Rlay: A Decentralized Information Network

Michael Hirn (m@rlay.com)

Draft v0.3.1 - June 20, 2018

Contents

1. Overview

The internet is in the middle of a revolution: centralized proprietary services are being replaced with decentralized open ones; trusted parties replaced with verifiable computation; inefficient monolithic services replaced with peer-to-peer algorithmic markets. Bitcoin, Ethereum, and other blockchain networks have proven the utility of decentralized transaction ledgers. These public ledgers process sophisticated smart contract applications and transact crypto-assets worth tens of billions of dollars. As first instances of internet-wide open services these blockchain networks are able to verify transactional information without central or trusted parties, providing useful payment services. Second instance blockchain networks like Filecoin emerged to verify storage information without a central or trusted party, providing useful storage services. We propose Rlay, a Decentralized Information Network to verify real-world or arbitrary information without a central or trusted party (1) providing direct monetization of information assets for many sectors like finance, health, transport, governance, education and (2) solving the decentralized information needs for Web 3.0 applications.

The Rlay protocol turns information verification into an algorithmic market. The market runs on a blockchain (for now Ethereum) with a native protocol token (called 'RLAY'), which clients earn by providing true propositions (i.e. information). Conversely, clients spend RLAY by submitting propositions. Similar to Bitcoin or Filecoin clients compete to submit true propositions first, which creates a powerful incentive for clients to submit as many true propositions as they can. The protocol weaves these amassed propositions into verified information network that anybody in the world can access. The network achieves robustness and security through a novel proof, called Proof-of-Coherence. Clients can query the information network and receive verified information, making Rlay work as a real-world information layer on top of smart contract or storage blockchains. It is especially useful for information systems to eliminate information asymmetries, building and running distributed applications that require information about the real-world, and implementing smart contracts.

This work:

1. Introduces the Rlay network, gives an overview of the protocol, and walks through several components in detail.
2. Formalizes Decentralized Information Network schemes and their properties, then constructs Rlay as a Decentralized Information Network.
3. Introduces a novel Proof-of-Justification scheme called Proof-of-Coherence, which allows proving a proposition (i.e. proving that it is information).
4. Introduces a novel Decentralized Ontology Language which allows clients to encode arbitrary information so that other human- and machine-clients can processes, verify, and unambiguously interpret them.
5. Formalizes verifiable markets and constructs a Public Information Market, which governs how information, access, exploration and research is managed.
6. Discusses use cases, connections to other systems, and how to use the protocol.

   Note: Rlay is a work in progress. Active research is under way, and new versions of this paper can be requested. For requests, comments, and suggestions contact us at m@rlay.com.

   Note: This paper embraces the structural design from previous work done by Filecoin and Zcash to further development of a concise standard for future work.

2. Introduction

RLAY is a protocol token whose protocol runs on a blockchain (for now Ethereum) and utilizes a novel proof, called Proof-of-Coherence, that enables clients to prove and validate propositions asserted by untrusted clients. Rlay protocol provides a proposition assert and query service that does not rely on a single coordinator, where: (1) clients assert arbitrary statements (i.e. propositions) about Reality, (2) clients earn tokens by asserting true propositions and forgo the opportunity to earn tokens when providing wrong information, (3) any client can become a validator, i.e. a Proof-of-Coherence Miner, and earn additional RLAY tokens by taking part in the asynchronous validation of propositions by running Proofs-of-Coherence.

Core Components

The Rlay protocol builds upon three novel components.

1. Decentralized Information Network: We provide an abstraction for a network of independent Suppliers to offer information retrieval and information validation services 3. Later, we present the Rlay as an incentivized, strategy-proof and verifiable Decentralized Information Network construction 5.

2. Decentralized Ontology Language: We present a decentralized formal language that allows clients to encode information and introduce Rlay Ontology as an universal Decentralized Ontology Language. In Rlay, Rlay

Ontology allows human- and machine-clients to processes, validate, and unambiguously interpret asserted propositions.

3. Proof-of-Coherence: We present a novel Proof-of-Coherence 4 that allows Suppliers to prove that asserted propositions are true (i.e. constitute as information). Proving asserted propositions enables Rlay to reward Suppliers only for true propositions and allows clients to verify asserted propositions.

Protocol Overview

- The Rlay protocol is a Decentralized Information Network construction built on a blockchain (for now Ethereum), with a native token called RLAY. Clients spend tokens for asserting propositions and earn tokens for true propositions.
- The Rlay network handles assert and query requests respectively via one verifiable market: the Public Information Market.
- The market is operated by the Rlay network which employs Proof-of-Coherence to guarantee that clients asserted propositions are true.

A sketch of the Rlay protocol, using nomenclature defined later within the paper, is provided in 3 and 5.

Paper Outline

The remainder of this paper is organized as follows. We present our definition of and requirements for a theoretical Decentralized Information Network scheme in 3. In 4 we motivate, define, and present our Proof-of-Coherence, used within Rlay to prove and verify propositions. In 5 we then describe the concrete instantiation of the Rlay protocol, describing data structures, processes, and the interactions between participants. 5 presents our `Payment` function, used within Rlay to incentive the honest participant of Suppliers. We conclude with a discussion of future work and areas of application in 7.

3. Definition of a Decentralized Information Network

We introduce the notion of a Decentralized Information Network (DIN) scheme. Decentralized Information Networks aggregate propositions submitted by multiple independent clients and self-coordinate to provide validated propositions (i.e. information) to clients. Coordination is decentralized and does not require trusted parties: the secure operation of theses systems is achieved through protocols that coordinate and verify operations carried out by individual parties. Decentralized Information Network can employ different strategies for coordination, including Byzantine Agreement, gossip protocols, or CRDTs, depending on the requirements of the system. Later, in 5, we provide a construction for the Rlay Decentralized Information Network.

Definition

A DIN scheme Π is a tuple of protocols run by clients:

(`Assert`, `Query`)

- `Assert(proposition) → key`: Clients execute the `Assert` protocol to submit propositions under a unique identifier key.
- `Query(key) → proposition`: Clients execute the `Query` protocol to retrieve propositions that are currently stored in the network.

A DIN Π must guarantee truth-tracking and its three requirements of competency, independence, and strategy-proofness. All properties are defined in more detail in the following sections.

Properties

Majority belief aggregation rules like the defined Decentralized Information Network, which decide a collective belief based on some concept of majority of its input propositions, have a unique and desirable property namely being truth-tracking (Condorcet 1785; Eckert and Pigozzi 2005; May 1952). Truth-tracking describes the behavior that the collective belief, produced by the aggregation rule, converges towards the truth as the amount of participants increases (related terms: wisdom-of-the-crowd, Condorcet's jury theorem). However, truth-tracking is only achieved under the following three assumptions: Competency, Independence, and Strategy-Proofness (List and Goodin 2001; Dietrich and Spiekermann 2013; Pivato 2017).

We describe the required properties for any Decentralized Information Network scheme as well as additional properties for our Rlay implementation.

Competency

Competency is the condition under which the average participant's likelihood of making no error, i.e. providing a true proposition, is more than 50% for binary propositions or more general has higher frequency than any wrong proposition.

Independence

Independence is the condition under which the errors made by participants, i.e. providing wrong propositions, are uncorrelated.

Strategy-Proofness

A DIN scheme `Π` is strategy-proof if it guarantees, that for any set of propositions, each participant weakly prefers the collective result of the aggregation rule that is obtained from expressing his or her own proposition truthfully over any collective result of the aggregation rule that would be obtained from misrepresenting his or her proposition. Game-theoretically this is equal to requiring that for a participant the expression of his or her true proposition must be a weakly dominant strategy.

Publicly Verifiable

A DIN scheme `Π` is publicly verifiable if for each successful `Assert` request, the network of Suppliers can generate a proof whether a proposition is true. The `Proof-of-Coherence` must convince any efficient verifier, which only knows key and does not know anything else about the proposition.

Auditable

A DIN scheme `Π` is auditable, if it generates a verifiable trace of operation that can be checked in the future to confirm a proposition was indeed true and rewarded the right amount of token.

Incentive-Compatible

A DIN scheme `Π` is incentive-compatible, if Suppliers are rewarded for successfully offering information, or penalized for misbehaving, such that the Suppliers dominant strategy is to provide true propositions (i.e. information).

Collusion-Strategy-Proofness

A DIN scheme `Π` is collusion-strategy-proof if it guarantees, that it is strategy-proof for any coalition of participants that controls no more than 50% of the voting power.

4. Proof-of-Coherence

For the Rlay DIN to be truth-tracking, Suppliers must satisfy the competence and independence requirements to prove to Consumers that a submitted proposition is true; in practice, Suppliers and other RLAY Token holders can generate Proofs-of-Coherence or other Proofs-of-Justification to verify a submitted proposition.

In this section we motivate, present and outline implementations for the Proof-of-Coherence schemes used in Rlay.

Motivation

Proof-of-Justification schemes such as Proof-of-Coherence provide necessary bounds on proposition Suppliers' competence and independence to guarantee Rlay Network's truth-tracking. Further, users can generate new information with it in a very efficient way, much more efficiently than analyzing, acquiring evidence, and/or reasoning about it themselves.

In the Rlay DIN, the presence of Proof-of-Coherence Miners and other Proof-of-Justification Miners, allows for a simple heuristic: the justification and truth-convergence of all propositions of the Rlay Network is higher the more propositions have been submitted.

Proof-of-Coherence

Proof-of-Coherence schemes are a subclass of Proof-of-Justification schemes that allow Proof-of-Coherence Miners to convince Consumers that some proposition is true.

Proof-of-Coherence together with the Rlay protocol forms an asynchronous, interactive scheme, where the Proof-of-Coherence Miner: (a) executes Proof-of-Coherence on his or her machine for some proposition given the Rlay Proposition Ledger, and then (b) asserts the mined proposition to the Rlay Network together with an optionally justification for the submitted proposition.

To the best of our knowledge, using Proof-of-Coherence schemes are the only practical algorithms to prove/justify propositions about the real-world (BonJour 1999; Landes and Williamson 2015), where the procedure is well illustrated by a crossword puzzle analogy, where entries are justified based on how well they cohere with the rest of the crossword puzzle.

Definition

A Proof-of-Coherence scheme is a tuple of polynomial-time algorithms, which can be approximated in better than polynomial-time:

(`Setup`, `Prove`, `Submit`)

- `PoC.Setup(PropTable) → PocStruc`, where `PocStruc` is a ordered and restructured subset of `PropTable`, that allows for efficient computation over its statements.
- `PoC.Prove(statement, PocStruc, PocStrucPrivate) → proposition`, where `PocStrucPrivate` is a ordered and restructured subset of some private information not public on the Rlay Network. `PoC.Prove` computes the justification of a given statement and returns a proposition based on the coherence with the relevant knowledge base.
- `PoC.Submit(proposition, utility) → integer`, where `proposition` is the previously computed justification and `utility` is determines the expected reward from submitting the `proposition` to the Rlay Network.

`PoC.Submit` computes if the proof in form of the proposition should be submitted to the Rlay Network or not.

Usage in Rlay

The Rlay Network allows for Proof-of-Coherence Miner and other Proof-of-Justification Miner to audit the propositions offered by Suppliers and ensure Competency and Independence. To use Proof-of-Coherence in Rlay, we modify our scheme to be non-interactive since there is no designated verifier, and we want any member of the network to be able to verify.

5. Rlay: A DIN Implementation

The Rlay DIN is a decentralized information network that is auditable, publicly verifiable, and designed on incentives. Suppliers pay a network of Blockchain Miners for aggregation of propositions; Payment Function Executors evaluate aggregation results. In this section, we present the Rlay DIN construction, based on the DIN definition and Proof-of-Coherence.

Frame

Participants

Any user can participate as one, some, or all of the following personas:

- *Suppliers* pay in RLAY to assert propositions (assigning them weight) in the DIN, via `Assert` requests. The decision to assert a specific proposition is generated by calculating the expected economic reward (in RLAY) for it.
- *Proof-of-Coherence Miners* are a special subclass of Suppliers who generate propositions via *Proofs-of-Coherence* and are thereby lending justification, i.e. validation, to aggregation outcomes over time.
- *Payment Function Executors* participate in Rlay by rewarding Suppliers, and Proof-of-Coherence Miners for their contributions of truthful propositions to the Rlay Network.

- *Consumers* can access and read propositions and their justification, without requiring RLAY to do so.

The Network, N

We personify all the users that run Rlay full nodes as one single abstract entity: *The Network*. The Network acts as an intermediary that runs the protocol; informally, at every new block in the Rlay blockchain, full nodes validate the well-formedness of propositions, and reward Suppliers for their truthful propositions.

The Ledger

Our protocol is applied on top of a ledger-based currency; for generality we refer to this as the Ledger, L. At any given time t (referred to as *epoch*), all users have access to Lt, the ledger at epoch t, which is a sequence of *transactions*. The ledger is append-only. The Rlay DIN protocol can be implemented on any ledger that allows for the verification of Rlay's proofs.

Data Objects

Statements

A *statement* is an ontological concept expressed in Decentralized Ontology Language, the Decentralized Ontology Language of the Network.

Propositions

A *proposition* is a statement, with an associated weight entered into the Network by Suppliers. Suppliers attach a weight in the form RLAY tokens to a proposition in order to earn tokens for the validation of the proposition.

Proposition Pools

A *proposition pool* is a set of *propositions*, whose *statement* are mutually exclusive. For each proposition pool, one aggregation result can be computed via the Aggregation Function. A proposition pool may be inferred either implicitly during the computation of the payment function, or defined explicitly in a smart-contract to allow other smart-contracts to interact with the aggregation result of the pool.

Proposition Ledger

The *PropLedger* (short for Proposition Ledger) is a data structure that keeps track of the propositions. The PropLedger is persisted in the state of the blockchain and updated when the *Assert* is called successfully. In practice, the PropLedger is used to keep the state of the DIN, allowing for quick look-ups during proof verification.

Aggregation Function

As mentioned in 5, each proposition pool has one associated *Aggregation Result*. This *Aggregation Result* can be seen as the consensus for that proposition pool. We construct an *Aggregation Function* in line with 8, to satisfy the requirements of Strategy-Proofness and Collusion-Strategy-Proofness as outlined in 3.

With the background of 8, we choose the weighted median over the propositions in a proposition pool as our aggregation function, as it satisfies those requirements. Together with Proof-of-Coherence and the Payment Function this guarantees the truth-tracking property for the Rlay DIN.

Since the specific *Aggregation Function* we use is the *weighted median over the value-restricted median-spaced domain*, this places some limitations on what the statements domains can be. (Nehring and Puppe 2006) goes into detail on what constitutes as median-spaced domain. The median spaces detailed there are sufficient for simple statements, but the domains of more complex interconnected statements may require some preprocessing to be fitted into a median-spaced domain.

Incentives & Payment Function

The Rlay is a Decentralized Information Network that incentives Suppliers to assert relevant, new, and true propositions. The rewards, computed by Payment Function Executors by running the Rlay Payment function, makes it a dominant strategy for Suppliers to submit true propositions and follow the Rlay protocol. The `Payment` function operates on four properties of propositions to assign rewards to Suppliers: `Weight`, `Pool`, `Chronology`, and `Distance`.

In addition to the reward of the `Payment`, Suppliers may receive incentive through the utility of having propositions verified to Consumers.

Weight

The `Weight` is the amount of RLAY tokens associated with a Supplier's proposition. The `Weight` of a proposition (1) determines the voting power in the propositions `Pool` (prevents Sybil attacks) and (2) makes the acquisition of as many RLAY tokens as they can a dominant strategy for honest information Suppliers.

Pool

The `Pool` is the set of mutually exclusive propositions asserted by Suppliers. The `Pool` associated with each asserted proposition (1) determines the reward allocation in each round and (2) makes it a dominant strategy for all Suppliers to contest and participate in existing proposition `Pools` so that at least 2 non-colluding Suppliers are present for each `Pool`.

Chronology

The `Chronology` is the temporal order in which identical propositions have been asserted. The `Chronology` associated with each proposition (1) determines the reward allocation inside a `Pool` for each round (2) makes it a dominant strategy for all Suppliers to contest and participate in existing proposition `Pools` as fast as they can and venture into new `Pools`.

Distance

The `Distance` is the distance between an asserted proposition and the determined weighted median by the Rlay protocol through the Payment Function Executors. The `Distance` associated with each asserted proposition (1) determines the reward allocation inside a `Pool` for each round and (2) makes it, with the presence of the other incentives and Proofs-of-Justification, a dominant strategy for all Suppliers to only assert true propositions (i.e. information).

Payouts

Every *epoch* a fixed amount of new RLAY tokens are minted. Every *epoch*, the Payment Function will be executed for all propositions recorded on the PropLedger, calculating the payout for the Supplier for that epoch. The Payment Function is constructed in a way that all the payouts for one epoch add up to the amount of newly minted tokens. Since all the propositions for a epoch are recorded on the blockchain and the Payment function is a deterministic algorithm, the calculated payouts are deterministic and auditable.

Additionally, it might also be required to mint tokens each epoch to incentivize the Payment Function Executors to execute the Payment Function.

6. Rlay Public Good Markets

Rlay Network is an ecosystem in which participants interact based on the predetermined incentive structure, with the goal of supplying a non-excludable, non-rival good: verified information. This section introduces the Rlay Information Monopsony and its first-layer as well as second-layer markets.

Public Good Information Market

Information, both verified and non-verified, constitutes as a club good (and not as a public good) because of its excludability property, as shown in 8 of the Appendix. However, Rlay, as a DIN implementation, transforms the property of excludability and therefore turns information into a public good, as the information and its consumption is open, public, and decentralized. To circumvent the issue of private provision to a public good, the protocol implements and enforces a private monopsony thereby circumventing the free-riding problem and any potentially arising market failure, problems outlined in more detail in 8 of the Appendix.

Rlay Information Monopsony

Rlay aims at transforming (verified) information into a public good, accessible, usable, and open to every individual or economic agent.

To resolve the problems associated with public goods, a monopsony on the demand side, i.e. a single “buyer” of the supplied information, is created. This single demand-side actor is the protocol itself, which is remunerating each individual supplier - given these suppliers follow their dominant strategy as laid out by the payment determinants. Moreover, through these determinants, the protocol ensures that the remunerated information suits the protocols demands, a property only feasible to be implemented when market power on the demand side is sufficiently large. As the monopsony is creating the rewards by itself, thereby ensuring liquidity, supply will always be met so long as it suffices the requirements and incentives set in its payment function.

Second-Layer Rlay Markets

Beyond the first-layer Rlay Public Good Information Market, which is at the center of the protocol, second-layer markets are enabled for the provision of verification services and information in more specific ways.

Information Bounty Market

Assuming that an agent derives utility from the output generated by the Rlay Protocol, which is in itself a public good, one may assume agents to derive utility not necessarily from the entirety of verified information gathered, but from information relevant to this respective agent.

Additionally, the protocol’s incentives may not under every circumstance suffice to meet each individual agent’s information needs, but rather the dominant information needs. Thus, a problem may occur for the individual agent with respect to selectively relevant information.

Given such cases, Information Bounty Markets as a second-layer market may come to existence. Assuming an agent (i), which derives a strong utility from having information (I) gathered and verified, would communicate a bounty on a specific proposition, which is backed by any sufficiently large number of the protocol’s native token, to be distributed to all suppliers participating in the referenced pool. Thereby, the individual agent (i) is able to steer incentives beyond the protocol’s inherent ones.

Nonetheless, one could argue that the Information Bounty Market may also be subject to free-riding, as the outcome of the bounty would subsequently be accessible to any other agent (j). Yet, the utility to be derived may be specific to the degree that no other agent would request verification, hence withholding from the market would most likely prove to be an inferior strategy.

7. Future Work & Acknowledgments

This work presents a clear and cohesive path toward the construction of the Rlay Decentralized Information Network; however, we also consider this work to be a starting point for future research on decentralized information systems. In this section we identify and populate three categories of future work. This includes work that has been completed and merely awaits description and publication, open questions for improving the current protocols, and formalization of the protocol.

On-going Work

The following topics represent ongoing work:

- Evaluation of the possibility of proposition weight withdrawal and associated effects on payment function.
- Detailed performance estimates and benchmarks for Rlay and its components.
- Evaluation of options for off-chain (zkSNARK, Truebit) or on-chain calculation of reward payments.
- A full implementable Rlay protocol specification.
- Rlay-in-Ethereum interface contracts and protocols.

- Formally prove the realizations of the Rlay Decentralized Information Network and the novel Proof-of-Coherence.

Open Questions

There are a number of open questions whose answers have the potential to substantially improve the network as a whole, despite the fact that none of them have to be solved before launch.

- Evaluation of interactions of smart-contract based mechanisms and communities (e.g. DAOs) with the protocol.
- Extension of the protocol to allow for confidential information on a public network. This could be used for e.g. revealing confidently information to a subset of Proof-of-Coherence Miner participants.
- New ontology primitives for use-cases that arise with usage of the protocol.
- Different Proof-of-Justification variants.

Proofs and Formal Verification

Because of the clear value of proofs and formal verification, we plan to prove many properties of the Rlay network and develop formally verified protocol specifications in the coming months and years. A few proofs are in progress and more in mind. But it will be hard, long-term work to prove many properties of Rlay (such as scaling, offline).

- Proofs of correctness for Expected Consensus and variants.
- Formally verify protocol descriptions (e.g. Verdi).
- Game theoretical analysis of Rlay's incentives and economics through Monte-Carlo and other evolutionary sampling algorithms.

Acknowledgments

We would like to recognize OST for providing funding as well as their utility chain and tooling to develop the Rlay and Proof-of-Coherence prototypes.

8. Appendix

Previous Work and (Im)possibility Results on Strategy-Proof Aggregation Functions

In 3 we described the core concepts of designing a *strategy-proof aggregation rule* which, given the input of private attitudes by all participating individuals, produces an outcome that is non-dictatorial and adheres to several other rationality constraints. We proceed by summarizing historic impossibility results and show the four escape-routes that lead to strategy-proof, non-dictatorial, rational aggregation rules.

(Condorcet 1785) chiefly demonstrated the potential effectiveness of collective aggregation functions (*Condorcet's jury theorem*) and the surprising problems that arise from them (*Condorcet's paradox*). (Arrow 1951) generalized this work by characterizing aggregation functions based on necessary and sufficient conditions, axiomatically describing them in terms of impossibilities (*Arrow's impossibility theorem*). (Gibbard 1973; Satterthwaite 1975) proceeded Arrow's work by showing further that there exists no aggregation rule that satisfies universal domain, non-dictatorship, the range constraint, resoluteness, and strategy-proofness (*Gibbard-Satterthwaite theorem*). Finally, (Dietrich and List 2007a), generalized Arrow's impossibility theorem by showing that there is no aggregation function that satisfies universal domain, collective rationality, independence, unanimity preservation, and non-dictatorship (*Dietrich-List theorem*).

Although these impossibility results mean that there are no aggregation functions that work under general assumptions, are strategy-proof, non-dictatorial, and produces collective rational outcomes, much research went into what conditions must be relaxed to recover from these impossibility theorems. Following, we will summarize that research and show that there are four possible escape-routes for strategy-proof and non-dictatorial aggregation functions.

(Dietrich and List 2007b) showed that *independence* and further *monotonicity* is necessary and sufficient for the non-manipulability of an aggregation rule by strategic voting (strategy-proof). From this follows, that relaxing the *independence* condition of the *Dietrich-List theorem* leads to a dead-end for us. Similarly, relaxing *collective rationality*, leading to aggregation functions that can not guarantee the desired property of *non-dictatorship* as well as most other relaxations (e.g. *range constraintness*), which lead to possible but unpractical solutions proved to be stalemates.

Although, (List 2003; Dietrich and List 2010) showed that relaxing *universal domain*, gives possibility results for non-dictatorial, strategy-proof, and collective rational aggregation functions. Further, relaxing the assumption of *universal domain* is practically sound, as it's assumption of having a peaking attitude and being indifferent to some extent about other states, covers many real-world scenarios.

Many possibility results fall under this category of domain relaxed aggregation functions, but three other possible escape-routes exist.

(Sen 1970) extended Arrow's theorem, which so far implied that attitudes are expressed ordinally and are not interpersonally comparable, to cardinal and interpersonally comparable attitudes that individuals would submit. It was later shown that these richter inputs make aggregation functions that satisfy the conditions of universal domain, non-dictatorship, collective rationality, and others required for strategy-proofness possible. Although feasible, the implications of such an assumption of interpersonally comparable attitudes remain controversial for many settings especially for preferences in a social choice/welfare setting.

Whereas the previous two escape-routes remained largely inside the classical framework of aggregation functions, the following two extend them in such a way, that strategy-proofness is achieved through other assumptions.

(Bartholdi, Tovey, and Trick 1989) described how the implied and unrealistic assumption of strategy-proofness, namely that individuals are computationally unbound, could be used to design aggregation functions which make it computationally hard for individuals to find a non-true attitude that would lead to a preferred outcome of the aggregation function. Although, no sufficient aggregation function was proposed yet.

Finally, *mechanism design theory*, a conjunction of game theory and economics, assumes that utility is transferable, e.g. via monetary payments, and that individuals have incomplete and imperfect information about what attitudes other individuals will reveal, which it uses to try to find mechanisms, i.e. game forms, that induce an equilibrium that is strategy-proof. Solutions of mechanism design theory usually suffer from being monetarily expensive to the collective.

The different approaches in all of the four escape-routes can be applied simultaneously. In fact, (Nakamoto 2008; Buterin and others 2014) borrow concepts that can be traced back to all four of these lines. In our work, by explicitly outlining these approaches and separating them from each other, we try to show that we can combine them in a new way so as to extend the applicability of strategy-proof, anonymous, peer-to-peer attitude aggregation functions.

We proceed in line with relaxing universal domain and leave the other three escape-routes for future research.

(Gruber 1993; Peleg and Sudhölter 1999; Nehring and Puppe 2002; May 1952; List 2003; Dietrich and List 2010, 2007b, 2007a; Eckert and Pigozzi 2005; Konieczny and Pérez 2005)

Information as a Club Good

Public economics traditionally classifies goods along two fundamental criteria, “rivalry of consumption” and “excludability”. A good is defined as non-rival (or non-subtractable of use), if the consumption by one economic agent (i) is not limiting the extent to which another economic agent (j) is able to consume the good. A good is, moreover, defined as non-excludable, if any economic agent (i) cannot be excluded from use and property rights can thus not be clearly assigned to any such agent individually.

Based on these two criteria, economics distinguish four types of goods:

1. Private goods (rival, excludable)
2. Club/toll goods (non-rival, excludable)
3. Common/”allmende” goods (rival, non-excludable)
4. Public goods (non-rival, non-excludable)

Beyond these two dimensions, public goods are further distinguished by their extent to which the defining properties apply. Pure public goods exhibit constant non-rivalry and non-excludability, whereas impure public goods may only be temporarily or to a limited extent non-rival and non-excludable.

Defining Information and Verified Information as a Club Good

Information can be defined as an entity resolving uncertainty. Whether it be information on the location of an object or information on an agent’s decision making.

Thus, applying the economic methodology and terminology to the subject of information, it can be inferred that information may suffice the requirements to be considered a club good in reality, limited mostly by excludability to be considered a public good. However, one should bear in mind the ongoing discussions in economic research on the nature of information and the dispute over its potential classification as an economic good.

One may argue at this point that there is a differentiation between private (confidential) information and public (non-confidential) information. This differentiation can be considered as artificially enforced by sovereign entities, hoping to safeguard misuse of information of economic relevance by a limited set of agents. Irrespective of that, the very nature and properties of both remain identical to the classification along rivalry and excludability, hence this artificial distinction is neglected.

Information, in general, can be classified as mostly non-rival. The information on an object's location may be reproduced to an arbitrary number of agents without impairing later agent's ability to also absorb this information and reproduce it again. Nonetheless, this does only hold true to a certain extent, as any agent's action resulting from the piece information may impair the consumption of information by any other agent (e.g. by removing the said object from it's stated location).

Assuming no contractually (i.e. by a sovereign entity) enforced excludability of information, one can still identify information as only partially excludable. This is mostly due to the medium over which information, subsequently knowledge, is propagated to agents. Two media can be distinguished: messages and observation, whereas both suffer from excludability. Any agents access to books, articles, speeches, the internet, and their kind is effectively excludable, despite many efforts to the contrary.

More specifically, verified information suffers likewise from excludability, despite being non-rival. Hence, both information and verified information can be defined as club goods.

(Fischbacher, Gächter, and Fehr 2001; Ostrom 2005; Samuelson 1954)

Private Provision of Public Goods & Market Failure

Public goods are commonly associated with market failure caused by the so-called "free-riding problem". Most strongly driven by the excludability property, no private market will emerge and any potentially emerging private suppliers will be driven out of the market quickly.

The following shall illustrate the free-riding problem and the resulting failure of private markets for public goods, given the example of a hypothetically existing market for an exemplary public good:

Assuming the case of national defense, one could question the absence of a private market for such, in which agents pay fees to any service provider from a pool of suppliers.

In such a private national defense market, assuming perfect competition, homogeneity, and absence of transaction costs, any agent would be able to choose and switch easily between the set of equivalent offerings. Hence, all suppliers in the market would produce at the same marginal costs (mc) = X.

Any increase in prices above marginal costs of production, would result in agents switching to the next best alternatives, as the price increase is by definition not justified (homogeneity). public increases in prices above marginal costs, on the other, are not sustainable as any individual supplier's dominant strategy would be to set prices marginally below the collusion price, in order to reap the benefits of attracting the entire market's demand. Moreover, any supplier in the market faces same, minimized marginal costs of production, as any single supplier with higher marginal costs is driven out of the market through demand switching to less costly alternatives.

Following this general setup of a competitive market, we now introduce the specific characteristics of public goods, behavioral assumptions on the demand-side agents, and thereby derive the resulting implications for the private national defense market:

National defense is neither rival nor excludable. Any supplier of national defense would protect all demand-side agents equally, as defense from any potential attacks, wars, and their kind would not be diminished by the number of agents that are to be safeguarded. Moreover, no single individual can be excluded from being protected as a collective that makes up the "nation".

Further, assuming any individual obtains a utility u(nd) from protection, it would be reasonable to assume that any agent would participate in the market as long as u(nd)>u(mc), unless the fundamental properties of national defense as a good are taken into consideration. For any single individual, the dominant strategy arising from the non-excludability from national defense would, hence, be to not participate in the market at all, as the utility from national defense is still obtained even when it is not explicitly purchased or paid for (free-riding). This then leads each agent to not purchase the services from any supplier. Further, as suppliers produce at marginal costs, no price reduction can be made to any level which could potentially alter the strategy for the agents. Hence, each supplier will be driven out of the market due to both short- and long-term losses to be expected from participation.

(Bergstrom, Blume, and Varian 1986; Buchanan 1968; Demsetz 1970)

9. Glossary

- *Decentralized Information Network*: A decentralized network of nodes that store, communicate, and process justified statements about Reality, i.e. information. Our Rlay protocol is a construction of a Decentralized Information Network.
- *Decentralized Ontology Language*: A decentralized network of nodes that store a set of concepts and categories in a subject area or domain that denotes their properties and the relations between them, which itself is denoted in the language of a description logic. Our Rlay Ontology protocol is a construction of a Decentralized Ontology Language.
- *Proof-of-Justification*: The general class of proofs that justify beliefs, observations, or information in general. Our Proof-of-Coherence is a member of the Proof-of-Justification class.
- *Public Information Market*: A market, potentially electronic, at which information, here defined as a public good, can be exchanged.
- *Byzantine Agreement*: A set of concurrent processes that can achieve coordination in spite of faulty behavior of some participants. A classical coordination problem in distributed computing which was introduced in two seminal papers by (Lamport, Shostak, and Pease 1982; Pease, Shostak, and Lamport 1980).
- *Reality*: The past, present, and future state of the root system that includes all existing particles, whether or not observable or comprehensible.
- *SNARK/STARK*: A set of instances of non-interactive zero-knowledge

proofs, a variant of zero-knowledge proofs, in which no interaction is necessary between prover and verifier.

10. Disclaimer

As of the date of publication of this whitepaper, RLAY Tokens have no known potential uses outside of the Rlay platform ecosystem and are not permitted to be sold or otherwise traded on third-party exchanges. This whitepaper does not constitute advice nor a recommendation by Project T Limited, its officers, directors, managers, employees, agents, advisors or consultants, or any other person to any recipient of this document on the merits of the participation in the TGE Sale. Participation in the Token Generation Event ("TGE") carries substantial risk and may involve special risks that could lead to a loss of all or a substantial portion of such an investment. Do not participate in the TGE unless you are prepared to lose the entire amount you allocated to purchasing RLAY Tokens. RLAY Tokens should not be acquired for speculative or investment purposes with the expectation of making a profit or immediate re-sale. No promises of future performance or value are or will be made with respect to RLAY Tokens, including no promise of inherent value, no promise of continuing payments, and no guarantee that RLAY Tokens will hold any particular value. Unless prospective participants fully understand and accept the nature of RLAY Tokens and the potential risks inherent in RLAY Tokens, they should not participate in the TGE. RLAY Tokens are not being structured or sold as securities. RLAY Tokens are sold as a functional good and all proceeds received by Project T Limited may be spent freely by Project T Limited, absent any conditions set out in this whitepaper. This whitepaper is not a prospectus or disclosure document and is not an offer to sell, nor the solicitation of any offer to buy any investment or financial instrument in any jurisdiction and should not be treated or relied upon as one. This whitepaper is for information only. Written authorisation is required for distribution of any or all parts contained herein.

All information here that is forward looking is speculative in nature and may change in response to numerous outside forces, including technological innovations, regulatory factors, and/or currency fluctuations, including but not limited to the market value of cryptocurrencies.

This whitepaper is for information purposes only and is subject to change. Project T Limited cannot guarantee the accuracy of the statements made or conclusions reached in this document. Project T Limited does not make and expressly disclaims all representations and warranties (whether express or implied by statute or otherwise) whatsoever, including but not limited to:

- any representations or warranties relating to merchantability, fitness for a particular purpose, suitability, wage, title or non-infringement;
- that the contents of this document are accurate and free from any errors; and
- that such contents do not infringe any third party rights. Project T Limited shall have no liability for damages of any kind arising out of the use, reference to or reliance on the contents of this document, even if advised of the possibility of such damages.

This whitepaper includes references to third party data and industry publications. Project T Limited believes that this industry data is accurate and that its estimates and assumptions are reasonable; however, there are no assurances as to the accuracy or completeness of this data. Third party sources generally state the information contained therein has been obtained from sources believed to be reliable; however, there are no assurances as to the accuracy or completeness of included information. Although the data are believed to be reliable, Project T Limited has not independently verified any of the data from third party sources referred to in this whitepaper or ascertained the underlying assumptions relied upon by such sources.

Please note that Project T Limited is in the process of undertaking a legal and regulatory analysis of the functionality of its RLAY Tokens. Following the conclusion of this analysis, Project T Limited may decide to amend the intended functionality of its RLAY Tokens in order to ensure compliance with any legal or regulatory requirements to which we are subject. In the event that Project T Limited decide to amend the intended functionality of its RLAY Tokens, Project T Limited will update the relevant contents of this whitepaper and upload the latest version of this to its website.

Any RLAY Tokens could be impacted by regulatory action, including potential restrictions on the ownership, use, or possession of such tokens. Regulators or other circumstances may demand that the mechanics of the RLAY Tokens be altered, all or in part. Project T Limited may revise mechanics to comply with regulatory requirements or other governmental or business obligations. Nevertheless, Project T Limited believe they have taken all commercially reasonable steps to ensure that its planned mechanics are proper and in compliance with currently considered regulations.

CAUTION REGARDING FORWARD-LOOKING STATEMENTS

This whitepaper contains forward-looking statements or information (collectively “forward-looking statements”) that relate to Project T Limited’s current expectations and views of future events. In some cases, these forward-looking statements can be identified by words or phrases such as “may”, “will”, “expect”, “anticipate”, “aim”, “estimate”, “intend”, “plan”, “seek”, “believe”, “potential”, “continue”, “is/are likely to” or the negative of these terms, or other similar expressions intended to identify forward-looking statements. Project T Limited has based these forward-looking statements on its current expectations and projections about future events and financial trends that it believes may affect its financial condition, results of operations, business strategy, financial needs, or the results of the TGE or the value or price stability of the RLAY Tokens.

In addition to statements relating to the matters set out here, this whitepaper contains forward-looking statements related to Project T Limited’s proposed operating model. The model speaks to its objectives only, and is not a forecast, projection or prediction of future results of operations.

Forward-looking statements are based on certain assumptions and analysis made by Project T Limited in light of its experience and perception of historical trends, current conditions and expected future developments and other factors it believes are appropriate, and are subject to risks and uncertainties. Although the forward-looking statements contained in this whitepaper are based upon what Project T Limited believes are reasonable assumptions, these risks, uncertainties, assumptions, and other factors could cause Project T Limited’s actual results, performance, achievements, and experience to differ materially from its expectations expressed, implied, or perceived in forward-looking statements. Given such risks, prospective participants in a TGE should not place undue reliance on these forward-looking statements. Risks and uncertainties include, but are not limited to those identified in the TGE's T&Cs. These are not a definitive list of all factors associated with a making a contribution to Project T Limited, in connection with its operations.

Project T Limited undertakes no obligation to update any forward-looking statement to reflect events or circumstances after the date of this whitepaper.

Project T Limited's business is subject to various laws and regulations in the countries where it operates or intends to operate. There is a risk that certain activities of Project T Limited may be deemed in violation of any such law or regulation. Penalties for any such potential violation would be unknown. Additionally, changes in applicable laws or regulations or evolving interpretations of existing law could, in certain circumstances, result in increased compliance costs or capital expenditures, which could affect Project T Limited's profitability, or impede Project T Limited's ability to carry on the business model and the RLAY Tokens model proposed in this whitepaper.

10. Bibliography

Arrow, Kenneth J. 1951. *Social Choice and Individual Values*.

Bartholdi, John J, Craig A Tovey, and Michael A Trick. 1989. "The Computational Difficulty of Manipulating an Election." *Social Choice and Welfare* 6 (3). Springer:227–41.

Bergstrom, Theodore, Lawrence Blume, and Hal Varian. 1986. "On the Private Provision of Public Goods." *Journal of Public Economics* 29 (1). Elsevier:25–49.

BonJour, Laurence. 1999. *The Dialectic of Foundationalism and Coherentism*. Wiley Online Library.

Buchanan, James M. 1968. "Demand and Supply of Public Goods." Rand McNally & Company.

Buterin, Vitalik, and others. 2014. "A Next-Generation Smart Contract and Decentralized Application Platform." *White Paper*.

Condorcet, M de. 1785. "Essay on the Application of Analysis to the Probability of Majority Decisions." *Paris: Imprimerie Royale*.

Demsetz, Harold. 1970. "The Private Production of Public Goods." *The Journal of Law and Economics* 13 (2). The University of Chicago Law School:293–306.

Dietrich, Franz, and Christian List. 2007a. "Arrow's Theorem in Judgment Aggregation." *Social Choice and Welfare* 29 (1). Springer:19–33.

———. 2007b. "Strategy-Proof Judgment Aggregation." *Economics & Philosophy* 23 (3). Cambridge University Press:269–300.

———. 2010. "Majority Voting on Restricted Domains." *Journal of Economic Theory* 145 (2). Elsevier:512–43.

Dietrich, Franz, and Kai Spiekermann. 2013. “Independent Opinions? On the Causal Foundations of Belief Formation and Jury Theorems.” *Mind* 122 (487). Mind Association:655–85.

Eckert, Daniel, and Gabriella Pigozzi. 2005. “Belief Merging, Judgment Aggregation and Some Links with Social Choice Theory.” In *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr. 2001. “Are People Conditionally Cooperative? Evidence from a Public Goods Experiment.” *Economics Letters* 71 (3). Elsevier:397–404.

Gibbard, Allan. 1973. “Manipulation of Voting Schemes: A General Result.” *Econometrica: Journal of the Econometric Society*. JSTOR, 587–601.

Gruber, Thomas R. 1993. “A Translation Approach to Portable Ontology Specifications.” *Knowledge Acquisition* 5 (2). Elsevier:199–220.

Konieczny, Sébastien, and Ramón Pino Pérez. 2005. “Propositional Belief Base Merging or How to Merge Beliefs/Goals Coming from Several Sources and Some Links with Social Choice Theory.” *European Journal of Operational Research* 160 (3). Elsevier:785–802.

Lamport, Leslie, Robert Shostak, and Marshall Pease. 1982. “The Byzantine Generals Problem.” *ACM Transactions on Programming Languages and Systems (TOPLAS)* 4 (3). ACM:382–401.

Landes, Jürgen, and Jon Williamson. 2015. “Justifying Objective Bayesianism on Predicate Languages.” *Entropy* 17 (4). Multidisciplinary Digital Publishing Institute:2459–2543.

List, Christian. 2003. “A Possibility Theorem on Aggregation over Multiple Interconnected Propositions.” *Mathematical Social Sciences* 45 (1). Elsevier:1–13.

List, Christian, and Robert E Goodin. 2001. “Epistemic Democracy: Generalizing the Condorcet Jury Theorem.” *Journal of Political Philosophy* 9 (3). Wiley Online Library:277–306.

May, Kenneth O. 1952. “A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision.” *Econometrica: Journal of the Econometric Society*. JSTOR, 680–84.

Nakamoto, Satoshi. 2008. “Bitcoin: A Peer-to-Peer Electronic Cash System.”

Nehring, Klaus, and Clemens Puppe. 2002. "Strategy-Proof Social Choice on Single-Peaked Domains: Possibility, Impossibility and the Space Between." *Unpublished Manuscript, Department of Economics, University of California at Davis.*

———. 2006. "The Structure of Strategy-Proof Social Choice - Part I: General Charactarization and Possibility Results on Median Spaces." *Journal of Economic Theory*.

Ostrom, Elinor. 2005. *Understanding Institutional Diversity*. Vol. 241. Princeton University Press Princeton.

Pease, Marshall, Robert Shostak, and Leslie Lamport. 1980. "Reaching Agreement in the Presence of Faults." *Journal of the ACM (JACM)* 27 (2). ACM:228–34.

Peleg, Bezalel, and Peter Sudhölter. 1999. "Single-Peakedness and Coalition-Proofness." *Review of Economic Design* 4 (4). Springer:381–87.

Pivato, Marcus. 2017. "Epistemic Democracy with Correlated Voters." *Journal of Mathematical Economics* 72. Elsevier:51–69.

Samuelson, Paul A. 1954. "The Pure Theory of Public Expenditure." *The Review of Economics and Statistics*. JSTOR, 387–89.

Satterthwaite, Mark Allen. 1975. "Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions." *Journal of Economic Theory* 10 (2). Elsevier:187–217.

Sen, Amartya Kumar. 1970. *Collective Choice and Social Welfare*. Vol. 11. Elsevier.