

Separ: Towards Regulating Future of Work Multi-Platform Crowdfunding Environments with Privacy Guarantees

Mohammad Javad Amiri
University of Pennsylvania
mjamiri@seas.upenn.edu

Joris Duguépéroux
Univ Rennes, CNRS, IRISA
joris.dugueperoux@irisa.fr

Tristan Allard
Univ Rennes, CNRS, IRISA
tristan.allard@irisa.fr

Divyakant Agrawal
University of California Santa Barbara
agrawal@cs.ucsb.edu

Amr El Abbadi
University of California Santa Barbara
amr@cs.ucsb.edu

ABSTRACT

Crowdfunding platforms provide the opportunity for diverse workers to execute tasks for different requesters. The popularity of the "gig" economy has given rise to independent platforms that provide competing and complementary services. Workers as well as requesters with specific tasks may need to work for or avail from the services of multiple platforms resulting in the rise of *multi-platform* crowdfunding systems. Recently, there has been increasing interest by governmental, legal and social institutions to enforce regulations, such as minimal and maximal work hours, on crowdfunding platforms. Platforms within multi-platform crowdfunding systems, therefore, need to collaborate to enforce cross-platform regulations. While collaborating to enforce global regulations requires the *transparent* sharing of information about tasks and their participants, the *privacy* of all participants needs to be preserved. In this paper, we propose an overall vision exploring the regulation, privacy, and architecture dimensions for the future of work multi-platform crowdfunding environments. We then present *Separ*, a multi-platform crowdfunding system that enforces a large sub-space of practical global regulations on a set of distributed independent platforms in a privacy-preserving manner. *Separ*, enforces *privacy* using *lightweight* and *anonymous tokens*, while *transparency* is achieved using fault-tolerant *blockchain ledgers* shared among multiple platforms. The privacy guarantees of *Separ* against covert adversaries are formalized and thoroughly demonstrated, while the experiments reveal the efficiency of *Separ* in terms of performance and scalability.

CCS CONCEPTS

• **Information systems** → **Crowdsourcing**; **Distributed database transactions**; • **Social and professional topics** → *Governmental regulations*; • **Security and privacy** → **Security services**.

KEYWORDS

Crowdfunding, Future of Work, Privacy, Regulation, Blockchain

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3449858>

ACM Reference Format:

Mohammad Javad Amiri, Joris Duguépéroux, Tristan Allard, Divyakant Agrawal, and Amr El Abbadi. 2021. *Separ: Towards Regulating Future of Work Multi-Platform Crowdfunding Environments with Privacy Guarantees*. In *Proceedings of the Web Conference 2021 (WWW '21), April 19–23, 2021, Ljubljana, Slovenia*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3442381.3449858>

1 INTRODUCTION

The rise of the "gig" or *platform economy* [17, 21] is reshaping work all around the world. Crowdsourcing platforms dedicated to work (also called *crowdfunding platforms* [10]) are online intermediaries between *requesters* and *workers*, where requesters propose *tasks* while *workers* propose skills and time. By providing requesters (resp. workers) 24/7 access to a worldwide workforce (resp. worldwide task market), crowdfunding platforms have grown in numbers, diversity, and adoption. Today, crowdworkers come from countries spread all over the world, and work on several, possibly competing, platforms [10]. The use of crowdfunding platforms is expected to continue growing [30], and in fact, they are envisioned as key technological components of the future of work [2, 11, 18].

Crowdfunding platforms, however, challenge national boundaries and weaken the formal relationships between the platforms, workers and task requesters. Guaranteeing the compliance of crowdfunding platforms with national or regional labor laws is hard¹ [30] despite the stringent need for regulating work. For example, the total work hours of a worker per week may not exceed 40 hours to follow *Fair Labor Standards Act*² (FLSA). In California, Assembly Bill 5 (AB5)³ entitles workers to greater labor protections, such as minimum wage laws, sick leave, and unemployment and workers' compensation benefits. AB5 is recently being challenged by *California Proposition 22*⁴, which also imposes its own set of regulations on minimal hours worked for health benefits. The global regulation of the work hours represents the *minimal* and *maximal* number of hours that participants, i.e., worker, requester, and platform, can spend on crowdfunding platforms. While legal tools are currently being investigated [29, 30], there is a stringent need for technical tools allowing official institutions to enforce regulations.

¹See, e.g., the *Otey v CrowdfLower* class action against a famous microtask platform for "substandard wages and oppressive working hours" (casetext.com/case/otey-v-crowdfLower-1).

²<https://www.dol.gov/agencies/whd/flsa>

³https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201920200AB5

⁴[https://ballotpedia.org/California_Proposition_22_App-Based_Drivers_as_Contractors_and_Labor_Policies_Initiative_\(2020\)](https://ballotpedia.org/California_Proposition_22_App-Based_Drivers_as_Contractors_and_Labor_Policies_Initiative_(2020))

Some platforms have already started implementing self-defined *local regulations*. For example, Uber⁵ and Lyft⁶ force drivers to rest at least 6 hours for every 12 hours in driver mode, or Wirk.io⁷ prevent micro-workers from earning more than €3000 per year. However, since workers and requesters can simply switch platforms when a local limit is reached, no global cross-platform regulation can be enforced. Moreover, participants in a crowdworking task may also behave maliciously or act as adversaries, e.g., violate the privacy of participants and the regulations for their benefits. The privacy of participants and the global consistency of regulations are critical for future of work crowdworking environments [11].

Most current crowdworking platforms are independent of each other. However, the emergence of complex tasks that may need multiple contributions from possibly different platforms, on one hand, and more importantly, the enforcement of legal regulations, on the other hand, highlight the need for collaboration between crowdworking platforms, resulting in *multi-platform crowdworking systems*. For example, many drivers work for both Uber and Lyft concurrently⁸, while a requester may also request multiple rides from both Uber and Lyft in parallel. The observation holds for microtask platforms [10] as well where a requester who has registered on *Amazon Mechanical Turk* and *Prolific* might need hundreds of contributions for a single microtask and accept these contributions from workers regardless of the platforms the microtasks are performed on. Since workers from different platforms might want to perform contributions, the system needs to establish consensus among different microtask platforms to assign workers and provide the specified number of contributions while ensuring minimal and maximum hourly regulations on participants without revealing any private information to the competing platforms.

Our overall vision for multi-platform crowdworking environments needs to address three main dimensions: *regulations*, *privacy*, and *architecture*. First, a multi-platform crowdworking system must clearly define the *types of regulations* supported in terms of the *complexity* of the regulation (e.g., a simple or a chain of interactions) as well as the *aggregate* requirements of the regulation (e.g., no aggregation or SUM aggregations). For example, California Proposition 22, which states "if a driver works at least 25 hours per week, companies (i.e., platforms) require to provide healthcare subsidies ...", is a simple, SUM-aggregate regulation. Second, the threat model (e.g., *honest-but-curious*, *covert*, *fully malicious*) as well as the privacy guarantees provided to each entity must be clearly specified. Finally, from an architecture design point of view, any multi-platform crowdworking environment consists of two critical components: *regulation management* that models and enforces regulations and *global state management* that manages the global states of all participants as well as tasks. Each of these components can be implemented in a centralized or a decentralized approach.

In this paper, we present our overall vision for future of work multi-platform crowdworking environments together with *Separ*, a possible instance of a precise point in the design space of multi-platform crowdworking systems. *Separ* results from choices, guided

by cutting-edge real-life regulation proposals, on all three regulations, privacy, and architecture dimensions of the design space. First, *Separ* focuses on managing lower and upper bounds on aggregate-based regulations. Second, *Separ* considers that any participant may act as a *covert adversary* [9] and ensures that no participant obtains or infers any information about a crowdworking task beyond what is strictly needed for accomplishing its local task and for the distributed enforcement of regulations. Finally, in *Separ*, we opted for simplicity and rapid prototyping by using a centralized (but fault-tolerant) component to manage regulations. However, we use a decentralized component to manage the global state of the system. In particular, *Separ* uses a *permissioned blockchain* as an underlying infrastructure that is shared among all involved platforms.

The complexity of the conjunction of the required properties makes the problem non-trivial. First, the regulations need to be expressed in a *simple* and *non-ambiguous* manner. Second, while enforcing regulations over multiple crowdworking platforms requires the global state of the system to be *transparent*, the *privacy* of participants needs to be preserved, hence *Separ* needs to reconcile *transparency* with *privacy*. Finally, the decentralized management of the global state among a distributed set of crowdworking platforms requires distributed consensus protocols.

Separ is a *two-level* solution consisting of a *privacy-preserving token-based system* (i.e., the *application level*) on top of a *blockchain ledger* shared across platforms (i.e., the *infrastructure level*). First, at the application level, global regulations are modeled using *light-weight* and *anonymous tokens* distributed to workers, platforms, and requesters. The information shared among participants is limited to the minimum necessary for performing the tasks against *adversarial participants acting as covert adversaries*. Second, at the infrastructure level, the blockchain ledger allows *Separ* to provide transparency across platforms. Nonetheless, for the sake of privacy and to improve performance, the ledger is *not maintained* by any platform and each platform maintains only a view of the ledger. We then design a suite of consensus protocols for coping with the concurrency issues inherent to a multi-platform context. Salient features of *Separ* include the simplicity of its building blocks (e.g., usual signature schemes) and its compatibility with today's platforms (e.g., it does not jeopardize their privacy requirements of requesters and workers for enforcing the regulation).

In a nutshell, the contributions of this paper are as follows:

- (1) A vision for the design space of regulation systems for future of work multi-platform crowdworking environments. In particular, we (1) express regulations as SQL constraints and categorize them according to their SQL expression, (2) propose a formal privacy model for multi-platform regulated crowdworking systems based on the well-known simulatability paradigm, and (3) discuss critical components of their architecture.
- (2) *Separ*, a two-level *privacy-preserving transparent* multi-platform proof-of-concept crowdworking system that enforces a precise point in the design space guided by cutting edge practical regulations currently discussed by societal organizations, legal entities and enterprises. *Separ* uses a simple language for expressing *global regulations* and mapping them to SQL constraints to ensure semantic clarity. It

⁵<https://www.uber.com/en-ZA/blog/driving-hours-limit/>

⁶<https://help.lyft.com/hc/articles/115012926787-Taking-breaks-and-time-limits>

⁷<https://www.wirk.io/50k-freelances-en-france/>

⁸For example, rideshareapps.com provides tutorials to help drivers manage apps to optimize their earnings <https://rideshareapps.com/drive-for-uber-and-lyft-at-the-same-time/>.

ensures privacy using *lightweight and anonymous tokens*, while transparency is achieved using a *blockchain* shared across platforms for both crash-only and Byzantine nodes.

- (3) A formal security analysis of Separ and thorough experimental evaluations.

The paper is organized as follows. The overall vision of the design space for regulation systems is presented in Section 2. Section 3 presents the Separ model within the design space. The application level including the implementation of regulations in Separ is discussed in Section 4. The infrastructure level consisting of the blockchain ledger and consensus protocols is presented in Section 5. Section 6 details an experimental evaluation, Section 7 discusses the related work, and Section 8 concludes the paper.

2 DESIGN SPACE OF REGULATION SYSTEMS

Designing a system for regulating crowdfunding platforms mandates three main choices. First, given the large variety of regulations that apply to working environments, any given regulation system must clearly define the *types of regulations* it supports. Second, a crowdfunding environment involves a set of distributed entities (platforms, workers, requesters) that cannot be fully trusted. Hence, the *privacy guarantees* provided to each entity must be rigorously stated, e.g., computational guarantees. Third, the distributed nature of a crowdfunding environment requires various *architectural choices* that range from fully decentralized approaches, e.g., peer-to-peer architectures, to the traditional centralized architecture. In this section, we first define the crowdfunding environment that we consider and then characterize the design space for crowdfunding regulation systems.

2.1 Crowdfunding Environment

A *crowdfunding environment* consists of a set of workers \mathcal{W} interacting with a set of requesters \mathcal{R} through a set of competing platforms \mathcal{P} . We refer to the workers, platforms, and requesters of a crowdfunding environment as *participants*. Each worker $w \in \mathcal{W}$ (1) registers to one or more platforms $\mathcal{P}_w \subset \mathcal{P}$ according to her preferences and, through the latter, (2) accesses the set of tasks available on \mathcal{P}_w , (3) submits a *contribution* to the platform $p \in \mathcal{P}_w$ she elects, and (4) obtains a *reward* for her work. On the other hand, each requester $r \in \mathcal{R}$ similarly (1) registers to one or more platforms $\mathcal{P}_r \subset \mathcal{P}$, (2) issues a *submission* which contains her tasks \mathcal{T}_r to one or more platforms $p \in \mathcal{P}_r$, (3) receives the contributions of each worker w registered to $\mathcal{P}_r \cap \mathcal{P}_w$ having elected a task $t \in \mathcal{T}_r$, and (4) launches the distribution of rewards. Platforms are thus in charge of facilitating the interactions between workers and requesters. A *crowdfunding process* π connects three participants – a workers w , a platform p , and a requester r – and aims to facilitate the execution of a task $t \in \mathcal{T}_r$ through platform p via a worker w . For simplicity and without loss of generality, we assume that each process corresponds to a time unit of work (e.g., 1 hour).

2.2 Types of Regulation

The space of possible regulations can be structured based on two orthogonal dimensions: the *complexity* of the regulation, e.g., constraints on a single process versus constraints that apply to multiple processes that are transitively related, and the *aggregate* nature of

the regulation, e.g., no aggregation versus restrictions involving SUM aggregate.

In order to give a clear semantics to each dimension and to rigorously map regulations to points in this space, we express regulations by SQL constraints over a *universal virtual table* storing information about all crowdfunding processes having been performed. We denote this table by U-TABLE and emphasize that it is *virtual* (we use it only for clarifying the various types of regulations, it is never instantiated). The attributes of U-TABLE refer to meta-data about the interactions (i.e., at least the worker, requester, and platform involved in a process, and also possibly additional metadata such as begin and end timestamps) and information about the contents (e.g., the time estimate for the task⁹ - the TIMECOST attribute below -, the proposed wage, the worker's contribution). For simplicity, we focus below on the attributes of the U-TABLE relevant to our illustrative examples: (1) TS_BEGIN, TS_END, WORKER, PLATFORM, and REQUESTER, and (2) TIMECOST, WAGE, and CONTRIBUTION. Finally, a regulation is simply a Boolean SQL expression nested within the usual CHECK clause: ALTER TABLE U ADD CONSTRAINT r CHECK (...).

The complexity and aggregate dimensions of a regulation can both be deduced from its SQL expression. The complexity is given by the presence of *join* operations while the aggregate dimension is given by the presence of *aggregate* function(s), possibly with GROUP BY and HAVING clause(s). For simplicity, we consider a coarse grain characterization of these two dimensions. A regulation is *simple* if there is no join and *complex* otherwise. A regulation is *row-only* if it does not involve any aggregate function, *aggregate-only* if it involves only comparison(s) over aggregate(s), and *mixed* if it involves comparisons over rows and aggregates.

We illustrate the possible types of regulations based on simple examples extracted from real-life crowdfunding regulations or from real-life proposals of regulation. First, we consider a regulation r_1 requiring the wage proposed by each task to be at least a given amount θ . This regulation is similar to CA Proposition 22. It illustrates the simple, row-only type of regulation.

```
ALTER TABLE U-TABLE ADD CONSTRAINT r1 CHECK (
    NOT EXISTS (
        SELECT * FROM U
        WHERE TIMECOST ≤ θ
    ) );
```

Second, we consider a regulation r_2 requiring each worker to work at most a given amount of time units θ per time period ρ . It illustrates a simple, mixed with SUM-aggregate regulation. It is similar to the regulation of *wirk.io* platform that limits the gains of any worker on the platform to €3000 per year. The following SQL constraint expresses r_2 , assuming that *current_time()* gives the current time in the same unit as the period ρ .

```
ALTER TABLE U-TABLE ADD CONSTRAINT r2 CHECK (
    NOT EXISTS (
        SELECT * FROM U
        WHERE WORKER=w AND current_time()-TS_BEGIN ≤ ρ
        GROUP BY WORKER
        HAVING SUM(TIMECOST) ≥ θ
    ) );
```

⁹Future regulation systems will need to design technical means to guarantee the reliability of the time estimates for tasks (e.g., privacy-preserving feedback systems from workers, automatic time estimation by analyzing task descriptions).

Finally, we complete our illustrations by considering a regulation r_3 that prevents any worker to submit two similar contributions to the same requester (even through two distinct platforms). We assume that the `sim` function computes the similarity between two contributions and that θ is the threshold above which we consider that two contributions are too similar. This illustrates the complex, row-only type of regulation.

```
ALTER TABLE U-TABLE ADD CONSTRAINT r3 CHECK (
  NOT EXISTS (
    SELECT *
    FROM U U1 JOIN U U2 ON
      U1.WORKER=U2.WORKER
    AND U1.REQUESTER=U2.REQUESTER
    AND sim(U1.CONTRIBUTION, U2.CONTRIBUTION) ≥ θ
  ) );
```

Although most regulations must always hold, e.g., a lower than constraint on a simple, Mixed with SUM-aggregate regulation, or a complex, row-only regulation, some regulations, inherently, *cannot* always hold. Similar to deferred SQL constraints, they must only hold after a given time period. For example, a periodic *greater-than* constraint on a simple, Mixed regulation cannot hold initially, but must hold at the end of a given period, e.g., CA Proposition 22 that requires a worker to work at least 25 hours per week to qualify for healthcare subsidies. We call *enforceable* the regulations that must always hold and *verifiable* the regulations that eventually hold. The verifiable/enforceable property of a regulation is only due to its nature not to its implementation (but it impacts it directly). Future crowdworking regulation systems need to determine the enforceable/verifiable properties of the regulations they support.

2.3 Threat Model and Privacy Model

Crowdworking environments are open environments that connect possibly adversarial participants. Any regulation system must clearly specify both the threat model (e.g., *honest-but-curious*, *covert*, *fully malicious*) and the privacy model. While the threat model only depends on the underlying system, the privacy model can be based on a common formal requirement parameterized for each system by the leaks it tolerates. The possible leakage ranges from no information leaked to any participant (similar to usual *secure multi-party computation* algorithms) to full disclosure (of all the information discussed above) to all participants (e.g., in current crowdworking platforms, the underlying system often requires full disclosure to the platform). The disclosures tolerated by future regulation systems will fall within this range.

We formalize below a common privacy model based on the well-known simulatability paradigm often used by secure multi-party computation algorithms. The proposed model guarantees that nothing leaks, with computational guarantees¹⁰, except the *pluggable* system-dependent tolerated *disclosures*.

Consider a crowdworking process π between worker w , platform p , and requester r for solving a task t . The task may have been sent to several platforms and, depending on the underlying crowdworking platforms, might have been accessed by several workers

¹⁰The majority of data protection techniques used in real-life are based on encryption schemes that provide guarantees against computationally-bounded adversaries but the model can be easily adapted to information theoretic attackers.

before being picked and solved. The information generated by the execution of π consists, similarly to the information captured by the U-TABLE, of information about interactions (e.g., at least w , r , and p , the participants directly involved in π) and information about contents (e.g., description of t , contribution, proposed wage). We propose to define the sets of disclosure according to the involvement of a participant in π . Indeed, the platforms not involved in π (i.e., different from p) but that received t may need to learn that t has been completed (e.g., to manage their local copy of the task). Similarly, workers who are not involved in task t might still need to know that it has been executed, while potentially preserving the privacy of worker w who executed the task. The three sets of disclosures defined below cover all these cases¹¹. Future regulation systems have to specify clearly the content of each disclosure set.

- Disclosures to the participants that are **not involved** in π and that have **not received** task t from requester r : δ_{-R-I}^π
- Disclosures to the platforms and workers that have **received** the task t from r but that are **not involved** in π : δ_{R-I}^π
- Disclosures to the participants that are directly **involved** in π (and have thus **received** task t): δ_{RI}^π

Definition 2.1. Let Π be a set of crowdworking processes executed by ζ a regulation system over a set of participants. We say that ζ is δ^Π -Private if, for all $\pi \in \Pi$, for all computationally-bounded adversaries A , the sets of disclosures $(\delta_{-R-I}^\pi, \delta_{R-I}^\pi, \delta_{RI}^\pi)$, assuming arbitrary background knowledge $\chi \in \{0, 1\}^*$, the distribution representing the adversarial knowledge over the input dataset in the real setting is *computationally indistinguishable* from the distribution representing the adversarial knowledge in an ideal setting in which a trusted third party cp executes the crowdworking process π of ζ : $\text{REAL}_{\zeta, A(\chi, \delta_i^\pi)}(\mathcal{W}, \mathcal{P}, \mathcal{R}, \mathcal{T}) \stackrel{c}{\equiv} \text{IDEAL}_{cp, A(\chi, \delta_i^\pi)}(\mathcal{W}, \mathcal{P}, \mathcal{R}, \mathcal{T})$ where $i \in \{-R-I, R-I, RI\}$, and REAL denotes the adversarial knowledge in the real and IDEAL its counterpart in the ideal setting.

2.4 Architecture

The architecture of a multi-platform crowdworking system consists of two main building blocks: *Regulation Management* and *Global State Management*. Regulation management models the regulations among the participants and ensures that the modeled regulations are adhered to by all participants. Global state management, on the other hand, stores the global states of the system including all information related to the participants and tasks. To implement these two components, similar to all distributed systems, either a centralized or a decentralized approach can be employed. Centralization is typically easier to rapid prototype, while requiring additional technologies to ensure fault-tolerance, privacy, and trustworthiness. A decentralized approach, on the other hand, is more compatible with the multi-platform settings, while resulting in more overhead and complex communication protocols among entities.

The verifiable/enforceable property of a regulation is another architectural challenge. While enforceable regulations could be enforced within the multi-platform system, verifiable regulations might be of interest to an outside entity, e.g., legal courts or insurance companies. In the latter case, the system needs to provide

¹¹This can be extended by tuning the disclosure sets (e.g., by distinguishing the requesters from the platforms in the set of involved participants).

evidence to an outside entity demonstrating that the regulation was adhered to and hence resolve any disputes that may arise.

3 SEPAR: DESIGN CHOICES

The vision that we present in Section 2 covers a broad space. We now propose Separ as a possible instance of a precise point in the design space of regulation systems. Separ results from choices, guided by cutting-edge real-life regulation proposals, on the three dimensions of the design space: supported regulations (i.e., simple, mixed with SUM-aggregate), disclosures tolerated, and architecture.

3.1 Supported Regulations

Separ focuses on enforcing lower and upper bounds on the aggregated working time spent on crowdworking platforms¹², a consensual societal need. The necessary information are the participants to crowdworking processes and the (discrete) time estimation of tasks¹³: for Separ, the U-TABLE thus consists of the WORKER, PLATFORM, REQUESTER, and TIMECOST attributes. Separ does not consider joins. As a result the kind of regulations supported by Separ is simple, mixed with SUM-aggregate. The enforceable/verifiable nature of the regulations supported by Separ are easy to determine: lower-than regulations are enforceable because the upper-bound guarantee must always hold, and greater-than regulations are verifiable because the lower-bound guarantee cannot always hold (but will have to hold eventually, e.g., at the end of each time period).

For simplicity, we propose to express such regulations by (1) a triple (w, p, r) that associates a worker w , a platform p , and a requester r , (2) a comparison operator $<$ or $>$, and (3) a threshold value θ (an integer) that defines the lower/upper bound that needs to hold. Intuitively, a regulation $((w, p, r), <, \theta)$ states that there must not be more than θ time spent by worker w on platform p for requester r . We also allow two wildcards to be written in any position of a triple: $*$ and \forall . First, the $*$ wildcard allows ignoring one or more elements of a triple¹⁴. For example, $(*, p, r)$ means that the regulation applies to the pair (p, r) . A triple may contain up to three $*$ wildcards. An element of a triple that is not a $*$ wildcard is called a *target* of the regulation. Second, the \forall wildcard allows a regulation to express a constraint that must hold for all participants in the same group of participants¹⁵. For example, (\forall, p, r) represents the following set of triples: $\{(w, p, r)\}, \forall w \in \mathcal{W}$. As a result, enforceable regulations are expressed by $((w, p, r), <, \theta)$ tuples (possibly with wildcards), and verifiable regulations by $((w, p, r), >, \theta)$ tuples (possibly with wildcards).

Examples. The semantics of an enforceable regulation without any wildcard, e.g., $e \leftarrow ((w, p, r), <, 26)$ expressing a higher bound on the number of time units spent by worker w for requester r on platform p , is the same as the following SQL query:

```
ALTER TABLE U-TABLE ADD CONSTRAINT e CHECK (
  NOT EXISTS (
    SELECT * FROM U-TABLE
    WHERE WORKER=w AND PLATFORM=p AND REQUESTER=r
    GROUP BY WORKER, PLATFORM, REQUESTER
```

¹²Regulating the wages earned through crowdworking platforms can be dealt with similarly.

¹³Extending regulations with validity periods (e.g., "one week"), is straightforward.

¹⁴Intuitively, the $*$ wildcard means "whatever".

¹⁵Intuitively, the \forall wildcard means "for each".

```
HAVING SUM(TIMECOST) ≥ 26
) );
```

The weekly FLSA limit on the total work hours per worker can easily be expressed as $((\forall, *, *), <, 40)$. A social security institution can request each worker w applying for insurance coverage to prove that she worked more than 5 hours: $((w, *, *), >, 5)$ is both necessary and sufficient. Similarly, the regulation $((*, p, *), >, 1000)$ allows a tax institution to require from each platform p applying for a tax refund that the total work hours of all its workers are at least 1000 hours.

3.2 Threat Model and Disclosure Sets

Separ considers that any participant in a crowdworking environment (worker, requester, platform) may act as a *covert adversary* [9] that, loosely speaking, aims at inferring anything that can be inferred from the execution sequence and that is able to deviate from the protocol if no other participant detects it. For simplicity, we assume in Separ that adversarial participants do not collude (although extending Separ to cope with colluding covert adversaries is easy (see Section 4)).

Consider a crowdworking process π between worker w , platform p , and requester r for solving task t . The information generated by the execution of π consists of the relationship between the three participants (w, p , and r) with task t . Additionally, we also consider the information that π is starting or ending through a starting event BEGIN and an ending event END. They may be determined, for example, by exchanging messages among the participants of π , and may include additional concrete information, e.g., timestamps, IP address. $(\text{BEGIN}, \text{END}, w, p, r, t)$ denotes the information generated by π . Separ does not leak any information about the worker and the requester involved in π when it is not needed by π . It tolerates the disclosure of the $\{\text{BEGIN}, \text{END}\}$ events and of the platform p to all participants, whatever their involvement in π . This allows platforms to share information for enforcing regulations (e.g., check that all participants satisfy the regulations before executing π), and to collaborate for correctly managing cross-platform tasks. Note that for simplicity we use the same notation δ for disclosures concerning sets of crowdworking processes as well.

The resulting disclosure sets of Separ are instantiated as follows:

- The participants that are **not involved** in π and that have **not received** task t from requester r must not learn anything about the worker, the task, and the requester involved in π : $\delta_{-R-I}^{\pi} = (\text{BEGIN}, \text{END}, p)$
- The platforms and workers that have **received** task t from r but that are **not involved** in π must be aware that t has been performed (e.g., for not contributing to t) but must not know that it has been performed by worker w : $\delta_{R-I}^{\pi} = (\text{BEGIN}, \text{END}, p, r, t)$
- The participants that are directly **involved** in π (and have thus **received** task t) learn the complete 6-tuple: $\delta_{RI}^{\pi} = (\text{BEGIN}, \text{END}, w, p, r, t)$

3.3 Architecture

Separ has two main components. The centralized *Registration Authority (RA)* and the decentralized *Multi-Platform Infrastructure (MPI)*. Although centralized, the RA can be made fault-tolerant

using standard replication [23] techniques. RA registers the participants to the crowdsourcing environment, models the regulations, and distributes to participants the cryptographic material necessary for enforcing or verifying regulations in a secure manner.

MPI, on the other hand, is a decentralized component that maintains the global state of the system including all operations performed by the participants. This state is maintained within a distributed persistent transparent *ledger*. MPI consists of a set of collaborating crowdsourcing platforms connected by an asynchronous distributed network. Due to the unique features of blockchains such as transparency, provenance, and fault tolerance, the MPI is implemented as a *permissioned blockchain*.

MPI processes two types of tasks: *internal* and *cross-platform*. Internal tasks are submitted to the ledger of a single platform, whereas cross-platform tasks are submitted to the ledgers of multiple platforms (i.e. involved platforms). Separ does not make any assumptions on the implementation of crowdsourcing processes by platforms, e.g., task assignment algorithm and workers contribution delivery. However, processing a task (either internal or cross-platform) requires agreement from the involved platform(s). To establish agreement among the nodes within or across platforms, Separ uses *local* and *cross-platform* consensus protocols. Furthermore, in both internal and cross-platform tasks, Separ enables all platforms to check the fulfillment of regulations using a *global* consensus protocol among *all* platforms (irrespective of whether they are involved in the execution of the task).

4 SEPAR: APPLICATION LEVEL

Separ can be viewed as a *two-level* system consisting of an application level, i.e., *privacy-preserving token-based system*, on top of an infrastructure level, i.e., *blockchain ledger* shared across platforms. This section presents the application level of Separ and Section 5 presents the infrastructure level. Inspired by e-cash systems, Separ implements both enforceable and verifiable regulations by managing two *budgets* per participant while guaranteeing both privacy and correctness. The overall system is conceptually simple and only relies on the correct use of individual and group signatures. The registration authority (RA, see Section 3.3) bootstraps Separ, its participants, and refreshes their budgets periodically, but is not involved in the continuous execution of crowdsourcing processes and their regulations.

4.1 Individual and Group Signatures

Workers, requesters, and platforms are all equipped by the registration authority with the following cryptographic material: a pair of public/private asymmetric *individual* keys (e.g., RSA) and a pair of public/private asymmetric *group* keys (e.g., [12]) where the union of all workers forms a group (in the sense of group signatures), the union of all requesters forms another group, and the union of all platforms forms the last group. A group signature scheme respects three main properties [16]: (1) only members of the group can sign messages, (2) the receiver of the signature can verify that it is a valid signature for the group but cannot discover which member of the group computed it, and (3) in case of dispute later on, the signature can be "opened" (with or without the help of the group members) to reveal the identity of the signer. A common way to enforce the third

property is to rely on a *group manager* that can add new members to the group or revoke the anonymity of a signature. Instances of such schemes are proposed in [16], but also in [8, 12].

In Separ, we use the protocol proposed in [12] and denote $\sigma_w^g(m)$ the group signature of the worker w (with her group private key) for message m . We use equivalent notations for requester r and platform p (i.e., respectively $\sigma_r^g(m)$ and $\sigma_p^g(m)$). The notation $\sigma_w^i(m)$ is used to refer to an individual asymmetric signature (e.g., RSA) of worker w (with her individual (non-group) private key) for the message m . We use equivalent notations for requester r and platform p (i.e., respectively $\sigma_r^i(m)$ and $\sigma_p^i(m)$). The registration authority is the group manager for the three groups (workers \mathcal{W} , requesters \mathcal{R} , platforms \mathcal{P}) and is also equipped with her own individual cryptographic keys for signing her messages: $\sigma_{RA}^i(m)$.

4.2 A Simple Token-Based System

To regulate crowdsourcing processes, our token-based system is defined by five functions: GENERATE for initializing the budgets with the correct number of tokens and refilling them, SPEND for spending portions of the budgets, PROVE for providing proofs for verifying verifiable regulations (e.g., to a third party), CHECK for checking whether a given spending is allowed or not, and ALERT for reporting dubious spending. Since the execution of these functions changes the global state of the system, the data involved in the execution (e.g., tasks, tokens, signatures) must be appended to the distributed ledger of the platforms.

The GENERATE Function. The registration authority uses the GENERATE function to create tokens for all participants (i.e., workers, platforms, requesters) according to the set of regulations. We call *e-tokens* the tokens implementing enforceable regulations and *v-tokens* the tokens implementing verifiable regulations¹⁶.

For each enforceable regulation $((w, p, r), <, \theta + 1)$, the registration authority generates θ *e-tokens* and sends a copy of each token to each target of the regulation. An *e-token* consists of a *public component*, i.e., a pair made of a *number used only once* (referred to as a *nonce* below) generated by the registration authority and a signature of the nonce by the registration authority¹⁷ and a *private component*: an index for selecting the correct set of tokens given the other targets, i.e., their public keys¹⁸. Let $e\text{-}\tau$ be an *e-token*, $e\text{-}\tau_{pub}$ be its public component, $e\text{-}\tau_{priv}$ be its private component, N be a nonce and λ the list of public keys of the targets of the corresponding regulation. The *e-token* is thus the pair $(e\text{-}\tau_{pub}, e\text{-}\tau_{priv})$ where $e\text{-}\tau_{pub} = (N, \sigma_{RA}^i(N))$ and $e\text{-}\tau_{priv} = \lambda$.

Similar to an *e-token*, a *v-token* consists of a public and a private component. The public component is a nonce together with its signature from the registration authority. The private component is simply a triplet of signatures, from the registration authority, binding the recipient of the *v-token* (called *owner* below) to each

¹⁶Although it is technically possible to use a unified token structure for both enforceable and verifiable regulations, using two different tokens reduces computation and communication costs.

¹⁷Extending tokens with labels and timestamps to support validity periods is straightforward.

¹⁸The use of a public key generated by the registration authority is important here because (1) it can be shared among participants without disclosing their identities, i.e., it is a pseudonym, (2) the corresponding private key can be used by participants for mutual authentication in order to guarantee the correctness of the index and consequently of the choice of tokens.

of the other targets of the verifiable regulation¹⁹. More formally, let v - τ be a v -token, v - τ_{pub} be its public component, v - τ_{priv} be its private component, N be a nonce, o be the identity of the participant owner of the token, and (w, p, r) be the related triplet. The v -token is thus the pair $(v$ - τ_{pub}, v - $\tau_{priv})$ where v - $\tau_{pub} = (N, \sigma_{RA}^i(N))$ and v - $\tau_{priv} = (\sigma_{RA}^i(N, o, w), \sigma_{RA}^i(N, o, p), \sigma_{RA}^i(N, o, r))$. The number of v -tokens to produce initially can be easily deduced from the lowest higher bound in enforceable regulations.

The SPEND Function. Requesters create and send their tasks to a platform. The tasks are appended to either the ledger of the platform (for internal tasks) or the ledgers of all involved platforms (for cross-platform tasks). Once the task t is published, the workers can indicate their intent to perform the task by sending a *contribution intent* to their platforms. If a contribution is still needed for the task, the crowdfunding process π starts with the SPEND function, performed as follows. Without loss of generality, we assume below that the given task has a time cost equal to 1.

The SPEND function essentially (1) gathers a sufficient number of e -tokens and v -tokens from the participants involved in π , (2) generates the group signatures of the public components of tokens by each participant, and (3) commits the public components of tokens (τ_{pub} for both e -tokens and v -tokens) with the corresponding group signatures $((\sigma_w^g(\tau_{pub}||t), \sigma_r^g(\tau_{pub}||t), \sigma_p^g(\tau_{pub}||t)))$ for both e -tokens and v -tokens to the ledgers of all platforms. Note that since only the public component of tokens is maintained in the ledgers, the number of tokens that a participant has spent cannot be learned from the ledgers. During this process, participants (e.g., requester) only learn one bit of information, i.e., whether the other participants have enough tokens to participate. Additional signatures are exchanged to guarantee that all the participants involved in π behave honestly (see the extended version [6] for details).

The PROVE function. Participants use the PROVE function to provide proofs for guaranteeing verifiable regulations (e.g., to a third party). The use of v -tokens is relatively straightforward. During the crowdfunding processes, participants simply store the private components of v -tokens and deliver them at the end of the validity periods of verifiable regulations. As an example, for a $((w, p, r), >, 4)$ verifiable regulation, the worker w sends the private components of 5 v -tokens²⁰. The entity in charge of guaranteeing the verifiable regulation checks the signature of the registration authority - to verify that the participant was involved in the task - and the nonce stored in the ledger - to ensure that the token has been indeed spent and committed to the ledgers of all platforms.

The CHECK and ALERT Functions. These functions are used to detect and report either the malicious behavior of participants resulting in an invalid consumption of tokens or the failure of a platform. The complete set of verifications protects against (1) the forgery of tokens (verification of the signatures), (2) the replay of tokens (verification of the absence of double-spending), (3) the relay of tokens (verification of the absence of usurpation), and (4)




Figure 1: Sequence chart (references to targets include all participants for v -tokens)

the illegitimate invalidation of tokens (timeout against malicious platform failures).

The first two verifications are straightforward and performed during global consensus. Verifications (3) and (4) are similar, and we explain here case (3). When a token is appended to the ledgers of all platforms, any participant (whether involved in the corresponding crowdfunding process or not) can CHECK its nonce. If a participant detects a nonce that was received from the registration authority but not spent (or spent on a wrong task), she ALERTs the registration authority. The registration authority will de-anonymize the group signature of the corresponding participant (e.g., the worker’s group signature if the alert comes from a worker) and check whether it has been signed by the same participant that sent the token. Depending on the result, the registration authority will take adequate sanctions against the fraudulent participant (true positive) or the alert-riser (false positive).

4.3 Task Processing Sequence

The processing of a crowdfunding task involves the following five main phases, as depicted in Figure 1.

Initialization. The registration authority provides all parties with their keys and tokens.

Publication. Requesters create and send their tasks to platforms. If a requester wants to publish its task on more than one platform (i.e., a cross-platform task), the involved platforms collaborate with each other to create a common instance of the task (e.g., a common task identifier). The involved platforms then append the task to their ledgers through submission transactions and inform their workers in their preferred manner for accessing tasks.

Assertion. After a worker has retrieved a task, the worker sends a *contribution intent* message to the platform without revealing the actual contribution. The platform then updates the number of required contributions for the task and appends the contribution intent to its ledger through a claim transaction. For cross-platform tasks, the platform informs other involved platforms about the received contribution intent, so that all involved platforms agree with the number (and order) of the received contribution intents (i.e., claim transactions) and append the claim transaction to their ledgers. If the desired number of contributions for the task has been achieved, the contribution intent is refused.

¹⁹Binding recipients to their v -tokens allows to trace possible malicious leakages of v -tokens (they strongly bind participant to a crowdfunding process, contrary to e -tokens) and consequently to prevent them (our participants act as covert adversaries).

²⁰For minimal disclosure purposes, participants are allowed to send only the parts of the private components of v -tokens that are relevant to the verifiable regulations being verified: verifiable regulations that do not specify all the targets (e.g., $((w, *, *), >, \theta)$ only require the private components involving the worker w (e.g., $\theta + 1$ distinct $\sigma_{RA}^i(N, o, w)$ signatures where $o = w$ here).

Verification. Once the contribution intent has been accepted by the platform(s), the platform asks the corresponding requester and worker to send the required tokens and signatures, through the SPEND function (see above). Upon receiving all tokens and signatures, the platform shares them with all platforms and the e -tokens and their signatures are appended to the ledgers of all platforms through verification transactions. From this point, anyone can check the validity of requirements with the CHECK function (and ALERT if required), as developed above.

Execution. Once all parties have checked the validity of the task, the tokens, and the group signatures, the contribution can be delivered to the requester and the reward to the worker.

4.4 Privacy Analysis

We show below in the suite of Theorem 1, Lemma 1, Lemma 2, and Theorem 2 that the global execution of Separ satisfies the δ^Π -privacy model against covert adversaries (where the disclosure sets are defined in Section 3). All proofs are included in the extended version of the paper [6].

First, Theorem 1 restricts the adversarial behavior to inferences (i.e., similar to a *honest-but-curious* adversary) and shows that the execution of Separ satisfies δ^Π -privacy (where the disclosure sets are defined in Section 3).

Theorem 1. (*Privacy (inferences)*) For all sets of crowdworking processes Π executed over participants \mathcal{W} , \mathcal{P} , and \mathcal{R} by an instance of Separ ζ , then it holds that ζ is δ^Π -Private against covert adversaries restricted to inferences (where the disclosure sets are defined in Section 3).

Second, we extend possible behavior to malicious behaviors aiming at jeopardizing regulations and show that they are systematically detected by Separ (Lemma 1 focuses on enforceable regulations and Lemma 2 on verifiable regulations). This prevents covert adversaries from performing malicious actions, limiting them to inferences.

Lemma 1. (*Detection of malicious behavior (enforceable regulations)*) A crowdworking process π executed over participants \mathcal{W} , \mathcal{P} , and \mathcal{R} by an instance of Separ ζ , completes successfully without raising a legitimate alert if and only if π does not jeopardize any enforceable regulation.

Lemma 2. (*Detection of malicious behavior (verifiable regulations)*) Participant P can produce a proof about process π executed over participants \mathcal{W} , \mathcal{P} , and \mathcal{R} by an instance of Separ ζ if and only if P was involved in π and π completed successfully.

Since Theorem 1 shows that Separ is δ^Π -private against adversaries restricted to inferences, and Lemma 1 and Lemma 2 show that malicious behaviors are prevented, it follows that Separ is δ^Π -private against covert adversaries (Theorem 2).

Theorem 2. (*Privacy (inferences and malicious behavior)*) For all sets of crowdworking processes Π executed over participants \mathcal{W} , \mathcal{P} , and \mathcal{R} by an instance of Separ ζ , then ζ is δ^Π -private against covert adversaries (where the disclosure sets are defined in Section 3).




Figure 2: (a): The ledger of Separ with 4 platforms, (b), (c), (d), and (e): The views of the ledger from different platforms

5 SEPAR: INFRASTRUCTURE LEVEL

Section 4 provides an abstract design for implementing the application level of Separ. In this section, we present the infrastructure level and show how Separ supports the execution of transactions on multiple globally distributed platforms that do not necessarily trust each other. In Separ and in order to provide fault tolerance, each platform consists of a set of nodes (i.e., replicas) that store copies of the platform’s ledger. Separ uses a *permissioned blockchain* as its underlying infrastructure (i.e., MPI). The unique features of blockchain such as transparency, provenance, fault tolerance, and authenticity are used by many systems to deploy a wide range of distributed applications in permissioned settings. In particular and for a crowdworking system, the *transparency* of blockchains is useful for checking integrity constraints, *provenance* enables Separ to trace how data is transformed, *fault tolerance* helps to enhance reliability and availability, and finally, *authenticity* guarantees that signatures and transactions are valid.

5.1 Blockchain Ledger

The blockchain ledger in Separ, as mentioned before, maintains the global state of the system and includes all submission, claim, and verification transactions of all internal as well as cross-platform tasks. To ensure data consistency, an ordering among transactions in which a platform is involved is needed. The order of transactions in the blockchain ledger is captured by *chaining* transaction blocks together, i.e., each transaction block includes the sequence number or the cryptographic hash of the previous transaction block. Since Separ supports both internal and cross-platform tasks and more than one platform is involved in each cross-platform transaction, the ledger (similar to [3, 5]) is formed as a *directed acyclic graph (DAG)* where the *nodes* of the graph are transaction blocks (each block includes a single transaction) and *edges* enforce the order among transaction blocks.

Fig. 2(a) shows a blockchain ledger created in the Separ model for a blockchain infrastructure consisting of four platforms p_1 , p_2 , p_3 , and p_4 . In this figure, λ is the unique initialization (*genesis*) block of the blockchain, t_i^s are submission transactions, $t_i^c_j$ is the j -th

claim transaction of task t_i , and t_iv is the verification transaction of task t_i . In Fig. 2(a), t_{10} , t_{20} , t_{30} , and t_{40} are internal submission transactions of different platforms that can be appended to the ledger in parallel. As shown, t_{10} requires 3 contributions (thus 3 claim transactions t_{10c_1} , t_{10c_2} , and t_{10c_3}) whereas each of t_{20} , t_{30} , and t_{40} needs two contributions. t_{10v} , t_{20v} , t_{30v} , and t_{40v} are the corresponding verification transactions. $t_{11,21}$ is a cross-platform submission among platforms p_1 and p_2 . Similarly, $t_{31,41}$ is a cross-platform submission among platforms p_3 and p_4 . Here, $t_{11,21}$ and $t_{31,41}$ require one and two contributions respectively. Note that the claim transactions of a cross-platform task might be initiated by different platforms and as mentioned earlier, the order of these claim transactions is important (to recognize the n first claims).

This global directed acyclic graph blockchain ledger includes all transactions of internal as well as cross-platform tasks initiated by all platforms. However, to ensure data privacy, each platform should only access the transactions in which the platform is involved. As a result, in Separ, the entire blockchain ledger is *not maintained* by any specific platform and each platform p_i , as shown in Fig. 2(b)-(e), only maintains its own *view* of the blockchain ledger including (1) all submission and claim transactions of its internal tasks, (2) all submission and claim transactions of the cross-platform tasks involving the platform, and (3) verification transactions of all tasks. Note that verification transactions are replicated on every platform to enable all platforms to check the satisfaction of global regulations. The global DAG ledger (e.g., Fig. 2(a)) is indeed the union of all these physical views (e.g., Fig. 2(b)-(e)). Note that, since there is no data dependency between the tasks that platform p is involved in and the verification transactions of the tasks that platform p is *not* involved in, the verification transactions might be appended to the ledgers in different orders, e.g., t_{20v} (of p_2) and t_{40v} (of p_4) are appended to the ledger of platforms p_1 and p_3 in different orders.

5.2 Consensus in Separ

In Separ, each platform consists of a (disjoint) set of nodes (i.e., replicas) where the platform replicates its own view of the blockchain ledger on its nodes to achieve fault tolerance. Nodes follow either the crash or the Byzantine failure model. In the crash failure model, nodes may fail by stopping and may restart, however, in the Byzantine failure model, faulty nodes may exhibit malicious behavior. Nodes of the same or different platforms need to establish consensus on a unique order in which entries are appended to the blockchain ledger. To establish consensus among the nodes, asynchronous fault-tolerant protocols have been used. Crash fault-tolerant protocols, e.g., Paxos [24], guarantee safety in an asynchronous network using $2f + 1$ nodes to overcome the simultaneous failure of any f nodes while in Byzantine fault-tolerant protocols, e.g., PBFT [13], $3f + 1$ nodes are usually needed to provide safety in the presence of f malicious nodes.

Completion of a crowdfunding task, as discussed earlier, requires a single submission transaction, one or more claim transactions (depending on the requested number of contributions), and a verification transaction. For an internal task of a platform, submission and claim transactions are replicated only on the nodes of the platform, hence, *local consensus* among nodes of the platform on the order of the transaction is needed. For a cross-platform task, on the other hand, submission and claim transactions are replicated on every node of all

involved platforms. As a result, *cross-platform consensus* among the nodes of only the *involved* platforms is needed. Finally, verification transactions are appended to the ledger of all platforms, therefore, all nodes of *every* platform participate in a *global consensus* protocol. In this section, we show how local, cross-platform, and global consensus are established with crash-only or Byzantine nodes.

5.2.1 Local Consensus. Processing a submission or a claim transaction of an internal task requires local consensus where nodes of a single platform, *independent* of other platforms, establish agreement on the order of the transaction. The local consensus protocol in Separ is pluggable and depending on the failure model of nodes, a crash, e.g., Paxos [24], or a Byzantine, e.g., PBFT [13] consensus protocol can be used.

5.2.2 Cross-Platform and Global Consensus. Both cross-platform consensus and global consensus require collaboration between multiple platforms. Since platforms do not trust each other and the primary node that initiates the transaction might behave maliciously, Separ uses *Byzantine* fault-tolerant protocols in both cross-platform and global consensus. Cross-platform and global consensus, however, are different in two aspects. First, in cross-platform consensus, only the involved platforms participate, whereas global consensus is established among all platforms, and second, at the platform level, while cross-platform consensus requires agreement from *every* involved platform, in global consensus, agreement from *two-thirds* of all platforms is sufficient. Cross-platform consensus requires agreement from every involved platform to ensure data consistency due to the possible data dependency between the cross-platform transaction and other transactions of an involved platform. Note that if an involved platform (as a set of nodes) behaves maliciously by not sending an agreement for a cross-platform transaction initiated by another platform, e.g., a claim transaction, its malicious behavior can be detected and penalties imposed. In global consensus, however, the goal is only to check the correctness of the transaction. To provide safety for global consensus at the platform level, we assume that at most $\lfloor \frac{|P|-1}{3} \rfloor$ platforms might behave maliciously. As a result, to commit a transaction, by a similar argument as in PBFT [13], at least two-thirds ($\lfloor \frac{2|P|}{3} \rfloor + 1$) of the platforms must agree on the order of the transaction.

Cross-Platform Consensus. Processing Submission and claim transactions of a cross-platform task requires cross-platform consensus among *all* involved platforms where due to the untrustworthiness of platforms, a Byzantine fault-tolerant protocol is used. Since the number of nodes within each platform depends on the failure model of nodes of a platform (i.e. $2f + 1$ crash-only or $3f + 1$ Byzantine nodes), the required number of matching replies from each platform, i.e., the quorum size, to ensure the safety of protocol is different for different platforms. We define *local-majority* as the required number of matching replies from the nodes of a platform. For a platform with crash-only nodes, local-majority is $f + 1$ (from the total $2f + 1$ nodes), whereas for a platform with Byzantine nodes, local-majority is $2f + 1$ (from the total $3f + 1$ nodes).

Separ establishes consensus on cross-platform transactions in four phases: (1) prepare, (2) propose, (3) accept, and (4) commit, and then (5) appends transactions to the ledger.

(1) Prepare Phase. Upon receiving a cross-platform (submission or claim) transaction m , the (pre-elected) node of the (recipient) platform p_i (called the *initiator primary*) initiates the consensus protocol by assigning a sequence number h_i to the transaction and multicasting a signed prepare message μ to the primary node of all involved platforms. The prepare message includes the received transaction m (either submission or claim), its digest d (cryptographic hash) and the sequence number h_i . The sequence number represents the correct order of the transaction block in the initiator platform. If the transaction is a claim transaction, the initiator primary includes the hash of the corresponding submission transaction as well.

(2) Propose Phase. Once the primary node of some platform p_j receives a prepare message μ for some transaction m from the initiator primary, it first validates the message. If the node is currently waiting for a commit message of some cross-platform transaction m' where the involved platforms of the two requests m and m' intersect in p_j as well as some other cluster p_k , the node does not process the new transaction m before the earlier transaction m' gets committed. This ensures that requests are committed in the same order on overlapping clusters (*consistency*), e.g., m and m' are committed in the same order on both p_j and p_k . Separ addresses deadlock situations, i.e., where overlapping clusters receive prepare messages in a different order, in the same way as SharPer [5]). If the primary is not waiting for an uncommitted transaction, it assigns sequence number h_j (that represents the order of m in platform p_j) to the message and multicasts a *signed propose* message to the nodes of its platform including both sequence numbers h_i and h_j , digest d and piggybacked prepare message μ . The initiator primary node, similarly, multicasts a signed propose message (including h_i , d , and piggybacked μ) to the nodes of its platform.

(3) Accept Phase. Once a node of an involved platform p_j receives a valid propose message from the primary node of its cluster which has no overlap with any uncommitted cross-platform requests, the node multicasts a signed accept message including both sequence numbers h_i and h_j , and digest d to every node of all involved platforms. The primary nodes of all involved platforms also multicast accept messages (with the same structure) to every node of all involved platforms. Note that if a platform behaves maliciously by not sending accept messages to other involved platforms, the initiator platform can report the malicious behavior of the platform by sending an ALERT message (as explained in Section 4.2) to the registration authority resulting in imposing penalties.

(4) Commit Phase. Upon receiving valid matching accept messages from a local-majority (i.e., either $f+1$ or $2f+1$ depending on the failure model) of every involved platform with h_i and d that match the propose message which was sent by primary, every node multicasts a signed commit message including all valid sequence numbers, e.g., h_i, h_j, \dots, h_k , to all nodes of every involved platform. The prepare, propose and accept phases of the algorithm guarantee that non-faulty nodes agree on a total order for the transactions.

(5) Append to Ledger. Finally, each node waits for valid matching commit messages from a local-majority of every involved platform that match its commit message before committing the transaction. If all transactions with lower sequence numbers than h_j have already been committed, the node appends a transaction block including the transaction as well as the corresponding commit messages to its copy of the ledger. Note that since commit messages include the digest

of the corresponding transactions, appending valid signed commit messages to the blockchain ledger in addition to the transactions, provides the same level of *immutability* guarantee as including the cryptographic hash of the previous transaction in the transaction block, i.e., any attempt to alter the block data can easily be detected.

In terms of message complexity, prepare phase consists of $|P|$ messages, propose phase needs n messages, and accept and commit phases, each requires n^2 messages where n is the number of nodes.

In addition to the normal case operation, Separ has to deal with two other scenarios. First, when the primary node fails. Second, when nodes have not received a quorum of *matching* accept messages from the local-majority of every involved platform due to conflicting accept messages. Indeed, the primary nodes of different platforms might multicast their propose messages in parallel, hence, different overlapping platforms might receive the messages in a different order. We use techniques similar to SharPer [5] to address these two situations. Due to space limitation, the detailed explanation of the techniques is omitted and are provided in the extended version of the paper [6].

Global Consensus. The verification transactions include group signatures and all tokens that are consumed by participants to perform a particular task. In Separ and in order to enable all platforms to check regulations, verification transactions are appended to the ledger of all platforms. To do so, a Byzantine fault-tolerant protocol, similar to cross-platform consensus, is run among all nodes of every platform where in each phase, the protocol needs agreement from the *local-majority* of *two-thirds* of the platforms.

The correctness of both cross-platform and global consensus protocols is proven in the extended version of the paper [6].

6 EXPERIMENTAL EVALUATIONS

In this section, we conduct several experiments to evaluate Separ. We consider a complex heavily loaded setting with 4 platforms, 20000 requesters and 20000 workers where requesters and workers are randomly registered to one or more platforms. Once a task is submitted to a platform, the platform randomly assigns the task to one or more (depending on the required number of contributions) idle workers (to avoid any delay). The experiments consist of two main parts. In the first part (Sections 6.1 and 6.2), the privacy costs of Separ (i.e., tokens and regulations) is evaluated, whereas in the second part (Sections 6.3 and 6.4), the scalability of Separ is evaluated. For the purpose of this evaluation, and as explained earlier, we do not focus on the description of tasks and contributions (both are modeled as arbitrary bitstrings). In addition, v -tokens, as explained earlier, are very similar to e -tokens except for the private part that has no significant impact on the performance and the number of interaction phases which is even less than e -tokens. Therefore, we only focus on e -tokens (i.e., enforceable regulations) in the experiments. To implement group signatures, we use the protocol proposed in [12]. The experiments were conducted on the Amazon EC2 platform. Each VM is a c4.2xlarge instance with 8 vCPUs and 15GB RAM, Intel Xeon E5-2666 v3 processor clocked at 3.50 GHz. When reporting throughput measurements, we use an increasing number of tasks submitted by requesters running on a single VM, until the end-to-end throughput is saturated, and state the throughput and latency just below saturation.




Figure 3: (a), (b) Class of regulations and (c) Geo-Scalability

6.1 Token Generation

In the first set of experiments, we measure the performance of token generation (performed by RA) in Separ. We consider different classes of regulations, i.e., single-target (e.g., $((w, *, *), <, \theta)$), two-target (e.g., $((*, p, r), <, \theta)$), and three-target (e.g., $((w, p, r), <, \theta)$) enforceable regulations. Our experiments show that Separ is able to generate tokens in linear time. Separ generates each token in $0.07ms$, hence, generating 1 million tokens in ~ 76 seconds. This clearly demonstrates the scalability of the token generation especially since token generation is executed periodically, e.g., every week or every month. Note that, we use a single machine to generate tokens, however, tokens related to different regulations can be generated in parallel. Hence, the throughput of Separ can linearly increase by running the token generation routine on multiple machines, e.g., with 10 machines, Separ is able to generate 1 million tokens in ~ 7.6 seconds. Moreover, to provide fault tolerance, the tokens can be replicated on multiple machines following the replication techniques. Furthermore, the class of regulations (i.e., the number of targets) does not affect the performance, i.e., the throughput and latency of token generation are constant in terms of the number of targets. It should, however, be noted that a regulations with more targets requires more tokens to be generated, e.g., $((w, *, *), <, \theta)$ requires $|W| * \theta$ tokens, whereas $((w, p, r), <, \theta)$ requires $|W| * |P| * |R| * \theta$ tokens to be generated.

6.2 Classes of Regulations

In the second set of experiments, we measure the overhead of privacy-preserving techniques, e.g., group signatures and tokens, used in Separ. We consider the basic scenario with no regulation (i.e., no need to exchange and validate tokens and signatures) and compare it with three different scenarios where each task has to satisfy a single target, a two-target, and a three-target regulation. The system consists of four platforms and the workload includes 90% internal and 10% cross-platform tasks (the typical settings in partitioned databases [32]) where two (randomly chosen) platforms are involved in submission and claim transactions of cross-platform tasks. Note that all platforms are still involved in the verification transaction of each task. We also assume that completion of each task requires a single contribution, i.e., claim transaction, (and obviously a submission and a verification transaction).

When nodes follow the crash failure model and the system includes no regulations, as shown in Figure 3(a), Separ processes 7000 tasks with 390 ms latency (the penultimate point). Adding regulations, results in more communication between participants to exchange tokens and signatures, however, Separ still processes

6200 tasks with 450 ms latency. In fact, all privacy-preserving techniques that are used in Separ result in only 11% and 15% overhead in terms of throughput and latency respectively. Moreover, the class of regulations does not significantly affect the performance of Separ. This is expected because more targets result in only increasing the number of (parallel) tokens and signature exchanges while the number of communication phases is not affected.

Similarly, in the presence of Byzantine nodes and as shown in Figure 3(b), Separ is able to process 6140 tasks with 409 ms latency with no regulations and 5331 tasks (13% overhead) with 467 ms (14% overhead) latency with single-target regulations. As before, the class of regulations does not affect the performance.

It should be noted that by increasing the number of regulations, Separ still demonstrates similar behavior as shown in this experiment. Indeed, adding more regulations, while it results in adding more tokens and possibly more participants and signatures, it does not affect the consensus protocols and other communication phases.

6.3 Scalability Over Spatial Domains

In the third set of experiments, we demonstrate that Separ scales over multiple spatial domains that are globally distributed in a realistic global setting. In particular, instead of placing all four platforms in the same data center, the platforms are placed in four different AWS regions, i.e., Tokyo, Hong Kong, Virginia, and Ohio data centers²¹. Since, as shown in Figure 3(a)-(b), the class of regulations (number of targets) does not affect performance, we assume all regulations as single-target. When nodes follow the crash failure model and the system has no regulations, as shown in Figure 3(c), Separ processes 2870 tasks with 891 ms latency before the end-to-end throughput is saturated. Adding regulations, however, Separ is still able to process 2590 tasks with 1012 ms latency. Hence, all privacy-preserving techniques used in Separ result in only 10% and 13% overhead in terms of throughput and latency respectively. Similarly, in the presence of Byzantine nodes, Separ is able to process 2523 tasks with 1105 ms latency with no regulations and 2239 tasks (12% overhead) with 1242 ms (13% overhead) latency with single-target regulations. These experiments demonstrated that as expected in a geo-distributed setting, the performance of Separ will be reduced, however, it is still able to process 2870 and 2523 tasks with crash-only and Byzantine nodes. More interestingly, in a geo-distributed setting, privacy-preserving techniques used in Separ have lower overhead in comparison to a setting with a single data center because the overhead of privacy-preserving techniques is covered by the latency of the geo-distributed setting.

6.4 Cross-Platform Tasks

In the next set of experiments, we measure the performance of Separ for workloads with different percentages of cross-platform tasks. We consider four different workloads with 0%, 20%, 80%, and 100% cross-platform tasks. Completion of each task requires a single contribution, two (randomly chosen) platforms are involved in each cross-platform task, and each task has to satisfy two randomly chosen regulations. When all nodes are crash-only, as presented in Figure 4(a), Separ processes 8600 tasks with 400 ms latency if

²¹The average measured Round-Trip Time (RTT) between every pair of Amazon data centers can be found at <https://www.cloudping.co/grid>




Figure 4: Varying Number of Cross-Platform Tasks

all tasks are internal. Note that even when all tasks are internal, the verification transaction of each task still needs global consensus among all platforms. Increasing the percentage of cross-platform tasks to 20%, reduces the overall throughput to 5800 (67%) with 400 ms latency since processing cross-platform tasks requires cross-platform consensus. By increasing the percentage of cross-platform tasks to 80% and then 100%, the throughput of Separ will reduce to 1900 and 700 with the same latency. This is expected because when most tasks are cross-platform ones, more nodes are involved in processing a task and more messages are exchanged. In addition, the possibility of parallel processing of tasks will be significantly reduced. In the presence of Byzantine nodes, as shown in Figure 4(b), Separ demonstrates a similar behavior as the crash-only case.

We also measure the scalability of Separ in crowdsourcing environments with varying number of platforms. The results (presented in the extended version [6]) demonstrate that the performance of Separ actually improves as the number of participating platforms increases.

7 RELATED WORK

Enhancing privacy in the context of crowdsourcing has been addressed by several recent studies with various kinds of guarantees, from differential privacy [33, 34] to cryptography [25–27], mostly focusing on spatial crowdsourcing and the use of geolocation to perform assignments. In ZebraLancer [28] and ZKCrowd [36], blockchain is also used to add transparency guarantees on top of privacy. However, all these studies consider a single-platform context. In the context of multi-platform crowdsourcing, Fluid [19] proposes to share workers' profiles between multiple platforms, and Wang et al. [35] provides an interesting insight on federated recommendation systems for workers and the balance of surplus of tasks or workers in a cross-platform context. However, none of them provides tools to manage external regulations: Separ is the first to support a multi-platform crowdsourcing context, with external regulations, transparency, and privacy expectations at the same time.

In the context of permissioned blockchains, Hyperledger Fabric [7] ensures data confidentiality using Private Data Collections [1]. Private Data Collections manage confidential data that two or more entities want to keep private from others. Quorum [15] supports public and private transactions and ensures the confidentiality of private transactions using the Zero-knowledge proof technique. Both Fabric and Quorum, however, order all transactions using a

single consensus protocol resulting in low throughput. Separ is inspired by various permissioned blockchain systems, and more specifically by SharPer [4][5]. SharPer is designed for environments with a single enterprise and uses sharding to improve scalability. SharPer, in contrast to Separ, does not deal with the privacy of participants and processes all transactions in the same way.

Providing anonymity as well as untraceability has been addressed by ZCash [20] which is restricted to the management of crypto-currency issues. Hawk [22] and Raziol [31] manage wider issues and include general smart contracts. However, these solutions do not incorporate infrastructures with multiple platforms, nor implement regulations (let alone anonymized ones). Finally, Solidus [14] proposes to privately manage a multi-platform banking system, with individual banks managing their own clients while allowing cross-platform transactions. While Solidus may be sufficient for banking systems, it does not consider users that subscribe to multiple platforms, nor envisions global profiles or regulations.

8 CONCLUSION

In this paper, we present an overall vision for future of work multi-platform crowdsourcing environments consisting of three main dimensions: regulations, security and architecture. We then introduce Separ, the first, to the best of our knowledge, to address the problem of enforcing global regulations over multi-platform crowdsourcing environments in a privacy preserving manner. Separ enables official institutions to express global regulations in simple and unambiguous terms, guarantees the satisfaction of global regulations by construction, and allows participants to prove to external entities their involvement in crowdsourcing tasks, all in a privacy-preserving manner. Separ also uses transparent blockchain ledgers shared across multiple platforms and enables collaboration among platforms through a suite of distributed consensus protocols. We prove Separ's privacy requirements and conduct extensive experiments to demonstrate Separ's performance and scalability.

Separ supports simple, mixed with SUM-aggregate regulations. Other kinds of regulations, e.g., complex, row-only may need additional verification mechanisms. We will extend Separ to support all kinds of regulations in future work.

ACKNOWLEDGMENTS

This work is funded by NSF grants CNS-1703560 and CNS-1815733 and by the ANR grant ANR-16-CE23-0004.

REFERENCES

- [1] [n.d.]. Private Data Collections: A High-Level Overview. <https://hyperledger-fabric.readthedocs.io/en/release-2.2/private-data/private-data.html>.
- [2] S. Amer-Yahia, S. Basu Roy, L. Chen, A. Morishima, J. Abello Monedero, P. Bourhis, F. Charoy, M. Danilevsky, G. Das, Gianluca Demartini, et al. 2020. Making ai machines work for humans in fow. *SIGMOD Record* 49, 2 (2020), 30–35.
- [3] M. J. Amiri, D. Agrawal, and A. El Abbadi. 2019. CAPER: a cross-application permissioned blockchain. *VLDB* 12, 11 (2019), 1385–1398.
- [4] M. J. Amiri, D. Agrawal, and A. El Abbadi. 2019. On Sharding Permissioned Blockchains. In *International Conference on Blockchain*. IEEE, 282–285.
- [5] M. J. Amiri, D. Agrawal, and A. El Abbadi. 2021. SharPer: Sharding Permissioned Blockchains Over Network Clusters. In *SIGMOD*. ACM.
- [6] M. J. Amiri, J. Duguéproux, T. Allard, D. Agrawal, and A. El Abbadi. 2020. Separ: Towards Regulating Future of Work Multi-Platform Crowdsourcing Environments with Privacy Guarantees. *arXiv preprint arXiv:2005.01038* (2020).
- [7] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, et al. 2018. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *EuroSys*. ACM.

- [8] G. Ateniese, J. Camenisch, M. Joye, and G. Tsudik. 2000. A practical and provably secure coalition-resistant group signature scheme. In *CRYPTO*. Springer, 255–270.
- [9] Y. Aumann and Y. Lindell. 2007. Security against covert adversaries: Efficient protocols for realistic adversaries. In *TCC*. Springer, 137–156.
- [10] J. Berg, M. Furrer, E. Harmon, U. Rani, and M. S. Silberman. 2018. *Digital labour platforms and the future of work : Towards decent work in the online world*. Technical Report ISBN 978-92-2-031025-0. International Labour Organization.
- [11] P. Bourhis, G. Demartini, S. Elbassuoni, E. Hoareau, and H. R. Rao. 2019. Ethical Challenges in the Future of Work. *IEEE Data Eng. Bull.* 42, 4 (2019), 55–64.
- [12] J. Camenisch and J. Groth. 2004. Group signatures: Better efficiency and new theoretical aspects. In *SCN*. Springer, 120–133.
- [13] Miguel Castro, Barbara Liskov, et al. 1999. Practical Byzantine fault tolerance. In *OSDI*, Vol. 99, 173–186.
- [14] E. Cecchetti, F. Zhang, Y. Ji, A. Kosba, A. Juels, and E. Shi. 2017. Solidus: Confidential distributed ledger transactions via PVORM. In *ACM CCS*, 701–717.
- [15] JP Morgan Chase. 2016. Quorum white paper.
- [16] D. Chaum and E. Van Heyst. 1991. Group signatures. In *Workshop on the Theory and Application of Cryptographic Techniques*. Springer, 257–265.
- [17] J. E. Cohen. 2017. Law for the platform economy. *UCDL Rev.* 51 (2017), 133.
- [18] D. Gross-Amblard, A. Morishima, S. Thirumuruganathan, M. Tommasi, and K. Yoshida. 2019. Platform Design for Crowdsourcing and Future of Work. *IEEE Data Eng. Bull.* 42, 4 (2019), 35–45.
- [19] S. Han, Z. Xu, Y. Zeng, and L. Chen. 2019. Fluid: A blockchain based framework for crowdsourcing. In *SIGMOD*. 1921–1924.
- [20] D. Hopwood, S. Bowe, T. Hornby, and N. Wilcox. 2016. Zcash protocol specification. *GitHub: San Francisco, CA, USA* (2016).
- [21] M. Kenney and J. Zysman. 2016. The rise of the platform economy. *Issues in science and technology* 32, 3 (2016), 61.
- [22] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou. 2016. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *S&P*. IEEE, 839–858.
- [23] Leslie Lamport. 1978. The implementation of reliable distributed multiprocess systems. *Computer Networks (1976)* 2, 2 (1978), 95–114.
- [24] L. Lamport. 2001. Paxos made simple. *ACM Sigact News* 32, 4 (2001), 18–25.
- [25] A. Liu, Z. Li, G. Liu, K. Zheng, M. Zhang, Q. Li, and X. Zhang. 2017. Privacy-preserving task assignment in spatial crowdsourcing. (2017).
- [26] A. Liu, W. Wang, S. Shang, Q. Li, and X. Zhang. 2018. Efficient task assignment in spatial crowdsourcing with worker and task privacy protection. *Geoinformatica* 22, 2 (2018), 335–362.
- [27] B. Liu, L. Chen, X. Zhu, Y. Zhang, C. Zhang, and W. Qiu. 2017. Protecting location privacy in spatial crowdsourcing using encrypted data. *EDBT* (2017).
- [28] Y. Lu, Q. Tang, and G. Wang. 2018. Zebalancer: Private and anonymous crowdsourcing system atop open blockchain. In *ICDCS*. IEEE, 853–865.
- [29] Conseil national du numérique. 2020. Travail à l'ère des plateformes. Mise à jour requise. <https://cnumerique.fr/publication-du-rapport-travail-leres-des-plateformes-mise-jour-requise-en-presence-de-cedric-o> (in French).
- [30] Global Commission on the Future of Work. 2019. *Work for a brighter future*. Technical Report ISBN 978-92-2-132796-7. International Labour Organization.
- [31] D. C. Sánchez. 2018. Raziol: Private and verifiable smart contracts on blockchains. *arXiv preprint arXiv:1807.09484* (2018).
- [32] A. Thomson, T. Diamond, S. Weng, K. Ren, P. Shao, and D. J. Abadi. 2012. Calvin: fast distributed transactions for partitioned database systems. In *SIGMOD*. 1–12.
- [33] H. To, G. Ghinita, L. Fan, and C. Shahabi. 2016. Differentially private location protection for worker datasets in spatial crowdsourcing. *IEEE Transactions on Mobile Computing* 16, 4 (2016), 934–949.
- [34] H. To, C. Shahabi, and L. Xiong. 2018. Privacy-preserving online task assignment in spatial crowdsourcing with untrusted server. In *ICDE*. IEEE, 833–844.
- [35] Y. Wang, T. Song, Q. Tao, Y. Zeng, Z. Zhou, Y. Xu, Y. Tong, and L. Chen. 2019. Interaction Management in Crowdsourcing. *Data Engineering* (2019), 23.
- [36] S. Zhu, Z. Cai, H. Hu, Y. Li, and W. Li. 2019. zkCrowd: a hybrid blockchain-based crowdsourcing platform. *Transactions on Industrial Informatics* (2019).