

Adversarial Level Agreements for Two-Party Protocols

Marilyn George
marilyn_george@brown.edu
Brown University
Providence, RI, USA

Seny Kamara
seny@brown.edu
Brown University
Providence, RI, USA

ABSTRACT

Adversaries in cryptography have traditionally been modeled as either semi-honest or malicious. Over the years, however, several works have investigated the design of cryptographic protocols against rational adversaries. The most well-known example are covert adversaries in secure computation (Aumann & Lindell, *TCC '07*) which are adversaries that wish to deviate from the protocol but without being detected. Protocols secure against such covert adversaries guarantee that deviations are detected with probability at least ϵ which is known as the deterrence factor.

In this work, we initiate the study of contracts in cryptographic protocol design. We show how to design, use and analyze contracts between parties for the purpose of incentivizing honest behavior from rational adversaries. We refer to such contracts as adversarial level agreements (ALA). The framework we propose can result in more efficient protocols and can enforce deterrence in covert protocols; meaning that one can guarantee that a given deterrence factor will deter the adversary instead of assuming it.

We show how to apply our framework to two-party protocols, including secure two-party computation (2PC) and proofs of storage (PoS). In the 2PC case, we integrate ALAs to publicly-verifiable covert protocols and show, through a game-theoretic analysis, how to set the parameters of the ALA to guarantee honest behavior. We do the same for PoS which are two-party protocols that allow a client to efficiently verify the integrity of a file stored in the cloud.

CCS CONCEPTS

• Security and privacy → Cryptography; Trust frameworks; Economics of security and privacy;

KEYWORDS

Adversarial Level Agreements, 2PC, Covert Security, Game Theory, Rational Cryptography, Proofs of Storage

ACM Reference Format:

Marilyn George and Seny Kamara. 2022. Adversarial Level Agreements for Two-Party Protocols. In *Proceedings of the 2022 ACM Asia Conference on Computer and Communications Security (ASIA CCS '22)*, May 30-June 3, 2022, Nagasaki, Japan. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3488932.3517385>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ASIA CCS '22, May 30-June 3, 2022, Nagasaki, Japan

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9140-5/22/05...\$15.00
<https://doi.org/10.1145/3488932.3517385>

1 INTRODUCTION

Adversaries in cryptographic protocols have traditionally been modeled as either semi-honest or malicious. Semi-honest adversaries are assumed to follow the protocol while trying to learn as much as possible whereas malicious adversaries can behave arbitrarily. In 2004, Halpern and Teague first considered adversaries that were *rational* in the sense that their behavior was governed by some utility function they sought to maximize [17]. While Halpern and Teague focused on proving and circumventing impossibility results in secret sharing and secure multi-party computation (MPC), follow-up work [10, 16] has shown that protocols secure against rational adversaries can be more efficient than protocols secure against malicious adversaries. The most well-known example is the work of Aumann and Lindell, which showed how to construct two-party computation (2PC) protocols secure against *covert* adversaries [8].

Covert adversaries. An adversary is covert if it deviates from the protocol but wishes to do so without being detected. A protocol secure against such an adversary guarantees that if the adversary is able to cheat then it will be detected with high enough probability. The covert model seems to capture many real-world adversaries, for example companies that cannot afford to be caught cheating because they have a reputation to protect. But to use a covert protocol in practice, one needs to set its deterrence factor which is a lower bound on the probability that cheating is detected. The idea is that if the adversary's utility function decreases when it gets caught and if the deterrence factor is set appropriately, then it is rational for a covert adversary to behave honestly.

In this work, we propose a new way to design cryptographic protocols against covert adversaries. Instead of focusing on deterrence, we focus on punishment. That is, our approach is to design protocols that explicitly punish adversaries when they cheat. It is well known in game theory that, when applied appropriately, punishment can lead to good outcomes. We will show this is the case in the setting of cryptographic protocols as well.

Contracts. In our framework, we integrate punishment into cryptographic protocols through *contracts*. Contracts are not a standard cryptographic primitive but they are powerful against rational adversaries. Furthermore, contracts are easy to design and well-understood by individuals and institutions. Many companies, including cloud providers, telecommunication companies and ISPs routinely use contracts called *service level agreements* (SLA) to guarantee a certain level of performance to their customers. In our setting, we use contracts to guarantee a certain level of "adversarialness" so we refer to them as *adversarial level agreements* (ALA). An ALA is a contract that specifies the damages a party has to pay if it is caught cheating.

Optimistic enforcement. One of the challenges in using contracts is that they require an enforcement mechanism. For our purposes, we assume the existence of a third party called the judge, that can enforce the damages stipulated in the ALAs. Of course, this requires our protocols to be auditable/verifiable by the judge. We stress, however, that our framework does not require the judge to be online except if there is a dispute, and as such our enforcement mechanism is *optimistic*. Since the judge is only required when a dispute occurs, this role could be played by either the existing legal system, or an appropriate smart contract, or another trusted entity depending on the application setting.

Advantages. The use of ALAs in cryptographic protocols provides several advantages. The first is to “enforce deterrence” in the covert model. In covert protocols, it is assumed that a high deterrence factor will prevent dishonest behavior but, in practice, there is no reason to believe this is true. For example, when faced with a rational garbler, we can set the deterrence factor of the two-party protocol to some value but we have no way of knowing if it is high enough to incentivize the garbler to behave honestly. Using ALAs, however, allows us to integrate punishment which can be used to incentivize honest behavior at a given deterrence level.

The second advantage is in improving efficiency against covert adversaries in a “cost-free” manner. Cryptographic protocols secure against deviating adversaries (i.e., either covert or malicious) have to guarantee some form of soundness error which bounds the probability that the adversary can cheat successfully. Interestingly, we will see that enhancing a protocol with an ALA creates a tradeoff between the damages in the ALA and the soundness error of the protocol: the higher we make the damages the higher the soundness error we can tolerate. Higher soundness error, in turn, implies a smaller security parameter which implies increased efficiency. This efficiency vs. damages tradeoff is particularly interesting because increasing damages is “free” in the sense that it does not impose any financial or computational costs on either party (unless cheating occurs). In other words, by using ALAs we can increase the efficiency of a protocol by decreasing its security parameter and increasing the damages but, for the most part, this increase in damages does not cost the parties anything.

Application to two-party protocols. In this work we explore how ALAs can be integrated into cryptographic protocols. We focus on two-party protocols and, specifically, the cases of 2PC and proofs of storage (PoS). We stress that the goal of this work is to initiate and motivate the use of contracts/ALAs in protocol design. The specific applications to 2PC and PoS are to illustrate how our framework can be applied to non-trivial cryptographic protocols and to motivate further study of ALAs.

In the setting of 2PC, we will be particularly interested in publicly-verifiable covert (PVC) protocols which were introduced by Asharov and Orlandi [5]. These protocols enhance covert protocols by providing honest parties with evidence of cheating that can be publicly verified. PoS, which were introduced by Juels and Kaliski [19] and by Ateniese et al. [6], are two-party protocols that allow a client to efficiently verify the integrity of a file stored at a remote server. Roughly, the client encodes its file before outsourcing it to the server. From that point on, it can verify the integrity of the file by sending a constant-size challenge to the server. The server then uses

the challenge and the file to compute a constant-size proof which it returns to the client. If the proof verifies, the client is convinced the server is storing the file.

1.1 Our Contributions

In this work, we introduce the notion of ALAs and show how to integrate them into the design of two-party cryptographic protocols; including secure two-party computation and proof of storage (PoS) protocols. We show that ALA-enhanced protocols are not only more efficient than standard protocols but that they also exhibit a “cost-free” tradeoff. Our work describes a new paradigm in cryptographic protocol design. More precisely, we make the following contributions.

Cryptographic inspection games. The introduction of contracts and punishment into cryptographic protocols results in new strategic interactions between the parties. We formalize these interactions as *cryptographic inspection games* (CIG) which are strategic games played between an inspector and an inspectee which correspond to the verifier/evaluator and prover/garbler, respectively. CIGs are a variant of *inspection games* which were introduced in the 1960s in the context of nuclear disarmament (we refer to [9] for an overview). In an inspection game the inspector’s goal is to detect deviation from some “good” behavior while the inspectee’s goals are to deviate without being detected.

One difference between CIGs and traditional inspection games is that, for CIGs, we also consider cases where the inspector is dishonest. This can occur in ALA-enhanced protocols because ALAs introduce an incentive for one of the parties to “frame” the other to recover the damages. To address this, we use a *dual ALA* which is a second contract; this time between the “inspector” party and a judge that, *if required*, can “inspect the inspector” to ensure honest behavior. In this manner, we find a desirable equilibrium for our CIGs where the parties behave honestly.

Secure two-party computation. We show how to integrate ALAs into 2PC protocols in order to get efficiency improvements against rational adversaries. More precisely, we augment PVC protocols [5] which guarantee *accountability* (i.e., deviations from the protocol can be publicly-verified) and *defamation-freeness* (i.e., no party can generate evidence that an honest party deviated) by having the garbler and evaluator first agree to an ALA that stipulates damages if the garbler deviates. At high level, the ALA induces a CIG where the evaluator is the inspector and the garbler is the inspectee. We analyze this game and, based on the accountability of the underlying protocol, find conditions under which the ALA guarantees honest behavior from the garbler. However, this result assumes an honest evaluator/inspector. We address this by adding a dual ALA between the inspector/evaluator and the judge that stipulates the damages the former must pay (to the judge) if it attempts to frame the inspectee/garbler. The introduction of the dual ALA gives rise to a more complex CIG whose analysis, in part, relies on the defamation-freeness of the PVC protocol.

Proofs of storage. To add ALAs to a PoS we must first ensure that it is auditable by a judge. Towards this, we introduce and formalize the notion of an auditable PoS (APoS) which is a PoS that includes an audit algorithm that a judge can use to determine if the server

has the file or not. From a security perspective, an APoS has to satisfy the same soundness requirement as a PoS (i.e., if the proof verifies the server holds the file) as well as a form of accountability and defamation-freeness. The former guarantees that if the server does not have the file then the judge will determine that it cheated; while the latter guarantees that if the server holds the file then the judge will determine that it did not cheat. This is analogous to accountability and defamation-freeness in the setting of 2PC [5]. We then describe a general transformation that turns any PoS into an APoS. The transformation relies on simple and efficient building blocks: collision-resistant hash functions and digital signatures.

We then consider how to extend an APoS with ALAs. We use an ALA between the client and server that stipulates the damages the server must pay if the APoS verification fails. This introduces strategic incentives which we model and analyze as a CIG. However, similar to the 2PC case, a rational client (as opposed to an honest one) has an incentive to frame the server to receive the damages. To address this we use a dual ALA between the client and the judge. We analyze the equilibria in both of the induced CIGs and show the conditions under which they guarantee that both parties will behave honestly.

Concrete analysis. We show how to apply our framework to concrete PVC and PoS protocols. Specifically, we design concrete ALAs for the 2PC protocols of Asharov and Orlandi[5], Kolesnikov and Mazolemoff[20] and Hong et al[18], and the APoS protocols that result from applying our transformation to the PoS protocols of Juels and Kaliski[19], Ateniese et al[6] and Shacham and Waters[21].

2 RELATED WORK

Secure two-party computation. 2PC was introduced by Yao [24] and the two main adversarial models are the semi-honest model where the adversary follows the protocol but tries to learn additional information and the malicious model where the adversary deviates arbitrarily. Covert adversaries, which were introduced by Aumann and Lindell [8], wish to deviate from the protocol as long as they are not detected. While the covert model is clearly and explicitly meant to capture rational adversaries, the strategic behavior of covert adversaries has, as far as we know, never been studied explicitly. In other words, protocols secure against covert adversaries guarantee that deviation is detectable with at least some probability ϵ , but the impact of this guarantee on a rational adversary’s strategy has not been formally studied. Our work provides such a study and, in addition, extends the covert model to consider what happens not only when deviations are detected but when deviations are punished. As discussed in Section 1, ALAs require the underlying cryptographic protocol to be auditable in the sense that a third party can verify whether one of the participants deviated from the protocol. Auditability has received some attention in the setting of secure computation. For example, Asharov and Orlandi [5] introduce publicly-verifiable covert secure computation (PVC) which augments covert protocols with a certificate that allows honest parties to prove their innocence. They proposed a garbled circuit-based protocol which was later improved by Kolesnikov and Malozemoff [20] and then by Hong et al. [18]. Baum, Damgard and Orlandi [11] show how to design PVC protocols whose correctness can be

verified by a trusted third party after the execution of the protocol. Cunningham, Fuller and Yakoubov [12] propose completely identifiable auditability, which allows the third party to identify the cheating parties in the protocol. In work more closely related to our own, Zhu, Ding and Huang [25] design and analyze PVC protocols with lightweight audit algorithms that can be implemented using smart contracts. The judge is optimistic and is only invoked if a dispute arises. There are several differences between our approaches. First, our framework relies on legal contracts and not blockchains/smart contracts—though the audits of the underlying auditable protocols could be implemented using smart contracts. Also, the approach taken in [25] requires that the parties deposit funds up front which is not the case with ALAs. More fundamentally, though, the framework of Zhu et al. assumes that the strategic interaction between the parties is a zero-sum game in the sense that the loss to the honest party is exactly the gain to the adversary. While this might apply in some settings, it is not always the case. Furthermore, the game-theoretic analysis is independent of blockchain transactions fees and assumes the deposits are high enough to deter framing.

Proofs of storage. Proofs of storage were introduced by Juels and Kaliski [19] and Ateniese et al. [6]. The former considered privately-verifiable proofs of retrievability whereas the latter considered publicly-verifiable proofs of data possession. PoS were further studied by Shacham and Waters [21] who gave both privately- and publicly-verifiable constructions. Dodis, Vadhan and Wichs then showed a connection between privately-verifiable PoS and hardness amplification. They present the idea of PoS codes and show improved constructions[14]. Ateniese, Kamara and Katz described a general compiler that transforms any identification protocol into a publicly-verifiable PoS [7]. PoS have been used in rational settings such as cryptocurrencies. Filecoin[3] is a decentralized storage network based on PoS that allows its miners to rent out storage to customers for tokens. In this work, we construct highly-efficient PoS against rational adversaries. As a building block, we also describe an auditable PoS against standard malicious adversaries that may be of independent interest.

3 PRELIMINARIES AND NOTATION

Notation. The security parameter is denoted $k \in \mathbb{N}$. A function is negligible in k if it is dominated by any inverse polynomial. The set of all binary strings of length n is denoted $\{0, 1\}^n$ and the set of all binary strings of arbitrary lengths is denoted $\{0, 1\}^*$. We write $x \stackrel{\$}{\leftarrow} X$ to denote that x is sampled uniformly at random from a set X . When an algorithm Alg_1 has oracle access to an algorithm Alg_2 we write $\text{Alg}_1^{\text{Alg}_2}(\cdot)$. The output of an interactive protocol Π executed between parties P_1 and P_2 with inputs x_1 and x_2 , respectively, is denoted by $(y_1, y_2) \leftarrow \Pi_{P_1, P_2}(x_1, x_2)$, where y_1 is P_1 ’s output and y_2 is P_2 ’s output. We denote the space of valid inputs to the protocol as $\mathbb{I}_\Pi = \mathbb{I}_1 \times \mathbb{I}_2$ where \mathbb{I}_i is the space of valid inputs for party i . Similarly, we define the output space $\mathbb{O}_\Pi = \mathbb{O}_1 \times \mathbb{O}_2$. Then for the protocol $(y_1, y_2) \leftarrow \Pi_{P_1, P_2}(x_1, x_2)$ we say that $x_i \in \mathbb{I}_i$ and $y_i \in \mathbb{O}_i$. Additionally let \mathbb{T}_Π be the set of valid transcripts of the protocol where $\mathbb{T}_\Pi = \mathbb{T}_1 \times \mathbb{T}_2$ and \mathbb{T}_i is the set of all messages sent by party i during the protocol. Whenever the protocol is obvious from context

we drop the subscript for convenience. We denote the expected utility (in a game-theoretic sense) to a player P due to playing a strategy S as $U_P[S]$. This expectation is calculated over the outcomes due to the distribution of the other players' strategies and any randomness in the game being played. The expected utility of playing S when other players are playing the strategy vector S_- is denoted as $U_P[S, S_-]$. A player's *best response* is the strategy that gives maximum expected utility given the game and the other players' strategies.

Collision resistant hash functions. A hash function $\text{Hash} = (\text{Gen}, \text{Hash})$ consists of two polynomial-time algorithms that work as follows. Gen takes as input a security parameter k and outputs a (non-secret) key α . Hash takes as input a key α and a message $m \in \{0, 1\}^*$ and outputs a hash $h \in \{0, 1\}^k$. The security of a hash function family is formalized using the following randomized experiment where \mathcal{A} is an adversary:

- $\text{Coll}_{\mathcal{A}}(k)$:
 - (1) compute $\alpha \leftarrow \text{Gen}(1^k)$;
 - (2) $(m_1, m_2) \leftarrow \mathcal{A}(\alpha)$;
 - (3) if $m_1 \neq m_2$ and $\text{Hash}(\alpha, m_1) = \text{Hash}(\alpha, m_2)$ output 1 else output 0.

We say that Hash is collision-resistant if for all PPT adversaries \mathcal{A} ,

$$\Pr [\text{Coll}_{\mathcal{A}}(k) = 1] \leq \text{negl}(k),$$

where the probability is over the coins of Gen and \mathcal{A} .

Digital signatures. A digital signature scheme $\text{Sig} = (\text{Gen}, \text{Sign}, \text{Vrfy})$ consists of three polynomial-time algorithms that work as follows. Gen takes as input a security parameter k and outputs a signing key sk and a verification key vk . Sign takes as input a signing key sk and a message m and outputs a signature σ . Vrfy takes as input a verification key vk , a message m and a signature σ and outputs a bit b . A signature scheme is correct if for all $k \in \mathbb{N}$, for all (vk, sk) output by $\text{Gen}(1^k)$, for all $m \in \{0, 1\}^*$, $\text{Vrfy}(\text{vk}, m, \text{Sign}(\text{sk}, m)) = 1$.

The security of a digital signature scheme is formalized using the following randomized experiment:

- $\text{Forge}_{\mathcal{A}}(k)$:
 - (1) compute $(\text{sk}, \text{vk}) \leftarrow \text{Gen}(1^k)$;
 - (2) $(m, \sigma) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(\text{vk})$;
 - (3) let Q be the set of messages on which the Sign oracle was queried
 - (4) if $\text{Vrfy}(\text{vk}, m, \sigma) = 1$ and $m \notin Q$, output 1 else output 0.

We say that a signature scheme Sig is existentially unforgeable if for all PPT adversaries \mathcal{A}

$$\Pr [\text{Forge}_{\mathcal{A}}(k) = 1] \leq \text{negl}(k),$$

where the probability is over the coins of Gen and \mathcal{A} .

PVC protocols. We assume the reader is familiar with the definitions of 2PC. Covert protocols include an additional parameter ϵ called the deterrence factor which is a lower bound on the probability that cheating is detected. PVC protocols further allow the honest parties to prove that cheating occurred to an external observer.

Definition 3.1. A publicly-verifiable covert protocol $\text{PVC} = (\Pi, \text{Blame}, \text{Judge})$ consists of a covert 2PC protocol Π and two polynomial-time algorithms Blame and Judge as follows:

- $\text{cert} \leftarrow \text{Blame}(\text{id}, \text{view})$: is a deterministic algorithm that takes as input the identifier of the cheating party id and a view of the protocol and outputs a certificate cert .
- $b \leftarrow \text{Judge}(\text{cert})$: is deterministic algorithm that takes as input a certificate cert and outputs 1 if the accused party is determined to have cheated and 0 otherwise.

We say that a PVC protocol is (ϵ, μ, ϕ) -secure if it satisfies *covert security*, *accountability* and *defamation-freeness*. Covert security guarantees that a cheating party is detected with probability at least ϵ . Accountability guarantees that if an honest party outputs a certificate of cheating for a dishonest party, $\text{Judge}(\text{cert})$ outputs 0 with probability at most μ . Defamation-freeness guarantees that if a party outputs a certificate cert that blames an honest party, $\text{Judge}(\text{cert})$ outputs 1 with probability at most ϕ .

Proofs of storage. A PoS is a protocol executed between a server that stores a dataset and a client that wishes to verify the integrity of the data. A publicly-verifiable PoS has the property that anyone in possession of the client's public key can run the verification of integrity. We recall the syntax and security definitions for a publicly-verifiable PoS.

Definition 3.2 (Proof of storage). A publicly-verifiable proof of storage $\text{PoS} = (\text{Gen}, \text{Encode}, \text{Chall}, \text{Prove}, \text{Vrfy})$ consists of five polynomial-time algorithms that work as follows:

- $(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^k)$: is a probabilistic algorithm that takes as input the security parameter k and outputs a secret key sk and a public key pk .
- $(\mathbf{e}, \text{st}) \leftarrow \text{Encode}(\text{sk}, \mathbf{f})$: is a probabilistic algorithm that takes as input a secret key sk and a file \mathbf{f} and outputs an encoded file \mathbf{e} and a state st .
- $\text{ch} \leftarrow \text{Chall}(\text{pk}, \text{st})$: is a probabilistic algorithm that takes as input a public-key pk and a state st and outputs a challenge ch .
- $\pi \leftarrow \text{Prove}(\text{pk}, \mathbf{e}, \text{ch})$: is a deterministic algorithm that takes as input a public key pk , an encoded file \mathbf{e} and a challenge ch and outputs a proof π .
- $b \leftarrow \text{Vrfy}(\text{pk}, \text{st}, \text{ch}, \pi)$: is a deterministic algorithm that takes as input a public key pk , a state st , a challenge ch and a proof π and outputs a bit b .

We say that PoS is complete if for all $k \in \mathbb{N}$, for all (sk, pk) output by $\text{Gen}(1^k)$, for all $\mathbf{f} \in \{0, 1\}^{\text{poly}(k)}$, for all (\mathbf{e}, st) output by $\text{Encode}(\text{sk}, \mathbf{f})$, for all ch output by $\text{Chall}(\text{pk}, \text{st})$, for all π output by $\text{Prove}(\text{pk}, \mathbf{e}, \text{ch})$, $\text{Vrfy}(\text{pk}, \text{st}, \text{ch}, \pi)$ outputs 1.

A publicly-verifiable PoS is typically used as follows. The client runs $(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^k)$ to generate a key pair. It keeps sk secret and sends pk to the verifier (which could be the client itself of course). The client then encodes its file by computing $(\mathbf{e}, \text{st}) \leftarrow \text{Encode}(\text{sk}, \mathbf{f})$ and sends \mathbf{e} to the server and st to the verifier. To verify the integrity of the file, the verifier sends a challenge $\text{ch} \leftarrow \text{Chall}(\text{pk}, \text{st})$ to the server who returns a proof $\pi \leftarrow \text{Prove}(\text{pk}, \mathbf{e}, \text{ch})$. Finally, the verifier verifies the proof by computing $b \leftarrow \text{Vrfy}(\text{pk}, \text{st}, \text{ch}, \pi)$. The security property we require is the following: if the proof verifies then the server must be storing the file (or some efficiently extractable form of the file). We refer the reader to the definitions of Ateniese et al. [7] for the formalization of this property.

4 ADVERSARIAL LEVEL AGREEMENTS

We augment cryptographic protocols with contracts that stipulate the punishments parties must incur if they deviate from the protocol. We call such contracts adversarial level agreements. Note that it is arguably the *avoidance* of punishment rather than the evasion of detection that is the true motivation of a covert adversary; i.e., it is the *consequence* of detection that makes detection undesirable, not detection on its own. As such, modeling and integrating punishment—here through contracts—is natural. Also, in practice, many cryptographic protocols operate in environments where there are (at least) implicit consequences for deviating.

We show that, using our framework, one can design cryptographic protocols that achieve new tradeoffs: namely, one can achieve better efficiency effectively *for free*, i.e., without any loss of security (against rational adversaries). To do so, it suffices to increase the damages in the ALA. Note that such a tradeoff is not possible with standard covert protocols because the covert model does not capture punishment. In the covert model, the only parameter we can tune is the deterrence factor (i.e., the probability of being detected) and there is a tight dependency between deterrence and efficiency.

ALAs. For a cryptographic protocol Π , an adversarial level agreement $\text{ALA}_{\Pi,k}(\mathbf{d})$ includes: (1) a specification of Π , defining each party’s prescribed behavior; (2) damages d due from each party if found deviating from Π , where $\mathbf{d} = (d_1, \dots, d_n)$ is a vector that describes the damages for each party. The inclusion of contracts and punishment in cryptographic protocols introduces new strategic interactions between the parties. We formalize these interactions as games we refer to as *cryptographic inspection games* which are a variant of inspection games.

Inspection games. An inspection game is two-player game between an inspectee who wishes to cheat and an inspector who wishes to deter cheating. Inspection games were introduced in the 1960s in the context of nuclear disarmament and are used to design nuclear inspections under nuclear Non-Proliferation Treaties. It is well-known that inspection games have no pure strategy Nash equilibria [9]. Informally if the inspector always inspects, the inspectee never cheats but then the inspector would never inspect so the inspectee would always cheat etc. Though pure equilibria do not exist, an inspection game could have mixed equilibria but they often have a positive probability of cheating and are therefore undesirable. To overcome this, an inspection game is usually converted to a *leadership* variant by having the inspector commit to a mixed strategy for which the inspectee’s best response is not to cheat. This is an application of the *leadership principle* in game theory [22] where the first player (the leader) declares and commits to a strategy which forces the second player (the follower) to play its the best response. The leader then has the power to choose a favorable equilibrium in this sequential game. We refer to [9] for a comprehensive summary of the applications of inspection games and the consequences of inspector leadership.

Cryptographic inspection games. A CIG is a two-party game between an inspectee that wishes to deviate from a protocol and an inspector that wishes to enforce honest behavior. There are two

main differences between inspection games and CIGs: (1) inspection games typically have both false-positives and false-negatives whereas CIGs have only false-negatives; and (2) in a CIG the inspector can be dishonest. We represent CIGs as extensive form games which are trees that represent a player’s possible actions as nodes and outcomes as leaves. Each ALA-enhanced cryptographic protocol results in a CIG. We then analyze the CIG to find ALA parameters that will guarantee the desirable equilibrium where the parties follow the protocol.

5 AUGMENTING TWO-PARTY COMPUTATION WITH ALAS

We now show how to use ALAs in secure two-party computation. As discussed above, we start with a PVC protocol Π and have the parties sign an ALA that describes Π and the damages they must pay if they deviate.

PVC protocols. Most two-party PVC protocols have a similar structure which we now describe. The garbler G creates many garbled circuits and the evaluator E opens and checks all but one of them. If any of the checks fail, the evaluator receives a publicly-verifiable certificate of the garbler’s deviation. If all the checks pass, the evaluator evaluates the unopened circuit and receives the output which it then sends to the garbler. The protocol guarantees accountability and defamation-freeness (see Section 3 for definitions).

Adversarial settings. We consider two adversarial settings:

- (1) Rational garbler vs. honest evaluator: the evaluator follows the protocol but the garbler can deviate to increase its utility;
- (2) Rational garbler vs. framing evaluator: when damages are introduced, a rational evaluator has an incentive to frame an honest garbler.

For each of these settings, we first describe the CIGs that result from the strategic interactions between the parties. We then analyze these games to find the damages that guarantee that the rational parties behave honestly.

5.1 Rational Garbler vs. Honest Evaluator

When an (ϵ, μ, ϕ) -secure PVC protocol is enhanced with an ALA between an honest evaluator and a rational garbler, it results in the CIG depicted in Figure 1. The garbler has a choice between deviating which we denote *Deviation* and behaving honestly which we denote *NoDeviation*. When the garbler behaves honestly (the left sub-tree) the evaluator’s circuit checks pass, it evaluates the circuit and both parties receive the output and their utilities u_G^o and u_E^o . On the other hand, when the garbler deviates (the right sub-tree), the deviation is detected by the evaluator with probability at least ϵ (from the deterrence of Π). This is a probabilistic event, represented by a hollow *chance* node in the tree. When the deviation is detected, the Judge is invoked and its audit gives the same output as the honest evaluator except with probability μ (from the accountability Π). If the audit outputs 1 the garbler is punished and pays d_G to the evaluator. If the audit outputs 0, which occurs with probability μ , the evaluator is punished and pays d_E to the Judge.

We denote the garbler’s utility for deviating (and potentially learning information about the evaluator’s input) as u_G^i and the

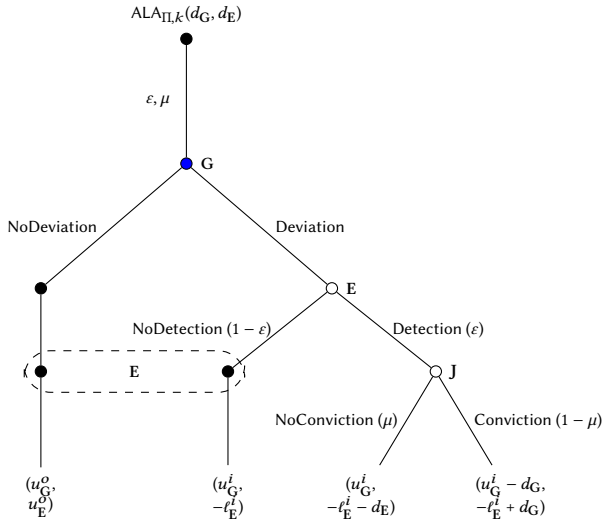


Figure 1: CIG between a rational garbler and an honest evaluator.

loss to the evaluator as ℓ_E^i . If a deviation occurs, the garbler (resp. evaluator) receives this utility (resp. loss) regardless of the Judge’s decision. We analyze the game in the following Lemma.

LEMMA 5.1. *Let Π be an (ϵ, μ, ϕ) -secure PVC protocol between a rational garbler and an honest evaluator that is augmented with an ALA with damages d_G and d_E . If*

$$d_G \geq \frac{u_G^i - u_G^o}{\epsilon(1 - \mu)},$$

then the dominant strategy for a rational garbler is to follow the protocol.

PROOF. To find the garbler’s dominant strategy, we compare its expected utilities from the actions available at the blue node of the game: namely, deviating from the protocol and not deviating. Clearly, we have $U_G[\text{NoDeviation}] = u_G^o$. Furthermore,

$$\begin{aligned} U_G[\text{Deviation}] &= (1 - \epsilon)u_G^i + \epsilon \left(\mu u_G^i + (1 - \mu)(u_G^i - d_G) \right) \\ &= u_G^i - \epsilon(1 - \mu)d_G \\ &\leq u_G^i - \epsilon(1 - \mu) \frac{u_G^i - u_G^o}{\epsilon(1 - \mu)} \\ &= u_G^o, \end{aligned}$$

where the inequality holds by the condition in the Lemma. The garbler’s expected utility from deviating is then at most its expected utility from behaving honestly, from which the Lemma follows. ■

5.2 Rational Garbler vs. Framing Evaluator

When an (ϵ, μ, ϕ) -secure PVC protocol is enhanced with an ALA between a rational garbler and a framing evaluator, it results in the CIG depicted in Figure 2. In this game, the evaluator has new strategies which account for framing. The garbler’s strategies remain the

same as in the previous game. The garbler chooses to deviate or to behave honestly. If the garbler is honest (the left sub-tree) then the evaluator’s checks pass. If the garbler deviates (the right sub-tree) the evaluator’s checks pass with probability $1 - \epsilon$. These two blue nodes represent an information set. The evaluator now has a choice between framing which we denote by Frame and not framing which we denote by NoFrame. If it chooses to frame the garbler, the audit will output 0 except with probability ϕ . This is a chance node whose probability is derived from the defamation-freeness of Π . If the audit outputs 1, the evaluator receives its utility u_E^f . If the audit outputs 0, the evaluator is punished for attempted framing and has to pay damages d_E . The garbler’s utility for deviating is u_G^i and the loss to the evaluator is ℓ_E^i and they remain regardless of the Judge’s decision. The evaluator’s utility for framing is u_E^f . This is equal to the damages imposed on the garbler, d_G , since the damages are paid to the evaluator. The damages paid by the evaluator for framing is d_E . Our next lemma, whose proof is in the full version, shows how to set these damages so that both parties behave honestly [15].

LEMMA 5.2. *Let Π be an (ϵ, μ, ϕ) -secure PVC protocol between a rational garbler and a framing evaluator that is augmented with an ALA with damages d_G and d_E . If,*

$$d_E \geq \frac{\phi d_G}{1 - \phi},$$

then the dominant strategy for a framing evaluator is to follow the protocol.

Notice that when we set the damages to rule out the evaluator’s framing strategy, the game reduces to the CIG of Figure 1 and the garbler has the choice to deviate or not. But if we also set the damages according to the conditions of Lemma 5.1, then we know that the garbler will not deviate in this game. So by setting damages according to both Lemmas, we can ensure that both the garbler and the evaluator will behave honestly.

Concrete improvements. Our framework can be applied to PVC protocols to achieve two kinds of improvements: (1) setting concrete deterrence factors; and (2) improving efficiency. Concrete instantiations for the PVC protocols of Hong et al. [18], Koleznikov and Malozemoff [20] and Asharov and Orlandi [5] are presented in Supplement B. At a high level, our framework allows the PVC protocols to use the lowest possible deterrence factor, with appropriate damages according to Lemmas 5.1 and 5.2. We note here that the setting of appropriate damages will depend on both the setting and the participating parties, and is an interesting question in its own right.

6 AUDITABLE PROOFS OF STORAGE

In order to apply our framework to proof of storage (PoS) protocols, we first introduce the notion of auditability in PoS. Once we have an auditable proof of storage (APoS), we can then augment the protocol using ALAs similar to our approach for PVC protocols. In this section we present (1) our definitions for auditable proofs of storage; (2) a general-purpose transformation to convert any PoS to an auditable PoS protocol. In the next section, we will analyze the CIGs that result from applying ALAs to APoS protocols, and design ALAs to improve the efficiency of concrete PoS protocols.

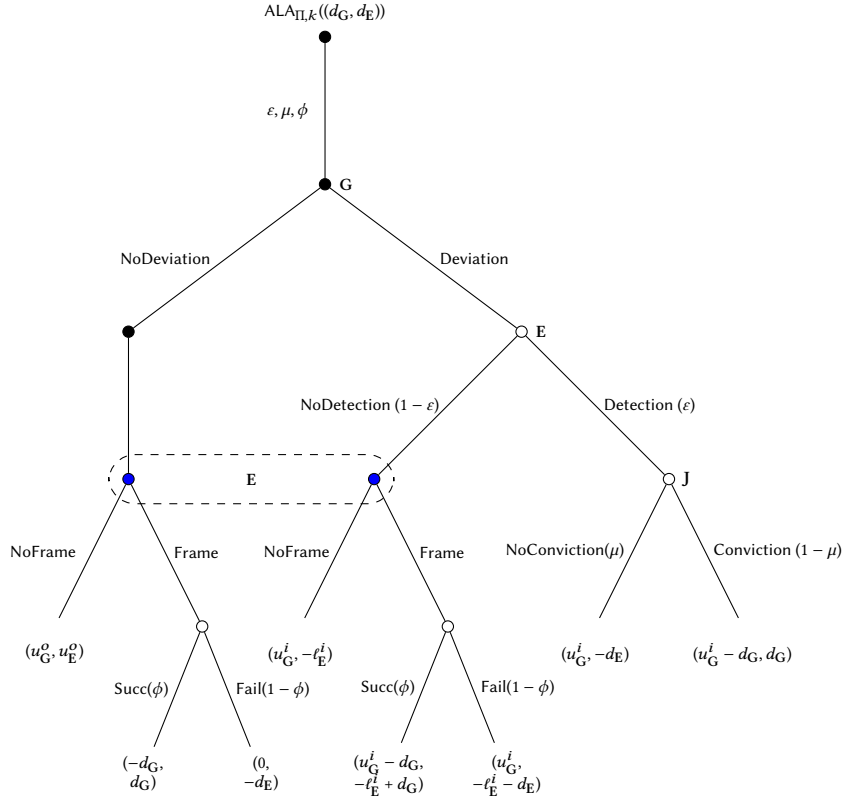


Figure 2: CIG between an rational garbler and a framing evaluator.

6.1 Definitions for APoS

We now introduce and formalize the notion of an *auditable* proof of storage which is a PoS whose execution can be audited by a trusted third party called the Judge. In addition to soundness, an APoS has to satisfy two properties: (1) *accountability* which, roughly, guarantees that the Judge will correctly detect when the file has been tampered with; and (2) *defamation-freeness* which, roughly, guarantees that the Judge cannot be fooled (by the client) into believing that the file has been tampered with if it has not.

Definition 6.1 (Auditable PoS). An auditable proof of storage APoS = (Gen, Π , Chall, Prove, Vrfy, Audit) is composed of five polynomial-time algorithms and one two-party protocol. Gen, Chall, Prove and Vrfy are as in a standard PoS (Definition 3.2) and Π , Receipt and Audit work as follows:

- $(st_J, clk, srk) \leftarrow \text{Gen}(1^k)$: is a probabilistic algorithm that takes as input the security parameter k and outputs a state st , a client key $clk = (clk_p, clk_s)$ and a server key $srk = (srpk, srsk)$ where $clk_p, srpk$ are public keys for the client and the server respectively.
- $((st_C, rec); \mathbf{e}) \leftarrow \Pi_{C,S}((clk, \mathbf{f}); srk)$: is a two-party protocol executed between the client and the server. The client inputs a client key clk and a file \mathbf{f} , and the server inputs a server key srk . The client receives as output a state st_C and a receipt rec whereas the server receives an encoded file \mathbf{e} .

- $b \leftarrow \text{Audit}(st_J, rec^*, \mathbf{e}^*)$: is a deterministic algorithm that takes as input a state st , a receipt rec^* from the client. It also receives an encoded file \mathbf{e}^* from the server—which may not be the original file and could possibly be empty if the server doesn't respond. It outputs a bit b which is 1 if the server is found guilty and 0 otherwise

We say that APoS is complete if for all $k \in \mathbb{N}$, for all $(st_J, clk, srk) \leftarrow \text{Gen}(1^k)$, for all $\mathbf{f} \in \{0, 1\}^{\text{poly}(k)}$, for all $((st_C, rec); \mathbf{e}) \leftarrow \Pi_{C,S}((clk, \mathbf{f}); srk)$, for all ch output by $\text{Chall}(clk, st_C)$, for all π output by $\text{Prove}(srk, \mathbf{e}, ch)$, $\text{Vrfy}(clk, st_C, ch, \pi)$ outputs 1.

An auditable PoS is used as follows: The Judge runs Gen to generate a state st_J , a client key clk and a server key srk . This step could be replaced by an interactive protocol that generates the keys—we only require that the Judge be able to trust the keys. The client and server then execute the Π protocol on a file \mathbf{f} . From this, the client receives a receipt rec and the server receives an encoded file \mathbf{e} . Chall, Prove and Vrfy are used as in a standard publicly-verifiable PoS. If at any time, the client claims the server is not storing the file, the Judge runs $b \leftarrow \text{Audit}(st_J, rec^*, \mathbf{e}^*)$ where rec^* is the receipt provided by the client and \mathbf{e}^* is the encoded file provided by the server.

Soundness. The notion of soundness of auditable PoS is similar to that of standard PoS. We formalize this with the following randomized experiments where \mathcal{K} is a stateful extractor which interacts

with the stateful adversary \mathcal{A} and attempts to output the original file \mathbf{f} :

- **OutputAdv $_{\mathcal{A}}(k)$:**
 - (1) generate $(st_j, clk, srk) \leftarrow \text{Gen}(1^k)$;
 - (2) repeat $p(k)$ times where p is a polynomial;
 - (a) the adversary computes $\mathbf{f} \leftarrow \mathcal{A}$;
 - (b) compute $((st_C, rec); \mathbf{e}) \leftarrow \Pi_{C, \mathcal{A}}((clk, \mathbf{f}); srk)$
 - (3) the adversary computes $\mathbf{f}^* \leftarrow \mathcal{A}$;
 - (4) compute $((st_C, rec); \mathbf{e}^*) \leftarrow \Pi_{C, \mathcal{A}}((clk, \mathbf{f}^*); srk)$
 - (5) compute $ch \leftarrow \text{Chall}(clk, st_C)$;
 - (6) the adversary computes $\pi^* \leftarrow \mathcal{A}(\mathbf{e}^*, ch)$;
 - (7) output (ch, π^*) .
- **OutputExt $_{\mathcal{A}, \mathcal{K}}(k)$:**
 - (1) generate $(st_j, clk, srk) \leftarrow \text{Gen}(1^k)$;
 - (2) repeat $p(k)$ times where p is a polynomial;
 - (a) the adversary computes $\mathbf{f} \leftarrow \mathcal{A}$;
 - (b) compute $((st_C, rec); \mathbf{e}) \leftarrow \Pi_{C, \mathcal{A}}((clk, \mathbf{f}); srk)$
 - (3) the adversary computes $\mathbf{f}^* \leftarrow \mathcal{A}$;
 - (4) compute $((st_C, rec); \mathbf{e}^*) \leftarrow \Pi_{C, \mathcal{A}}((clk, \mathbf{f}^*); srk)$
 - (5) repeat $q(k)$ times where q is a polynomial;
 - (a) the extractor computes $ch_i \leftarrow \text{Chall}(clk, st_C)$;
 - (b) the adversary computes $\pi_i \leftarrow \mathcal{A}(\mathbf{e}^*, ch_i)$;
 - (6) the extractor computes $((ch^e, \pi^e), \mathbf{f}^e) \leftarrow \mathcal{K}(clk, st_C, rec, \{ch_i, \pi_i\})$;
 - (7) output (ch^e, π^e) .
- **Sound $_{\mathcal{A}, \mathcal{K}}(k)$:**
 - (1) run **OutputExt $_{\mathcal{A}, \mathcal{K}}(k)$** ;
 - (2) obtain $\mathbf{f}^* \leftarrow \mathcal{A}$ and $((ch^e, \pi^e), \mathbf{f}^e) \leftarrow \mathcal{K}(clk, st_C^*, rec^*, \{ch_i, \pi_i\})$ from the above run;
 - (3) if $1 \leftarrow \text{Vrfy}(clk, st_C, ch^e, \pi^e)$ and $\mathbf{f}^e \neq \mathbf{f}^*$ output 1 else output 0.

Definition 6.2 (Soundness). An auditable proof of storage APoS = (Gen, Π , Chall, Prove, Vrfy, Audit) is ϵ -sound if there exists an expected polynomial-time extractor \mathcal{K} such that for all PPT adversaries \mathcal{A} , **OutputAdv $_{\mathcal{A}}(k)$** and **OutputExt $_{\mathcal{A}, \mathcal{K}}(k)$** are identically distributed and,

$$\Pr \left[\text{Sound}_{\mathcal{A}, \mathcal{K}}(k) = 1 \right] \leq \epsilon,$$

where the probabilities are over the coins of Gen, Π , \mathcal{A} .

Accountability. As discussed above, accountability guarantees that the Judge will correctly determine if the file has been tampered with. Note that since standard publicly-verifiable PoS allow for a third party to verify the integrity of the file, they naturally achieve accountability if we simply let the Judge run the verification algorithm using the client's public key and state. We formalize accountability with the following randomized experiment where \mathcal{A} is a stateful adversary, \mathbf{f} is a file and APoS = (Gen, Π , Chall, Prove, Vrfy, Audit) is an auditable PoS:

- **Accnt $_{\mathcal{A}, \mathbf{f}}(k)$:**
 - (1) generate $(st_j, clk, srk) \leftarrow \text{Gen}(1^k)$;
 - (2) compute $((st_C, rec); \mathbf{e}) \leftarrow \Pi_{C, \mathcal{A}}((clk, \mathbf{f}); srk)$
 - (3) the adversary computes $\mathbf{e}^* \leftarrow \mathcal{A}$;
 - (4) if $\text{Audit}(st_j, rec, \mathbf{e}^*) = 0$ and $\mathbf{e}^* \neq \mathbf{e}$ output 1 else output 0.

Definition 6.3 (Accountability). An auditable proof of storage APoS = (Gen, Π , Chall, Prove, Vrfy, Audit) is α -accountable if for all PPT adversaries \mathcal{A} , for all files \mathbf{f} ,

$$\Pr \left[\text{Accnt}_{\mathcal{A}, \mathbf{f}}(k) = 1 \right] \leq \alpha$$

where the probabilities are over the coins of Gen, Π , \mathcal{A} .

Defamation-freeness. The third security property we require of an APoS is defamation-freeness which guarantees that the Judge cannot be fooled into believing that the file has been tampered with when it has not. We formalize defamation-freeness with the following randomized experiment where \mathcal{A} is a stateful adversary, \mathbf{f} is a file and APoS = (Gen, Π , Chall, Prove, Vrfy, Audit) is an auditable PoS:

- **Defame-Free $_{\mathcal{A}, \mathbf{f}}(k)$:**
 - (1) generate $(st_j, clk, srk) \leftarrow \text{Gen}(1^k)$;
 - (2) compute $((st_C, rec); \mathbf{e}) \leftarrow \Pi_{\mathcal{A}, S}((clk, \mathbf{f}); srk)$
 - (3) the adversary computes $rec^* \leftarrow \mathcal{A}$;
 - (4) if $\text{Audit}(st_j, rec^*, \mathbf{e}) = 1$ output 1 else output 0.

Definition 6.4 (Defamation-freeness). An auditable proof of storage APoS = (Gen, Π , Chall, Prove, Vrfy, Audit) is δ -defamation-free if for all PPT adversaries \mathcal{A} , for all files \mathbf{f} ,

$$\Pr \left[\text{Defame-Free}_{\mathcal{A}, \mathbf{f}}(k) = 1 \right] \leq \delta$$

where the probabilities are over the coins of Gen, Π , \mathcal{A} .

Notation. Throughout this work, we will refer to an APoS that is ϵ -sound, α -accountable and δ -defamation-free simply as being $(\epsilon, \alpha, \delta)$ -secure.

6.2 A PoS-to-APoS Transformation

In this section we present a transformation that converts any PoS to an auditable PoS. Our transformation makes use of a hash function family $\text{Hash} = (\text{Gen}, \text{Hash})$ and of a signature scheme $\text{Sig} = (\text{Gen}, \text{Sign}, \text{Vrfy})$. We assume authenticated channels between all the parties.

Overview. At a high level, our transformation allows the client and server to agree on the file that is to be stored, by both signing a hashed digest of the file. If a dispute arises during the protocol and the client asks the Judge to perform an audit, it checks the signed digest. If both the signatures verify, the client is honest, and if the file stored matches the signed digest, then the server is honest. The transformation works as follows: the Gen algorithm creates a public/private key pair for the underlying PoS, a key w for the hash function family, and two signing/verification key pairs: one for the client and one for the server. The two-party protocol Π is executed between the client C and the server S. The client encodes the file \mathbf{f} and signs its hash h . It then sends the encoded file \mathbf{e} , its hash h and its signature σ_C on the hash to the server. The server checks if h is indeed the hash of the encoded file \mathbf{e} and if so signs it as well. It then returns the hash h , the signature σ_C and its own signature σ_S on the hash. Finally, the client verifies that σ_S is indeed a signature on h . The receipt rec consists of the hash h and the two signatures σ_C and σ_S . The Chall, Prove and Vrfy algorithms are the same as the Chall and Prove and Vrfy algorithms of the underlying PoS. The Audit algorithm takes as input a receipt $rec = (h, \sigma_C, \sigma_S)$ from the client

and some encoded file e^* from the server. The Judge first checks whether σ_C and σ_S are indeed signatures over h . If not, it determines that the client is dishonest. If both signatures are valid and if the encoded file e^* hashes to the hash h , then the server is considered honest. On the other hand, if the two signatures are valid but e^* does not hash to h then the server is considered dishonest since e^* could not be the original encoded file e . The transformation and security proof are detailed in Appendix A.

7 AUGMENTING APOS WITH ALAS

We now show how our techniques are applicable to PoS protocols. In this setting, the client C and the server S agree on the ALA. We assume an out-of-band “traditional” contract that specifies the hash of the file and the damages the server will pay to the client if the file has been tampered with. To enhance an APoS with an ALA, the ALA has to be signed by both parties after they agree on the file. We denote by $ALA(d_S)$ an adversarial level agreement with server damages d_S . The strategic interaction between a rational server and client in an ALA-enhanced APoS protocol can be viewed as a CIG where the client plays the role of the inspector and the server plays the role of the inspectee.

Overview of CIG. The client’s actions are the set of possible soundness error parameters for the APoS and the server’s actions are to either keep the file or lose the file. Here, the client “inspects” the server by using the APoS. If the inspection fails the client invokes the Judge and the server is punished by having to compensate the client for losing the file. The client’s strategies can then be viewed as a mixed strategy over the actions “inspection passes” which we denote Pass and “inspection fails” which we denote Fail, where the former is played with probability $1 - \epsilon$ and the latter is played with probability ϵ . More precisely, the client picks a soundness error ϵ which gives it a probabilistic payoff when the server tampers with the file. When the server does not tamper with the file, however, the client will always have the same deterministic payoff (which is the cost of the APoS or the “inspection”). We model these payoffs as mixing over two strategies “pass” and “fail” which: (1) have different payoffs when the server tampers with the file; (2) have the same payoffs when the server does not tamper with the file. In other words, when the server tampers with the file, the APoS can “pass” or “fail” and the client gets the appropriate payoffs but when the server does not tamper, the actions “pass” and “fail” give the same payoff.

Given solely this choice, a client would want to pick a low ϵ , or even $\epsilon = 0$ in order to increase the probability of the inspection passing. However, the technical difficulty is that all practical constructions of proofs of storage protocols have non-zero soundness error. Additionally, in order to reduce the soundness error, the security parameter of these schemes must be increased. With a greater security parameter comes greater computation costs both for the server and the client. These costs can also depend on other factors, such as the cost of computation or storage. In order to simplify the model, we assume that these costs of computation for both the server and the client are fixed constant values.

On the other hand, the server strategies are mixed strategies over the set of actions “tamper with file” which we denote Tamp and “do not tamper with file” which we denote NoTamp, where the

former is played with probability t_S and the latter is played with probability $1 - t_S$. Similar to the 2PC case, the client has an incentive to frame the honest server in order to receive compensation.

In all the games that follow, the Judge runs Audit to identify the guilty party. If Audit finds the server guilty the Judge enforces the damages d_S to the server and compensates the client. On the other hand, if the audit finds the server innocent the Judge enforces the damages d_C to the client for attempting to frame the honest server.

Adversarial settings. We consider two adversarial settings:

- (1) Rational server vs. honest client: the client follows the protocol but the server can deviate if it increases its utility.
- (2) Rational server vs. framing client: the server is rational (and may even choose not to participate) and the client is rational and may try to frame the server.

7.1 Rational Server vs. Honest Client

In this section, we describe how to enhance an APoS with an ALA and, specifically, how to set the damages in the ALA as a function of the parameters of underlying APoS. Here, we focus on the case of an honest client and a rational server. Our goal is to show that if the damages in the ALA are set appropriately, a rational server will never tamper with the file.

APoS for honest clients. The simplest APoS for the setting of honest clients works as follows. Given a standard PoS, one simply defines the Audit algorithm to always output 1. Since the client is honest, Audit will only be invoked when PoS verification fails. This simple construction is $(\epsilon, 0, 1)$ -secure, where ϵ is the soundness error of the underlying PoS, $\alpha = 0$ because Audit never outputs 0 and $\delta = 1$. Note that the defamation-freeness of this APoS is the worst possible but this does not matter as an honest client will never invoke the Judge when PoS verification succeeds. We describe the corresponding CI game and prove the following theorem about the equilibrium strategies of a rational server.

Game tree. When an $(\epsilon, 0, 1)$ -secure APoS is enhanced with an ALA between a rational server and an honest client, it results in the CIG in Figure 3. The server has a choice between Tamp and NoTamp. When the server chooses not to tamper with the file, the client’s check always passes. Since the client is constrained to honest actions, it does nothing. The utilities to both the server and the client are just the costs of running the check c_S, c_C . On the other hand, when the server tampers with the file (right sub-tree) it is detected by the client with probability $(1 - \epsilon)$. This probabilistic event is denoted by the hollow chance node. When the tampering is detected, the client invokes the (trivial) Judge, who always punishes the server. The damages are paid as compensation to the client. The gain to the server from tampering with the file is denoted as g_S and the corresponding loss to the client as ℓ_C . If the server tampers with the file, this loss to the client and the gain to the server is part of their utilities – even if the tampering is not detected by the client. We analyze the game in the following Theorem, whose proof is in the full version of our paper [15].

THEOREM 7.1. *Let Π be an $(\epsilon, 0, 1)$ -secure APoS protocol between a rational server and an honest client that is augmented with an ALA*

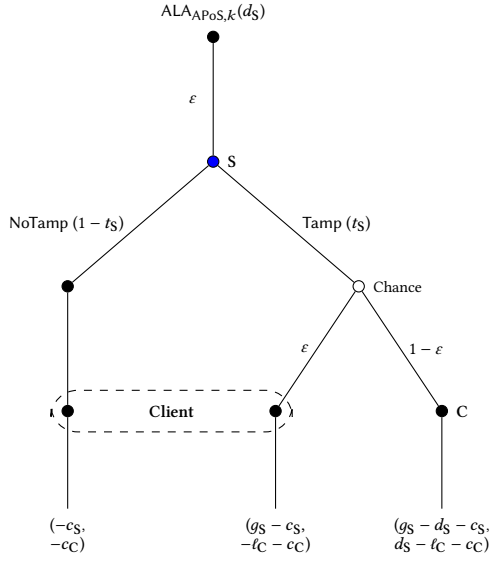


Figure 3: CIG between a rational server and an honest client.

with damages d_S . If,

$$d_S \geq \frac{g_S}{(1 - \epsilon)},$$

where g_S is the server’s gain when tampering with the file, then the dominant strategy for a rational server is to not tamper.

7.2 Rational Server vs. Framing Client

We now consider dishonest clients and, specifically, clients that may want to frame the server. As in the 2PC case, this is an important setting because the use of an ALA introduces an incentive for the client to frame the server. As before, we handle framing clients with a dual ALA between the client and the Judge that specifies the damages payable to the Judge if the client is caught deviating. In addition, however, we also include a server fee φ that the client pays to the server when a proof verifies. Now, the ALA, the dual ALA and the server fee together specify a contract parameterized by (d_S, d_C, φ) .

Equilibrium conditions. In this setting, we have no guarantee on client behavior. In particular, the client might invoke the Judge even when the APoS verification passes. This, in turn, affects the game from the previous section since the server is no longer sure of the client’s behavior. In fact, a rational server may not even want to participate in the protocol. Our first challenge then is to prove that, even in this setting, there exists an equilibrium where the client invokes the Judge *if and only if* the APoS fails to verify and the server participates in the protocol (in the sense that it responds to the challenges issued by the client). Notice how this equilibrium behavior is identical to the honest client setting. At this equilibrium, we can then derive the soundness error ϵ that will force the server to play NoTamp. We prove the existence of the equilibrium in Lemma 7.2 and Lemma 7.3 and show the final result about soundness error in Theorem 7.4. We start by proving that an equilibrium exists where both the server and the client behave as in the honest client setting.

Server fee. In order to have a rational server respond to a challenge, even to a possibly framing client, we need to include a server fee so that the server has some utility for completing a proof. Note that in the context of PoS the server fees occur naturally; that is, most cloud services would charge a fee to store a client’s file. Given that the honest server always receives φ for completing a proof, the honest server will always respond to a challenge. Then if the client does not receive a response to its challenge it assumes that the server has played Tamp and invokes the Judge. This forces the server to respond to all challenges *even if it has tampered with the file*. This may seem surprising but the intuition is as follows. If the server does not respond to the challenge, then the client will invoke the Judge which increases its cost due to the damages it has to pay. On the other hand, if the server always responds the soundness error of the APoS creates the possibility that the verification might succeed even if the file was tampered with. Then we have the first part of the equilibrium condition: *If the client invokes the Judge only when the APoS verification fails (or) when it does not receive a response, then the server will always respond to challenges.*

Game tree. The game and the payoffs for the server are shown in Figure 4. After the server chooses whether to play Tamp or NoTamp, the server can choose whether to “respond to the APoS challenge” which we denote Resp or “not to respond” which we denote NoResp. Since the client actions are fixed in this game, the client invokes the Judge only when either receiving no response from the server (or) when the APoS verification fails. For clarity, in Figure 4, the outcomes when the Judge is invoked are compressed into a single outcome and marked with the server’s expected utility as follows: (1) if the server plays NoTamp and responds to the challenge, the Judge is never invoked; (2) if the server plays Tamp and responds, the Judge is invoked when the proof fails (with probability $1 - \epsilon$); (3) If the server does not respond, the client’s view is the same in both the left and the right sub-tree— it does not know if the server has played Tamp or NoTamp first, it only knows that NoResp was played second. Then the green nodes from the game tree are in the same *information set* and the actions that are played from either must be identical. Since the action available to the client is to invoke or not invoke the Judge, the probability of the Judge being invoked (p) is the same at both the green nodes. If the server gets away with tampering, it receives the gain g_S and if the server is found guilty of tampering (even in error) it pays the damages d_S . We now state Lemma 7.2, whose proof is in the full version, about the server response to a challenge using the server’s expected payoffs [15].

LEMMA 7.2. *Let Π be an $(\epsilon, \alpha, \delta)$ -secure APoS protocol between a rational server and a framing client augmented with an ALA with damages d_S and d_C and a fee φ . If,*

$$\varphi \geq \frac{c_S}{\epsilon}$$

where c_S is the server’s cost of participation, and if the client invokes the Judge only if APoS verification fails or the server does not respond, then the dominant strategy for a rational server is to always respond to a challenge.

We have now shown the first part of the equilibrium condition: *If the client invokes the Judge only when the APoS verification fails (or) when the server does not respond then the server will*

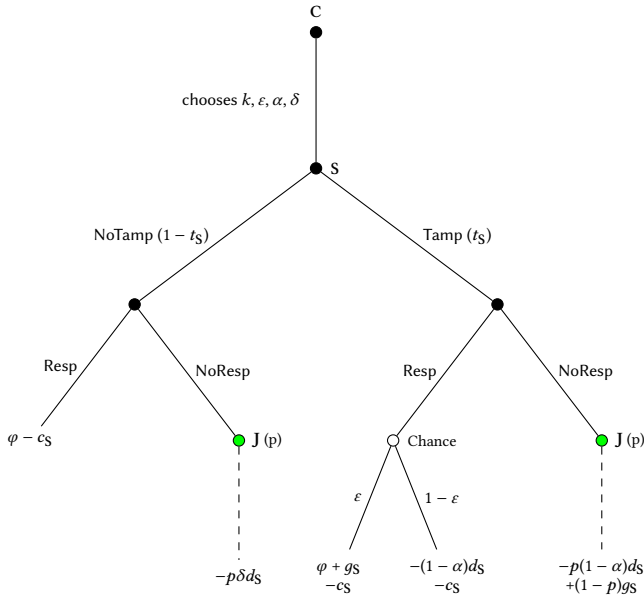


Figure 4: Server responses in CI game.

always respond to challenges. It only remains to show that *if* the server always responds to challenges *then* the client invokes the Judge only when the APoS verification fails.

Game tree. If the server responds to all challenges, we have the game in Figure 5 where the leaves show the client payoffs. The server makes a choice to play Tamp or NoTamp. The client receives the server response to its challenge. If the server plays NoTamp the verification always passes. However, if the server plays Tamp, the verification passes with the probability ϵ – the soundness of the APoS. In both these outcomes the client has the same information (at the blue nodes). After the verification, the client has a choice to invoke (I) or not invoke (NI) the Judge. If the verification fails, the Judge awards the server damages to the client except with probability α . If the verification passes, the Judge’s audit depends on if the server is playing Tamp or NoTamp. If the server is honest, the client pays damages to the Judge else the server pays damages to the client. We now state the following lemma, whose proof is in the full version [15], which completes the proof that an equilibrium exists where the server responds to all challenges and the client only invokes the Judge if the verification fails.

LEMMA 7.3. *Let Π be an $(\epsilon, \alpha, \delta)$ -secure APoS between a rational server and a framing client that is augmented with an ALA with damages d_S and d_C . If the server always responds to the APoS challenges and if*

$$\alpha \leq \frac{d_S}{d_S + d_C}, \quad \delta \leq \frac{d_C}{d_S + d_C}$$

$$\text{and } t_S \leq \frac{-(\delta d_S - (1 - \delta)d_C)}{-(\delta d_S - (1 - \delta)d_C) + \epsilon((1 - \alpha)d_S - \alpha d_C)}$$

where t_S is the probability that the server tampers with the file, then the dominant strategy for a rational client is to invoke the Judge if and only if the server’s proof does not verify.




Figure 5: Rational Client actions in APoS

It follows from Lemmas 7.2 and 7.3 that there exists an equilibrium where a rational client will honestly execute an $(\epsilon, \alpha, \delta)$ -secure APoS and a rational server will always respond to challenges with a proof. This game now simplifies to the tree in Figure 6, which is the same as that of the honest client and rational server in Section 7.1. Then under the conditions for Lemmas 7.2 and 7.3, we can use the insight from the honest client game and derive the values of the parameters ϵ , α , and δ such that the server will not deviate from the protocol. Our final theorem, whose proof is in the full version, concerns the choice of parameters that a client needs to declare to incentivize a rational server to not tamper with the file [15].

THEOREM 7.4. *Let Π be an $(\epsilon, \alpha, \delta)$ -secure APoS between a rational server and an honest client that is augmented with an ALA with damages d_S and d_C and fee ϕ . If,*

$$\phi \geq c_S/\epsilon, \quad \alpha \leq \frac{d_S}{d_S + d_C}, \quad \delta \leq \frac{d_C}{d_S + d_C}, \quad \phi + (1 - \alpha)d_S \geq g_S/(1 - \epsilon)$$

$$\text{and } t_S \leq \frac{-(\delta d_S - (1 - \delta)d_C)}{\epsilon((1 - \alpha)d_S - \alpha d_C) - (\delta d_S - (1 - \delta)d_C)},$$

then the dominant strategy for a rational server is not to tamper with the file.

Concrete instantiations. In summary, our framework can be applied to general PoS protocols as follows: (1) use our general PoS-to-APoS transformation to obtain the corresponding APoS protocol; (2) augment the resulting APoS protocol with an ALA; (3) set concrete deterrence factors according to our theorems. We present applications of our framework to the PoS of Juels and Kaliski [19], Ateniese et al. [6] and Shacham and Waters [21] in Appendix C.

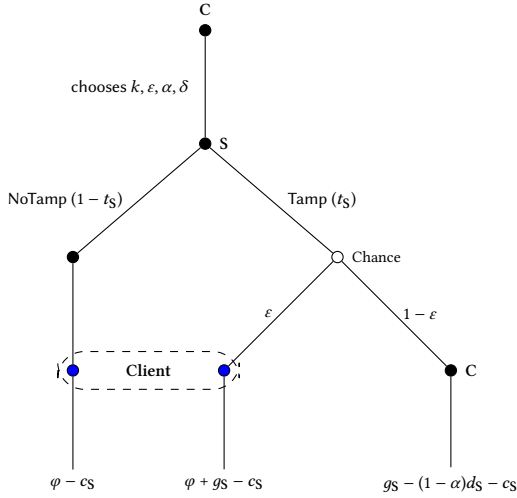


Figure 6: Server-Client Game Tree

8 CONCLUSIONS AND FUTURE WORK

We introduced the notion of an ALA and showed how to apply it to non-trivial two-party protocols including 2PC and PoS. We showed that ALA-enhanced protocols can both enforce deterrence and be more efficient than previous protocols. Our work motivates new and interesting directions for future exploration. The most immediate is to design ALAs for other two-party protocols like zero-knowledge proofs and verifiable computation. Also, extending our framework to multi-party protocols, including to secure multi-party computation, would be interesting and non-trivial since the resulting CIGs would be substantially more complex to define and analyze. Another interesting direction would be to design and study new and more complex kinds of ALAs, for example, using techniques from financial engineering (i.e., stocks, derivatives etc.) and microeconomics (e.g., contract theory). Yet another interesting, but orthogonal direction, would be to study how to set the parameters of the ALA for different application settings, depending on the the parties involved in the secure protocol.

REFERENCES

- [1] [n. d.]. Amazon S3 Pricing. <https://aws.amazon.com/s3/pricing/>. Accessed: 2019-05-29.
- [2] [n. d.]. Average Energy Prices, Seattle-Tacoma-Bellevue – April 2019. https://www.bls.gov/regions/west/news-release/averageenergyprices_seattle.htm. Accessed: 2019-05-29.
- [3] [n. d.]. Filecoin: A Decentralized Storage Network. <https://filecoin.io/filecoin.pdf>. Accessed: 2019-03-06.
- [4] [n. d.]. Once You're in the Cloud, How Expensive Is It to Get Out? <https://www.nefiber.com/blog/cloud-egress-charges/>. Accessed: 2019-05-29.
- [5] Gilad Asharov and Claudio Orlandi. 2012. Calling out cheaters: Covert security with public verifiability. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 681–698.
- [6] Giuseppe Ateniese, Randal Burns, Reza Curtmola, Joseph Herring, Lea Kissner, Zachary Peterson, and Dawn Song. 2007. Provable data possession at untrusted stores. In *Proceedings of the 14th ACM conference on Computer and communications security*. Acm, 598–609.
- [7] Giuseppe Ateniese, Seny Kamara, and Jonathan Katz. 2009. Proofs of storage from homomorphic identification protocols. In *Asiacrypt*. Springer, 319–333.
- [8] Yonatan Aumann and Yehuda Lindell. 2007. Security against covert adversaries: Efficient protocols for realistic adversaries. In *Theory of Cryptography Conference*. Springer, 137–156.

- [9] Rudolf Avenhaus, Bernhard Von Stengel, and Shmuel Zamir. 2002. Inspection games. *Handbook of game theory with economic applications* 3 (2002), 1947–1987.
- [10] Pablo Daniel Azar and Silvio Micali. 2012. Rational proofs. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 1017–1028.
- [11] Carsten Baum, Ivan Damgård, and Claudio Orlandi. 2014. Publicly auditable secure multi-party computation. In *Security and Cryptography for Networks*. Springer, 175–196.
- [12] Robert Cunningham, Benjamin Fuller, and Sophia Yakubov. 2017. Catching MPC cheaters: Identification and openness. In *International Conference on Information Theoretic Security*. Springer, 110–134.
- [13] R Damasevicius, G Ziberkas, V Stukys, and J Toldinas. [n. d.]. Energy consumption of hash functions. *Elektronika ir elektrotechnika* 18, 10 ([n. d.]), 81–84.
- [14] Yevgeniy Dodis, Salil Vadhan, and Daniel Wichs. 2009. Proofs of retrievability via hardness amplification. In *Theory of Cryptography Conference*. Springer, 109–127.
- [15] Marilyn George and Seny Kamara. 2020. Adversarial Level Agreements for Two-Party Protocols. Cryptology ePrint Archive, Report 2020/1249. <https://ia.cr/2020/1249>.
- [16] S Dov Gordon and Jonathan Katz. 2006. Rational secret sharing, revisited. In *International Conference on Security and Cryptography for Networks*. Springer, 229–241.
- [17] Joseph Halpern and Vanessa Teague. 2004. Rational secret sharing and multiparty computation. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. ACM, 623–632.
- [18] Cheng Hong, Jonathan Katz, Vladimir Kolesnikov, Wen-jie Lu, and Xiao Wang. 2019. Covert Security with Public Verifiability: Faster, Leaner, and Simpler. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 97–121.
- [19] Ari Juels and Burton S Kaliski Jr. 2007. PORs: Proofs of retrievability for large files. In *Proceedings of the 14th ACM conference on Computer and communications security*. Acm, 584–597.
- [20] Vladimir Kolesnikov and Alex J. Malozemoff. 2015. Public Verifiability in the Covert Model (Almost) for Free. In *Advances in Cryptology – ASIACRYPT 2015*, Tetsu Iwata and Jung Hee Cheon (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 210–235.
- [21] Hovav Shacham and Brent Waters. 2008. Compact proofs of retrievability. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 90–107.
- [22] Bernhard Von Stengel and Shmuel Zamir. 2010. Leadership games with convex strategy sets. *Games and Economic Behavior* 69, 2 (2010), 446–457.
- [23] Arvinderpal S Wander, Nils Gura, Hans Eberle, Vipul Gupta, and Sheueling Chang Shantz. 2005. Energy analysis of public-key cryptography for wireless sensor networks. In *Third IEEE international conference on pervasive computing and communications*. IEEE, 324–328.
- [24] Andrew C Yao. 1982. Protocols for secure computations. In *23rd annual symposium on foundations of computer science (sfcs 1982)*. IEEE, 160–164.
- [25] Ruiyu Zhu, Changchang Ding, and Yan Huang. 2019. Efficient Publicly Verifiable 2PC over a Blockchain with Applications to Financially-Secure Computations. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019*, Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz (Eds.). ACM, 633–650. <https://doi.org/10.1145/3319535.3363215>

A THE POS TO APOS TRANSFORMATION

THEOREM A.1. *If PoS is sound, Sig is existentially unforgeable and Hash is collision-resistant with security parameter k , then APos as described in Figure 7 is $(\epsilon, \alpha, \delta)$ -secure where ϵ, α , and δ are negligible in k .*

PROOF. The soundness of APos follows directly from the soundness of PoS so we focus here on auditability.

To show accountability, we show that if there exists a PPT adversary \mathcal{A} and a file f such that

$$\Pr \left[\text{Accnt}_{\mathcal{A}, \mathcal{F}}(k) = 1 \right] = \epsilon(k),$$

for a non-negligible function $\epsilon(k)$, then there exists a PPT adversary \mathcal{F} such that $\Pr \left[\text{Forge}_{\mathcal{F}}(k) \right] = \epsilon_1(k)$ or $\Pr \left[\text{Coll}_{\mathcal{F}}(k) \right] = \epsilon_2(k)$, where ϵ_1 and ϵ_2 are non-negligible in k .

Given hash key w and access to a signing oracle $\text{Sign}(\text{sk}, \cdot)$, \mathcal{F} starts by computing $(\text{sk}, \text{pk}) \leftarrow \text{PoS.Gen}(1^k)$, $(\text{sk}_S, \text{vk}_S) \leftarrow \text{Sig.Gen}(1^k)$ and setting $\text{st}_j = \text{vk}_S$, $\text{srk} = (\text{pk}, \text{sk}_S, \text{vk}, w)$. It then simulates \mathcal{A} and

Let $\text{PoS} = (\text{Gen}, \text{Encode}, \text{Chall}, \text{Prove}, \text{Vrfy})$ be a publicly-verifiable proof of storage, $\text{Sig} = (\text{Gen}, \text{Sign}, \text{Vrfy})$ be a signature scheme and $\text{Hash} = (\text{Gen}, \text{Hash})$ be a hash function family. Consider the auditable proof of storage $\text{APoS} = (\text{Gen}, \Pi, \text{Chall}, \text{Prove}, \text{Vrfy}, \text{Audit})$ defined as follows:

- $\text{Gen}(1^k)$:
 - (1) compute $(\text{sk}, \text{pk}) \leftarrow \text{PoS.Gen}(1^k)$;
 - (2) compute $w \xleftarrow{\$} \text{Hash.Gen}(1^k)$;
 - (3) compute $(\text{sk}_C, \text{vk}_C) \leftarrow \text{Sig.Gen}(1^k)$ and $(\text{sk}_S, \text{vk}_S) \leftarrow \text{Sig.Gen}(1^k)$;
 - (4) output $\text{st}_J = \text{vk}_S, \text{clk} = (\text{clpk}, \text{clsk}) = ((\text{pk}, \text{vk}_C, w), (\text{sk}, \text{sk}_C))$ and $\text{srk} = (\text{srpk}, \text{srsk}) = ((\text{pk}, \text{vk}_S), (\text{sk}_S, w))$
- $\Pi((\text{clk}, \mathbf{f}); \text{srk})$:
 - (1) the client:
 - (a) computes $(\mathbf{e}, \text{st}) \leftarrow \text{PoS.Encode}(\text{sk}, \mathbf{f})$;
 - (b) computes $h \leftarrow \text{Hash}(w, \mathbf{e})$;
 - (c) computes $\sigma_C \leftarrow \text{Sign}(\text{sk}_C, h)$;
 - (d) sends \mathbf{e} and $\tau = (h, \sigma_C)$ to the server;
 - (2) the server:
 - (a) if $h = \text{Hash}(w, \mathbf{e})$
 - (i) computes $\sigma_S \leftarrow \text{Sign}(\text{sk}_S, h)$;
 - (ii) sends $\text{rec} = (h, \sigma_C, \sigma_S)$ to the client
 - (b) else outputs \perp and aborts;
 - (3) the client:
 - (a) if $\text{Vrfy}(\text{vk}_S, h, \sigma_S) = 1$ outputs rec otherwise outputs \perp ;
- $\text{Chall}(\text{clpk}, \text{st})$: output $\text{ch} \leftarrow \text{PoS.Chall}(\text{pk}, \text{st})$;
- $\text{Prove}(\text{srk}, \mathbf{e}, \text{ch})$: output $\pi \leftarrow \text{PoS.Prove}(\text{pk}, \mathbf{e}, \text{ch})$;
- $\text{Vrfy}(\text{clk}, \text{st}, \text{ch}, \pi)$: output $b \leftarrow \text{PoS.Vrfy}(\text{pk}, \text{st}, \text{ch}, \pi)$;
- $\text{Audit}(\text{st}_J, \text{rec}^*, \mathbf{e}^*)$:
 - (1) parse rec^* as $(\text{rec}_1, \text{rec}_2, \text{rec}_3)$ i.e. (h, σ_C, σ_S) ;
 - (2) compute $b_1 \leftarrow \text{Sig.Vrfy}(\text{vk}_C, h, \text{rec}_2)$;
 - (3) compute $b_2 \leftarrow \text{Sig.Vrfy}(\text{vk}_S, h, \text{rec}_3)$;
 - (4) if $h = \text{Hash}(w, \mathbf{e}^*)$ set $b_3 = 1$ else set $b_3 = 0$;
 - (5) if $b_1 = 1 \wedge b_2 = 1 \wedge b_3 = 1$ output 0
 - (6) if $b_1 = 1 \wedge b_2 = 1 \wedge b_3 = 0$ output 1
 - (7) otherwise output 0.

Figure 7: A PoS-to-APoS transformation.

executes $((\text{st}_C, \text{rec}), \mathbf{e}) \leftarrow \Pi_{\mathcal{A}, \mathcal{S}}((\text{clk}, \mathbf{f}); \text{srk})$ playing the role of the client C. During this execution, \mathcal{F} generates the signature σ_S by querying its Sign oracle on h . It then sends $\tau = (h, \sigma_C)$ to \mathcal{A} and records (h, σ_C) as a pair queried to the Sign oracle. \mathcal{A} sends back $\text{rec}^* = (h^*, \sigma_C^*, \sigma_S^*)$, then \mathcal{F} aborts unless $h^* = \text{Hash}(w, \mathbf{e})$ and the signatures verify. If the client signature verifies but (h^*, σ_C^*) has not been queried to the oracle, \mathcal{F} produces this as a forgery and terminates the experiment. Notice that \mathcal{A} 's view up to this point when simulated by \mathcal{F} is distributed exactly as its view during an $\text{Accnt}_{\mathcal{A}, \mathcal{F}}(k)$ experiment. If the experiment has not been terminated by \mathcal{F} , then \mathcal{A} 's view is still consistent and the receipt $\text{rec} = (h, \sigma_C, \sigma_S)$ is produced honestly. Then by our initial assumption and the definition of the Accnt experiment, this implies that

$$\Pr [\text{Audit}(\text{st}_J, \text{rec}, \mathbf{e}^*) = 0 \wedge \mathbf{e} \neq \mathbf{e}^*] \geq \varepsilon(k).$$

But note that $\text{Audit}(\text{st}_J, \text{rec}, \mathbf{e}^*)$ outputs 0 if and only if h is a valid hash of \mathbf{e}^* . Then \mathcal{F} produces $(\mathbf{e}, \mathbf{e}^*)$ as the hash collision. Since the total probability of \mathcal{A} winning is non-negligible it follows that it either produces a forged signature in the receipt rec with some $\varepsilon_1(k)$ or an \mathbf{e}^* with the same hash value as \mathbf{e} with $\varepsilon_2(k)$ where $\varepsilon_1, \varepsilon_2$ are non-negligible in k . But since Sig is existentially unforgeable and Hash is collision-resistant we have a contradiction. Hence we know that

$$\Pr [\text{Accnt}_{\mathcal{A}, \mathcal{F}}(k) = 1] \leq \alpha(k),$$

where $\alpha(k)$ is negligible. Towards showing defamation-freeness, we show that if there exists a PPT adversary \mathcal{A} and a file \mathbf{f} such that

$$\Pr [\text{Defame-Free}_{\mathcal{A}, \mathcal{F}}(k) = 1] = \varepsilon(k),$$

for a non-negligible function $\varepsilon(k)$, then there exists a PPT adversary \mathcal{F} such that $\Pr [\text{Forge}_{\mathcal{F}}(k)] = \varepsilon(k)$.

Given vk and access to a signing oracle $\text{Sign}(\text{sk}, \cdot)$, \mathcal{F} starts by computing $(\text{sk}, \text{pk}) \leftarrow \text{PoS.Gen}(1^k)$, $w \leftarrow \text{Hash.Gen}(1^k)$ and $(\text{sk}_C, \text{vk}_C) \leftarrow \text{Sig.Gen}(1^k)$ and setting $\text{st}_J = \text{vk}$ and $\text{clk} = (\text{sk}, \text{pk}, \text{sk}_C, \text{vk}_C, w)$. It then simulates \mathcal{A} and executes $((\text{st}_C, \text{rec}), \mathbf{e}) \leftarrow \Pi_{\mathcal{A}, \mathcal{S}}((\text{clk}, \mathbf{f}); \text{srk})$ playing the role of the server S. During this execution, \mathcal{F} generates the signature σ_S by querying its Sign oracle on h . When the execution finishes and \mathcal{A} outputs $\text{rec}^* = (h^*, \sigma_C^*, \sigma_S^*)$, \mathcal{F} outputs the pair (h^*, σ_S^*) as its forgery. Notice that \mathcal{A} 's view when simulated by \mathcal{F} is distributed exactly as its view during an $\text{Defame-Free}_{\mathcal{A}, \mathcal{F}}(k)$ experiment. By our initial assumption and the definition of the Defame-Free experiment, this implies that

$$\Pr [\text{Audit}(\text{st}_J, \text{rec}^*, \mathbf{e}) = 1] \geq \varepsilon(k).$$

But note that $\text{Audit}(\text{st}_J, \text{rec}^*, \mathbf{e})$ outputs 1 if and only if σ_S^* is a valid signature on h^* and if $\text{Hash}(w, \mathbf{e}) \neq h^*$. Since \mathcal{F} only queries its Sign oracle on $\text{Hash}(w, \mathbf{e})$ (if $h \neq \text{Hash}(w, \mathbf{e})$ it aborts), it follows that (h^*, σ_S^*) is a new and valid message/signature pair. This is a contradiction to the existential unforgeability of Sig . Therefore we

have

$$\Pr \left[\text{Defame-Free}_{\mathcal{A}, \mathcal{F}}(k) = 1 \right] \leq \delta(k),$$

where $\delta(k)$ is negligible. ■

B CONCRETE IMPROVEMENTS FOR PVC PROTOCOLS

In this section, we apply our framework of Section 5 to three concrete PVC protocols to achieve two kinds of improvements: (1) setting concrete deterrence factors; and (2) improving efficiency.

The Hong et al. protocol [18]. The most efficient PVC protocol is by Hong et al. [18]. Its deterrence factor is $\varepsilon = 1 - (1/\lambda)$, where λ is the number of garbled circuits generated by the garbler so the smallest deterrence factor that can be achieved is $1/2$ (with two circuits). Also, it achieves negligible accountability μ and defamation-freeness ϕ . For concreteness, we assume that the security parameter k is set such that μ and ϕ are at most $1/64$ and that $u_G^i \leq \$100$, $u_G^o \geq \$50$ and $u_E^w \leq \$100$. By Lemmas 5.1 and 5.2, we can incentivize honest behavior by using an ALA with damages

$$d_G \geq \frac{u_G^i - u_G^o}{\varepsilon(1 - \mu)} = \$101.60 \quad \text{and} \quad d_E \geq \frac{\phi d_G}{1 - \phi} = \$1.61.$$

The Koleznikov and Malozemoff protocol [20]. Another PVC protocol was proposed by Koleznikov and Malozemoff (KM). Its deterrence factor is $\varepsilon = (1 - 1/\lambda)(1 - 2^{-\eta+1})$, where λ is the number of garbled circuits and η is the XOR-tree replication factor. The smallest deterrence factor possible is $1/4$ with $\lambda = 2$ and $\eta = 2$ but Koleznikov and Malozemoff set ε to $1/2$. Using ALAs, we can guarantee honest behavior even with the minimal deterrence factor $\varepsilon = 1/4$. As above, we assume the security parameter is set so that μ and ϕ are at most $1/64$ and that $u_G^i \leq \$100$, $u_G^o \geq \$50$ and $u_E^w \leq \$100$. Again, by Lemmas 5.1 and 5.2, we can incentivize honest behavior by using an ALA with damages

$$d_G \geq \frac{u_G^i - u_G^o}{\varepsilon(1 - \mu)} = \$203.20 \quad \text{and} \quad d_E \geq \frac{\phi d_G}{1 - \phi} = \$3.22.$$

The same analysis holds for the protocol Asharov and Orlandi (AO) since the deterrence factor is the same. Notice, however, that our use of ALAs also improves the computation and communication complexity of the protocols since we can use the minimal deterrence factor. These protocols have several parameters which include the bit-length of their underlying field τ , the length of the inputs n , and the size of their commitments κ , the number of replicated circuits λ , and the XOR-tree replication factor ν . In addition, for the KM protocol let t be the number of OTs and assume that $3\kappa > \tau$. Based on the analysis of [20], the communication complexity of the KM protocol when using signed OTs is,

$$\tau(7t + 11) + 2\lambda\kappa t + \lambda\kappa(2n + 1) + \tau(13\lambda - 4) + \lambda(2\kappa(\lambda - 1) + n\kappa) + \log(\lambda) + 2\kappa|G_C|,$$

where $|G_C|$ is the number of non-XOR gates in the circuit. The communication cost of the AO protocol with signed OTs is,

$$\tau(7\nu n + 11) + 2\lambda\kappa\nu n + \lambda(2\kappa|G_C| + \tau) + 2n\lambda(\kappa + \tau) + \tau(13\lambda - 8) + \lambda\kappa(2(n + \nu n)(\lambda - 1) + 2n(\lambda - 1) + n).$$

For $\kappa = 128$, $\tau = 256$ and $\lambda = \nu = 3$, which leads to $\varepsilon = 1/2$, the KM protocol needs 3.47Mbit to compute a 128-bit AES circuit with 9100 non-XOR gates. For these same parameters, the AO protocol requires 9.3Mbit. Using the ALA described above, however, we can set $\lambda = \nu = 2$ which leads to the minimal deterrence factor of $1/4$

and a communication complexity of 3.03Mbit for the KM protocol and 5.75Mbit for the AO protocol.

C CONCRETE INSTANTIATIONS FOR POS

C.0.1 The Juels-Kaliski PoS. In the paper, Juels and Kaliski analyze a concrete example - the file to be stored is encoded such that it can be retrieved for upto $\delta = 0.01$ fraction of corruption with high probability. Sentinel blocks are then inserted such that a corruption rate of $\delta = 0.005$ can be detected with probability at least 0.71 [19]. The guarantee provided by their scheme is that of high enough detection of deviation i.e. corruption of the file. We show how to augment this scheme with an appropriate ALA and damages in order to guarantee that a rational server *will not deviate*. We show that we need fewer sentinel blocks in order to provide this stronger guarantee due to the presence of the ALA and the damages which come ‘for free’.

Example from Juels and Kaliski [19]. We consider a 2GB file f with 128-bit block size and a (255, 223, 32) Reed Solomon code over $GF[2^{128}]$. This results in a 2^{27} blocks before erasure encoding and 153, 477, 870 blocks after erasure encoding. Adding $s = 1, 000, 000$ sentinels, the file now has 154, 477, 870 blocks which amounts to about 2.3GB. Extending the analysis, consider a corruption rate of $\delta = 0.01$ the adversary will corrupt the file beyond retrievability with probability at most 0.089. To detect a corruption rate of $\delta = 0.005$, they use 1000 sentinels per challenge which guarantees detection with probability at least 0.71. Keeping the storage constant we show that we can enhance the Juels-Kaliski PoS with an ALA and guarantee that the rational server will not corrupt the file above $\delta = 0.005$ using only 200 sentinels per challenge. Although the detection probability is now lower (~ 0.23) we can adjust the damages to get this result. We now show how to calculate the appropriate parameters for this ALA.

Estimating server costs. First we estimate the cost incurred by the server to execute the protocol. During the encode protocol Π , the server needs to compute a hash, verify a signature and sign a hash. SHA-256 is estimated to consume $100\mu\text{Wh}$ per MB[13] so the energy to hash a 2.3GB file is $2.355 \times 10^{-4}\text{kWh}$. The cost of verifying plus signing a 32-bit SHA-256 hash using RSA-1024 is about $350\text{mJ} \approx 10^{-7}\text{kWh}$ [23]. The average energy cost in Seattle is 12 cents per month per kWh[2] so the cost of the hash, the signature and the verification amount to 30×10^{-4} cents. To answer a challenge, the server needs to return 200 blocks to the client. This amounts to sending $200 \times 128\text{bits}$ which is 25.6kB back to the client. The AWS egress prices (i.e., the cost to move data out of a cloud AWS) is 0.09\$ per GB[4], so sending the proof will cost about 21.6×10^{-5} cents. So the total cost for the server is $\sim 32 \times 10^{-4}$ cents.

Setting the parameters of the ALA. From this analysis, we should set the server fee to $\varphi \geq c_S/\varepsilon = 32 \times 10^{-4}/0.77 \approx 0.004$ cents. Next, we have to ensure that $\varphi + (1 - \alpha)d_S \geq \frac{g_S}{1 - \varepsilon}$, where α is the upper bound on the accountability of the APoS, d_S are the damages the server must pay if caught cheating and g_S is the gain the server incurs by losing the file. To estimate g_S we use the price of storing 2.3GB for a year on Amazon S3 which is $\$0.6348 \approx \0.7 [1]. Using RSA-1024 and SHA-256 implies that δ and α are approximately 0. So if we set the server damages in the ALA to be $d_S = \$4$ then the condition above is satisfied. For the equilibrium to exist, we also

need the server’s probability of violation to be bounded above as

$$t_S \leq \frac{-(\delta d_S - (1-\delta)d_C)}{\varepsilon((1-\alpha)d_S - \alpha d_C) - (\delta d_S - (1-\delta)d_C)} \approx \frac{d_C}{\varepsilon d_S + d_C}.$$

To get the guarantee for $t_S \leq 0.75$ we need the client damages in the dual ALA to be $d_C = \$10$. We can always adjust the bound for t_S to be arbitrarily close to 1 by increasing the client’s damages. Finally we need both $\alpha \leq d_S/(d_S + d_C) = 2/7$, and $\delta \leq d_C/(d_S + d_C) = 5/7$ which is also satisfied. Then by using an ALA with damages of \$4 and a dual ALA with damages of \$10, we can guarantee that a rational server and a framing client will behave honestly using only 200 sentinels per challenge as opposed to the 1000 sentinels per challenge needed originally. This will also allow us to run five times as many challenges. We note that we could trade off storage or the parameters of the error-correcting code similarly, as long as we ensure that the conditions for Theorem 7.4 hold.

C.0.2 The Ateniese et al. PoS. We now apply our framework to the publicly-verifiable proofs of data possession S-PDP by Ateniese et al [6]. In their paper they show that for a file of n blocks where the server corrupts t of them and the client asks for c blocks per challenge then the probability that the client detects this corruption p is given by: $1 - \left(\frac{n-t}{n}\right)^c \leq p \leq 1 - \left(\frac{n-c+1-t}{n-c+1}\right)^c$.

They then say that when $t = 1\%$ of n the client must ask for $c = 460, c = 300$ blocks in order to achieve p of 99% and 95% respectively [6]. In our setting the above relationship implies for the soundness error ε which is $1 - p$: $\left(\frac{n-t}{n}\right)^c \geq \varepsilon \geq \left(\frac{n-c+1-t}{n-c+1}\right)^c$.

Now we show that we can provide the guarantee that the rational server will not deviate with $c = 50$ blocks and d_S, d_C equal to \$2, \$7 respectively. In order to calculate the parameters we choose a 2GB file and assume that with tags it expands to at most 2.5GB. Let each block be 4KB as per the optimal size in the paper. Then $n = 524288, t = 5243$ and from the above we have: $0.60499 \geq \varepsilon \geq 0.60497$ *Estimating server costs.* First we estimate the cost incurred by the server to execute the protocol. During the encode protocol Π , the server needs to compute a hash, verify a signature and sign a hash. SHA-256 is estimated to consume $100\mu\text{Wh}$ per MB[13] so the energy to hash a 2.5GB file is $2.56 \times 10^{-4}\text{kWh}$. The cost of verifying plus signing a 32-bit SHA-256 hash using RSA-1024 is about $350\text{mJ} \approx 10^{-7}\text{kWh}$ [23]. The average energy cost in Seattle is 12 cents per month per kWh[2] so the cost of the hash, the signature and the verification amount to 30.7×10^{-4} cents. To answer a challenge in the publicly-verifiable scheme, the server needs to return ~ 1 block. This amounts to sending at most 5kB back to the client. The AWS egress prices (i.e., the cost to move data out of a cloud AWS) is 0.09\$ per GB[4], so sending the proof will cost about 4.29×10^{-5} cents. So the total cost for the server is $\sim 32 \times 10^{-4}$ cents.

Setting the parameters of the ALA. From this analysis, we should set the server fee to $\varphi \geq c_S/\varepsilon = 32 \times 10^{-4}/0.61 \approx 0.005$ cents. Next, we have to ensure that $\varphi + (1 - \alpha)d_S \geq \frac{g_S}{1-\varepsilon}$, where α is the upper bound on the accountability of the APoS, d_S are the damages the server must pay if caught cheating and g_S is the gain the server incurs by losing the file. To estimate g_S we use the price of storing 2.5GB for a year on Amazon S3 which is $\$0.69 \approx \0.7 [1]. Using RSA-1024 and SHA-256 implies that δ and α are approximately 0. So if we set the server damages in the ALA to be $d_S = \$2$ then the condition above is satisfied. For the equilibrium to exist, we also

need the server’s probability of violation to be bounded above as

$$t_S \leq \frac{-(\delta d_S - (1-\delta)d_C)}{\varepsilon((1-\alpha)d_S - \alpha d_C) - (\delta d_S - (1-\delta)d_C)} \approx \frac{d_C}{\varepsilon d_S + d_C}.$$

To get the guarantee for $t_S \leq 0.85$ we need the client damages in the dual ALA to be $d_C = \$7$. We can always adjust the bound for t_S to be arbitrarily close to 1 by increasing the client’s damages. Finally we need both $\alpha \leq d_S/(d_S + d_C) = 2/9$, and $\delta \leq d_C/(d_S + d_C) = 7/9$ which is also satisfied. We then see that by using an ALA with damages of \$2 and a dual ALA with damages of \$7, we guarantee that a rational server and a framing client will behave honestly using only 50 blocks per challenge as opposed to the ~ 300 blocks needed to guarantee a reasonable chance of detection. We have hence reduced the overall computation required for the server and have shown a stronger guarantee by adding the ALA.

C.0.3 The Shacham-Waters PoS. We consider the example parameterization mentioned in the paper [21]. They take the security parameter $k = 80$ in order to have a small enough soundness error ε . We assume their (negligible) soundness error is $1/2^{k1}$. With our framework we can assume a much smaller security parameter while guaranteeing that the rational server will not cheat. Let the security parameter $k = 10$. Then by our simplification the soundness $\varepsilon = 1/1024 \sim 0.001$. We assume a 2.5GB file as in the previous example and the communication costs are similar. From the above analysis the server’s cost of participation is $\sim 32 \times 10^{-4}$ cents.

Setting the parameters of the ALA. From this analysis, we should set the server fee to $\varphi \geq c_S/\varepsilon = 32 \times 10^{-4}/0.001 \approx 3.2$ cents. Next, we have to ensure that $\varphi + (1 - \alpha)d_S \geq \frac{g_S}{1-\varepsilon}$, where α is the upper bound on the accountability of the APoS, d_S are the damages the server must pay if caught cheating and g_S is the gain the server incurs by losing the file. To estimate g_S we use the price of storing 2.5GB for a year on Amazon S3 which is $\$0.69 \approx \0.7 [1]. Using RSA-1024 and SHA-256 implies that δ and α are approximately 0. So if we set the server damages in the ALA to be $d_S = \$1$ then the condition above is satisfied. For the equilibrium to exist, we also need that the server’s probability of violation is bounded as $t_S \leq \frac{-(\delta d_S - (1-\delta)d_C)}{\varepsilon((1-\alpha)d_S - \alpha d_C) - (\delta d_S - (1-\delta)d_C)} \approx \frac{d_C}{\varepsilon d_S + d_C}$, which is $t_S \leq 0.999$ for the client damages in the dual ALA $d_C = \$1$. Finally we need both $\alpha \leq d_S/(d_S + d_C) = 1/2$, and $\delta \leq d_C/(d_S + d_C) = 1/2$ which is also satisfied. Then by using an ALA and a dual ALA with damages of \$1, we can guarantee that a rational server and a framing client will behave honestly using only $k = 10$. From the paper this now implies that the prime p is 20 bits (down from 160), the challenge space is $B = \{0, 1\}^{10}$ since the challenge length is only 10 elements as opposed to 80 in their example parameterization. In the paper, they attempt to reduce this challenge size by sending a pseudorandom seed instead to the server. The server evaluates the pseudorandom function at specific points (usually a counter) to generate the challenge. However, even with the modification; our framework reduces the amount of computation overall for both the server and the client while providing a stronger guarantee with the addition of the ALA.

¹We note that this is in order to have a concrete value for the example, and any other negligible (or even non-negligible) function can be used similarly.