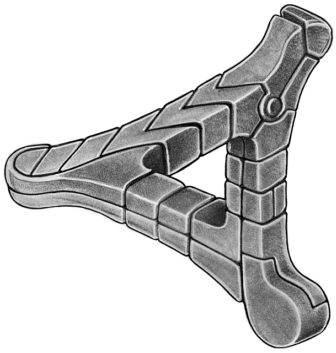

BACKFEED

DECENTRALIZED VALUE DISTRIBUTION SYSTEM

FOR BLOCKCHAIN-BASED APPLICATIONS

Cooperation is a powerful value creating force. Well-coordinated groups of individuals can accomplish together much more than the sum of what they could do on an individual basis. Yet, collective action is a difficult task that cannot be achieved without proper coordination. Many mechanisms for social coordination have been suggested thus far —each with their own benefits and drawbacks. The goal is to organise individuals into a larger structure so as to reach a common objective that would be otherwise unachievable by individuals working disjointly. But when it comes to the management of social relationships and interactions, most organisational structures known to date suffer from two main drawbacks:

1. Large-scale organizations often rely on a rigid hierarchical structure to operate. Such hierarchical structure serves as an effective coordination mechanism, but it comes at the expense of operational agility and flexibility.
2. In the absence of a hierarchical structure or central authority in charge of managing and evaluating the work of others, it is difficult to measure the value that every individual contributes to an organisation and reward them accordingly.

The **Back***feed* protocol offers a practical solution to these two challenges. It consists of a decentralized value distribution system to support *large-scale*, *free*, and *systematic* cooperation without giving up on agility or flexibility. The protocol can be used to govern the process of value distribution in the context of both formal and informal social organisations (*e.g.* social groups, communities, associations, non government organizations, companies, corporations, etc.)
In lieu of the traditional top-down and centralized coordination mechanisms, the **Back***feed* protocol is based on bottom-up decentralized collaboration (DC) mechanisms that indirectly incentivize individuals to synchronize their actions around collective goals. At the same time, these mechanisms ensure a fair distribution of the generated value to each individual, according to the perceived value of their respective contribution to the organization as a whole.
Inspired by and built on the emerging blockchain ecosystem, the protocol provides:

- the basic economic incentives for rewarding spontaneous contributions in the context of both legacy institutions and alternative, open source, decentralized organisations;
- a generalized model for the distribution of influence (or reputation) within a community, to those contributors who are the most aligned with the community's value system.

GLOSSARY & BUILDING BLOCKS

Decentralized Collaboration (DC) refers to a new type of organizational structure that brings multiple *agents* to collaborate. Each DC is defined by the following three key components:

1. a transferrable token, or a DC-issued digital asset that can be transferred to third parties
2. a non-transferrable reputation score, establishing the influence of an agent within a DC
3. a protocol for the dynamic distribution or reallocation of tokens and reputation in a DC

Agent can be an individual, one facet of an individual (as an individual can be split into multiple agents), a group of individuals, or any other entity that can act as an independent unit (*e.g.*, a DC can be an agent in another DC). Each agent in a DC has a unique account that tracks the record of actions (*i.e.*, a historical log of contributions and evaluations) and record of equity (*i.e.*, her balance of tokens and her reputation score over time). An agent may participate in multiple DCs.

Contribution can be any action with a potential value —tangible or intangible— provided by an agent to the collective. A contribution may consist of new information, an idea, a code snippet or design, providing a service, or any other sort of contribution that may be found valuable by other agents in a DC (including emotional labour).

Evaluation is the act of appraisal of a contribution —*i.e.* the determination of its value to a particular DC. Evaluation can be done on a variety of scales. Agents can evaluate a contribution (including their own) only if they have *reputation* in the DC where the contribution is evaluated.

Reputation Score (or reputation) is a *non-transferrable* unit of account that is earned by helping the DCs goals and determines the influence each agent has in the *evaluation* process for a specific DC. It can change in two ways: (a) by making a valuable contribution to the DC —reputation is gained additively, also by a new agent with no reputation; and (b) by evaluating a contribution —this can only be done by agents with non-negligible reputation, whose reputation score may increase or decrease depending on the degree of *retrospective alignment* that each agent has with the rest of the evaluating agents.

Token is a *transferrable,* value-carrying unit that serves as currency in the DC. Accordingly, tokens are used as an instrument of reward, medium of exchange, means of payment, and measure of wealth. Tokens are disbursed as a reward to an agent who made a contribution to the DC according to the specific evaluation system that is defined by the *Backfeed protocol*.

Backfeed Protocol is the technology that provides the governance mechanism for the DC. It defines the evaluation process, disbursement of tokens and reputation scores, as well as the secured record keeping. The Backfeed protocol is the "secret sauce" of the DC. The key mechanisms that drive the protocol and its inherent economic model are the tokens distribution system and the reputation scoring system.

PROTOCOL REQUIREMENTS

The **Back*feed*** protocol is actually a family of protocols, whose different parameters can be configured to meet any collaborative style and organizational need. However, regardless of its application, the protocol should always satisfy the following basic requirements:

- **Distributed governance -** There should be no centralized evaluation and reward distribution. The protocol should be able to operate in a decentralized peer-to-peer network, in a way that ensures accountability and auditability.

- **Scalability -** The protocol should support the operations of a variety of DC, regardless of their nature, objectives, and size. The protocol must accommodate organizations of any scale, from a handful of individuals to hundreds, thousands, or millions of people.

- **Adaptability -** The protocol must be able to adapt to changes within a DC, *e.g.* changes in size (growth or contraction) or changes in value systems.

- **Resilience -** The protocol should sustain any potential operational failure, fraudulent activities, or malicious attack. In particular, the protocol needs to be resilient against Sybil attacks —a particular typology of attack whereby a reputation system is subverted by forging identities in peer-to-peer networks.

- **Incentive for beneficial collaboration -** The protocol should provide agents with incentives to contribute to the DC. The protocol motivates participation by ensuring that agents will be rewarded in DC tokens as well as a higher reputation score according to the community's evaluation of their respective contributions. The protocol also motivates agents to evaluate genuinely each other's contributions by ensuring that agents' reputation in the DC will be adjusted, upward or downward, according to the quality of their respective evaluations.

- **Internal alignment -** The protocol should promote the alignment of agents within the DC. The reputation scoring mechanism must be designed to enhance the reputation (and thus the impact) of agents that further the interests of the community. By adjusting the reputation according to the quality of evaluations, the agents whose evaluations are the most aligned with the community's values acquire a stronger level of influence within the DC's evaluation system.

- **Diversification -** In addition to encouraging the internal alignment of agents within a DC, the protocol should also encourage diversification at the external level, by facilitating the exploration of alternatives features (*e.g.* spin-offs or 'forks'), which might prove to be the most suitable to a particular context.

THE BACKFEED PROTOCOL

TOKENS DISTRIBUTION

The Backfeed protocol can be used to encourage more individuals to contribute to a DC's project or initiative, by rewarding them with a number of transferable (economic) tokens.
Whenever a contribution is evaluated positively within the DC community (*i.e*. it reaches a positive median value), the DC might decide that —along with the new reputation— a specific amount of tokens should also be issued to the contributor(s), according to a particular value function that must be collectively agreed upon by the DC.
Just like currency, DC tokens are merely an indication that value has been created at some point, by someone. Therefore they are not inherently linked to the person to which they have been issued in the first place —and they can, as such, be freely exchanged on the market as any other currency or commodity.
We define here two key elements needed for the rigorous quantitative specification of the token distribution mechanism within a DC:

Evaluation set
The evaluation set is the set of possible values which a person can evaluate a contribution with. It can be continuous (*i.e.* all real numbers between 0 and 1) or discrete (*i.e*. any natural number from 1 to 5); and it can be either bounded or unbounded. The Backfeed protocol addresses all of these cases.
It is important that the evaluation set be DC-specific, as it might be desirable to have different types of evaluation sets for different applications. For example, a simple DApp that merely collects a list of contributions into a joint repository might only need to rely on binary evaluations (*e.g*. up or down votes), whereas a more general purpose DC which accepts contributions of many different kinds might require an evaluation set which is continuous and/or unbounded, so that users can evaluate contributions without any limitation of sort.

Value function
The role of the value function is to specify the worth (in DC tokens) of an evaluation in the evaluation set. It is the bridge between the evaluation set and the actual amount of tokens that will be issued and distributed to the relevant contributor(s) to a DC. It is important to note that this function does not need to be linear, although for consistency it must be monotonically increasing ($a < b \rightarrow f(a) < f(b)$).
E.g., given an evaluation set of $E = \{1, 2, 3\}$ and a value function $v : E \rightarrow R$ equal to $v(1) = 0$, $v(2) = 10$, $v(3) = 50$ ⇒ a vote of 3 would be equivalent to 50 DC tokens.
Tokens issuance

The amount of tokens distributed to the contributor depends on the ***median*** value of all weighted evaluations —accounting for the total reputation of the DC (and not just that of the voters). This means that tokens will only be issued after there is at least 50% of the overall reputation of the DC community engaged in evaluating a particular contribution.

Note that using the median, rather than the weighted average, renders the protocol more resilient to (adversarial) agents making indefinitely high evaluations so as to falsely inflate the weighted average of a contribution (and therefore the contribution reward in DC tokens). Indeed, as opposed to the average which is directly affected by every evaluation, the median is basically a 50% threshold: as long as no more than half of the reputation engaged in evaluations is contaminated, the median will not reach an arbitrarily large result.

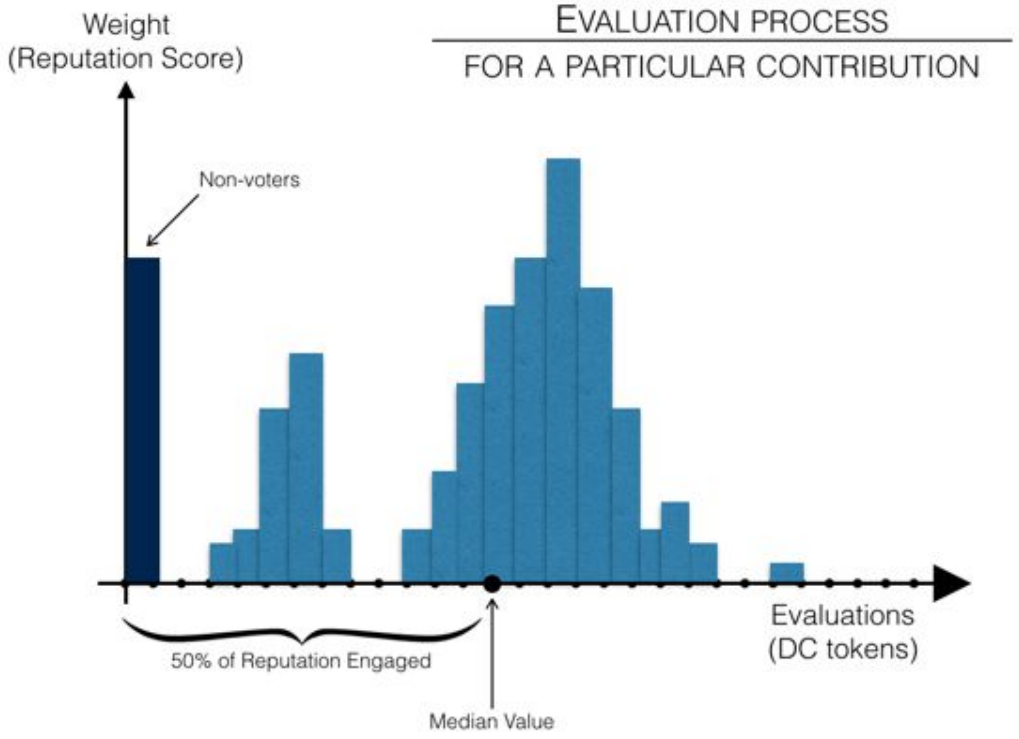

Fig. 1: Illustration of the Backfeed protocol for token distribution based on the median value

Note that all DC members who did not evaluate a contribution count as having made an evaluation of 0 (but their reputation is not accounted for in the total sum of engaged reputation).

Anyone with a non-negligible amount of reputation within the DC can evaluate a contribution at any point in time over the lifetime of the DC. If they so wish, people can also change their mind over time. Changing the value of an evaluation will qualify as a new evaluation (with all the necessary costs this might entail) —whereas the former evaluation will simply be deleted from the system (*i.e.* it will no longer count towards the median value).

Tokens are issued on an on-going basis, with respect to the current median evaluation. New tokens are issued every time the median goes over its previous maximal value, and the differential value is issued to the contributor. Hence, a contribution that initially remains unnoticed and eventually becomes more relevant over time will reward the contributors with a flow of new tokens as time passes. Conversely, if the median value reaches a peak and then decreases, issuance of new tokens is frozen until the value grows back above that peak.

REPUTATION FLOW

While token distribution can be useful to incentivize contributions, the Backfeed protocol is, first and foremost, a **public reputation system** which determines the rules according to which new reputation is both issued and distributed within a particular DC.

A reputation system is public if the reputation of an agent in the system is the result of information which has been publicly authenticated —or, in other words, information which does not need to be authenticated by every individual member of the network. More precisely, the Backfeed protocol relies on *objectively subjective* information (*i.e.* information that can be regarded as being subjective to a network, but objective within that network). This makes it possible for people to judge the worth of any given agent according to that specific value system that characterises a particular network.

As such, the Backfeed protocol distinguishes itself from most personal reputation systems which only reflect the subjective perception of a particular agent within the network. These systems rely on subjective information —*i.e.* information which has been authenticated by the agent herself. They are generally much more resilient than public reputation systems, but they are also more limited in that they require agents to authenticate every bit of information (although many of such systems rely on a web-of-trust mechanism).

Backfeed is the first public reputation system known to date that is fully decentralized and resilient against Sybil attacks.

As opposed to DC tokens, which can be freely transferred to third parties, reputation in a DC is an indication of how much an individual is aligned to a particular DC's value system. As such, reputation is inherently linked to the specific agent who earned it, and cannot therefore be transferred to anyone. Reputation can be earned in two distinct ways:

- **Contributions**:
  Anyone can acquire reputation within a particular DC by making contributions that will be evaluated positively within the DC. Each contribution is evaluated by all (or a portion of) DC members, and the evaluation is weighted by their respective reputation within the

DC. When the median reaches a positive value, new reputation is issued and distributed to the corresponding contributor(s) based on the aggregated (weighted) evaluation.

- **Evaluations**:
  Anyone can gain or lose reputation by making evaluations of others' contributions within a particular DC. The evaluation process is based on a positive-feedback loop, whereby evaluators gain or lose reputation according to the retrospective alignment of their evaluations with regard to the rest of the community. Note that, given the evaluation process is costly in terms of reputation (*i.e.* evaluators pay with a portion of their reputation in order to make an evaluation), it can only be done by agents who have a non-negligible reputation within the DC.

Reputation issuance for new contributions:
The reputation gain when a contribution is evaluated by a DC is quite straightforward: Reputation is issued to contributors, on an on-going basis, whenever the median value of their respective contributions reaches a positive value, *i.e.* when more than 50% of the DC reputation considers that a contribution is valuable. Hence, new reputation cannot be issued without a consensus within the network. The precise amount of reputation to be issued for each evaluation is a function of the evaluation set that is specific to any given DC. It needs, therefore, to be defined, on a case-by-case basis, by each individual DC.

Reputation flow upon evaluation:
In the following section, we focus on the reputation flow algorithm which is triggered by new evaluations. As mentioned, different DC's might adopt different evaluation sets; however, for the sake of clarity, we present here the simplest case of binary evaluation, namely where the evaluation set is $\{0,1\}$. In the appendix, we cover the generalization of this scheme both for discrete and continuous evaluation sets (bounded and unbounded).
In order to better explain the reputation flow Table 1 lists the variables relevant to the protocol.

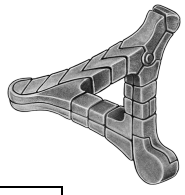

N	The current number of users in the system.
r_i	The current reputation of the i'th user.
R	The current sum of all the reputation in the system $\left(R = \sum_{i=1}^{N} r_i\right)$
e_i	The evaluation of the i'th user of contribution C (in {0,1})
U_k	The set of users which have evaluated contribution C (at the time the k'th user evaluates, inclusive of the k'th user)
V_k	The sum of reputation voted on contribution C. $\left(V_k = \sum_{j \in U_k} r_j\right)$
W_k	The sum of reputation voted on contribution C with equal value as the k'th user's vote (inclusive of the k'th user). $\left(W_k = \sum_{j \in U_k} \delta_{e_i,e_j} r_j\right)$ note: $\delta_{i,j} = 1$ when $i = j$ and $\delta_{i,j} = 0$ otherwise.
s	The stake fraction (in [0,1])
d	The stake distribution fraction
$r'_i(e_k)$	The new reputation of user i once user k has evaluated contribution C.
α	Reputation payment skewness

Table 1: Listing of relevant variables for the reputation flow triggered by new evaluations.

Binary evaluation reputation flow

The Backfeed protocol determines the change in reputation of DC members during and following an evaluation of a contribution. Overall the reputation dynamics consists of two complementary elements.

- The first element is the **evaluation cost** (or *reputation stake*) which every evaluator pays upon making an evaluation. This cost is levied multiplicatively —*i.e*. a certain fraction of the evaluator's reputation is deducted from its overall reputation. The rationale behind the evaluation cost is to encourage genuine participation and effort, as only those who are

willing to put their reputation 'at stake' will be able to make an evaluation. In order that the evaluation cost of an evaluation be proportional to the potential reputation gain, we introduce a cost function which depends on the amount of reputation which has already voted on the contribution. Specifically, the cost decreases with the voted reputation so that last evaluators break even with non evaluators. The evaluation costs is also a mechanism to avoid spam evaluations, since any agent doing bad evaluations will eventually lose most (if not all) of its reputation, and will thus no longer be able to evaluate contributions.

- The second element is **reputation reallocation**, which consists of reputation *stake distribution* to all previous aligned evaluators. The distribution mechanism is a means to incentivize early evaluators —since the earlier the evaluation is made, the greater the potential gain from later evaluations is.

These two stages combined constitute the basic logic of the Backfeed reputation flow algorithm. Each of these parameters are more rigorously specified and expanded upon in the following sections.

Evaluation cost function

As mentioned the evaluation cost decreases with the amount of previously voted reputation. So the cost for the k'th evaluator is given by:

$$Cost(s, r_k, V_k, R, \alpha) = s \cdot r_k \cdot \left(1 - \left(\frac{V_k}{R}\right)^{\alpha}\right)$$

The relevant variables are defined in Table 1. The formula consists of several elements. First it should be noted that $s \ll 1$ is a small fraction and that the right most parenthesis are always less than Note that the last evaluator, i.e. the one which completes the 100% of reputation voted, does not need to pay which is in agreement because he can not gain anything either, as no one can vote after him.

Reputation updates

We can now put together everything and write down how the reputation of all the users changes following a new evaluation by the k'th user.

$$r'_k = r_k - Cost(s, r_k, V_k, R, \alpha)$$

For the previous evaluators we have:

For $i < k$

$$r'_i = r_i + \delta_{e_i,e_k} \cdot d \cdot r_k \cdot \frac{r_i}{W_k}$$

The delta ensures that only users which have evaluated like the current evaluator gain reputation. The amount gained is a product of a fraction d multiplied by the current evaluators reputation and the reputation of the user divided by the total reputation which has voted like the current evaluator. So everybody gets in proportion to how much reputation they have voted with (their stake).

Figure 2 illustrated the reputation flow in a simple scenario in which 10 users evaluate a single contribution consecutively by the same value. The x-axis is the the index of the evaluations and the y-axis is the normalized reputation of the users. The first user (orange) evaluates and therefore loses a fraction of his reputation. When the other users vote, he starts to gain reputation. The second user (yellow) evaluates and again loses some reputation, though a smaller amount, according to the cost function, and gains when the rest evaluate and so on. Each evaluator is taking a risk that the next evaluators might not evaluate like him or might not evaluate at all - In that case he might lose a portion of his reputation.

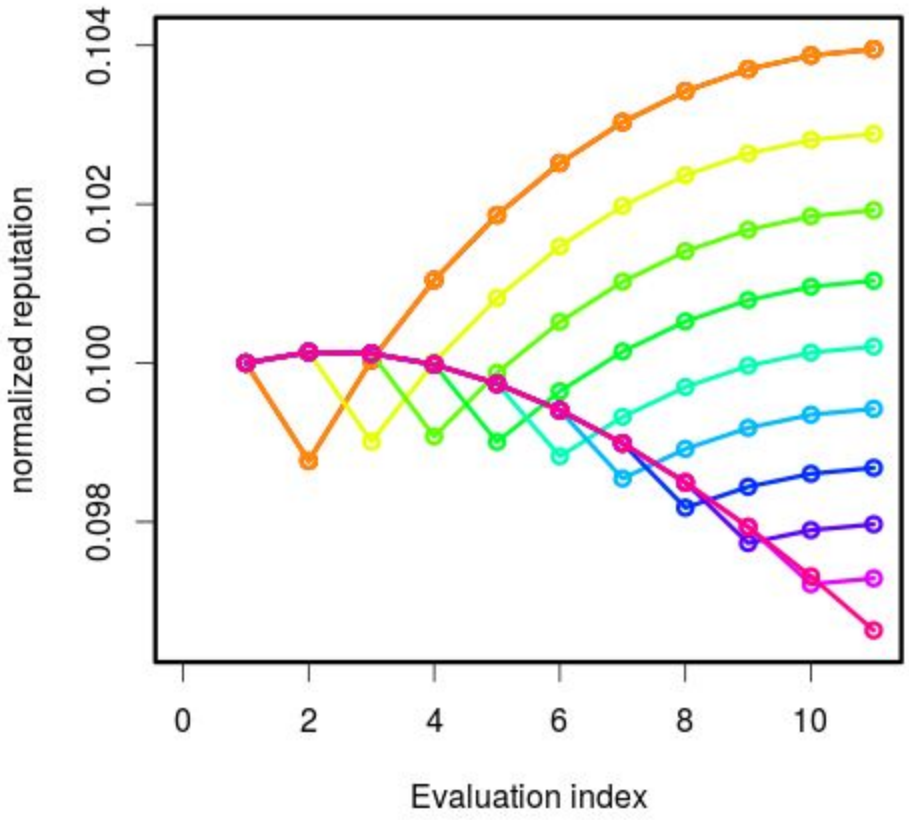

Fig. 2: Illustration of the reputation flow for a simple scenario in which 10 users evaluate the same contribution consecutively by the same vote.

Overall, the protocol encourages people to evaluate contributions at an early stage, since the earlier an evaluation is made, the greater are the potential rewards to be earned. Eventually, as others evaluate the same contribution with a similar evaluation, those who are the most in line with the overall community's evaluation will be able to retrieve the reputation they lost, and often gain more reputation than they initially had.