

Argus: A Fully Transparent Incentive System for Anti-Piracy Campaigns

Xian Zhang*, Xiaobing Guo[†], Zixuan Zeng[‡], Wenyan Liu[§], Zhongxin Guo*, Yang Chen*, Shuo Chen*,
Qiufeng Yin*, Mao Yang*

*Microsoft Research Asia

[†]Alibaba Group

[‡]Carnegie Mellon University

[§]Shanghai Key Laboratory of Trustworthy Computing, East China Normal University, Shanghai, China

{zhxian, zhogu, yachen, shuochen, qfyin, maoyang}@microsoft.com

xiaobing.gxb@alibaba-inc.com

zixuanze@andrew.cmu.edu

wyliu@stu.ecnu.edu.cn

Abstract—Anti-piracy is fundamentally a procedure that relies on collecting data from the open anonymous population, so how to incentivize credible reporting is a question at the center of the problem. Industrial alliances and companies are running anti-piracy incentive campaigns, but their effectiveness is publicly questioned due to the lack of transparency. We believe that full transparency of a campaign is necessary to truly incentivize people. It means that every role, e.g., content owner, licensee of the content, or every person in the open population, can understand the mechanism and be assured about its execution without trusting any single role.

We see this as a distributed system problem. In this paper, we present Argus, a fully transparent incentive system for anti-piracy campaigns. The groundwork of Argus is to formulate the objectives for fully transparent incentive mechanisms, which securely and comprehensively consolidate the different interests of all roles. These objectives form the core of the Argus design, highlighted by our innovations about a Sybil-proof incentive function, a commit-and-reveal scheme, and an oblivious transfer scheme. In the implementation, we overcome a set of unavoidable obstacles to ensure security despite full transparency. Moreover, we effectively optimize several cryptographic operations so that the cost for a piracy reporting is reduced to an equivalent cost of sending about 14 ETH-transfer transactions to run on the public Ethereum network, which would otherwise correspond to thousands of transactions. With the security and practicality of Argus, we hope real-world anti-piracy campaigns will be truly effective by shifting to a fully transparent incentive mechanism.

I. INTRODUCTION

Intellectual property is one of the most valuable assets for present-day companies, especially in the software, movie, gaming and digital publishing industries. Anti-piracy is a long-lasting and heavily invested effort, because piracy impacts the fundamental business models of these industries. Anti-piracy has the legal aspect and the technological aspect. The former crucially depends on the latter to collect credible and undeniable evidences, so that appropriate legal actions can be taken against the infringers. For example, if evidences prove that the number or the retail value of the pirated copies exceeds a certain threshold during any 180-day period, according to the Title 17 of the United States Code, infringers shall be imprisoned for a maximum of 5

years, or fined a maximum \$250,000, or both [1].

Since piracy is fundamentally about disseminating copyrighted contents outside legitimate distribution channels, a central question about anti-piracy is how to incentivize people in the open population to report pirated copies. Industrial alliances (e.g., BSA [2], FACT [3], SIIA [4]) and companies (e.g., Custos [5], Veredictum [6]) have offered big amounts of bounties for piracy reporting. For example, the Business Software Alliance (i.e., BSA [2]), whose members include Apple, IBM, Microsoft, Symantec and many others, posted a \$1-million bounty for reporting. However, the approach is not yet effective and is questioned/criticized by the public, mainly due to the lack of transparency [7]. For example, it is unclear whether the \$1-million total bounty is simply a marketing gimmick, as BSA had only rewarded a small fraction of it in a long period of time. Also, BSA’s neutrality is really problematic, as its members are copyright holders, who may not represent the best interest of the informers in the public. Moreover, it is unclear how BSA evaluates the credibility of the piracy reports or whether they are strong enough against an infringer’s repudiation. Obviously, opaqueness about incentive, fairness, and credibility-criteria seriously limit the effectiveness of these anti-piracy campaigns.

We envision that a methodological progress for anti-piracy campaigns can be made once they are formulated as a decentralized computing problem. It is promising to advance the status quo by distributed-system technologies, especially those about incentive model (NF-Crowd [8], Arbitrum [9], Hydra [10]), consensus mechanism (PBFT-Hyperledger [11], Algorand [12]), secure messaging (Decentralized release [13], Hyperpubsub [14], Zerocash [15], [16]) and Sybil resistance (Arbitrum [9], TrueBit [17]). Toward the vision, we have built a concrete system named Argus. The design is based on a clear problem statement and a set of properties as objectives, which are explained next.

Problem statement. We describe the anti-piracy problem using the following terminology. An *owner* is the one who owns the copyrighted content (e.g., a film maker). The content is distributed through a controlled channel to a set of *licensees* (e.g., cinemas and film critics). Some licensees

^{†‡§}Work done during employment (Xiaobing Guo) and internship (Zixuan Zeng and Wenyan Liu) at Microsoft Research Asia.

may leak their copies of the content, which leads to many pirated copies in the open population. These licensees are called the *infringers*. The anti-piracy system’s goal is to incentivize people in the open population to report the pirated copies to the system. We refer to these people as the *informers*.

The challenge in the anti-piracy problem is that the interests of these roles are different or even conflicting. Specifically, the owner’s goal is to identify infringers and assess the severity of the infringement. The owner wishes that, by giving a financial incentive (e.g., a bounty) to the open population, as many good-faith reports as possible can be received. However, the motivation of informers is not always aligned with the owner. It is only reasonable to assume that informers are financially motivated [8], [9], like black-hat security researchers motivated by bug bounties [10]. Not surprisingly, an infringer’s best interest is to refute the credibility of an evidence by arguing that it could be fabricated by an informer or the owner.

Because of the conflict of interests, an unbiased solution would require a contract that was agreed upon by all these roles. But who should be the executor of the contract? One possibility is to introduce an “executor” role into the problem, as in the BSA situation described earlier, but the role is really undesirable because its neutrality is hard to be assured in reality (e.g., even big companies like Facebook and Google had controversial practices that put their neutrality in doubt [18]). Our work is to explore the feasibility of an open contract of which the execution is fully transparent to the public, without an additional role as the trust basis.

The Argus system. In this paper, we present the design, implementation and evaluation of the Argus system. To the best of our knowledge, it is the first public anti-piracy system which (1) does not hinge on any “trusted” role; (2) treats every participant fairly (in particular, it is resilient to greed and abuse, and resolves conclusively every foreseeable conflict); and (3) is efficient and economically practical to run on a public blockchain (e.g. it achieves an impressive off-chain throughput of 82.6 data-trades per second per machine, and incurs only a negligible on-chain cost equivalent to sending 14 ETH-transfer transactions per report on the public Ethereum blockchain).

The four pillars in the Argus design are *full transparency*, *incentive*, *information hiding* and *optimization*. These are the main focuses to be elaborated on in this paper. It is worth noting that they are not four problems to be solved individually, but integral aspects in one coherent design. We highlight some properties of Argus, which represent some of our core innovations:

Incentivizing good-faith informers. A fundamental challenge is about the interest of informers, who are anonymous people in the open population. The owner’s interest is to

collect good-faith reports so that the severity of the infringement can be accurately estimated. However, each individual informer’s interest is to maximize his own reward. What prevents an informer from creating multiple identities to make multiple reports, so that he gets multiple rewards but causes the owner’s estimation to be inflated? Note that an attack using multiple forged identities is often referred to as the Sybil attack [9], [17]. In Argus, the incentive model ensures that the total reward of the informer and all his Sybils is less than the reward he would get without forging the Sybils. In other words, our model disincentivizes Sybil attacks, so the informers’ interest is aligned with the owner’s. In addition, our model is superior to previous models because of several other properties for better incentives (Section III).

Information hiding for report submission. Because Argus runs on a public ledger, its execution is fully transparent to everybody [15], [16]. It is crucial that an informer is unable to resubmit any report previously submitted by somebody else. For this reason, Argus’ report submission protocol is based on Multi-period Commitment Scheme, which gives a “zero-knowledge” style guarantee, i.e., a submission only proves that the informer has a copy of the content without disclosing other information. Compared to traditional commitment schemes, our scheme leaks no useful information even in the reveal phase while avoiding the heavy cost of zero-knowledge proof (Section IV).

Strong accusation against infringer. An owner’s accusation against a licensee is always subject to an inherent paradox – since both have the leaked copy, why can’t the licensee refute the accusation by arguing that the infringer may be the owner himself? We believe that the only solution to circumvent this paradox is to resort to a probabilistic argument. Argus uses Oblivious Transfer (OT) to ensure that the false accusation is bounded by a probability ϕ , which can be arbitrarily small. Hence, the accusation is very hard to refute. Moreover, we improve the efficiency of leveraging OT, which carefully considers the scalability limitation of distributed ledgers [11], [12] (Section V).

Contributions. Our contributions are as follows:

- We formulate anti-piracy as a problem about consolidating different interests of multiple roles, including informers in the open population. We also clearly state the design objectives of anti-piracy solutions, which give a foundation for this work and future research.
- The Argus contract is fully transparent — no role is considered as the trust base. This is a significant advancement. In addition, our approach is systematic. Because of the clarity on the design objectives, we are able to deduce the general form of the incentive model and identify all the unavoidable technical challenges. Because these challenges are general, solving them in Argus will have a far-reaching impact in the broad problem space.
- To achieve full transparency and a number of novel ob-

jectives, Argus needs to implement sophisticated cryptographic operations as contract code, rather than native code. Optimizations are a vital effort in the design of Argus. We show that, if existing cryptographic operations were adopted without optimization, the cost would equal sending thousands of transactions (as opposed to 11 transactions in Argus), which would make the solution economically unreasonable.

Roadmap. The rest of the paper is organized as follows. Section II gives an overview of Argus to address how we leverage a transparent contract to achieve the mutual trust and fairness between owner, licensees and informers, along with related primitives. In Section III, Section IV and Section V, we further figure out that those primitives should be optimized to achieve four pillars with better practicality. Then, we brief our implementation with optimizations in Section VI. In Section VII, the security analysis and performance evaluation of Argus are provided, followed by related work in Section VIII and a conclusion. Due to the page limit, we move details of mathematical deduction, protocol implementations and security analysis to an extended version of this paper [19].

II. OVERVIEW OF ARGUS

In this section, we give an overview of the anti-piracy solution and Argus contract. We assume familiarity with contract [20] and cryptographic primitives such as oblivious transfer [21], zero-knowledge proof [22] and commitment scheme [23].

A. The Argus Contract

Figure 1 illustrates the important elements in the Argus contract and how different roles interact with it. An instance of the Argus contract is created by the owner of a copy-righted content. We denote M as the number of licensees and list the data fields used in the Argus system in Table I.

When the owner distributes the content to the licensees, she generates a large number of watermarked copies, e.g., $10000 * M$ copies. In other words, each copy is embedded with a unique secret string. In this work, we assume that nobody can remove the watermark from a copy without badly deteriorating the content quality. Despite this assumption, we will describe in Section II-B our insights about improving the robustness of watermark.

A licensee retrieves a copy from the owner through oblivious transfer (OT), which ensures: (1) the owner does

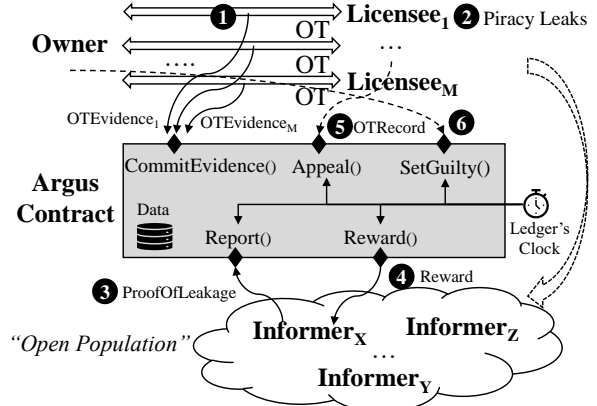


Figure 1: Overview of the Argus system.

not know which of the watermarked copies is retrieved; (2) the licensee knows nothing about other copies except the retrieved one. The OT procedure is performed by private computations on the owner side and the licensee side. One of our important features added to the existing OT protocol is to produce a data record called *OTRecord* on the licensee side, which is needed in case the licensee appeals against an accusation. In addition, the OT scheme produces another piece of data called *OTEvidence*, which opaquely represents the existence of the current OT procedure. Every *OTEvidence* is submitted to the contract via function `CommitEvidence(OTEvidence)`. The set `OTEList[M]` defined in the contract stores all *OTEvidences*.

Suppose a copy of the content is leaked out into the open population. In Figure 1, we assume that informer_X gets the copy and wants to report. The informer should first extract the secret string from the watermarked copy. The reporting is to fulfill an information-hiding procedure to show informer’s acquaintance of the secret string, which is via function `Report(ProofOfLeakage)` of the contract and increases an element of `ReportNumber[M]`. Function `Report()` can be either implemented via commitment scheme or zero knowledge proof. Because `Report()` is an information-hiding procedure, other people in the population cannot learn anything about the watermarked copy reported by informer_X, thus cannot report the same copy unless they actually have it. Nevertheless, the information-hiding submission does not prevent informer_X from creating multiple Sybil identities to submit multiple reports about one watermarked copy. Function `Reward()` of the Argus contract implements an incentive model that causes the total reward obtained through the Sybil attack to be less than what the informer would normally receive. We will explain the incentive model in Section III. Function `Reward()` is periodically invoked according to the ledger’s clock so that informers can get timely rewards.

Once a report is received by the contract, the reported copy will reveal which licensee is accused and update `LicenseeStatus[M]`. The status of the licensee

Table I: Data fields used in the Argus system

Data fields	Description
<code>LicenseeStatus[M]</code>	To store the status of every licensee. Each status belongs to set {NORMAL, ACCUSED, GUILTY, EXONERATED}
<code>OTEList[M]</code>	To store the <i>OTEvidence</i> of every OT between owner and licensee
<code>ReportNumber[M]</code>	To indicate the number of leaked copies from every licensee by counting informers’ reports

is changed from NORMAL to ACCUSED. The licensee has a time period for appeal. If he does not call function `Appeal(OTRecord)` within the period, function `SetGuilty()`, also invoked by the owner according to the ledger’s clock, changes the licensee’s status to GUILTY. If he calls `appeal(OTRecord)` within the period, the `OTRecord` will reveal which watermarked copy the licensee received via OT. If it is the same copy revealed by the report, the licensee’s status is changed to GUILTY. Otherwise, it is changed to EXONERATED. Once exonerated, this licensee is exonerated for this copyrighted content, because the identity of his copy has been disclosed.

There are different ways to implement the Argus contract, but every solution should include three elements: (1) an *incentive model*, (2) an *information-hiding submission scheme*, and (3) an *OT scheme*. Before presenting these elements in the following sections, we introduce the trust assumptions of Argus below.

B. Trust Assumptions

We make three trust assumptions for the design and evaluation of Argus.

- **Robust watermarking.** We assume watermarking cannot be compromised without considerable degradation of the content’s value. Existing work has shown success of watermark robustness in certain domains. For example, the image watermark can protect against attacks such as splitting, sampling, filtering, image compression [24]. Although novel attacks such as oracle attack [25] and collusion attack [26] emerge, watermark techniques keep evolving with new countermeasures [27], [28]. We see this as a separate research concern.
- **Financially motivated informers.** We assume that a financial reward is the only motivation for every informer [9], [17]. This implies that, if an action results in a loss of reward, no informer will take action.
- **Trusted blockchain.** We assume the trustworthiness of the blockchain. There are technologies to enhance security at the ledger layer [29] and the smart contract layer [10], [30]. We also see this as a topic complementary to Argus.

III. INCENTIVE MODEL

As mentioned in the introduction, an incentive-compatible mechanism is required in Argus. It is a challenge because of the open population — the owner does not know the real-life identities of informers. The interests of the owner and the informers are different: the owner wants to receive good-faith reports, and the informers (both honest ones and greedy ones) want to get financial rewards. The goal of the incentive model is to consolidate these different interests. The design of our incentive model is inspired by some models in the literature [9], [17]. However, because we have the target problem clearly formulated (as in the introduction), we are able to deduce a general form of the incentive model. Our

approach is superior for three reasons: (1) it is unclear how the incentive models in the literature were obtained. They seem more of a “creative art” than a result of a disciplined design; (2) our general form encompasses all models that satisfy the incentive objectives of all roles. The existing models in the literature are simplified special cases of our general form; (3) our model ensures extra desirable properties, such as *timely payout* and *guaranteed amount*, which we see as very important aspects of informers’ incentive.

A. The Objectives of Incentive Models

Before introducing our general-form model, it is worth making explicit the objectives of a desirable incentive model.

- First, the model must *disincentivize Sybil attacks*. It is easy to understand that, due to the open population, any informer can create multiple identities (i.e., Sybils), so it is not possible for the owner to detect Sybil attacks. Hence, an objective of the incentive model is to disincentivize them, so that the total reward of all Sybils of a duplicate report is lower than the reward of a single unique report. With this property, an informer (who is assumed to be financially motivated) will not inflate the owner’s counts in `ReportNumber[]`.
- Second, it is the owner’s interest to *incentivize timely reports*, so an informer reporting earlier should be rewarded more than another one reporting later. Essentially, under the incentive model, informers are competitors racing against time. Nothing can be gained by delaying a report.
- Third, it is the informers’ interest to *get timely payouts and guaranteed amounts*. In all existing models, informers need to wait until the end of the campaign to know the amounts and get the rewards. We believe that a good incentive should reward an informer shortly after the report is confirmed valid. The time to reward should be independent of the campaign’s duration.

B. Deducing the Reward Function from the Objectives

We formulate the incentive model as a reward function $B(I_i, n)$, which denotes the bounty value for the i -th successful informer I_i (informers are chronologically ordered) when the total number of informers is n . Previous proposals have given a special form of $B(I_i, n)$ as $c * 2^{-n+1}$ (c is a constant denoting the total bounty value) under certain constraints [9], but they have not formalized the properties of $B(I_i, n)$ in a general form. For example, the work [9] only gives the special form with the assumption that $B(I_i, n) = B(I_j, n) (i \neq j)$, i.e., every informer’s reward is equal.

In this paper, we go through the process of deducing the reward function from the objectives. The mathematical definitions of the properties corresponding to aforementioned objectives are as follows:

- **Sybil-proofness.** Arbitrary subset of informers get a re-

duced total reward if the total number of submissions increases: $\sum_{i \in S_m} B(I_i, m) \geq \sum_{i \in S_k} B(I_i, k)$, where $m \leq k$ are two positive integers and $S_m \subseteq S_k$ are two arbitrary subsets of $\{1, \dots, m\}$ and $\{1, \dots, k\}$, respectively.

- **Order-awareness.** The earlier the informer reports, the more bounty he/she obtains: $B(I_i, n) \geq B(I_{i+1}, n)$.
- **Timely payout.** Each informer is rewarded in an amortized style: $B(I_i, n) = B_1(i) + B_2(n)$, where $B_1(i)$ is only related to i , so the reward can be paid immediately.
- **Guaranteed amount.** An informer, upon a confirmed reporting, is guaranteed a minimal reward amount $c_i > 0$, i.e., his/her total reward will not be under c_i as the number of informers increases: $\lim_{n \rightarrow \infty} B(I_i, n) \geq c_i$.

Note that except for the first property, the others are not achieved by reward function $B(I_i, n) = c * 2^{-n+1}$ from previous work. We ascribe the limitations of previous incentive models to the lack of a general-form deduction, which is addressed in our work. Due to the page limit, we omit the mathematical deduction for the general form and the process of enriching properties (details in the Appendix A of [19]). We only highlight the final expression of $B(I_i, n)$, which can be formalized as Theorem III.1:

Theorem III.1 (The expression of $B(I_i, n)$). To satisfy Sybil-proofness, Order-awareness, Timely Payout and Guaranteed Amount, $B(I_i, n)$ should satisfy the equation:

$$B(I_i, n) = -\xi_i + \sum_{j=i+1}^n \xi_j + c * 2^{-n+1} \quad (1)$$

where $\xi_{i+1} = \sum_{j=1}^i 2^{j-i} * \Delta_j$, $\sum_{j=1}^n 2^j * \Delta_j \leq c$, $\{\Delta_i\} \in \mathbb{R}_{\geq 0}^*$

C. Plugging in real-world numbers

In the introduction, we describe the \$1-million campaign opaquely run by the BSA [7]. After a long period of time, only a small percentage of the total amount was paid out to informers. This is obviously a low incentive to the public. We now analyze the outcome if Argus is used for the \$1-million campaign. We instantiate the parameters $\Delta_i = 2^{-i} * c/l$ (for $i = 1, \dots, l$) or 0 (for $i > l$) in Theorem III.1 where l can be an arbitrary positive integer. For simplicity, we set $l = 20$ in this work, which is a typical boundary to classify copyright infringement [1]. With this setting, we compare Argus with previous work [9], [17], which are also Sybil-proof models.

Figure 2 shows the reward amount (in the log scale) that each of the n informers will get in our incentive model (the upper diagram) and the previous model (the lower diagram). Every line in the lower diagram is horizontal, because the previous model does not have order-awareness. As a result, all informers get the same reward, and the reward is exponentially decreased with n , in order to ensure Sybil-proofness. There is no guaranteed amount when a report is confirmed. When the campaign ends, even an early informer

may find the reward almost zero if there are many later informers. It is a problematic model to incentivize people.

The upper diagram shows our model. The line of $n = \infty$ corresponds to the guaranteed amounts for the informers. Our lines are also affected by n , but not as drastically as in the previous work. As n increases, the lines become closer to the $n = \infty$ line (note that the $n=100$ line is visually overlapped with it), suggesting that every existing informer loses some reward when a new informer joins. The design of our reward function ensures that the loss of reward is big enough so that no informer wants to fake a Sybil identity to get another reward.

Compared to the (problematic) incentive shown in the lower diagram of Figure 2, the reward function of Argus is superior in all objectives we set at the beginning of this section.

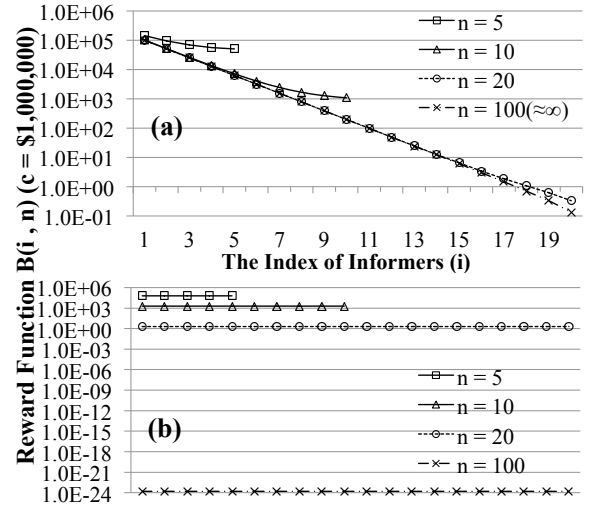


Figure 2: The reward of first twenty informers in (a) our incentive model and (b) previous incentive model [9], [17], respectively (when total number of informers $n = 5, 10, 20, 100$).

IV. INFORMATION HIDING FOR REPORT SUBMISSION

As mentioned in the introduction, full transparency is a major advancement of Argus. To achieve it, our design requires an effective strategy for information hiding. Specifically for the procedure of report submission, Argus needs to ensure that, although everybody in the open population can see the interactions between an informer and the Argus contract, nobody other than the informer can replay the interactions effectively. (Note that the informer himself will not replay due to the Sybil-proof property of the incentive model discussed in Section III.)

A. The objectives of the information-hiding submission

Before designing the information-hiding submission procedure, we first consider the interests of different roles, and specify the objectives to be achieved:

- First, the owner’s interest is to get an accurate count of the number of piracy copies. The count should not be inflated, meaning that no replay attack as described above should be possible.
- Second, the owner wants to run a bounty campaign for a long duration, e.g., 3-6 months, so that the magnitude of the piracy infringement can be evaluated more thoroughly. However, an informer’s interest is to get a timely payout after a report is submitted successfully.
- Third, both the owner and the informers want the submission to be concluded in a short period of time. Also, the submission should be efficient. Since Argus runs on a public blockchain, the gas consumption of the procedure is required to be very low.

B. Previous proposals and their limitations

Previous work have proposed commitment scheme [10], [23] and zero-knowledge proof [9], [31] to eliminate the replay attack. However, in Argus’s scenarios, they may either conflict with the interests of owner/informer or encounter serious performance problems:

- In a traditional commitment scheme, the informer submits a commitment of the report (e.g., Hash) in the first period (“commitment phase”) and reveals the report during a following period (“reveal phase”). Informers’ reports are not accepted in the second phase unless a corresponding commitment exists in the first phase. Therefore, the replay attack cannot succeed since the attacker cannot generate valid commitments in the first phase. However, setting the length of the commitment phase has a dilemma because of the different interests of the owner and the informers (definitions in Section IV-A): (1) if the phase is too long, it is not timely for the owner to validate infringers and take subsequent actions, and the informer needs to wait for a long time to get the bounty; (2) if it is too short, the number of the reported pirated copies cannot accurately reflect the severity of the infringement.
- In a Zero-Knowledge Proof (ZKP) scheme, the informer can generate a proof to show the possession of a pirated copy. In the ZKP scheme, the bounty can be paid as soon as the contract verifies the proof, which is what we desire. However, the ZKP scheme has a prohibitively high performance overhead and gas cost (see Section VII).

C. Multi-period Commitment Scheme

To achieve the objectives with good performance, we propose a novel technique called multi-period commitment. The scheme can be considered as an extension of the traditional commitment scheme, but has a good performance. Meanwhile, it has the advantage similar to the zero-knowledge proof — the owner can still set a long bounty campaign period, but confirm every report almost in real time.

Our scheme allows multiple commit-and-reveal phases so that there are sufficient time windows for the informers to

submit piracy reports. For a desired length of collection period T , the owner can divide T into K **sub-periods** $\{T_1, \dots, T_K\}$. Each sub-period T_i ($1 < i < K$) is the i -th commitment phase and also the $(i-1)$ -th reveal phase. In other words, an informer can claim a bounty in T_{i+1} (i -th reveal phase) by revealing if the corresponding commitment is submitted in T_i (i -th commitment phase).

However, dividing into periods introduces a problem: informers can replay the process of commit-reveal in sub-periods T_i and T_{i+1} to later sub-periods T_j and T_{j+1} ($j > i$). To defend again this kind of replay attack, we introduce a time stamp into the formula of commitment and process of verification: if we denote the piracy report as X and hash function as \mathcal{H} , and there is a predefined list $L[\cdot] = \{\mathcal{H}(\mathcal{H}(X||1)), \dots, \mathcal{H}(\mathcal{H}(X||K))\}$ in the contract. Then, the commitment cm submitted in T_i can be $\mathcal{H}(\mathcal{H}(X||i)||n)$ where “||” denotes concatenation and “ n ” denotes a randomized nonce. In corresponding reveal phase T_{i+1} , $rv = \mathcal{H}(X||i)$ and n should be submitted to the contract for verification that $\mathcal{H}(rv||n) = cm$ and $\mathcal{H}(rv) = L[i]$. By this reinforcement, the aforementioned replay attack cannot pass the verification “ $\mathcal{H}(rv) = L[j]$ ” in later sub-periods T_{j+1} ($j > i$).

The multi-period commitment scheme achieves the objectives given in Section IV-A:

- The multi-period commitment scheme foils the replay attack, thus meets the first objective.
- The multi-period commitment scheme supports an arbitrary length of collection period T , thus meets the second objective.
- With a sufficiently large K , each sub-period can be short enough. Thus, informers can reveal reports and get their rewards within a short interval after commitments. The owner gets quick confirmations of the infringers. This achieves the third objective.

V. GUARDING AGAINST INFRINGER’S REPUDIATION

No role in Argus, even the owner, is assumed trusted. This presents a challenge: when the owner accuses a licensee of leaking a copy, the licensee can refute the accusation by arguing that the copy could have been leaked out by the owner himself. To resolve the dispute, Argus must make the evidence of the accusation so convincing that *the probability of the accused infringer being an innocent licensee is extremely small*. Hence, a true infringer’s attempt to repudiate will be unsubstantiated.

To approach the objective, we use a 1-out-of- N Oblivious Transfer (OT) [21] protocol to achieve this goal. The 1-out-of- N OT protocols were used for data sharing [32], [33]: the owner generates N different copies of data (e.g. via watermarking) and plays OT protocol with the licensee. Then, the licensee can obtain only one copy without owner’s knowing which one. Thus, owner can infer the chosen version with a pirated copy. When there is a dispute between

a licensee and the owner, they can submit messages occurred in OT to a *trusted judge* for resolving disputes. There is only a $\frac{1}{N}$ probability that the successfully accused licensee is innocent.

A goal of Argus is not to have any trusted role, so the Argus contract has to implement the functionalities of the “judge” on a public blockchain. However, this may introduce a big bandwidth overhead: messages incurred in OT are proportional to $O(N)$. To achieve a desirable security level with a large N (e.g. 10,000), existing solutions introduce enormous on-chain overhead (e.g. bandwidth, execution, storage).

To greatly reduce the overhead, we introduce $O(1)$ -Appeal which only incurs $O(1)$ on-chain messages and operations. $O(1)$ -Appeal has two properties:

- **Obliviousness.** It is the property of 1-out-of- N OT [21]:
 - (1) the licensee can arbitrarily choose and obtain 1 data from N candidate data but cannot know the unchosen data;
 - (2) the owner does not know which data are chosen by the licensee. This property guarantees that the probability to successfully incriminate an innocent licensee is $\frac{1}{N}$ and thus an infringer is hard to deny accusation with a large N (e.g. 10000).
- **Non-repudiation.** When the licensee is accused, the licensee can appeal by showing committed records. When there is a dispute, neither the owner nor licensees can deny which copy the licensee had chosen in the previous OT protocol. The contract is able to give a conclusive answer.

A. Constant-Size-OTRecord Appeal ($O(1)$ -Appeal)

The protocol of $O(1)$ -Appeal is shown in Figure 3, which includes four sub-protocols: Initialize, GenerateEvidence, TransferData, Appeal. The first three sub-protocols are very similar¹ to those in [21], while the fourth sub-protocol is our new invention that incurs only an $O(1)$ on-chain cost. Unlike in Section II, the owner does not have to commit OTEvidence (i.e. R) to the contract. Instead, R can be signed and kept locally. For simplicity, we assume that the owner and the licensee do not abort during the procedure (e.g. this can be ensured by using the state channel technology [34] or the fair exchange protocol [35]).

Different from previous work that utilizes transferred vectors (e.g. $\{E_1, \dots, E_N\}$) in TransferData for dispute resolving, we find that utilizing one transferred variable R in GenerateEvidence can be of the same effect. Our key discovery is that R has a one-to-one correspondence to the licensee’s chosen index l (see Theorem A.6 in [19]). Therefore, R can be used as the evidence to indicate the licensee’s chosen index in the dispute-resolving stage (i.e. the appeal stage).

As shown in Appeal in Figure 3, if R is signed by owner/licensee and indicates that the corresponding index

¹The only difference is that we add a step (step 2) in GenerateEvidence.

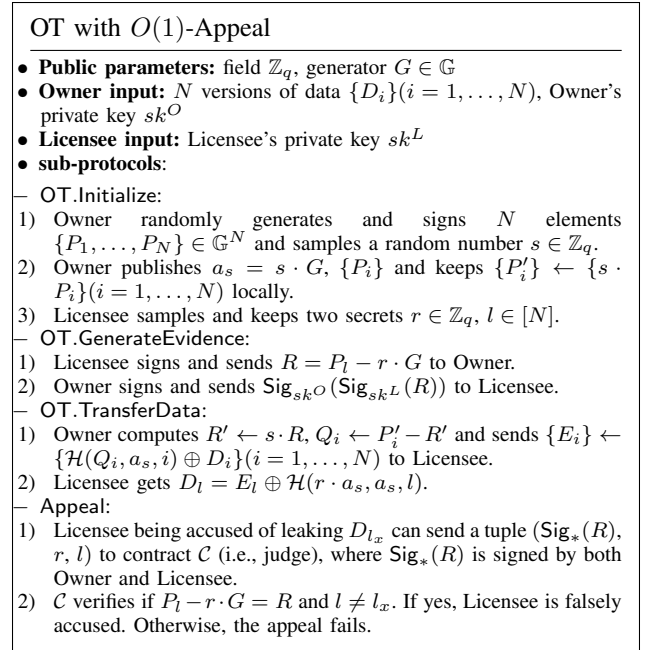


Figure 3: Licensee gets one data D_l from N data $\{D_1, \dots, D_N\}$ from Owner via $O(1)$ -Appeal with $OTEvidence = R$ and $OTRecord = (r, l)$.

l differs from accused index l_x , we can conclude that the licensee is wrongly accused. We prove the Obliviousness and Non-repudiation of OT with $O(1)$ -Appeal in Appendix B of [19].

Though $O(1)$ -Appeal has greatly reduced the on-chain overhead of the appeal stage, there is still a considerable cost of the off-chain bandwidth in TransferData. To reduce the off-chain bandwidth overhead, in Section VI, we will further leverage a PIR protocol [36] and slightly adapt $O(1)$ -Appeal, which guarantees that the size of the data transferred is about the size of data D_l . In addition, we will show how to integrate $O(1)$ -Appeal with the information-hiding report scheme in Appendix C of [19].

VI. IMPLEMENTING THE ARGUS SYSTEM

The previous three sections explain the main objectives and the core ideas of Argus. It is important to recognize that the objectives are not separate problems to solve individually. The Argus contract needs to achieve the objectives altogether in a coherent design. Due to the page limit, we only provide a sketch of our construction and implementation here. The holistic view and the implementation details of the Argus system are provided in Appendix C and Appendix D of [19], respectively.

With corresponding watermark algorithms, current Argus system supports three data types: image [24], audio [37] and software [38]. A Merkle tree structure is leveraged to reduce the on-chain storage: for any list of data, only the Merkle root of the list is uploaded to the blockchain. We also leverage Private Information Retrieval (PIR) [36] to

Table II: Interests/threats of participants in Argus

Participants	Interest if honest	Threat if malicious
Owner	To discover infringers and tally the number of copies	To falsely accuse innocent licensees
Licensee	To win in an appeal because of innocence	To appeal despite the guilt
Informer	To submit a honest report once	To submit a fake report, steal a report or submit a valid report multiple times.

reduce the bandwidth overhead of downloading data for the licensees.

VII. SECURITY ANALYSIS AND PERFORMANCE EVALUATION

In this section, we first analyze the security of Argus then describe the experimental setup for the performance evaluation. The evaluation results include the performance measurements and the cost of Argus transactions.

A. Security Analysis

The detailed security analysis is given in Appendix E of [19]. Without loss of generality, we only present an analysis which considers a game of five participants $\text{Game}_{O,L_1,L_2,I_1,I_2}^{\text{Argus}}$, where O, L_1, L_2, I_1, I_2 , representing the Owner, one Licensee, another Licensee, one Informer and another Informer, respectively. Their interests if honest and their threats if malicious are summarized in Table II. Based on $\text{Game}_{O,L_1,L_2,I_1,I_2}^{\text{Argus}}$, we can easily extend our security analysis to scenarios with multiple owners, informers and licensees. Since L_1 and I_1 are identical to L_2 and I_2 respectively, to demonstrate the security of Argus, we enumerate all cases that O, L_1, I_1 are individually honest (i.e. following the protocol) while other four participants may collude. For each case, we find that the interest of the honest participant will not be affected. In other words, we conclude that *if a participant (i.e. owner, licensee or informer) has no fault, the interest of this participant will not be hurt even when others collude.*

B. Performance Evaluation

Experimental Setup. Our testbed consists of relatively low-end Azure Virtual Machines (D2s_v3, 2 vCPUs, 8GB RAM, Linux) for the nodes of owner, licensees, informers and blockchain nodes. For blockchain nodes, we adopt the default PoW algorithm (i.e., Ethash [39]) and parameters (block interval, block Gaslimit, etc.) of current Ethereum (date: 2021-04-05) to simulate the public blockchain. The average block interval of Ethereum is set as 12 seconds. The bandwidth of uploading and downloading is tested as around 50 MB/s. To guarantee a sufficient OT security and a short confirmation of reports, we by default set the OT versions (i.e., N) in Section V as 10,000 while the number of periods (i.e., K) in Section IV as 1000. In other words, the probability of accusing an innocent licensee is $\frac{1}{10000}$ while the report confirmation time during a 180-day period is 4.3

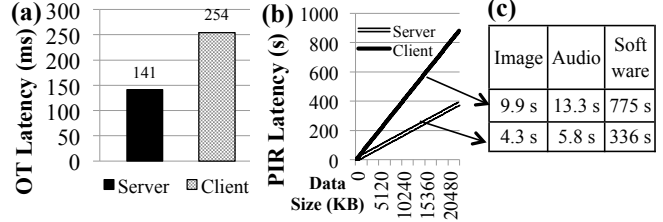


Figure 4: The performance of OT and PIR: (a) OT Latency (b) PIR Latency (c) PIR Latency for different data

hours.

Evaluation Results. We evaluate Argus system from various perspectives of practicality, such as the throughput of system, client latency, gas cost on Ethereum, etc. All protocols of Argus are tested end-to-end. We also give the comparison between Argus and previous work from aforementioned perspectives. Due to the page limit, we only present evaluation results of Initiate, ShareData, ReportPiracy and Appeal, which are the most resource-consuming protocols. In other words, these four protocols can introduce considerable overhead in throughput, latency, storage and cost to impede the Argus’s adoption in practice. For an intuitive understanding of the gas cost in Ethereum, we represent the gas cost in the number of sending simplest transactions².

We first evaluate Initiate, which is the setup phase of Argus. The main element of Initiate is to generate a Merkle tree and to deploy the contract. For every licensee, there is a time complexity of $O(N * K)$ for the owner to generate the Merkle tree, which needs about 10 minutes. This process is totally offline and can be parallelized and accelerated with high-end machines. Deploying the contract costs about 5.2×10^6 gas, which equals to the cost of sending ~ 248 simplest Ethereum transactions. In addition, we also evaluate the off-chain storage cost for Argus, which is about 960 KB per licensee, which has a space complexity of $O(N)$.

We evaluate the latency/bandwidth in data sharing. ShareData includes two phases, OT (Section V) and PIR (Section VI). While the PIR phase can be done offline, the OT phase can directly affect the throughput of the Argus system especially when there are a number of clients (licensees) concurrently communicating with the server (owner):

- For the OT phase, the evaluation result is shown in Figure 4 (a). It takes about 141 ms for the server to complete the phase. Therefore, a throughput of 7.1 OT requests per second per machine can be served. With a stronger server (Azure D32s_v3, 32 cores, 128GB RAM), the throughput can be further improved to 82.6 per second. Note that, the throughput can be linearly scaled up by increasing the number of machines.
- For the PIR phase, we list the latency of client/server in Figure 4 (b) and Figure 4 (c). As shown in Figure 4 (b),

²Ethereum’s simplest transaction only transfers ether and costs 21,000 gas.

Table III: High-level comparison between state-of-the-art work and Argus

Desired properties	Details	BSA [2]	Custos [5]	AWM [33]	ZKP [15]	Hydra [10]	Arbitrum [9]	This work
Trusted payments	Full transparency	×*	✓	N/A**	✓	✓	✓	✓
Better payments	Timely/guaranteed payout	×	✓	N/A	✓	×	×	✓
Identifying infringers	Strong accusation	✓	×	✓	N/A	N/A	N/A	✓
Assessing severity	Sybil-proofness	✓	×	N/A	N/A	N/A	✓	✓
	Information-hiding	✓	×	N/A	✓	✓	✓	✓
Scalability	High throughput	✓	×	×	×	✓	✓	✓

* Symbols of “✓” and “×” denote corresponding property is “achievable” and “hard to achieve”, respectively.

** Properties are marked as “not applicable (N/A)” if corresponding work is not designed for these properties.

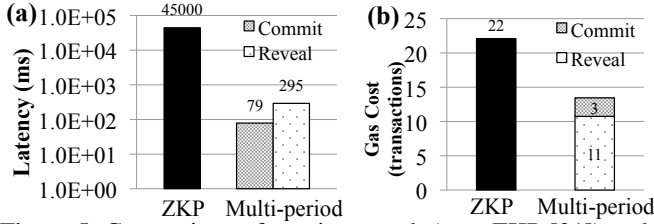


Figure 5: Comparison of previous work (e.g., ZKP [31]) and multi-period commitment scheme: (a) Latency (b) Gas Cost

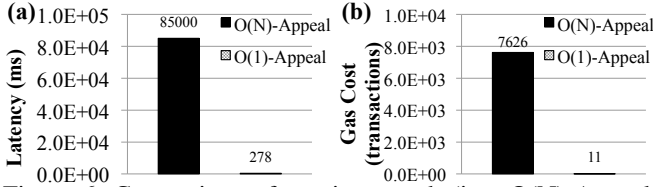


Figure 6: Comparison of previous work (i.e., $O(N)$ -Appeal [32], [33]) and $O(1)$ -Appeal: (a) Latency (b) Gas Cost

the PIR latency is proportional to the data size, which corresponds to a bandwidth of 206 Kbps and 476 Kbps for the client and server, respectively. And in Figure 4 (b), we have the PIR latency for data types in Section VI, which is shown in Figure 4 (c). Without PIR, the direct downloading time of the 10,000 copies for the licensee would be 51 s, 67 s and 3800 s, respectively. Similar to the OT phase, the latency and bandwidth of PIR phase can also be linearly and considerably improved with more and higher-performance (i.e., faster core and larger RAM) machines. Since PIR can be offline after the OT phase, the bandwidth of Argus system is determined by the OT phase.

ReportPiracy includes the commit phase and the reveal phase (Section IV). The evaluation results of ReportPiracy are shown in Figure 5. The latency in Figure 5 (a) denotes the time which is spent by Informer’s machine to submit a transaction to Argus contract until the transaction is executed by blockchain nodes. In other words, the consensus time is excluded in the latency result. The latencies of commit and reveal are negligible compared to the block time of Ethereum, respectively.

The on-chain cost of our multi-period scheme is also negligible. As shown in Figure 5 (b), the gas costs for commit and reveal are about 8×10^4 and 2×10^5 , which equal to ~ 3 and ~ 11 simplest Ethereum transactions, respectively. In other words, a total cost equivalent to sending 14 simplest

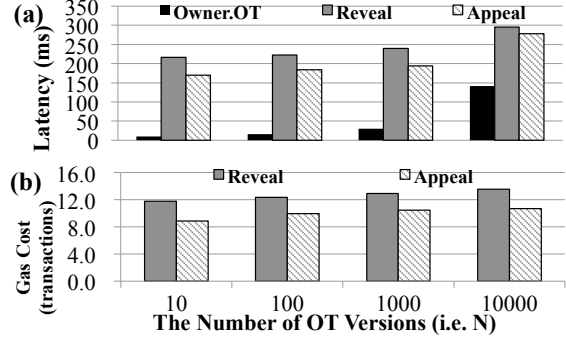


Figure 7: Sensitivity analysis of #OT versions: (a) Latency (b) Gas Cost

transactions is required for an informer to report piracy in our system. From the gas consumption, given that the maximum gas limit of every Ethereum block is around 12,000,000, we can conclude that the Transactions Per Second (tps) for ReportPiracy is about $12 \times 10^6 / (2.0 \times 10^5 \times 12) \approx 5.0$, which is 11% of the theoretically maximum Ethereum throughput³ (commit and reveal occur in different blocks and thus only the more expensive reveal is considered). We also compare our multi-period scheme with ZKP scheme [16], [31]. Compared to ZKP scheme, our scheme can reduce the informer client latency by 99.3% and the gas cost by 39%.

As in Figure 6, we also evaluate Appeal (Section V) and get similar results as ReportPiracy: a tps of 4.3 is achieved given the gas consumption of Appeal ($\sim 2 \times 10^5$, i.e., the cost of ~ 11 simplest Ethereum transactions). The appeal protocol proposed by previous work [32], [33] (denoted as “ $O(N)$ -Appeal”) introduces unacceptable on-chain operations and exceeds the maximum gas limit of an Ethereum block. Thus, we cannot evaluate corresponding latency and gas consumption in an end-to-end style. Instead, we estimate them using the number of on-chain operations. Results show that $O(1)$ -appeal can significantly reduce the client latency and the gas cost: compared to previous appeal scheme, the informer client latency can be reduced by 99.7% and the gas cost by a factor of 960X, respectively.

As shown in Figure 7, we also investigate the impact of choosing different numbers of OT versions (i.e., N), from 10 to 10,000. On the one hand, as addressed in Section V, the value of $\frac{1}{N}$ determines the probability ϕ of false accusation;

³The maximum tps of Ethereum is about 47.6 when the block only contains the simplest Ethereum transactions. By contrast, the average tps of Ethereum is 15.0 currently (2021-04-05).

On the other hand, the increase of N has a negative impact on the performance, gas cost and storage overhead of Argus system. We can see that the increase of OT affects the OT latency of owner significantly (Figure 7 (a)) while the gas cost increases logarithmically (Figure 7 (b)) since we use Merkle tree structure in our design ⁴.

To summarize the performance of Argus system: in the normal case, which does not involve piracy-reporting or appeal, only ShareData is involved. In this case, the throughput of Argus system is equal to the throughput of ShareData (82.6 off-chain transactions per second per machine). In the uncommon case of piracy reporting, the throughput of the reporting transactions is 5.0 tps. In the rare case in which the appeal procedure is performed (i.e., the owner is malicious), the throughput of the appeal transactions is 4.3 tps.

VIII. RELATED WORK

We summarize the comparison of Argus with previous work in Table III. The third and the fourth Columns (i.e., BSA and Custos) are two competitive solutions of Argus while the fifth to the eighth (AWM, ZKP, Hydra and Arbitrum) are primitives corresponding to $O(1)$ -appeal, multi-period commitment and incentive model, respectively. Details are listed below:

- **Comparison with previous solutions.** Centralized schemes such as BSA [2] can considerably disincentivize informers due to the opacity of payments. To increase the trust of payments, Custos [5] leverages blockchain. However, Custos does not consider strong accusations and cannot assess the severity of piracy, which is important in law enforcement. By contrast, Argus achieves full transparency along with all other properties.
- **Comparison with state-of-the-art primitives:**
 - Previous work of Asymmetric Watermarking (AWM) [32], [33] all rely on the existence of a “trusted judge”. In addition, their appeal protocols introduce $O(N)$ bandwidth cost, which is unacceptable in blockchain scheme.
 - Current Zero-Knowledge Proof (ZKP) [15], [16] can eliminate replay attack by enabling informers to prove their acquaintance of $id_{i,j}$ without revealing the answer. However, current ZKP can introduce considerable performance overhead and gas consumption which limits system scalability.
 - Commitment scheme (or so-called “commit-reveal”) [10] can address replay attack by dividing the single submission into phases of commitment and revealing, which is much more efficient than ZKP schemes. However, to fully evaluate piracy, the commitment phase should be relatively long, which delays owner’s confirmation of piracy and informers’ payments.
 - Sybil-proofness [9], [17] is introduced to disincentivize

⁴We omit the commit operation of informer in Figure 7 since the latency and gas cost of commit is unrelated to the value of N .

informers to report repeatedly: the more times an informer reports, the less bounty the informer can claim. However, existing sybil-proof incentive models merely depend on the total number of informers, which is known only when the collection period ends.

- By contrast, Argus overcomes the limitations of the above primitives. In addition, further integration and optimization are introduced in this work.

IX. CONCLUSIONS

Anti-piracy is fundamentally a procedure that relies on collecting data from the open anonymous population, so how to incentivize credible reports is a question at the center of the problem. Academic researchers and real-world companies have come up with various incentive mechanisms. However, without explicitly prescribing the interests of different roles and the objectives of an anti-piracy system, designing such a mechanism has been more of a “creative art” than a systematic and disciplined exploration. Currently, there is no good framework to evaluate these designs and actual systems.

The most essential value of our work is not the Argus system itself, but the approach leading to its design and implementation. We first state clearly the interests of different roles and the goal of full transparency without trusting any role. Once these are stated, all the design requirements naturally surface, such as Sybil-proofness, information-hiding submission, resistance to infringer’s repudiation, etc; once these design requirements are clear, we are able to *deduce*, rather than *invent*, the general form of valid solutions; the deduced general form then boils down to a set of unavoidable technical obstacles, which we overcome by adapting cryptographic schemes, building contract code and optimizing performance.

Argus exemplifies the outcome of this disciplined approach. It is superior to existing solutions in terms of the trust assumption and the assured properties. In particular, we draw the following conclusions: (1) it is feasible to build a fully transparent solution without introducing a trusted role. This could enable a paradigm shift for anti-piracy incentive solutions. Also, it is a compelling application scenario for public blockchains; (2) such a solution indeed consolidates all roles’ interests fairly, i.e., as long as a role is not at fault, his/her interest will not be impaired by other malicious or at-fault roles; (3) besides logic soundness, the solution is economically practical, as a result of our effective optimizations.

ACKNOWLEDGMENT

The authors would like to thank Lidong Zhou and anonymous reviewers for their valuable comments. Wenyan Liu was supported by NSFC grants (No. 61972155), the Science and Technology Commission of Shanghai Municipality (20DZ1100300) and the Open Project Fund from Shenzhen Institute of Artificial Intelligence and Robotics for Society.

REFERENCES

- [1] “Copyright law of the united states,” <https://www.copyright.gov/title17/title17.pdf>.
- [2] “Piracy bounty from business software alliance (bsa),” <https://reporting.bsa.org/>.
- [3] “The federation against copyright theft,” <https://www.fact-uk.org.uk/>.
- [4] “Siia’s corporate anti-piracy reward program,” https://www.siia.net/piracy/report/siia_reward_program.pdf.
- [5] “Custos,” <https://custostech.com/>.
- [6] “Veredictum,” <https://www.veredictum.io/>.
- [7] “The public doubt about bsa bounty,” <https://www.pcworld.com/article/147448/article.html>.
- [8] C. Li, B. Palanisamy, R. Xu, J. Wang, and J. Liu, “Nf-crowd: Nearly-free blockchain-based crowdsourcing,” in *2020 International Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 2020, pp. 41–50.
- [9] H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, and E. W. Felten, “Arbitrum: Scalable, private smart contracts,” in *Proceedings of the 27th USENIX Conference on Security Symposium*. USENIX Association, 2018, pp. 1353–1370.
- [10] L. Breidenbach, I. Cornell Tech, P. Daian, F. Tramèr, and A. Juels, “Enter the hydra: Towards principled bug bounties and exploit-resistant smart contracts,” in *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, 2018.
- [11] H. Sukhwani, J. M. Martínez, X. Chang, K. S. Trivedi, and A. Rindos, “Performance modeling of pbft consensus process for permissioned blockchain network (hyperledger fabric),” in *2017 IEEE 36th Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 2017, pp. 253–255.
- [12] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, “Algorand: Scaling byzantine agreements for cryptocurrencies,” in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 51–68.
- [13] C. Li and B. Palanisamy, “Decentralized release of self-emerging data using smart contracts,” in *2018 IEEE 37th Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 2018, pp. 213–220.
- [14] G. Bu, T. S. L. Nguyen, M. P. Butucaru, and K. L. Thai, “Hyperpub-sub: Blockchain based publish/subscribe,” in *2019 38th Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 2019, pp. 366–3662.
- [15] E. B. Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, “Zerocash: Decentralized anonymous payments from bitcoin,” in *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 459–474.
- [16] “Counteracting front-running with zero-knowledge proof,” <https://medium.com/@schor/on-zero-knowledge-proofs-in-blockchains\ -14c48cfd1dd1>.
- [17] J. Koch and C. Reitwiessner, “A predictable incentive mechanism for truebit,” *arXiv preprint arXiv:1806.11476*, 2018.
- [18] “Pallone: Google, facebook, twitter content treatment not ‘neutral’,” <https://www.nexttv.com/news/pallone-google-facebook-twitter-content-treatment-not-neutral-169576>.
- [19] X. Zhang, X. Guo, Z. Zeng, W. Liu, Z. Guo, Y. Chen, S. Chen, Q. Yin, M. Yang, and L. Zhou, “Argus: A fully transparent incentive system for anti-piracy campaigns (extended version),” 2021, <https://arxiv.org/abs/2107.06049>.
- [20] “The solidity contract-oriented programming language,” <https://github.com/ethereum/solidity>.
- [21] M. Naor and B. Pinkas, “Efficient oblivious transfer protocols,” in *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2001, pp. 448–457.
- [22] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza, “Succinct non-interactive zero knowledge for a von neumann architecture,” in *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, 2014, pp. 781–796.
- [23] G. Brassard, D. Chaum, and C. Crépeau, “Minimum disclosure proofs of knowledge,” *Journal of computer and system sciences*, vol. 37, no. 2, pp. 156–189, 1988.
- [24] I. J. Cox, J. Kilian, F. T. Leighton, and T. Shanon, “Secure spread spectrum watermarking for multimedia,” *IEEE transactions on image processing*, vol. 6, no. 12, pp. 1673–1687, 1997.
- [25] J.-P. M. Linnartz and M. Van Dijk, “Analysis of the sensitivity attack against electronic watermarks in images,” in *International Workshop on Information Hiding*. Springer, 1998, pp. 258–272.
- [26] D. Boneh and J. Shaw, “Collusion-secure fingerprinting for digital data,” *IEEE Transactions on Information Theory*, vol. 44, no. 5, pp. 1897–1905, 1998.
- [27] T. Y. Kim, H. Choi, K. Lee, and T. Kim, “An asymmetric watermarking system with many embedding watermarks corresponding to one detection watermark,” *IEEE signal processing letters*, vol. 11, no. 3, pp. 375–377, 2004.
- [28] G. Tardos, “Optimal probabilistic fingerprint codes,” *Journal of the ACM (JACM)*, vol. 55, no. 2, p. 10, 2008.
- [29] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, “On the security and performance of proof of work blockchains,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 3–16.
- [30] Y. Li, H. Liu, Z. Yang, B. Wang, Q. Ren, L. Wang, and B. Chen, “Protect your smart contract against unfair payment,” in *2020 International Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 2020, pp. 61–70.
- [31] R. Khalil, A. Gervais, and G. Felley, “Tex-a securely scalable trustless exchange,” *IACR Cryptology ePrint Archive*, vol. 2019, p. 265, 2019.
- [32] A. Kiayias, N. Leonardos, H. Lipmaa, K. Pavlyk, and Q. Tang, “Communication optimal tardos-based asymmetric fingerprinting,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2015, pp. 469–486.
- [33] L. Xu, F. Zhang, W. Susilo, and Y. Wen, “Solutions to the anti-piracy problem in oblivious transfer,” *Journal of Computer and System Sciences*, vol. 82, no. 3, pp. 466–476, 2016.
- [34] J. Poon and V. Buterin, “Plasma: Scalable autonomous smart contracts,” *White paper*, pp. 1–47, 2017.
- [35] S. Dziembowski, L. Eeckey, and S. Faust, “Fairswap: How to fairly exchange digital goods,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2018, pp. 967–984.
- [36] C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian, “Xpir: Private information retrieval for everyone,” *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 2, pp. 155–174, 2016.
- [37] D. Kirovski and H. S. Malvar, “Spread-spectrum watermarking of audio signals,” *IEEE transactions on signal processing*, vol. 51, no. 4, pp. 1020–1033, 2003.
- [38] C. Collberg, E. Carter, S. Debray, A. Huntwork, J. Kececiloglu, C. Linn, and M. Stepp, “Dynamic path-based software watermarking,” in *ACM Sigplan Notices*, vol. 39, no. 6. ACM, 2004, pp. 107–118.
- [39] “Ethereum,” <https://www.ethereum.org/>.