-inclusion. Finally, accountability follows from the provable non-inclusion of the DLC,

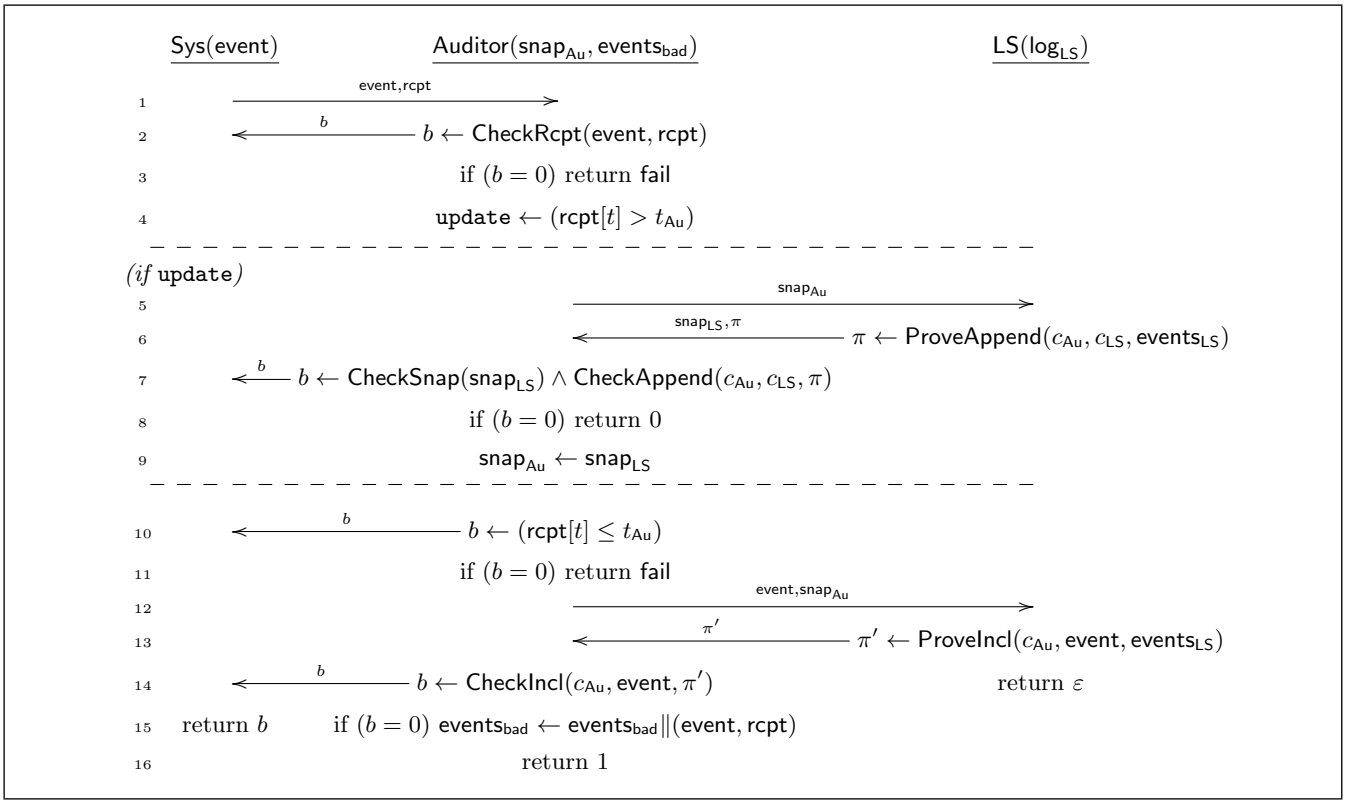

Figure 3: The `CheckEntry` protocol for pledged transparency overlays. The parts of the protocol that may not be carried out (depending on the 'if' clause) are marked with dashed lines.

as if an event is missing from the log then the auditor and monitor should be able to provide valid evidence of this (in the form of a receipt promising to include a given event and a proof of non-inclusion of that event).

4.4 A generic basic transparency overlay

A basic transparency overlay is essentially a simpler version of a pledged transparency overlay, so we do not give a full description of the protocols here, but instead describe the necessary modifications that must be made.

The most obvious modification is that all of the parts that involve receipts do not exist in the basic variant. Thus, the `Log` protocol for a basic transparency overlay omits lines 4-5 from Figure 2 but otherwise remains the same. Next, in the `CheckEntry` protocol, the auditor now cannot use the receipt to check if it needs to update, so it must use time(event) instead. The `Inspect` protocol contains no mention or usage of receipts, and thus is exactly the same in the basic variant. This leaves the `Gossip` protocol, in which the only significant modification is that basic transparency overlays cannot provide the second type of evidence (in which the auditor and monitor use the receipt to prove that the log server promised to include an event but then did not), so do not attempt to produce it (lines 9-12 of Figure 5). This also means that evidence consists only of the two snapshots and a proof.

As the basic transparency overlay thus involves only minor modifications to the pledged transparency overlay, we do not prove its security from scratch, but instead prove the following theorem as a special case of Theorem 4.3.

THEOREM 4.4. *If the DLC is secure in the augmented setting and the signature scheme is unforgeable (i.e., EUF-CMA secure), then the modified protocols and algorithms described above comprise a secure basic transparency overlay, as defined in Section 4.1.*

5. CERTIFICATE TRANSPARENCY

In this section, we describe how CT instantiates a pledged transparency overlay (as defined formally in Section 4.2), discuss how the formal notions of overlay security imply more intuitive notions of security specific to the setting of issuing certificates, and finally discuss the requirements of a practical deployment of CT.

5.1 CT is a secure pledged overlay

As depicted in Section 2, Certificate Transparency has three actors in the system Sys: a certificate authority CA, a website owner Site, and a client Client. One of the first two actors must participate in the `Log` protocol,[2] to ensure that the certificate issued by CA to Site ends up in the log, and the client participates in the `CheckEntry` protocol to check that the certificate presented to it by a website is in the log.

In the parlance of CT, an event is a (basic) certificate $\mathsf{cert} = (pk_{\mathsf{name}}, \sigma_{\mathsf{CA}})$, where σ_{CA} is the CA's signature on

[2]This means that either the website obtains the signed certificate from the CA and then goes on to enter it into the log, or the CA signs the certificate and enters it into the log before returning the extended certificate to the website.

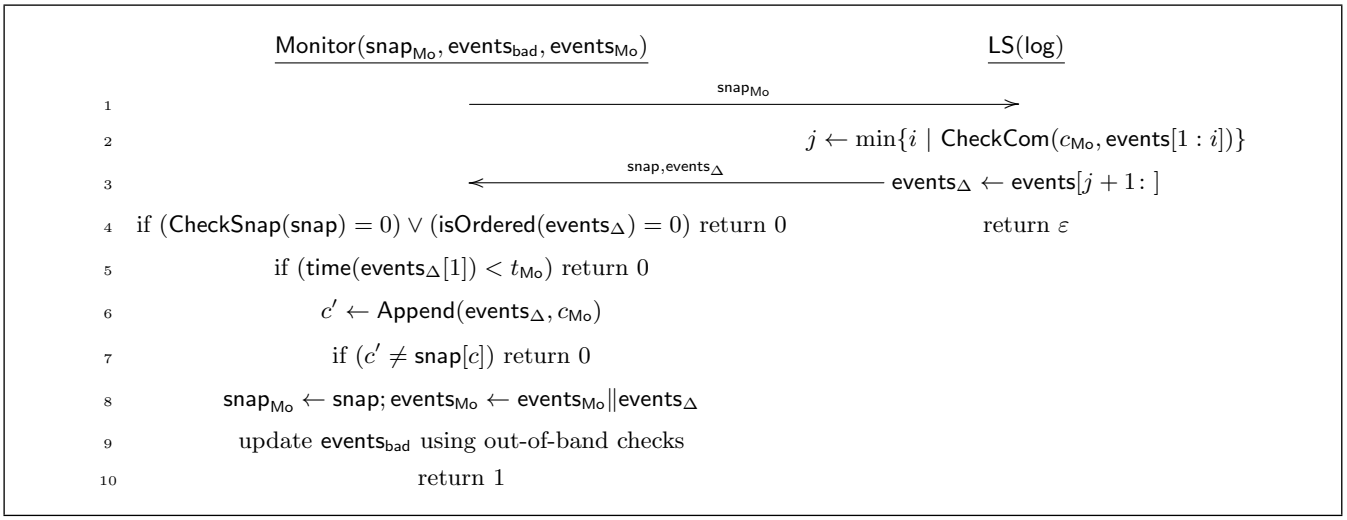

Figure 4: The `Inspect` protocol.

the site's public key pk_{name},[3] a receipt is a *signed certificate timestamp* (SCT), and a snapshot is a *signed tree head* (STH). For the notion of timing needed for snapshots and receipts, one could pick the current local time of the log server and either use this value directly as t or incorporate into it some buffer period, which is referred to in the CT documentation as the *maximum merge delay* (MMD). We discuss this further in Section 5.3. Finally, CT instantiates `GenEventSet` as follows:

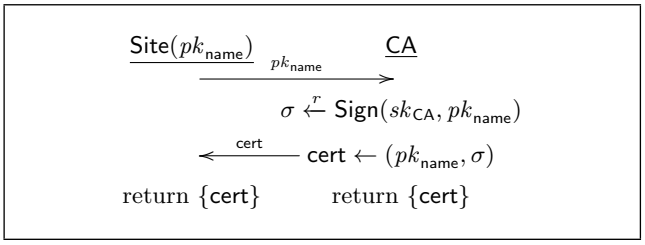

The rest of the protocols needed for the transparency overlay can be instantiated exactly as in Section 4.3, so Theorem 4.3 carries over directly and we can see that CT provides a secure pledged transparency overlay.

5.2 Further security implications

We have just demonstrated that CT provides a secure transparency overlay, but it is not clear what this means for the specific setting of certificate issuance. To explore this, we first remind ourselves of the security of the underlying system (i.e., the issuance of basic certificates), in which (1) it should be difficult to produce a basic certificate without contacting the CA, and (2) an honest client should accept only (basic) certificates that verify. These are clearly satisfied assuming the correctness and unforgeability of the signature scheme.

Combining the underlying issuance security with the security of the overlay, we can argue that three more intuitive security goals are largely satisfied. First, **an *extended* certificate (i.e., a certificate augmented with an SCT) should not pass verification if it has not been jointly produced by the CA and log server.** This holds because the underlying issuance security implies that it is difficult to produce cert without the CA, and non-frameability implies that it is difficult to produce rcpt without the log server, so it should be difficult to produce (cert, rcpt) without both the CA and the log server.

Second, **honest clients shouldn't accept "bad" certificates; i.e., certificates that are either improperly formatted or not being monitored.** The underlying issuance security says that if cert does not verify then the client won't accept. Following this, the honest client accepts only if the auditor indicates that the certificate is in the log. By consistency, the auditor's view of the log is consistent with the monitor's view from the last time they engaged in the `Gossip` protocol (unless evidence has been produced to the contrary, at which point we can assume the auditor ceases communication with the log server). If the certificate is older than this, then the certificate is definitely being monitored; if it is newer, then it is not guaranteed that the certificate is being monitored, but if it is not then the auditor can at least detect this during its next iteration of the `Gossip` protocol. Thus, honest clients never accept improperly formatted certificates, and are unlikely to accept unmonitored certificates provided that the auditor and monitor are engaging in the `Gossip` protocol with sufficient frequency.

Finally, **if a log server is misbehaving by omitting certificates from the log that it promised to include, it should be possible to blame it.** If a log server refuses to answer queries, then there is little we can do about this in the context of our overlay (although in a practical setting with more than one log server this problem could be mitigated). If a log server does answer, then it can be formally blamed by accountability, as the SCT acts as non-repudiable evidence that the log server has promised to include a certificate and the corresponding proof of non-inclusion demonstrates that it has not done so.

[3]For simplicity, we include here only the most basic version of the information that needs to be checked for and included in a certificate.

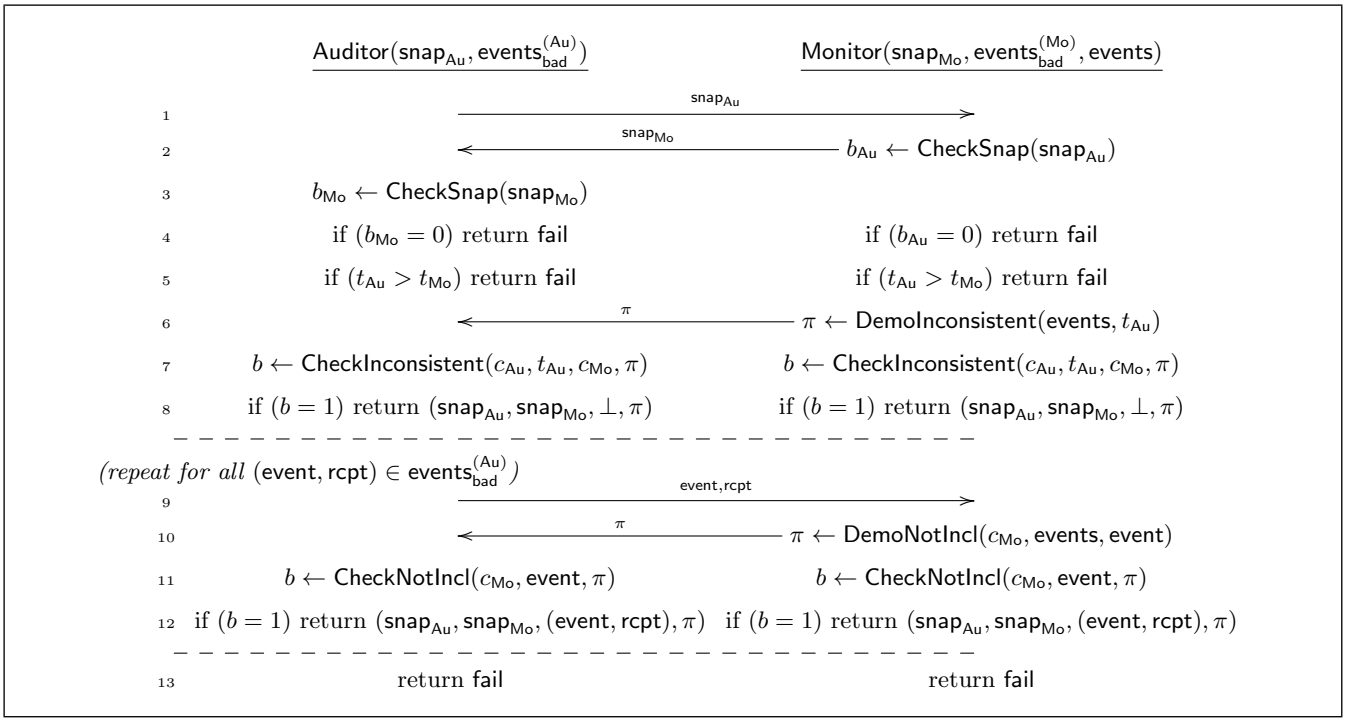

Figure 5: The `Gossip` protocol for pledged transparency overlays. The optional part of the protocol is marked with dashed lines.

5.3 Practical considerations

Finally, we discuss some necessary alterations that would be made to our protocol if used in a real deployment.

Batched additions to the log. In Figure 2, the log server currently updates the log during the `Log` protocol, and as a result includes the exact current time in the SCT. To avoid doing this operation every time, this process would be batched, so the time in the SCT would instead be some time in the near future (e.g., the end of the current day). This gap between the current and promised times is referred to in the CT documentation as the *maximum merge delay* (MMD).

Collapsing the overlay into the system. As discussed in the CT documentation, in a real deployment we expect auditors to interact with many different log servers (as the certificates seen by clients may be logged in many different places), but expect monitors to focus on one log and the certificates it contains. There are therefore two possible models: in one, the auditors and monitors are operated as separate services, and monitors can even be used as backup log servers. In the other, the role of the auditor could collapse into the client (e.g., it could be run as a browser extension and responses could be cached), and the role of the monitor could collapse (at least partially) into the website, who could monitor the log to at least keep track of its own certificates.

Privacy concerns. While SSL certificates are public and thus storing them in a public log presents no privacy concern, information might be revealed about individual users through the certificates queried by the auditor (to both the log server and monitor), as well as the choice of signed tree heads and SCTs. We view this as an interesting area for future research, but mention briefly that some of these concerns can be mitigated — with minimal effect on the security of the transparency overlay — by omitting the optional part of Figure 5, in which the auditor reveals to the monitor some of the certificates that it has seen.

6. AMPLIFYING BITCOIN'S SECURITY

Although Bitcoin already provides a large degree of transparency — as its transaction ledger, called the blockchain, is globally visible — it does not satisfy the requirements of a transparency overlay. In particular, the miners, who play a role analogous to the log server in producing the blockchain, are not known entities and thus cannot be held responsible; this in turn means that consistency and non-frameability cannot be satisfied. In this section, we thus begin by presenting in Section 6.1 a secure basic transparency overlay for Bitcoin.

One might naturally wonder whether such a distinction is purely pedantic; i.e., if overlaying transparency on top of a transparent system provides any actual benefits. To answer this question in the affirmative, we discuss in Section 6.2 the benefits (in terms of both security and efficiency) that are achieved by applying the transparency overlay. In particular, we show that the addition of a secure transparency overlay relieves regular Bitcoin users (i.e., users wishing only to spend and receive bitcoins) from having to store and verify the entire Bitcoin blockchain, which as of this writing is over 80GB.[4] To go even further, we argue that if one is willing to adopt a distributed rather than a fully decentralized solution (i.e., if one is willing to trust any set of named parties), then the entire Bitcoin system collapses into a CT-like

[4]https://blockchain.info/charts/blocks-size

transparency overlay and the need for hash-based mining is eliminated.

6.1 A transparency overlay for Bitcoin

As depicted in Section 2, Bitcoin has three actors in the system Sys: a sender Sender, a receiver Receiver, and a miner Miner. The sender and the miner must participate in the Log protocol to enter transactions into the log (although really this can be done by only the miner, after it has collected all relevant transactions), and the receiver participates in the `CheckEntry` protocol to check that the transaction in which it should be receiving bitcoins is in the log. Our transparency overlay for Bitcoin then instantiates `GenEventSet` as follows:

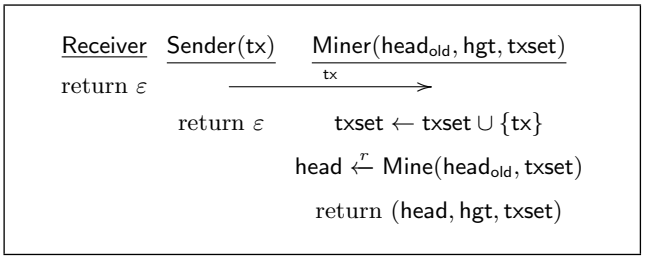

An event is a *transaction* tx, which must have a certain structure (i.e., lists of input and output addresses) and satisfy certain requirements (i.e., that is does not represent double-spending). A set of events eventset is a *block*, which contains not only a list of transactions txset but also a hash head, a pointer $\mathsf{head}_{\mathsf{prev}}$ to the previous block, and a height hgt; combining events in an event set also allows us to impose the required notion of timing, which is the block height hgt. By combining `GenEventSet` with the modified protocols described in Section 4.4, we can thus apply Theorem 4.4 to get a secure basic transparency overlay in the setting of Bitcoin.

6.2 Further security implications

By applying a transparency overlay to Bitcoin, we have provided a method for achieving provable transparency guarantees in this setting. We have also achieved (in a manner similarly observed by Miller et al. [27], although they did not provide any security guarantees) a much more efficient version of the system: senders and receivers now store nothing (or, if the auditor collapses into the users as discussed in Section 5.3 for CT, they store a snapshot), as compared to the entire blockchain or set of block headers that they were required to store previously. While this goal was of course already achievable by Bitcoin senders and receivers using web solutions (i.e., storing their bitcoins in an online wallet), our system is the first to achieve this goal with any provable security guarantees, thus minimizing the trust that such users must place in any third party.

Our analysis also has implications beyond users' storage of the blockchain. To go beyond our initial attempt at an overlay (which we dub the "naïve overlay" in Table 1), one might observe that the miner provides no additional value beyond that of the log server: whereas in CT the CA was necessary to provide a signature (and more generally is assumed to perform external functions such as verifying the owner of a website), here the miner just collates the transactions and sends them to the log server. By having senders contact log servers directly, one could therefore eliminate entirely the role of mining without any adverse effects on security. Thus,

	Bitcoin	Naïve overlay	CT-like overlay
Hashing	yes	yes	no
Set of miners	decentralized	hybrid*	distributed
Broadcast	yes	yes	no
Provable security	no	yes*	yes

Table 1: The different tradeoffs between Bitcoin, our naïve overlay, and a "CT-like" overlay in which log servers completely replace miners. Our naïve solution provides the same openness that Bitcoin has for miners but also provable security guarantees for those who make (optional) use of distributed log servers, while our CT-like solution requires trust in the set of log servers but achieves both provable security and significantly better efficiency.

if users are willing to make the trust assumptions necessary for our transparency overlay — namely, to assume that some honest majority of a distributed set of log servers provide the correct response about the inclusion of a transaction — then the system can collapse into a distributed structure (the "CT-like overlay" in Table 1) in which no energy is expended to produce the ledger, and users have minimal storage requirements. Moreover, if users communicate directly with the log server, then we could add a signed acknowledgment from the log server that would allow us to satisfy accountability. Interestingly, this solution closely resembles the recent RSCoin proposal [14] (but with our additional consistency and non-frameability guarantees), which achieves linear scaling in transaction throughput; this provides additional validation and suggests that this distributed approach presents an attractive compromise between the two settings.

7. CONCLUSIONS AND OPEN PROBLEMS

In this paper, we initiated a formal study of transparency overlays by providing definitions and a generic secure construction of this new primitive. To demonstrate the broad applicability of our generic formalization, we proved that Certificate Transparency (CT) is a secure transparency overlay, and presented a Bitcoin-based transparency overlay that achieves provable notions of security and significantly reduces the storage costs of regular Bitcoin users. Our comparison reveals that in any settings where distributed trust is possible (i.e., one is willing to trust any set of known participants), Bitcoin can collapse into CT and the need for both mining and the storage of the blockchain disappears. On the other hand, if one is not willing to trust anyone, then on a certain level these requirements seem inevitable.

While our constructions provide provably secure properties concerning integrity, it is not clear how our transparency overlay could provide this same value to any system in which a meaningful notion of privacy is required. It is thus an interesting open problem to explore the interaction between transparency and privacy, and in particular to provide a transparency overlay that preserves any privacy guarantees of the underlying system.

Acknowledgments

Sarah Meiklejohn is supported in part by EPSRC Grant EP/M029026/1.

8. REFERENCES

[1] A. Anagnostopoulos, M. T. Goodrich, and R. Tamassia. Persistent authenticated dictionaries and

their applications. In G. I. Davida and Y. Frankel, editors, *ISC 2001*, volume 2200 of *LNCS*, pages 379–393, Malaga, Spain, Oct. 1–3, 2001. Springer, Berlin, Germany.

[2] M. Andrychowicz and S. Dziembowski. Pow-based distributed cryptography with no trusted setup. In *Proceedings of Crypto 2015*, 2015.

[3] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek. Secure multiparty computations on Bitcoin. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2014.

[4] D. Basin, C. Cremers, T. H.-J. Kim, A. Perrig, R. Sasse, and P. Szalachowski. ARPKI: Attack Resilient Public-Key Infrastructure. In *Proceedings of ACM CCS 2014*, pages 382–393, 2014.

[5] M. Bellare and S. Keelveedhi. Interactive message-locked encryption and secure deduplication. In *Proceedings of PKC 2015*, volume 9020 of *LNCS*, pages 516–538, 2015.

[6] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2014.

[7] J. C. Benaloh and M. de Mare. One-way accumulators: A decentralized alternative to digital sinatures (extended abstract). In T. Helleseth, editor, *EUROCRYPT'93*, volume 765 of *LNCS*, pages 274–285, Lofthus, Norway, May 23–27, 1993. Springer, Berlin, Germany.

[8] I. Bentov and R. Kumaresan. How to use bitcoin to design fair protocols. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 421–439, Santa Barbara, CA, USA, Aug. 17–21, 2014. Springer, Berlin, Germany.

[9] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten. Research perspectives and challenges for Bitcoin and cryptocurrencies. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2015.

[10] P. Bright. Independent Iranian hacker claims responsibility for Comodo hack, Mar. 2011.

[11] J. Camenisch, M. Kohlweiss, and C. Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In S. Jarecki and G. Tsudik, editors, *PKC 2009*, volume 5443 of *LNCS*, pages 481–500, Irvine, CA, USA, Mar. 18–20, 2009. Springer, Berlin, Germany.

[12] J. Camenisch and A. Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 61–76, Santa Barbara, CA, USA, Aug. 18–22, 2002. Springer, Berlin, Germany.

[13] S. Crosby and D. Wallach. Efficient data structures for tamper-evident logging. In *Proceedings of the 18th USENIX Security Symposium*, 2009.

[14] G. Danezis and S. Meiklejohn. Centrally banked cryptocurrencies. In *Proceedings of NDSS 2016*, 2016.

[15] B. Dowling, F. Günther, U. Herath, and D. Stebila. Secure logging schemes and Certificate Transparency. In *Proceedings of ESORICS 2016*, 2016. To appear.

[16] C. Fromknecht, D. Velicanu, and S. Yakoubov. A decentralized public key infrastructure with identity retention. IACR Cryptology ePrint Archive, Report 2014/803, 2014. http://eprint.iacr.org/2014/803.pdf.

[17] J. Garay, A. Kiayias, and N. Leonardos. The Bitcoin backbone protocol: Analysis and applications. In *Proceedings of Eurocrypt 2015*, 2015.

[18] C. Garman, M. Green, and I. Miers. Decentralized anonymous credentials. In *Proceedings of the NDSS Symposium 2014*, 2014.

[19] D. Goodin. Fraudulent Google credential found in the wild, Aug. 2011.

[20] T. H.-J. Kim, L.-S. Huang, A. Perrig, C. Jackson, and V. Gligor. Accountable key infrastructure (AKI): a proposal for a public-key validation infrastructure. In *Proceedings of WWW 2013*, pages 679–690, 2013.

[21] B. Laurie, A. Langley, and E. Kasper. Certificate transparency, 2013.

[22] J. Leyden. Inside 'Operation Black Tulip': DigiNotar hack analysed, Sept. 2011.

[23] H. Lipmaa. Secure accumulators from euclidean rings without trusted setup. In F. Bao, P. Samarati, and J. Zhou, editors, *ACNS 12*, volume 7341 of *LNCS*, pages 224–240, Singapore, June 26–29, 2012. Springer, Berlin, Germany.

[24] M. S. Melara, A. Blankstein, J. Bonneau, E. W. Felten, and M. J. Freedman. CONIKS: Bringing key transparency to end users. In *Proceedings of USENIX Security 2015*, 2015.

[25] J. Menn. Key Internet operator VeriSign hit by hackers, Feb. 2012.

[26] R. C. Merkle. A certified digital signature. In G. Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 218–238, Santa Barbara, CA, USA, Aug. 20–24, 1989. Springer, Berlin, Germany.

[27] A. Miller, M. Hicks, J. Katz, and E. Shi. Authenticated data structures, generically. In *Proceedings of POPL 2014*, 2014.

[28] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. bitcoin.org/bitcoin.pdf.

[29] Nasdaq. Nasdaq launches enterprise-wide blockchain technology initiative, May 2015.

[30] D. O'Leary, V. D'Agostino, S. R. Re, J. Burney, and A. Hoffman. Method and system for processing Internet payments using the electronic funds transfer network, Nov. 2013.

[31] C. Papamanthou, E. Shi, R. Tamassia, and K. Yi. Streaming authenticated data structures. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 353–370, Athens, Greece, May 26–30, 2013. Springer, Berlin, Germany.

[32] C. Papamanthou, R. Tamassia, and N. Triandopoulos. Authenticated hash tables. In P. Ning, P. F. Syverson, and S. Jha, editors, *ACM CCS 08*, pages 437–448, Alexandria, Virginia, USA, Oct. 27–31, 2008. ACM Press.

[33] M. D. Ryan. Enhanced certificate transparency and end-to-end encrypted mail. In *Proceedings of NDSS 2014*, 2014.