

# ZebraLancer: Private and Anonymous Crowdsourcing System atop Open Blockchain

Yuan Lu, Qiang Tang, Guiling Wang  
 Department of Computer Science, New Jersey Institute of Technology  
 Email: {yl768, qiang, gwang}@njit.edu

**Abstract**—We design and implement the first *private and anonymous* decentralized crowdsourcing system ZebraLancer<sup>1</sup>, and overcome two fundamental challenges of decentralizing crowdsourcing, i.e. data leakage and identity breach.

First, our *outsource-then-prove* methodology resolves the tension between blockchain transparency and data confidentiality, which is critical in crowdsourcing use-case. ZebraLancer ensures: (i) a requester will not pay more than what data deserve, according to a policy announced when her task is published via the blockchain; (ii) each worker indeed gets a payment based on the policy, if he submits data to the blockchain; (iii) the above properties are realized not only without a central arbiter, but also without leaking the data to the open blockchain.

Furthermore, the transparency of blockchain allows one to infer private information about workers and requesters through their participation history. On the other hand, allowing anonymity will enable a malicious worker to submit multiple times to reap rewards. ZebraLancer overcomes this problem by allowing anonymous requests/submissions without sacrificing the accountability. The idea behind is a subtle linkability: if a worker submits twice to a task, anyone can link the submissions, or else he stays anonymous and unlinkable across tasks. To realize this delicate linkability, we put forward a novel cryptographic concept, i.e. the *common-prefix-linkable* anonymous authentication. We remark the new anonymous authentication scheme might be of independent interest.

Finally, we implement our protocol for a common image annotation task and deploy it in a test net of Ethereum. The experiment results show the applicability of our protocol atop the existing real-world blockchain.

## I. INTRODUCTION

Crowdsourcing empowers open collaboration over the Internet. One remarkable example is the solicitation of annotated data: the famous ImageNet [1] was created in Amazon's crowdsourcing marketplace, Mechanical Turk (MTurk) [2]. Another notable example is mobile crowdsensing [3] where one (called "requester") can request a group of individuals (called "workers") to use their mobile devices to gather information fostering a data-driven application [4]. Various monetary incentive mechanisms were introduced [5–11] to motivate workers to make real efforts. To facilitate these mechanisms, the state-of-the-art solution necessarily requires a trusted third-party to host crowdsourcing tasks to fulfill the fair exchange between the crowd-shared data and the rewards;

<sup>1</sup>Two popular hypotheses of zebra strips were: (i) camouflage used to confuse predators by motion dazzle, and (ii) visual cues used by herd peers to identify. The delicate anonymity in our system can be analog, as it overcomes the natural tension between anonymity and accountability. On the other hand, freelancer is a typical position enabled by crowdsourcing.

otherwise, the effectiveness of incentive mechanisms can be hindered by the so-called "free-riders" (i.e. dishonest workers reap rewards without making real efforts) and "false-reporters" (i.e. dishonest requesters try to repudiate the payment).

It is well-known that reducing the reliance on a trusted third-party is desirable in practice, and the same goes for the use-case of crowdsourcing. First, numerous real-world incidents reveal that the party might silently misbehave in-house for self-interests [12]; or, some of its employees [13] or attackers [14] can compromise its functionality. Second, the party often fails to resolve disputes. For instance, requesters have a good chance to collect data without paying when using MTurk, which is biased towards requesters over workers [15]. Third, a centralized platform inevitably inherits all the vulnerabilities of the single point failure. For example, Waze, a crowdsourcing map application, suffered from 3 unexpected server downs and 11 scheduled service outages during 2010 to 2013 [4]. Last but not least, a centralized platform hosting all tasks also increases the worry of massive privacy breach. One fresh lesson to us is the tremendous private data leakage of the leader in crowdsourcing economy, Uber [16].

In contrast, an (open) blockchain is a distributed, transparent and immutable public "bulletin board" organized as a chain of blocks. The blockchain is usually managed and replicated by a peer-to-peer (P2P) network collectively. Each block includes some messages committed by network peers, and is validated by the whole network according to a pre-defined consensus protocol. This ensures reliable delivery of messages via the untrusted Internet. More interestingly, the messages contained in each block can be program code, the execution of which is enforced and verified by all blockchain network peers; hence, a more exotic application of smart contract [17] is enabled. Essentially, the smart contract can be viewed as a "decentralized computer" that faithfully handles all computations and message deliveries (except the adversary can choose the order of messaging) related to a specified task. It becomes enticing to build a *decentralized* crowdsourcing platform atop.

Unfortunately, this new fascinating technology also brings about new challenges, which were never that severe in the centralized setting before [18]. One notable feature of the blockchain is its *transparency*. The whole chain is replicated by the whole network to ensure consistency, thus the data submitted to the blockchain will be visible to the public. This causes an immediate problem violating data privacy as many

of the crowdsourced data maybe sensitive. Sometimes, the data are simply valuable to the requester who paid to get them. What is worse, since the block confirmation (which corresponds to the time when the submitted answers are actually recorded in a block) normally takes some time after the data is submitted to the network, a malicious worker can simply copy the data committed by others, and submit the same data as his own to run the free-riding attack. Without *data confidentiality*, the incentive mechanisms could be rendered completely ineffective.

Furthermore, most crowdsourcing systems [2, 4] and incentive mechanisms [6–8] implicitly require participants to authenticate on requesting and submitting in order to prevent misbehaviors. When decentralizing crowdsourcing atop open blockchain, this basic requirement might cause the history of submitting/requesting to become public knowledge recorded in the blockchain (which was originally “protected” in a data center such as the breached one of Uber’s). This breached participation history could leak considerable amount of information about workers/requesters [19] to the *public*, and therefore seriously impairs their privacy. Notably, if a worker/requester frequently joins traffic monitoring tasks, then *anyone* can read the blockchain ledger and figure out location traces of them.

To address the above fundamental privacy challenges, we have to resolve two natural tensions: (i) the tension between blockchain transparency and data confidentiality, and (ii) the tension between anonymity and accountability. Simple solutions utilizing some standard cryptographic tools (e.g. encryption and/or group signature) to protect the data confidentiality and the anonymity do not work well: the encryption of data immediately prevents smart contracts from enforcing the rewards policy, as the decrypting key is not accessible to the blockchain peers; to allow fully anonymous participation will give a dishonest worker an opportunity of submitting multiple times to one crowdsourcing task, and thus he may claim more rewards than what is supposed, and a malicious requester can anonymously submit many colluded answers to herself to downgrade the actual workers.

**Our contributions.** In this paper, we construct a general blockchain based protocol to enable the first *private and anonymous* decentralized data crowdsourcing system<sup>2</sup>. Our protocol can realize a class of incentive mechanisms for crowdsourcing without relying on a trusted party. The rewards promised by a requester can be enforced according to pre-specified policies. More importantly, we also protect the data confidentiality and the worker/requester anonymity while still holding them accountable. Specifically,

- 1) A blockchain based protocol is proposed to realize a decentralized crowdsourcing system satisfying: (i) the fair exchange between data and rewards, i.e., a worker will be paid the correct amount according to the pre-defined

<sup>2</sup>A couple of recent attempts on decentralized crowdsourcing have been made, [20–22], however, none of them address the above privacy and anonymity issues. See section II for details.

policy, *if and only if* he submits data independent with other submissions; (ii) the data confidentiality, i.e., the submitted data is confidential to anyone other than the requester; and (iii) anonymity and accountability.

Intuition behind data confidentiality is an *outsource-then-prove* methodology that: (i) the requester is enforced to deposit the budget of her incentive policy to a smart contract; (ii) answers encrypted under the requester’s public key are collected by the smart contract; (iii) the evaluation of the reward is *outsourced* to the requester who then needs to send the blockchain an instruction to reward each answer. The instruction is ensured to follow the promised incentive policy, because the requester is also required to attach a valid succinct zero-knowledge proof to attest its correctness.

The worker anonymity of our protocol can ensure: (i) the public including the requester and the implicit registration authority is not able to tell whether a data comes from a given worker; (ii) if a worker joins multiple tasks announced via the blockchain, no one can link these tasks. More importantly, we also address the threat of multi-submission exacerbated by anonymity misuse. In particular, if a worker anonymously submits more than the allowed numbers in *one* task, our scheme allows the blockchain to tell and drop the invalid submissions. We similarly achieve the requester anonymity.

- 2) To achieve the above goal of anonymity while preserving accountability, we propose a new cryptographic primitive, called *common-prefix-linkable* anonymous authentication. In most of the time, a user can authenticate on messages and also the validity of his certificate without being linked. The only exception is when the same key holder authenticates two messages with the same prefix, everyone can link the authentications. In all other cases, authentications are unlinkable. Namely, the anonymity is strong, as everyone (including the registration authority) cannot tell the actual identities behind anonymous authentications. Such a primitive may be of independent interests for its balance between anonymity and accountability.

To utilize the new primitive in our protocol, a worker has to commit to the task via an anonymous authentication. The task will be the common prefix so that the special linkability will prevent multi-submission to the same task. A requester can also use it to authenticate in each task she publishes, and convince workers that she cannot maliciously submit to downgrade their rewards.

- 3) To show the feasibility of applying our protocol, we implement the system we call ZebraLancer for a common image annotation task on top of Ethereum, a real-world blockchain infrastructure. Intensive experiments and performance evaluations are conducted in an Ethereum test net. Since current smart contracts support only primitive operations, tailoring such protocols compatible with existing blockchain platforms is non-trivial.

## II. RELATED WORK

We thoroughly review related works, and briefly discuss the insufficiencies of the state-of-the-art solutions.

**Centralized data crowdsourcing systems.** MTurk [2] is the most commercially successful data crowdsourcing platform. But it has a well-know vulnerability allowing false-reporters gain short-term advantage [15]. Also, MTurk collects plain-texts of answers, which causes considerable worry of data leakage. Last, the pseudo IDs in MTurks can be trivially linked by a malicious requester. Dynamo [23] was designed as a privacy wrapper of MTurk. Its pseudo ID can only be linked by the pseudo ID issuer, but still it inherited all other weaknesses of MTurk. SPPEAR [24] considered a couple of privacy issues in data crowdsourcing, and thus introduced a couple more authorities, each of which handled a different functionality. Distributing one authority into multiple reduces the excessive trust, but, unfortunately, it is still not clear how to instantiate all those different authorities in practice.

**Decentralized data crowdsourcing.** We also note there are several attempts [20–22, 25] using blockchain to decentralize data crowdsourcing, but neither of them considers privacy and anonymity which are fundamentally arguable for basic utility.

**Anonymous data crowdsourcing.** Li and Cao [26] proposed a framework to allow workers generate their own pseudonyms based on their device IDs. But the protocol sacrificed the accountability of workers, because workers can forge pseudonyms without attesting that they are bound to real IDs, which gave a malicious worker chances to forge fake pseudo IDs and cheat for rewards. Rahaman et al. [27] proposed an anonymous-yet-accountable protocol for crowdsourcing based on group signature, and focused on how to revoke the anonymity of misbehaved workers. Misbehaved workers could be identified and further revoked by the group manager. The authors in [24] similarly relied on group signature but introduced a couple of separate authorities. Our solution can be considered as a proactive version that can prevent worker misbehavior, and without relying on a group manager.

**Accountable anonymous authentication.** The pioneering works in anonymous e-cash [28, 29] firstly proposed the notion of one-time anonymous authentication. The concept later was studied in the context of one-show anonymous credential [30]. Some works [31, 32] further extended the notion of one-time use to be  $k$ -time use, and therefore enabled a more general accountability for anonymous authentications. In [33], the authors considered a special flavor of accountability to periodically allow  $k$ -time anonymous authentications.

Conceptually similar to the linkability appeared in one-show credential [30], linkable group/ring signatures [34, 35] were proposed to allow a user to sign messages on behalf of his group unlinkably up to twice. In [36], a more general concept of event-oriented linkable group signature was formulated to realize more fine-grained trade-off between accountability and anonymity: a user can sign on behalf of his group unlinkably up to  $k$  times per *event*, where an *event* could be a common

reference string (e.g., the unique address to call a smart contract deployed in the blockchain). But its main disadvantage is that the group manager can reveal the actual identities of users, when they have no intention but occasionally submit twice per *event*.

Our new primitive can be considered as a special cryptographic notion to formalize the subtle balance between event-oriented linkability and irrevocable anonymity. Particularly, our scheme ensures that no authority can revoke users' anonymity (which is strictly stronger than [36]).

**Privacy-preserving smart contracts.** Privacy-preserving smart contract is a recent hot topic in blockchain research. Most of them are for general purpose consideration [37, 38], and thus deploy heavy cryptographic tools including general secure multi-party computation (MPC). Hawk [39] did provide a general framework for privacy-preserving smart contracts using light zk-SNARK, but mainly for reward receiver to prove to the contract. Our work can be considered as a very specially designed MPC protocol, and a lot of dedicated optimizations of zk-SNARK exist which can directly benefit our protocol. Last, cryptocurrencies like Zcash [40] and Ethereum [41] also leverage zk-SNARK to build a public ledger that supports anonymous transactions. We note that they consider more basic blockchain infrastructures, on top of which we may build our application for crowdsourcing.

## III. PRELIMINARIES

**Blockchain and smart contracts.** A blockchain is a global ledger maintained by a P2P network collectively following a pre-defined consensus protocol. Each block in the chain will aggregate some transactions containing use-case specific data (e.g., monetary transfers, or program codes).

In general, we can view the blockchain as an ideal public ledger [42] where one can write and read data in the clear. Moreover, it will faithfully carry out certain pre-defined functionalities. The last property captures the essence of smart contracts, that every blockchain node will run them as programs and update their local replicas according to the execution results, collectively. More specifically, the properties of the blockchain can be informally abstracted as the following ideal public ledger model [39]:

- 1) *Reliable delivery of messages.* The blockchain can be modeled as an ideal public ledger that ensures the *liveness and persistence* of messages committed to it [42]. Detailedly, a message sent to the blockchain is in the form of a validly signed transaction broadcasted to the whole blockchain network, and then it will be solicited by a block and written into the blockchain. Although the blockchain ensures persistence in asynchronous network (under honest majority assumption), its liveness still requires a-prior network delay [43]. We will assume synchronous network model as well as liveness through the paper, such that a valid message will be agreed by the whole blockchain network within a certain time, after it is sent. Also we remark that a network

adversary can reorder transactions that are broadcasted to the network but not yet written into a block.

- 2) *Correct computation.* The blockchain can be seen as a state machine driven by messages included in each block [44]. Specifically, miners and full nodes will persistently receive newly proposed blocks, and faithfully execute “programs” defined by current states with taking messages in new blocks as inputs. Moreover, the computing results can be written into the chain, and therefore be reliably delivered to the whole network.
- 3) *Transparency.* All internal states of the blockchain will be visible to the whole blockchain (intuitively, anyone). Therefore, all message deliveries and computations via the blockchain are in the clear. For example, when a smart contract decrypts some ciphertexts with using a decrypting key, the key will be learnt by everyone.
- 4) *Blockchain address (Pseudonym).* A message committed to the blockchain should be sent in a pseudonym, a.k.a. blockchain address. In practice, a blockchain address is usually bound to the hash of a public key. More importantly, the message should be correctly signed by using the corresponding secret key, and the security of digital signatures can further ensure that one cannot send messages in the name of the blockchain address, unless she has the secret key.

Also, the program code of a smart contract deployed in the blockchain can be referred by a unique blockchain address, such that one can call the blockchain network to execute the contract, through sending the blockchain a message that points to this address.

**zk-SNARK.** A zero-knowledge proof (zk-proof) allows a party (i.e. prover) to generate a cryptographic proof convincing another party (i.e. verifier) that some values are obtained by faithfully executing a pre-defined computation on some private inputs (i.e. witness) without revealing any information about the private state. The security guarantees are: (i) *soundness*, that no prover can convince a verifier if she did not compute the results correctly; sometimes, we require a stronger soundness that for any prover, there exists an extractor algorithm which interacts with the prover and can actually output the witness (a.k.a. *proof-of-knowledge*); (ii) *zero-knowledge*, that the proof distribution can be simulated without seeing any secret state, i.e., it leaks nothing about the witness. Both above will hold with an overwhelming probability.

The zero-knowledge succinct non-interactive argument of knowledge (zk-SNARK) further allows such a proof to be generated non-interactively. More importantly, the proof is *succinct*, i.e., the proof size is independent on the complexity of the statement to be proved, and is always a small constant. More precisely, zk-SNARK is a tuple of three algorithms. A setup algorithm can output the public parameters to establish a SNARK for a NP-complete language  $\mathcal{L} = \{\vec{x} \mid \exists \vec{w}, s.t., C(\vec{x}, \vec{w}) = 1\}$ . The Prover algorithm can leverage the established SNARK to generate a constant-size proof attesting

the trueness of a statement  $\vec{x} \in \mathcal{L}$  with witness  $\vec{w}$ . The Verifier algorithm can efficiently check the proof.

#### IV. PROBLEM FORMULATION

In this section, we will give more precise definitions about the problem and its security requirements.

**Data crowdsourcing model.** As illustrated in Fig.1, there are four roles in the model of data crowdsourcing, i.e., requesters, workers, a platform and a registration authority. A *requester*, uniquely identified by  $R$ , can post a task to collect a certain amount of answers from the crowd. When announcing the task, the requester promises a concrete reward policy to incentivize workers to contribute (see details about the definition of reward policy below). A *worker* with a unique ID  $W_j$ , submits his answer  $A_j$  and expects to receive the corresponding reward. The *platform*, a medium assisting the exchange between requesters and workers, is either a trusted party or emulated by a network of peers. The platform considered in this paper is jointly maintained by a collection of network peers, and in particular, we will build it atop an open blockchain network. The *registration authority* (RA), can play an important role of verifying and managing unique identities of workers/requesters, by binding each identity to a unique credential (e.g. a digital certificate).




Fig. 1. The model of data crowdsourcing: workers and requesters obtain unique credentials bound to their identities at a registration authority (RA); authenticated requesters and authenticated workers can make fair exchange between data and rewards through a third-party platform (arbiter).

RA is a necessary demand of real-world crowdsourcing systems such as MTurk and Waze to prevent malicious participants. Moreover, many auction-based incentive mechanisms [7, 8] are built upon the non-collusion game theory that implicitly requires RA to ensure one bid from one unique identity. In practice, RA can be instantiated by (i) the platform itself, (ii) the certificate authority who provides authentication service, or (iii) the hardware manufacturer who makes trusted devices that can faithfully sign on messages [45]. Our solution should be able to inherit these established RAs in real-world.

In this paper, w.l.o.g., we assume that each unique identity is only allowed to submit one answer to a task. Also, we consider settings that the value of crowd-shared answers can be evaluated by a well-defined process such as auctions or qualities, and also the corresponding rewards, c.f. [9–11] about

quality-aware rewards, and [7, 8] about auction-based incentives. These incentives share the similar essence as follows.

Suppose an authenticated requester publishes a task  $T$  with a budget  $\tau$  to collect  $n$  answers from  $n$  workers. An authenticated worker who is interested in it will then submit his answers. Then the *reward* of an answer  $A_j$  will be determined by a function parameterized by some variables, denoted by  $a_1, \dots, a_m$ , i.e., the reward of  $A_j$  can be defined as a function  $R_j := R(A_j; a_1, \dots, a_m, \tau)$  parameterized by some variables, denoted by  $a_1, \dots, a_m, \tau$ . Remark that  $\tau$  is the budget of the requester, and we will use  $R_j := R(A_j; \tau)$  for short.

Particularly, in some simple crowdsourcing tasks (e.g., multiple choice problems), the quality of an answer can be straightforwardly evaluated by all answers to the same task, i.e.  $R_j = R(A_j; A_1, \dots, A_n, \tau)$ , with using either majority voting or estimation maximization iterations [9–11]. More generally, [46] proposed a universal method to evaluate quality: (i) some workers are requested to answer a complex task; (ii) other different workers are then requested to grade each answer collected in the previous stage. What’s more, our model actually captures the essence of many auction-based incentive mechanism such as [7, 8], when the parameters  $a_1, \dots, a_m$  represent the bids of workers (and other necessary auxiliary inputs). So our protocol design and implementations will focus on instantiating quality-aware incentives, and it should be extendable to the scope of auction-based incentives trivially.

**Security models.** Next, we specify the basic security requirements for our (decentralized) crowdsourcing system.

*Data confidentiality.* This property requires that the communication transcripts (including the blocks in the blockchain) do not leak anything to anyone (except the requester) about the input parameters  $a_1, \dots, a_m$  of the incentive policy  $R$ . Because these parameters might actually be the valuable data submitted by workers. We can adapt the classical semantic security [47] style of definition from cryptography for this purpose: the distribution of the public communication can be simulated with only public knowledge.

*Anonymity.* Private information of worker/requester can be explored by linking tasks they join/publish [19]. Intuitively, we might require two properties for workers’ anonymity: (i) *unlinkability between a submission and a particular worker* and (ii) *unlinkability among all tasks joined by a particular worker*. However, (i) indeed can be implied by (ii), because the break of first one can obviously lead up to the break of the latter one. Similarly, the anonymity of requester can be understood as the unlinkability among all tasks published by her. The requirement of worker anonymity can be formulated via the following game. An adversary  $\mathcal{A}$  corrupts a requester, the registration authority (RA), and the platform (e.g. the blockchain); suppose there are only two honest workers,  $W_0$  and  $W_1$ . In the beginning, the adversary announces  $n$  tasks. For each task  $T_i$ , suppose there are a set of participating workers  $\mathbf{W}_{T_i}$ . After seeing all the communications, for any  $T_i \neq T_j$ ,  $\mathcal{A}$  cannot tell whether  $\mathbf{W}_{T_i} \cap \mathbf{W}_{T_j} = \emptyset$  better

than guessing. We note that the anonymity should hold even if all entities, including the requester and the platform (except  $W_0$  and  $W_1$ ), are corrupted. The requester anonymity can be defined via the above game similarly, and we omit the details.

*Security against a malicious requester.* A malicious requester may avoid paying rewards (defined by the pre-specified  $R$ ), e.g. launches the *false-reporting* attack. Security in this case can be formulated via the following security game: an adversary  $\mathcal{A}$  corrupts the requester and executes the protocol by specifying a task with a budget  $\tau$ . Let us define a bad event  $B_1$  to be that there exists a worker  $W_j$ , who receives payment smaller than  $R_j := R(A_j; \tau)$ . We require that for every polynomial time  $\mathcal{A}$ ,  $\Pr[B_1]$  is negligibly small.

*Security against malicious workers.* A dishonest worker may try to harvest more rewards than what he deserves. Security in this case can be formulated as follows: an adversary  $\mathcal{A}$  corrupts one worker,<sup>3</sup> and participates in the protocol interacting with a requester (and the platform).  $\mathcal{A}$  submits some answers  $\mathbf{A} := \{A_1, \dots, A_n\}$ ,  $n \geq 1$ . Let us define the bad event  $B_2$  as that  $\mathcal{A}$  receives a payment greater than  $\max_{\{A_i \in \mathbf{A}\}} R_j := R(A_j; \tau)$  from the requester. We require that for all polynomial time  $\mathcal{A}$ ,  $\Pr[B_2]$  is negligibly small.

We remark that the above securities against a malicious requester and malicious workers have capture the special fairness of the exchange between crowd-shared data and rewards.

## V. PRIVATE AND ANONYMOUS DECENTRALIZED DATA CROWDSOURCING: PROTOCOL

In this section, we will construct a private and anonymous protocol to address the critical challenges of decentralizing crowdsourcing, without sacrificing security against “free-riders” and “false-reports”. The crowdsourcing platform will be built upon an existing network of blockchain. More specifically, we will tackle the new privacy and anonymity challenges brought by the blockchain.

As we briefly mentioned in previous sections, the system will implicitly have a separate registration service that validates each participant’s unique identity before issuing a certificate. Such setup alleviates some basic problems that every worker is allowed to submit no more than a fixed number  $k$  of answers. For simplicity, we consider here  $k = 1$ .

**Intuitions.** Our basic strategy is to let the smart contract to enforce the fair exchange between submitted answers and the corresponding rewards, but without revealing data or identities (or certificates). Let us walk through the high level ideas first.

The requester firstly codifies a reward policy parameterized by her budget (i.e.  $R(\cdot; \tau)$ ) into a smart contract. She broadcasts a transaction containing the contract code and the budget. Once the smart contract is included in the blockchain, it can be referred by a unique blockchain address, and the

<sup>3</sup>We remark that we focus on resolving the *new* challenges introduced by blockchain, and put forth the best possible security, as if there is a fully trusted third-party serving as the crowdsourcing platform. For example, it is not clear how to handle worker collusion even in the centralized setting, thus such a problem is out of the scope of this paper.

budget should be deposited to this address (otherwise, no one would participate). After that, any worker who is interested in contributing could simply submit his answer to the blockchain, via a transaction pointing to the contract’s address.

As pointed out above, we have to protect the *confidentiality* of the answers, in order to ensure that answers from different workers are independent. So the workers encrypt the answers under the requester’s public key. Now the contract cannot see the answers so it cannot calculate the corresponding rewards. But the requester can retrieve all the encrypted answers and decrypt them off-chain, and further learn the rewards they deserve. It would be necessary that the requester will *correctly* instruct the smart contract how to proceed forward. Concretely, we will leverage the practical cryptographic tool of zk-SNARK to enforce the requester to prove: she indeed *followed the pre-specified reward policy* calculating the rewards. Detailedly, the requester should prove her instruction of rewards is computed as follows: (i) obtain all answers by decrypting all encrypted answers using a secret key corresponding to the public key contained in the smart contract; (ii) use all those answers and the announced  $R(\cdot; \tau)$  to compute the quality of each answer.

A more challenging issue arises regarding *anonymity-yet-accountability*. Also as briefly pointed before, we would like to achieve a balance between anonymity and accountability. Here we put forth a new cryptographic primitive to resolve the natural tension. A user can anonymously authenticate messages (which are composed of a fixed length prefix and the remaining part). But if the two authenticated messages share a common prefix, anyone can tell whether they are done by a same user or not. Moreover, no one can link any two message-authentication pairs, as long as these messages have different prefixes. Having this new primitive in hand, our protocol can be designed by letting each worker to anonymously authenticate on the message of  $\alpha_{\mathcal{C}} || c$ , when he submits the encrypted answer  $c = \text{Enc}(pk, a)$  to a contract  $\mathcal{C}$  (that has a unique address  $\alpha_{\mathcal{C}}$  and specifies the crowdsourcing task). This implies that the submissions from the same worker to one task can be linked and thus counted, but two submissions of any worker to two different tasks will not be linkable. Also, the number of maximum allowed submissions in each task can be easily tuned (by counting linked submissions).

Last, we also need to augment the smart contract by building the zk-SNARK verification algorithm in it. In particular, when the smart contract receives the signed instruction and proof, the verification algorithm will be executed. All inputs of the verification algorithm are immutable common knowledge stored in the blockchain, e.g., the budget, the encrypted answers and the requester’s public key. If a dishonest requester reports a false instruction, her proof cannot be verified and the contract will simply drop the instruction. What’s more, if the smart contract does not receive a *correct* instruction within a time limit, it can directly disseminate the budget to all workers evenly as punishment, since the budget has been deposited. In this way, the requester cannot gain any benefit by deviating from the protocol, and she will be self-enforced to respond

properly and timely, resulting in that each worker will receive the expected reward. On the other hand, a dishonest worker can never claim more rewards than that he is supposed to get, as the reward is calculated by the requester herself.

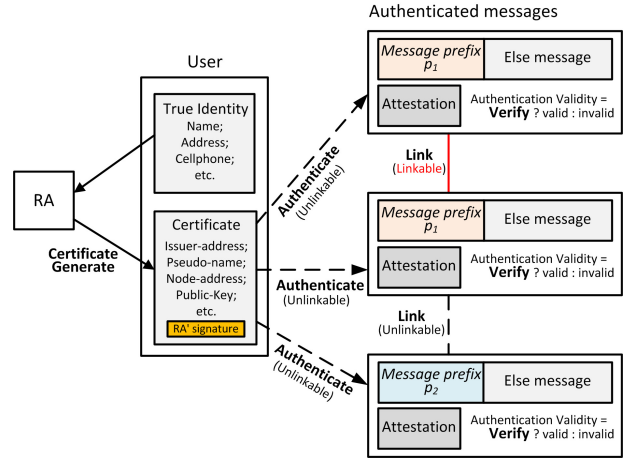


Fig. 2. Subtle linkability of the common-prefix-linkable anonymous authentication scheme. All involved algorithms except Setup are shown in bold.

#### A. Common-prefix-linkable anonymous authentication

Before the formal description of ZebraLancer’s protocol, let us introduce the new primitive for achieving the anonymous-yet-accountable authentication first. As briefly shown in Fig.2, our new primitive can be built atop any certification procedure, thus we include a certification generation procedure that can be inherited from any existing one. Also, we insist on *non-interactive* authentication, thus all the steps (including the authentication step) are described as algorithms instead of protocols. Formally, a common-prefix-linkable anonymous authentication scheme is composed of the following algorithms:

- **Setup**( $1^\lambda$ ). This algorithm outputs the system’s master public key  $mpk$ , and system’s master secret key  $msk$ , where  $\lambda$  is the security parameter.
- **CertGen**( $msk, pk$ ). This algorithm outputs a certificate  $cert$  to validate the public key.
- **Auth**( $m, sk, pk, cert, mpk$ ). This algorithm generates an attestation  $\pi$  on a message  $m$  that: the sender of  $m$  indeed owns a secret key corresponding to a valid certificate.
- **Verify**( $m, mpk, \pi$ ). This algorithm outputs 0/1 to decide whether the attestation is valid or not.
- **Link**( $mpk, m_1, \pi_1, m_2, \pi_2$ ). This algorithm takes two valid message-attestation pairs, i.e.  $(m_1, \pi_1), (m_2, \pi_2)$ , as inputs. If  $m_1, m_2$  have a common-prefix with length  $\lambda$ , and  $\pi_1, \pi_2$  are generated from the same certificate, it outputs 1; otherwise outputs 0.

The first one characterizes a special *accountability* requirement in anonymous authentication. It requires that no efficient adversary can authenticate two messages with a common-prefix without being linked, if using the same certificate. More generally, if an attacker corrupts  $q$  users, she cannot

authenticate  $q + 1$  messages sharing a common-prefix, without being noticed. Formally, consider the following cryptographic game between a challenger  $\mathcal{C}$  and an attacker  $\mathcal{A}$ :

- 1) Setup. The challenger  $\mathcal{C}$  runs the Setup algorithm and obtains the master keys.
- 2) CertGen queries. The adversary  $\mathcal{A}$  submits  $q$  public keys with different identities and obtains  $q$  different certificates:  $cert_1, \dots, cert_q$ .
- 3) Auth. The adversary  $\mathcal{A}$  chooses  $q + 1$  messages  $p||m_1, \dots, p||m_{q+1}$  sharing a common-prefix  $p$  (with  $|p| = \lambda$ ) and authenticates to the challenger  $\mathcal{C}$  by generating the corresponding attestations  $\pi_1, \dots, \pi_{q+1}$ .

Adversary  $\mathcal{A}$  wins if all  $q + 1$  authentications pass the verification, and no pair of those authentications were linked.

**Definition 1** (Common-prefix-linkability). *For all probabilistic polynomial time algorithm  $\mathcal{A}$ ,  $\Pr[\mathcal{A}$  wins in the above game] is negligible on the security parameter  $\lambda$ .*

Next is the *anonymity* guarantee in normal cases. We would like to ensure the anonymity against any party, including the public, the registration authority, and the verifier who can ask for multiple (and potentially correlated) authentication queries. Also, our strong anonymity requires that no one can even link whether the same people is authenticating for different messages, if these messages have different prefixes. The basic requirement for anonymity is that no one can recognize the real identity from the authentication transcript. But our *unlinkability* requirement is strictly stronger, as if one can recognize identity, obviously, she can link two authentications by first recovering the actual identities.

To capture the unlinkability (among the authentications of different-prefix messages), we can imagine the most stringent setting, where there are only two honest users in the system, the adversary still cannot properly link any of them from a sequence of authentications. Formally, consider the following game between the challenger  $\mathcal{C}$  and adversary  $\mathcal{A}$ .

- 1) Setup. The adversary  $\mathcal{A}$  generates the master key pair.
- 2) CertGen. The adversary  $\mathcal{A}$  runs the certificate generation procedure as a registration authority with the challenger. The challenger submits two public keys  $pk_0, pk_1$  and the adversary generates the corresponding certificates for them  $cert_0, cert_1$ .  $\mathcal{A}$  can always generate certificates for public keys generated by herself.
- 3) Auth-queries. The adversary  $\mathcal{A}$  asks the challenger to serially use  $(sk_0, pk_0, cert_0)$  and  $(sk_1, pk_1, cert_1)$  to do a sequence of authentications on messages chosen by her. Also, the number  $q$  of authentication queries is chosen by  $\mathcal{A}$ . The adversary obtains  $2q$  message-attestation pairs.
- 4) Challenge. The adversary  $\mathcal{A}$  chooses a new message  $m^*$  which does not have a common prefix with any of the messages asked in the Auth-queries, and asks the challenger to do one more authentication.  $\mathcal{C}$  picks a random bit  $b$  and authenticates on  $m^*$  using  $sk_b, pk_b, cert_b$ . After receiving the attestation  $\pi_b$ ,  $\mathcal{A}$  outputs her guess  $b'$ .

The adversary wins if  $b' = b$ .

**Definition 2** (Anonymity). *An authentication scheme is unlinkable, if  $\forall$  probabilistic polynomial-time algorithm  $\mathcal{A}$ ,  $|\Pr[\mathcal{A}$  wins in the game] -  $\frac{1}{2}|$  is negligible.*<sup>4</sup>

**Construction.** Now we proceed to construct such a primitive. Same as many anonymous authentication constructions, we will also use the zero-knowledge proof technique towards anonymity. In particular, we will leverage zk-SNARK to give an efficient construction. For the above concept of common-prefix-linkable anonymous authentication, we need to further support the special accountability requirement. The idea is as follows, since the condition that “breaks” the linkability is common-prefix, thus the authentication will do a special treatment on the prefix. In particular, the authentication shows a tag committing to the prefix together with the user’s secret key, and then presents zero-knowledge proof that such a tag is properly formed, i.e., computed by hashing the prefix and a secret key. To ensure other basic security notions, we will also compute the other tag that commits to the whole message. The user will further prove in zero-knowledge that the secret key corresponds to a certified public key.

We remark that our main goal is to construct a decentralized crowdsourcing, such a new anonymous authentication primitive could be further studied systematically in future works. Concretely, we present the detailed construction as follows:

- Setup( $\lambda$ ). This algorithm establishes the public parameters  $PP$  that will be needed for the zk-SNARK system. Also, the algorithm generates a key pair  $(msk, mpk)$  which is for a digital signature scheme.<sup>5</sup>
- CertGen( $msk, pk_i$ ): This algorithm runs a signing algorithm on  $pk_i$ ,<sup>6</sup> and obtains a signature  $\sigma_i$ . It outputs  $cert_i := \sigma_i$ .
- Auth( $p||m, sk_i, pk_i, cert_i, PP$ ): On inputting a message  $p||m$  having a prefix  $p$ . The algorithm first computes two tags (or interchangeably called headers later),  $t_1 = H(p, sk_i)$  and  $t_2 = H(p||m, sk_i)$ , where  $H$  is a secure hash function. Then, let  $\vec{w} = (sk_i, pk_i, cert_i)$  represent the private witness, and  $\vec{x} = (p||m, mpk)$  be all common knowledge, the algorithm runs zk-SNARK proving algorithm Prover( $\vec{x}, \vec{w}, PP$ ) for the following language  $\mathcal{L}_T := \{t_1, t_2, \vec{x} = (p||m, mpk) \mid \exists \vec{w} = (sk_i, pk_i, cert_i) \text{ s.t. } CertVrfy(cert_i, pk_i, mpk) = 1 \wedge pair(pk_i, sk_i) = 1 \wedge t_1 = H(p, sk_i) \wedge t_2 = H(p||m, sk_i)\}$ , where the CertVrfy algorithm checks the validity of the certificate using a signature verification, and *pair* algorithm verifies whether

<sup>4</sup>We remark that our definition of anonymity is strictly stronger than that definition of the event-oriented linkable group signature [36], in which the RA can revoke user anonymity under certain conditions.

<sup>5</sup>To be more precise, the public parameter generation could be from another algorithm, for simplicity, we put it here. In the security game for anonymity, the adversary only generates the  $msk, mpk$ , not the public parameter.

<sup>6</sup>Here we assume there is an external identification procedure to check the actual identity with the public key, and the user key pairs are generated using common algorithms, e.g., for a digital signature, here we ignore the details.

two keys are a consistent public-secret key pair. This prove algorithm yields a proof  $\eta$  for the statement  $\bar{x} \in \mathcal{L}$  (also for the proof-of-knowledge of  $\bar{w}$ ).

Finally, the algorithm outputs  $\pi := (t_1, t_2, \eta)$ .

- $\text{Verify}(p||m, \pi, mpk, PP)$ : this algorithm runs the verifying algorithm of zk-SNARK Verifier on  $\bar{x}, \pi$  and  $PP$ , and outputs the decision bit  $d \in \{0, 1\}$ .
- $\text{Link}(m_1, \pi_1, m_2, \pi_2)$ : On inputting two attestations  $\pi_1 := (t_1^1, t_2^1, \eta_1)$  and  $\pi_2 := (t_1^2, t_2^2, \eta_2)$ , the algorithm simply checks  $t_1^1 \stackrel{?}{=} t_1^2$ . If yes, output 1; otherwise, output 0. We also use  $\text{Link}(\pi_1, \pi_2)$  for short.

**Security analysis (sketch).** Here we briefly sketch the security analysis for the above construction.

Regarding *correctness*, it is trivial, because of the completeness of the underlying SNARK.

Regarding *unforgeability*, we require an uncertified attacker cannot authenticate. The only transcripts can be seen by the adversary are headers and the zero-knowledge attestation. Headers include one generated by hashing the concatenation of  $p||m, sk$ . In order to provide a header, the attacker has to know the corresponding  $sk$ , as it can be extracted in the random oracle queries. Thus there are only two different ways for the attacker: (i) the attacker generates forges the certificate, which clearly violates the signature security; (ii) the attacker forges the attestation using an invalid certificate, which clearly violates the proof-of-knowledge of the zk-SNARK.

Regarding the *common-prefix-linkability*, it is also fairly straightforward, as the final authentication transcript contains a header computed by  $H(p, sk)$  which is an invariant for a common prefix  $p$  using the same secret key  $sk$ .

Regarding the *anonymity/unlinkability*, we require that after seeing a bunch of authentication transcripts from one user, the attacker cannot figure out whether a new authentication comes from the same user. This holds even if the attacker can be the registration authority that issues all the certificates. To see this, as the attacker will not be able to figure the value of the  $sk$  from all public value, thus the headers/tags can be considered as random values. It follows that  $H(p, sk)$  and a random value  $r$  cannot be distinguished (similarly for  $H(p||m, sk)$ ). More importantly, due to the zero-knowledge property of zk-SNARK, given  $r$ , a simulator can simulate a valid proof  $\eta^*$  by controlling the common reference string of the zk-SNARK. That said, the public transcript  $t_1, t_2, \eta$  can be simulated by  $r_1, r_2, \eta^*$  where  $r_1, r_2$  are uniform values, and  $\eta^*$  is a simulated proof, all of which has nothing to do with the actual witness  $sk$ .

Summarizing the above intuitive analysis, we have the following theorem:

**Theorem 1.** *Conditioned on that the hash function to be modeled as a random oracle and the zk-SNARK is zero-knowledge, the construction of the common-prefix-linkable anonymous authentication satisfies anonymity. Conditioned on the underlying digital signature scheme used is secure, and the zk-SNARK satisfies proof-of-knowledge, our construction*

*of the common-prefix-linkable anonymous authentication will be unforgeable. It is also correct and common-prefix linkable.*

## B. The protocol for ZebraLancer

Now we are ready to present a general protocol for a class of crowdsourcing tasks having proper quality-aware incentives mechanisms defined as in Section IV. As zk-SNARK requires a setup phase, we consider that a setup algorithm generated the public parameters  $PP$  for this purpose, and published it as common knowledge.<sup>7</sup> Our descriptions focuses on the application atop the open blockchain, and therefore omits details of sending messages through the underlying blockchain infrastructure. For example, “one uses blockchain address  $\alpha$  to send a message  $m$  to the blockchain” will represent that he broadcasts a blockchain transaction containing the message  $m$ , the public key associated to  $\alpha$ , and the signature properly generated under the corresponding secret key.

Remark that here we let each worker/requester to generate a different blockchain address for each task (i.e. a *one-task-only* address) as a simple solution to avoid de-anonymization in the underlying blockchain.<sup>8</sup> For concrete instantiations of the underlying infrastructures, see our implementations discussed in Section VI.

**Protocol details.** The details of ZebraLancer protocol can be described as follows:

- **Register.** *Everyone registers at RA to get a certificate bound to his/her unique ID, which is done off-line only once for per each participant.*

A requester, having a unique ID denoted by  $R$ , creates a public-secret key pair  $(pk_R, sk_R)$ , and registers at the registration authority (RA) to obtain a certificate  $cert_R$  binding  $pk_R$  to  $R$ . Each worker, having a unique ID denoted by  $W_i$ , also generates his public and secret key pair  $(pk_i, sk_i)$ , and registers his public key at RA to obtain a certificate  $cert_i$  binding  $pk_i$  and  $W_i$ .

- **TaskPublish.** *A requester anonymously authenticates and publishes a task contract with a promised reward policy via the blockchain.*

When the requester  $R$  has a crowdsourcing task, she generates a *fresh* blockchain account address  $\alpha_R$ , and a key pair  $(epk, esk)$  (which will be used for workers to encrypt submissions) for this task only.

$R$  then prepares parameters Param, including the encryption key  $epk$ , the number of answers to collect (denoted by  $n$ ), the deadline, the budget  $\tau$ , the reward policy R,

<sup>7</sup>This in practice can be done via a secure multiparty computation protocol [48] to eliminate potential backdoors.

<sup>8</sup>Our anonymous protocol mainly focuses on the application layer such as the crowdsourcing functionality that is built on top of the blockchain infrastructure. If the underlying blockchain layer supports anonymous transaction, such as Zcash [40], the worker and the requester can re-use account addresses. We further remark that the anonymity in network layer are out the scope of this paper, we may deploy our protocol on existing infrastructure such as Tor.



SNARK's public parameters  $PP$ , RA's public key  $mpk$ , and also  $\pi_R = \text{Auth}(\alpha_{\mathcal{C}} \parallel \alpha_R, sk_R, pk_R, cert_R, PP)$ .<sup>9</sup>

The requester then codes a smart contract  $\mathcal{C}$  that contains all above information for her task. After compiling  $\mathcal{C}$ , she puts  $\mathcal{C}$ 's code and a transfer of the budget into a blockchain transaction, and uses the one-task-only address  $\alpha_R$  to send the transaction into the blockchain network. When a block containing  $\mathcal{C}$  is appended to the blockchain,  $\mathcal{C}$  gets an immutable blockchain address  $\alpha_{\mathcal{C}}$  to hold the budget and interact with anyone.<sup>10</sup>

See Algorithm 1 below for a concrete example of task contract. (The important component of verifying zk-proofs is done by calling a library `libsnark.Verifier` integrated into the blockchain infrastructure, and implementation details will be explained in Section VI).

- **AnswerCollection.** *The contract collects anonymously authenticated encrypted answers, only when an answer is sent from a worker who didn't submit before.*

If a registered worker  $W_i$  is interested in contributing, he first validates the contract content (e.g., checking the parameters), then generates a *one-task-only* blockchain address  $\alpha_i$ . He encrypts his answer  $A_i$  under the task's public key  $epk$  to obtain ciphertext  $C_i$ .

He then uses common-prefix-linkable anonymous authentication scheme to generate an attestation  $\pi_i = \text{Auth}(\alpha_{\mathcal{C}} \parallel \alpha_i \parallel C_i, sk_i, pk_i, cert_i, PP)$ .<sup>9</sup> Then he uses his *one-task-only* address  $\alpha_i$  to send  $C_i, \pi_i$  to the blockchain network (with a pointer to  $\alpha_{\mathcal{C}}$ , i.e. the unique address of the contract  $\mathcal{C}$ ).

Then,  $\mathcal{C}$  runs  $\text{Verify}(\alpha_{\mathcal{C}} \parallel \alpha_i \parallel C_i, \pi_i, mpk, PP)$ , and also executes  $\text{Link}(\pi_i, \pi_*)$  for each valid authentication attestation  $\pi_*$  that was received before (including requester's, namely  $\pi_R$ ). Such that,  $\mathcal{C}$  can ensure  $C_i$  is the first submission of a registered worker. For unauthenticated or double submissions,  $\mathcal{C}$  simply drops it.<sup>11</sup>

The contract  $\mathcal{C}$  will keep on collecting answers, until it receives  $n$  answers or the deadline (in unit of block) passes. It also records each address  $\alpha_i$  that sends  $C_i$ . Remark that  $\text{Link}$  algorithm will be executed  $O(n^2)$  times, but it is a simple equality check over a pair of

<sup>9</sup>We remark that the requester should authenticate her one-task-only blockchain address  $\alpha_R$  along with the task contract, and workers will join the task only if the task contract is indeed sent from a blockchain address as same as the authenticated  $\alpha_R$ . So a malicious requester cannot "authenticate" a task by copying other valid authentications. In addition, each worker has to authenticate his one-task-only blockchain address  $\alpha_i$  with his answer submission as well. The task contract will check the submission is indeed sent from an address same to the authenticated  $\alpha_i$ . Otherwise, a malicious worker can launch free-riding through copying and re-sending authenticated submissions that have been broadcasted but not yet confirmed by a block.

<sup>10</sup>We emphasize that  $\alpha_{\mathcal{C}}$  will be unique per each contract. In practice,  $\alpha_{\mathcal{C}}$  can be computed via  $H(\alpha_R \parallel \text{counter})$ , where  $H$  is a secure hash function, and  $\text{counter}$  is governed by the blockchain to be increased by exact one for each contract created by the blockchain address  $\alpha_R$ . It's also clear that the requester  $R$  can predicate  $\alpha_{\mathcal{C}}$  before  $\mathcal{C}$  is on-chain, such that she can compute  $\pi_R$  off-line and let it be a parameter of contract  $\mathcal{C}$ .

<sup>11</sup>We remark that our protocol can be extended trivially to allow each worker to submit some  $k$  answers in one task by modifying the checking condition programmed in the smart contract of crowdsourcing task.

---

### Algorithm 1: Example using quality-aware incentive

---

**Require :** This contract's address  $\alpha_{\mathcal{C}}$ ; requester's one-time blockchain address  $\alpha_R$ ; requester's authenticating attestation  $\pi_R$ ; RA's public key  $mpk$ ; budget  $\tau$ ; public key  $epk$  for encrypting answers; SNARK's public parameters  $PP$ ; number of requested answers  $n$ ; deadline of answering in unit of block  $T_A$ ; deadline of instructing reward in unit of block  $T_I$ .

```

1 List keeping answers' ciphertexts,  $C \leftarrow \emptyset$ ;
2 Map of anonymous attestations and authenticated one-time
  blockchain addresses of workers,  $W \leftarrow \emptyset$ ;
3 if  $\text{getBalance}(\alpha_{\mathcal{C}}) < \tau \vee \neg \text{Verify}(\alpha_{\mathcal{C}} \parallel \alpha_R, \pi_R, mpk, PP)$ 
  then
4    $\perp$  goto 21 ;  $\triangleright$  Budget not deposited or requester not identified.
5  $\text{timer}_A \leftarrow$  a timer expires after  $T_A$ ;
6 while  $\|C\| < n \wedge \text{timer}_A$  NOT expired do
7   if  $\alpha_i$  sends  $\pi_i, C_i$  then
8     if  $\neg \text{Link}(\pi_i, \pi_R) \wedge \forall \pi_* \in W.\text{keys}() \neg \text{Link}(\pi_i, \pi_*) \wedge$ 
        $\text{Verify}(\alpha_{\mathcal{C}} \parallel \alpha_i \parallel C_i, \pi_i, mpk, PP)$  then
9        $W.\text{add}(\pi_i \rightarrow \alpha_i)$ ;  $C.\text{add}(C_i)$ ;
10  $\text{timer}_I \leftarrow$  a timer expires after  $T_I$ ;  $\triangleright$  Start to wait instruction
11 while  $\text{timer}_I$  NOT expired do
12   if  $\alpha_R$  sends  $\bar{R} := (R_1, \dots, R_n)$  and  $\pi_{\text{reward}}$  then
13      $\bar{P} \leftarrow (epk, \tau, C_1, \dots, C_n)$ ;
14     if  $\text{libsnark.Verifier}((\bar{P}, \bar{R}), \pi_{\text{reward}}, PP)$  then
15       for each  $(\pi_i \rightarrow \alpha_i) \in W$  do
16          $\perp$   $\text{transfer}(\alpha_{\mathcal{C}}, \alpha_i, R_i)$ ;
17       goto 21;
18  $R \leftarrow \tau / \|W\|$ ;  $\triangleright$  Reward all if no correct instruction
19 for each  $(\pi_i \rightarrow \alpha_i) \in W$  do
20    $\perp$   $\text{transfer}(\alpha_{\mathcal{C}}, \alpha_i, R)$ ;
21  $\text{transfer}(\alpha_{\mathcal{C}}, \alpha_R, \text{getBalance}(\alpha_{\mathcal{C}}))$ ;  $\triangleright$  Refund the remaining
22 function  $\text{getBalance}(\text{addr})$ 
23    $\perp$  return the balance of  $\text{addr}$  in the blockchain ledger;
24 function  $\text{transfer}(\text{src}, \text{dst}, \text{value})$ 
25   if  $\text{getBalance}(\text{src}) < \text{value}$  then
26      $\perp$  return false;
27   the balance of  $\text{src}$  subtracts  $\text{value}$  in the blockchain ledger;
28   the balance of  $\text{dst}$  adds  $\text{value}$  in the blockchain ledger;
29   return true;
30    $\triangleright$  The contract program is driven by a "discrete" clock that
    increments with validating each newly proposed block
     $\triangleright$   $\text{libsnark.Verifier}$  is a library embedded in the runtime
    environment of smart contract such as EVM

```

---

hashes, such that the cost of running it for several times will be nearly nothing in practice.

- **Reward.** *The requester computes and prove to the smart contract how to reward each anonymous answer.*

The requester  $R$  keeps listening to the blockchain, and once  $\mathcal{C}$  collects  $n$  submissions, she retrieves and decrypts all of them to obtain the corresponding answers  $A_1, \dots, A_n$  (if there are not enough submissions when the deadline passes, the requester simply sets the remaining answers to be  $\perp$  which has been considered by the incentive mechanism  $R$ ).

Next, the requester computes the reward for each answer

$R_i = R(A_i; A_1, \dots, A_n, \tau)$  as specified by the policy codified in  $\mathcal{C}$ . More importantly, she generates a zero knowledge proof  $\pi_{reward}$ , with the secret key  $esk$  as witness to attest the validity of the instruction. In particular, the proof is for the following NP-language  $\mathcal{L} = \{\bar{R}, \bar{P} \mid \exists esk \text{ s.t. } \bigwedge_{j=1}^n A_j = \text{Dec}(esk, C_j) \wedge \bigwedge_{j=1}^n R_j = R(A_j; A_1, \dots, A_n, \tau) \wedge \text{pair}(esk, epk) = 1\}$ , where  $\bar{P}$  denotes Param together with ciphertexts  $C_1, \dots, C_n$ ; while  $\bar{R} := (R_1, \dots, R_n)$  is the instruction about how to reward each answer. After computing  $\bar{R}$  and  $\pi_{reward}$ ,  $R$  puts them into a blockchain transaction, and still use her one-task-only blockchain address  $\alpha_R$  to send the transaction to  $\mathcal{C}$  (by using a pointer to  $\alpha_{\mathcal{C}}$ ). This finishes the *outsource-then-prove* methodology.

Once a newly proposed block contains the reward instruction  $\bar{R}$  and its attestation  $\pi_{reward}$ , the contract  $\mathcal{C}$  first checks that they are indeed sent from  $\alpha_R$  (by verifying the digital signature of the underlying blockchain transaction). Then it leverages SNARK's Verifier algorithm to verify the proof  $\pi_{reward}$  regarding the correctness of  $\bar{R}$ . If the verification passes, it transfers each amount  $R_i$  to each account  $\alpha_i$ , and refunds the remaining balance to  $\alpha_R$ . Otherwise, pause. If receiving no valid instruction after a predefined time (in unit of block), the contract simply transfers  $\tau/n$  to each  $\alpha_i$  as part of the policy  $R$ .

### C. Analysis of the protocol

**Correctness and efficiency.** It is clear to see that the requester will obtain data and the workers would receive the right amount of payments. If they all follow the protocol, under the conditions that (i) the blockchain can be modeled as an ideal public ledger, (ii) the underlying zk-SNARK is of completeness, (iii) the public key encryption is correct, and (iv) common-prefix-linkable anonymous authentication satisfies correctness. Regarding *efficiency*, we note the *on-chain* computation (and storage which are two of the major obstacles for applying blockchain in general) is actually very light, as the contract essentially only carries a verification step. Thanks to zk-SNARK, the verification can be efficiently executed by checking only a few pairing equalities; moreover, the special library can be dedicatedly optimized in various ways [49].

**Security analysis (sketch).** Here we briefly discuss the security of ZebraLancer protocol. The underlying primitives including ours are well abstracted, which allows us to argue in a modular way.

Regarding the *data confidentiality* of answers, all related public transcripts are simply the ciphertexts  $C_1, \dots, C_n$ , and the zk-SNARK proof  $\pi$ . The ciphertexts are easily simulatable according to the semantic security of the public key encryption, and the proof  $\pi$  can also be simulated without seeing the secret witness because of the *zero-knowledge* property.

Regarding the *anonymity*, an adversary has two ways to break it: (i) link a worker/requester through his blockchain addresses; (ii) link answers/tasks of a worker/requester through his authenticating attestations. The first case is trivial, simply

because every worker/requester will interact with each task by a randomly generated one-task-only blockchain address (and the corresponding public key.). The second case is more involved, but the anonymity of workers and requesters can be derived through the anonymity of the common-prefix-linkable anonymous authentication scheme.

Regarding the *security against a malicious requester*, a malicious requester has three chances to gain advantage: (i) deny the policy announced in Publish phase; (ii) cheat in Reward phase; (iii) submit answers to intentionally downgrade others in AnswerCollection phase. The first threat is prevented because the smart contract is public, and the requester cannot deny it once it is posted in the immutable blockchain. The second threat is prohibited by the soundness of the underlying SNARK, since any incorrect instruction passing the verification in the smart contract, directly violates the proof-of-knowledge of zk-SNARK. The last threat is simply handled the unforgeability and common-prefix-linkability of our common-prefix-linkable anonymous authentication scheme.

*Security against malicious workers* is straightforward, the only ways that malicious workers can cheat are: (i) submitting more than one answers in AnswerCollection phase; (ii) sending the contract a fake instruction in the name of requester in Reward phase; (iii) altering the policy specified in the contract. The first threat is simply handled by the common-prefix-linkability and unforgeability of common-prefix-linkable anonymous authentication. The second threat can be approached by predicting others' answers, and it is prevented due to the semantical security of public key encryption. The third threat is simply handled by the security of digital signatures. The last issue is trivial, because the blockchain security ensures the announced policy is immutable.

**Theorem 2.** *Our protocol satisfies the data confidentiality, the anonymity, and the security against a malicious requester and workers, if all the underlying cryptographic primitives and the blockchain platform satisfy the corresponding securities.*

## VI. ZEBRALANCER: IMPLEMENTATION OF THE SYSTEM AND EXPERIMENTAL EVALUATION

We implement the protocol of ZebraLancer atop Ethereum, and instantiate a series of typical image annotation tasks [10] with using it. Furthermore, we conduct experiments of these tasks in an Ethereum test net to evaluate the applicability.

**System in a nutshell.** As shown in Fig.3, the decentralized application (DApp) of our system is composed of an on-chain part and an off-chain part. The on-chain part consists of crowdsourcing task contracts and an interface contract of the registration authority (RA). The RA's contract simply posits the system's master public key as a common knowledge stored in the blockchain. The off-chain part consists of requester clients and worker clients. These clients can be blockchain clients wrapped with functionalities required by our system. Specifically, a client of requester should codify a specific task with a given incentive mechanism and announces it as a smart contract. Note that we, as the designers of the DApp, can

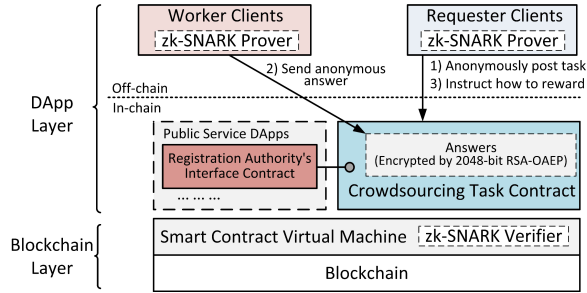


Fig. 3. The system-level view of ZebraLancer. Our DApp layer can be built on top of an existing blockchain, e.g. the Ethereum Byzantium release [41].

provide contract templates to requesters for easier instantiation of incentive mechanisms, c.f. [50]. The clients further need an integrated zk-SNARK prover to produce the anonymous authenticating attestations; moreover, a requester client should also leverage SNARK prover to generate proofs attesting the correct execution of incentive policies.

**Implementation challenges.** The main challenge of deploying smart contracts in general is that they can only support very light on-chain operations for both computing and storing.<sup>12</sup> Our protocol actually has taken this into consideration. In particular, our on-chain computation only consists SNARK verifications, while the heavy computation of SNARK proofs are all done off the blockchain. Even still, building an efficient privacy-preserving DApp compatible with existing blockchain platform such as Ethereum is not straightforward. For instance, in order to allow smart contracts to call a zk-SNARK verification library, a contract of this library should be thrown into a block, but this library is a general purpose tool that can be too complex to be executed in the smart contract runtime environment, e.g. Ethereum Virtual Machine (EVM). Alternatively, we modify the the runtime environment of smart contracts (i.e. Ethereum Homestead release), so that an optimized zk-SNARK verification library [49] is embedded in it as a primitive operation. The modified client is written in Java 1.8 with Spring framework, and is available at [github.com/maxilbert/ethereumj](https://github.com/maxilbert/ethereumj).

We remark that Ethereum project recently integrated some new cryptographic primitives into EVM to enable SNARK verification as well [41], which ensures our DApp can essentially inherit all Ethereum users to maintain the blockchain infrastructure to govern the faithful execution of the smart contracts in our protocol.

**Establishments of zk-SNARKs (off-line).** As the feasibility

<sup>12</sup>We remark the communication overhead is not a serious worry, because: (i) a blockchain network such as Ethereum does not require fully meshed connections, i.e. requesters and workers only connect a constant number of Ethereum peers, and to deliver one more message via the blockchain only brings about a constant cost of communication for each peer; (ii) if necessary, requesters and workers can even run on top of so-called light-weight nodes, which eventually allows them receive and send messages only related to crowdsourcing tasks; (iii) even if there is a trusted arbiter facilitating incentive mechanisms, the only saving in communication is just an instruction about how to reward answers (and its attestation).

of ZebraLancer highly depends on the tininess of SNARK proofs and the efficiency of SNARK verifications, it becomes critical to establish necessary zk-SNARKs off-line. As formally discussed, our anonymous authentication scheme and incentive mechanisms can be stated as well-defined deterministic constraint relationships. We first translate these mathematical statements into their corresponding boolean circuit satisfiability representations. Furthermore, we establish zk-SNARK for each boolean circuit, such that all required public parameters are generated. All the above steps are done off-line, as they are executed for only once when the system is launched.

**An image annotation crowdsourcing task.** To showcase the usability of our system, we implement a concrete crowdsourcing task of image annotation [10]. The task is to solicit labels for an image which can later be used to train a learning machine. The task requests  $n$  answers from  $n$  workers, and can be considered as a multi-choice problem. Majority voting is used to estimate the “truth”. An answer is seen as “correct”, if it equals to the “truth”. The reward amount of a worker is  $\tau/n$  if he answers correctly, otherwise, he receives nothing. In our terminology, the reward  $R_i := R(A_i; A_1, \dots, A_n, \tau) = \tau/n$ , if  $A_i$  equals the majority; otherwise,  $R_i = 0$ . Following [10], we implement and deploy 5 contracts in the test net to collect 3 answers, 5 answers, 7 answers, 9 answers and 11 answers from anonymous-yet-accountable workers, respectively.

The smart contracts are written in Solidity, a high-level scripting language translatable to smart contracts of Ethereum. We also modify Solidity compiler, such that a programmer can write a contract involving zk-SNARK verifications at high-level. We instantiate the encryption to be RSA-OAEP-2048, the DApp-layer hash function to be SHA-256, and the DApp-layer digital signature to be RSA signature. Moreover, for zk-SNARK, we choose the construction of *libsnark* from [49]. We deploy a test network consisting of four PCs: three PCs are equipped with Intel Xeon E3-1220V2 CPU (PC-As), and the other one is equipped with Intel i7-4790 CPU (PC-B); all PCs have 16 GB main memory and have Ubuntu 14.04 LTS installed. In the test net, a PC-A and a PC-B play the role of miners, and the other two PC-As only validate blocks (i.e. full nodes that do not mine). One full node plays the role of the requester, and anonymously publishes crowdsourcing tasks to the blockchain; and the other full node mimics workers, and sends each anonymously authenticated answer from a different blockchain address. Miners are only responsible to maintain the test net and do not involve in tasks.

**Performance evaluation.** As the main bottleneck is the on-chain computation of the smart contract, we first measure the time cost and the spatial cost of miners, regarding the executions of zk-SNARK verifications used in the above annotation tasks. These zk-SNARKs are established for common-prefix-linkable anonymous authentications and incentive mechanisms, respectively. The results of time cost are listed in Table I. It is clear that zk-SNARK verifications in our system can be efficiently executed in respect of verification time. Moreover, our experiment results also reveal that the spatial cost of zk-

TABLE I  
EXECUTION TIME OF IN-CONTRACT ZK-SNARK VERIFICATIONS.

Verification for	Operands Length			Time@ PC-A	Time@ PC-B
	Proof	Key	Inputs		
Anonymous authentication	729B	1.2KB	1.5KB	10.9ms	6.2ms
Majority (3-Worker)	729B	16.0KB	3.4KB	15.5ms	9.1ms
Majority (5-Worker)	730B	21.6KB	4.7KB	16.3ms	9.8ms
Majority (7-Worker)	731B	27.3KB	6.0KB	17.0ms	10.3ms
Majority (9-Worker)	729B	32.9KB	7.3KB	17.5ms	12.1ms
Majority (11-Worker)	730B	38.6KB	8.6KB	17.9ms	13.1ms

SNARK verifications is constant and tiny at both types of PCs (exactly 17MB main memory). Also, Table I indicates that the required on-chain storage for the task contracts (i.e. SNARK proofs, SNARK verification keys, and ciphertexts of answers) is on an acceptable order of kilobyte<sup>13</sup>. Therefore, the on-chain performance of the system can be clearly practical, regarding time and space costs of blockchain peers.

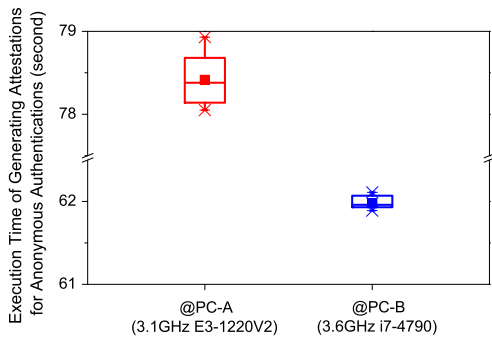


Fig. 4. The time of generating common-prefix-linkable anonymous authentications in two PCs. The box plot is derived from 12 different experiments.

We also consider the cost of anonymity, if one uses the common-prefix-linkable anonymous authentication. We measure the running time of generating the authenticating attestations at PCs. As shown in Fig.4, our experiment results clarify that about 78 seconds are spent on generating an anonymous attestation with using PC-A (3.1 GHz CPU). In PC-B (3.6GHz CPU), the running time can be shortened to about 62 seconds. Those are not ideal, but acceptable by the anonymity-sensitive workers. We remark that our protocol can be trivially extended to support non-anonymous mode, in case that one gives up the anonymity privilege: s/he can generate a public-private key pair (for digital signatures), and then registers the public key at RA to receive a certificate bound to the public key; to authenticate, s/he can simply show the certified public key, the certificate, along with a message properly signed under

<sup>13</sup>We note that there are many alternatives to minimize the on-chain storage in the implementations, e.g. to use off-chain storages [51, 52]. So, when a requester is publishing a data-intensive crowdsourcing task (e.g. image labeling, voice captioning) via the blockchain, it is not necessary for her to store all the data (e.g. images, audios) in the chain. These trivial techniques are beyond the scope of this paper, in which we only focus on the technical feasibility of decentralizing crowdsourcing instead.

the corresponding secret key, which essentially costs nearly nothing regarding the computational efficiency.

## VII. CONCLUSION & OPEN PROBLEMS

ZebraLancer can facilitate the fair exchange between the crowd-shared data and their corresponding rewards, without the involvement of any third-party arbiter. Moreover, it shows the practicability to resolve two natural tensions in decentralized crowdsourcing atop open blockchain: one between data confidentiality and transparency, and the other one between participants' anonymity and their accountability.

Along the way, we put forth a new anonymous authentication scheme. Besides the strong anonymity that cannot be revoked by any authority, it also supports a delicate linkability only for messages that share the common prefix and are authenticated by the same user. A concrete construction of the scheme is proposed, and it shows the compatibility to real-world blockchain infrastructure. We envision the scheme could be of independent interests. Also, we develop a general *outsource-then-prove* technique to use smart contracts in a privacy-preserving way. This technique can further extend the scope of applications atop some existing privacy-preserving blockchain infrastructures such as [39–41].

**Open questions.** Since this work is the first attempt of decentralizing crowdsourcing system atop the real-world blockchain in a privacy-preserving way, the area remains largely unexplored. Here we name a few open questions. First, there are many incentive mechanisms using reputation systems, can we further extend our implementations to support those incentives? Second, as the current smart contract technology is at an infant stage and can only allow tiny on-chain storage, can we further optimize our implementations with using off-chain storage [51, 52] or information oracle [53] to assist more large-scale tasks, e.g. to collect annotations for millions of images (i.e. the scale of ImageNet dataset)? Third, our anonymous protocol currently either relies on the underlying blockchain to support anonymous transaction, or requires workers/requesters use one-time blockchain account to submit data and receive reward. Can we design a (DApp-layer) protocol to solve the drawbacks? Last but not least, as all existing studies and platforms rely on trusted registration authorities (RAs) to establish identities for crowdsourcing, it is tempting to further remove RAs by leveraging blockchain in a more complex way.

## ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers of this paper for their valuable comments.

## REFERENCES

- [1] J. Deng, W. Dong, R. Socher *et al.*, “Imagenet: A large-scale hierarchical image database,” in *Proc. IEEE CVPR 2009*, pp. 248–255.
- [2] Mechanical Turk. [Online]. Available: <https://www.mturk.com/mturk/>
- [3] R. Ganti, F. Ye, and H. Lei, “Mobile crowdsensing: current state and future challenges,” *IEEE Commun. Mag.*, vol. 49, no. 11, pp. 32–39, 2011.
- [4] Waze. [Online]. Available: <https://status.waze.com>

- [5] Y. Wen, J. Shi, Q. Zhang *et al.*, "Quality-driven auction-based incentive mechanism for mobile crowd sensing," *IEEE Trans. Veh. Technol.*, vol. 64, no. 9, pp. 4203–4214, 2015.
- [6] Y. Zhang and M. Van der Schaar, "Reputation-based incentive protocols in crowdsourcing applications," in *Proc. IEEE INFOCOM 2012*, pp. 2140–2148.
- [7] D. Zhao, X.-Y. Li, and H. Ma, "How to crowdsource tasks truthfully without sacrificing utility: Online incentive mechanisms with budget constraint," in *Proc. IEEE INFOCOM 2014*, pp. 1213–1221.
- [8] D. Yang, G. Xue, X. Fang, and J. Tang, "Crowdsourcing to smartphones: Incentive mechanism design for mobile phone sensing," in *Proc. ACM MobiCom 2012*, pp. 173–184.
- [9] D. Peng, F. Wu, and G. Chen, "Pay as how well you do: A quality based incentive mechanism for crowdsensing," in *Proc. ACM MobiHoc 2015*, pp. 177–186.
- [10] N. Shah and D. Zhou, "Double or Nothing: multiplicative Incentive Mechanisms for Crowdsourcing," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 5725–5776, 2016.
- [11] H. Jin, L. Su, D. Chen, K. Nahrstedt, and J. Xu, "Quality of information aware incentive mechanisms for mobile crowd sensing systems," in *Proc. ACM MobiHoc 2015*, pp. 167–176.
- [12] P. Muncaster. China's Internet wunderkind in the dock over alleged fraud. [Online]. Available: [http://www.theregister.co.uk/2012/07/03/qihoo\\_fraud\\_traffic\\_comscore](http://www.theregister.co.uk/2012/07/03/qihoo_fraud_traffic_comscore)
- [13] H. Wei. Alipay apologizes for leak of personal info. [Online]. Available: [http://www.chinadaily.com.cn/china/2014-01/07/content\\_17219203.htm](http://www.chinadaily.com.cn/china/2014-01/07/content_17219203.htm)
- [14] H. Kelly. Apple account hack raises concern about cloud storage. [Online]. Available: <http://www.cnn.com/2012/08/06/tech/mobile/icloud-security-hack/>
- [15] B. McInnis, D. Cosley, C. Nam, and G. Leshed, "Taking a HIT: Designing around rejection, mistrust, risk, and workers' experiences in Amazon Mechanical Turk," in *Proc. ACM CHI 2016*, pp. 2271–2282.
- [16] K. B. Mike Isaac and S. Frenkel. Uber hid 2016 breach, paying hackers to delete stolen data. [Online]. Available: <https://www.nytimes.com/2017/11/21/technology/uber-hack.html>
- [17] N. Szabo, "Formalizing and securing relationships on public networks," *First Monday*, vol. 2, no. 9, 1997.
- [18] A. Kate. Blockchain privacy: Challenges, solutions, and unresolved issues. [Online]. Available: [http://www.isical.ac.in/~rcbse/blockchain2017/lecture/Kate\\_Slides.pdf](http://www.isical.ac.in/~rcbse/blockchain2017/lecture/Kate_Slides.pdf)
- [19] K. Yang, K. Zhang, J. Ren, and X. Shen, "Security and privacy in mobile crowdsourcing networks: challenges and opportunities," *IEEE Commun. Mag.*, vol. 53, no. 8, pp. 75–81, 2015.
- [20] F. Buccafurri, G. Lax, S. Nicolazzo, and A. Nocera, "Tweetchain: An alternative to blockchain for crowd-based applications," in *International Conference on Web Engineering 2017*. Springer, pp. 386–393.
- [21] C. Tanas, S. Delgado-Segura, and J. Herrera-Joancomart, "An integrated reward and reputation mechanism for MCS preserving users privacy," in *International Workshop on Data Privacy Management and Security Assurance 2015*, pp. 83–99.
- [22] M. Li, J. Weng, A. Yang, and W. Lu, "CrowdBC: A Blockchain-based Decentralized Framework for Crowdsourcing," 2017. [Online]. Available: <http://eprint.iacr.org/2017/444.pdf>
- [23] N. Salehi, L. C. Irani, M. S. Bernstein *et al.*, "We are dynamo: Overcoming stalling and friction in collective action for crowd workers," in *Proc. ACM CHI 2015*, pp. 1621–1630.
- [24] S. Gisdakis, T. Giannetos, and P. Papadimitratos, "Security, privacy, and incentive provision for mobile crowd sensing systems," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 839–853, 2016.
- [25] A. Muehleemann, "Sentiment protocol: A decentralized protocol leveraging crowd sourced wisdom," 2017. [Online]. Available: <https://eprint.iacr.org/2017/1133.pdf>
- [26] Q. Li and G. Cao, "Providing efficient privacy-aware incentives for mobile sensing," in *Proc. IEEE ICDCS 2014*, pp. 208–217.
- [27] S. Rahaman, L. Cheng, D. D. Yao, H. Li, and J.-M. J. Park, "Provably secure anonymous-yet-accountable crowdsensing with scalable sublinear revocation," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 384–403, 2017.
- [28] D. Chaum, "Blind signatures for untraceable payments," in *Advances in Cryptology 1983*, pp. 199–203.
- [29] D. Chaum, A. Fiat, and M. Naor, "Untraceable electronic cash," in *Advances in Cryptology – CRYPTO 1988*, pp. 319–327.
- [30] J. Camenisch and A. Lysyanskaya, "An efficient system for non-transferable anonymous credentials with optional anonymity revocation," in *Proc. EUROCRYPT 2011*, pp. 93–118.
- [31] S. Xu and M. Yung, "K-anonymous secret handshakes with reusable credentials," in *Proc. ACM CCS 2004*, pp. 158–167.
- [32] I. Teranishi, J. Furukawa, and K. Sako, "K-times anonymous authentication," in *International Conference on the Theory and Application of Cryptology and Information Security (Asiacrypt) 2004*, pp. 308–322.
- [33] J. Camenisch, S. Hohenberger, M. Kohlweiss, A. Lysyanskaya, and M. Meyerovich, "How to win the clonewars: efficient periodic n-times anonymous authentication," in *Proc. ACM CCS 2006*, pp. 201–210.
- [34] T. Nakanishi, T. Fujiwara, and H. Watanabe, "A linkable group signature and its application to secret voting," *Trans. of Information Processing Society of Japan*, vol. 40, no. 7, pp. 3085–3096, 1999.
- [35] J. K. Liu, V. K. Wei, and D. S. Wong, "Linkable spontaneous anonymous group signature for ad hoc groups," in *Australasian Conference on Information Security and Privacy 2004*, pp. 325–335.
- [36] M. H. Au, W. Susilo, and S.-M. Yiu, "Event-oriented k-times revocable-iff-linked group signatures," in *Australasian Conference on Information Security and Privacy 2006*, pp. 223–234.
- [37] A. Kiayias, H.-S. Zhou, and V. Zikas, "Fair and robust multi-party computation using a global transaction ledger," in *Proc. EUROCRYPT 2016*. Springer, pp. 705–734.
- [38] G. Zyskind, O. Nathan, and A. Pentland, "Enigma: Decentralized computation platform with guaranteed privacy," 2015. [Online]. Available: <https://arxiv.org/abs/1506.03471>
- [39] A. Kosba, A. Miller, E. Shi *et al.*, "Hawk: The blockchain model of cryptography and privacy-preserving smart contracts," in *Proc. IEEE S&P 2016*, pp. 839–858.
- [40] D. Hopwood, S. Bowe, T. Hornby, and N. Wilcox, "Zcash Protocol Specification." [Online]. Available: <https://github.com/zcash/zips/blob/master/protocol/protocol.pdf>
- [41] Ethereum Team. Byzantium HF Announcement. [Online]. Available: <https://blog.ethereum.org/2017/10/12/byzantium-hf-announcement/>
- [42] J. A. Garay, A. Kiayias, and N. Leonardos, "The bitcoin backbone protocol: Analysis and applications," in *Proc. EUROCRYPT 2015*. Springer, pp. 281–310.
- [43] R. Pass, L. Seeman, and A. Shelat, "Analysis of the blockchain protocol in asynchronous networks," in *Proc. EUROCRYPT 2017*. Springer, pp. 643–673.
- [44] V. Buterin, "A next-generation smart contract and decentralized application platform," 2014. [Online]. Available: <https://github.com/ethereum/wiki/wiki/White-Paper>
- [45] S. Saroiu and A. Wolman, "I am a sensor, and I approve this message," in *Proc. ACM HotMobile 2010*, pp. 37–42.
- [46] Y. Baba and H. Kashima, "Statistical quality estimation for general crowdsourcing tasks," in *Proc. ACM SIGKDD 2013*, pp. 554–562.
- [47] S. Goldwasser and S. Micali, "Probabilistic encryption & how to play mental poker keeping secret all partial information," in *Proc. ACM STOC 1982*, pp. 365–377.
- [48] E. Ben-Sasson, A. Chiesa, M. Green, E. Tromer, and M. Virza, "Secure sampling of public parameters for succinct zero knowledge proofs," in *Proc. IEEE S&P 2015*, pp. 287–304.
- [49] E. Ben-Sasson, A. Chiesa, D. Genkin *et al.*, "SNARKs for C: Verifying program executions succinctly and in zero knowledge," in *Advances in Cryptology – CRYPTO 2013*, pp. 90–108.
- [50] C. Frantz and M. Nowostawski, "From institutions to code: Towards automated generation of smart contracts," in *Proc. IEEE FAS\*W 2016*, pp. 210–215.
- [51] Swarm. [Online]. Available: <http://swarm-guide.readthedocs.io>
- [52] J. Benet, "IPFS: Content Addressed, Versioned, P2P File System." [Online]. Available: <https://github.com/ipfs/papers/raw/master/ipfs-cap2pfs/ipfs-p2p-file-system.pdf>
- [53] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi, "Town crier: An authenticated data feed for smart contracts," in *Proc. ACM CCS 2016*, pp. 270–282.