

A Scalable Voting System: Validation of Holographic Consensus in DAOstack

<p>Youssef Faqir-Rhazoui Dpt Software Engineering and Artificial Intelligence Universidad Complutense de Madrid (Spain) yelfaqir@ucm.es</p>	<p>Javier Arroyo Institute of Knowledge Technology Universidad Complutense de Madrid (Spain) javier.arroyo@fdi.ucm.es</p>	<p>Samer Hassan Berkman Klein Center Harvard University (USA) Institute of Knowledge Technology (Spain) shassan@cyber.harvard.edu</p>
--	---	--

Abstract

Blockchain technology allows for new forms of online governance relying on decentralized infrastructure. In particular, it has enabled the so called Decentralized Autonomous Organizations (DAOs): pieces of software deployed on a blockchain which mediate the interaction of groups of people. Allegedly, their main aim is to facilitate large-scale decentralized cooperation in online communities. In order to do that, they provide governance mechanisms for democratic decision-making in communities. One of them is "Holographic Consensus", a voting mechanism intended to filter relevant proposals for large communities. To which extent is Holographic Consensus working as intended, facilitating effective scalable cooperation? In order to validate this method, we analyze the 22 DAO communities and 6000 users from the DAOstack platform, which uses Holographic Consensus in all its DAOs. Our results show that this mechanisms seems to follow the expected behavior, facilitating scalable decision-making in large DAOs.

1. Introduction

Blockchain is an emergent technology that provides a distributed and immutable ledger, which can be thought as a distributed database with a synchronization mechanism. Like the Internet, (public)¹ blockchains constitute an open infrastructure, which is not owned or controlled by a central authority. Since it is replicated by each node in a network, the ledger is transparent to users and third-parties [1, 2]. It aims to provide a mechanism in which decentralized transactions and operations are secure, without needing to trust a third-party as in centralized systems. [2, 3].

¹Typically, referring to "blockchains" is understood as referring to the public blockchains like Bitcoin or Ethereum, which are transparent, have miners, and have a high degree of decentralization. However, "private" or "permissioned" blockchains also exist, which have access control, are not transparent, and typically rely on a group of entities controlling access. In this paper, we refer just to the first.

Blockchain emerged with the appearance of Bitcoin, the popular "cryptocurrency", but has grown enormously since then. Such growth has been promoted by the *Ethereum* project, which provides a blockchain-based distributed computing platform [4]. Ethereum provided a fully-fledged programming language, *Solidity*, to build decentralized applications, *Dapps*. In addition, it facilitated the creation of "smart contracts", computational agreements between parties which may be self-executed and self-enforced.

These Dapps have been applied in many fields [5] including financial applications such as banking services [6], or cryptocurrency payment [7], but also applications in fields such as Internet of Things [8]. More relevant to this paper, Ethereum boosted blockchain applications to governance, allowing the rise of Decentralized Autonomous Organizations, or simply DAOs [9].

In plain terms, DAOs are organizations where the interaction of members (humans or machines) is mediated by a blockchain application that is controlled by a set of rules embedded in its source code. DAOs are said to be able to autonomously hire people, provide services, gain money for their own aims, own smart property, coordinate with other autonomous software, or facilitate cooperation, to name a few [3, 10]. Still, while this is yet to be seen in this emerging paradigm, these promises have attracted both idealistic defenders [11, 12] and those warning of its dangers [13, 14].

DAOs are supposed to enable large-scale cooperation in online communities. Large-scale cooperation has been a matter of debate for decades, and the Internet has facilitated the emergence of online communities where this is possible [15]. Thus, we can observe large commons-based peer production projects such as Wikipedia or the Linux kernel [15]. Still, these communities face multiple challenges, such as the emergence of bureaucratization and elites [16]. Blockchain enthusiasts claim DAOs can reduce the cost of large-scale decentralized cooperation, incentivize cooperative behavior, increase participation, and facilitate open democratic organizations [17, 18].

One of the issues in large-scale online communities is attention: after reaching certain scale, participants cannot review (and vote, when possible) all submitted proposals. Thus, some sort of proposal filtering is required, so participants give their attention to the proposals which are most aligned to the community ends. This is the principle behind the mechanism of *holographic consensus* [19]. This mechanism establishes some sort of prediction market that acts as filter for the community, enabling them to make predictions about which proposals will be accepted, rewarding those that were accurate.

In this work, we will analyze the emergent DAO ecosystem, DAOstack, that implements *holographic consensus*, to empirically validate if such method facilitates large-scale proposal voting. With this method, high-quality proposals aligned with the purpose of the DAO are supposed to be selected and will require a simple majority vote instead of absolute majority. While the idea seems good in theory, does this innovative blockchain-enabled mechanism effectively facilitate large-scale voting? Is it actually used in practice? Does it work for DAO communities?

In order to respond to these questions, we analyze the activities of 6,000 DAOstack users organized in 22 DAO communities. Section 2 covers DAOs, the governance scalability challenges, and the governance solutions provided by DAOs. Section 3 explains holographic consensus, while Section 4 provides an analytical overview of DAOstack activity and user-base. Section 5 focuses on the empirical validation of holographic consensus based on the DAO communities operations, and Section 6 provides some concluding remarks.

2. DAOs in a nutshell

There are multiple definitions of DAO. As defined by the co-founder of the Aragon DAO platform: “a DAO is an internet-native entity with no central management which is regulated by a set of automatically enforceable rules on a public blockchain, and whose goal is to take a life of its own and incentives people to achieve a shared common mission”². In short, it would be a group of people with common goals that coordinate online relying on blockchain infrastructure.

In practice, DAOs are pieces of software which define the interactions of its participants. Their source code is typically deployed on a blockchain such as the aforementioned Ethereum. Thus, each DAO has a unique address it can be referred with. Following its most common implementation nowadays, DAOs make use of “governance” tokens, which participants need to have in

order to participate. Usually, DAOs achieve governance through a proposal system. It is also common for DAOs to manage resources, typically cryptocurrencies, and DAO members can decide how to allocate them through the decision-making system. For example, DAO members send proposals³ asking for resources (typically, cryptotokens) in exchange for some work (coding smart contracts, auditing some code, writing documentation, etc). Such proposals are supposed to help the community to reach its goals.

Nowadays, the vast majority of DAOs use Ethereum as a platform. In Ethereum, every operation or transaction a user wants to do, implies a small “commission”, a cost in crypto-currency (“gas”). As opposed to traditional centralized web apps which imply a running cost for server rental, Ethereum apps do not have running costs, but a cost per operation performed, typically paid by the user performing them. That is due to the fact operations are validated and confirmed through blockchain “miners” which are paid these commissions. The cost of an operation may be more or less expensive depending on the kind of activity to be performed. These facts condition the DAO activity and the kind of activity that is recorded in the blockchain.

2.1. The governance scalability problem

Typically, electronic voting systems focus on providing a secure infrastructure with encrypted and verifiable votes, which is trustworthy for voters [20]. Under a blockchain, those technical issues are allegedly reduced, due to its nature as a distributed ledger without central management, and with its transactions validated cryptographically [21]. We are not tackling here the technical scalability challenges that blockchain faces [22]. Instead we focus on how blockchain and DAOs deal with the classical challenges in political organization, on how to scale direct democratic participation to large groups of people [23, 24].

It is already common to find arguments in which this scalability is facilitated by modern technologies [25], and DAO proponents aim to take it a step further [18]. Within the DAO proponents, the co-founder of the DAOstack DAO platform, Matan Field, describes as the most prominent DAO governance problem the tension between scalability and resilience [26]. He defines resilience as the tolerance and even resistance of a governance system to faulty behaviour. And he sees scalability as the ability of a governance system to take a large number of decisions in a given period as more agents participate in the system. As a community

³Typically, the proposals and their outcome are available as transactions in a public blockchain, e.g. Ethereum. Additionally, DAOs have a web to show the proposals and their state.

²<https://twitter.com/licuende/status/1263511552709267456>

grows, if an absolute majority (>50%) is required to make a decision, more agents will need to vote. On the other hand, resilience in a decentralized community imply decisions not being hijacked by a small minority. Thus, growing a community increases the number of voters required to satisfy the majority, while growing the number of decisions demands more attention from voters in a time frame. Hence, scaling an organization, either in terms of members or decisions, hinders the governance of the organization. Reducing the number of voters required to take decisions (e.g. >30%) is no solution, because it compromises the organization resilience that could be hijacked by a minority. The duality between scalability and resilience is addressed in different ways by the DAO platforms.

2.2. Scalability solutions in DAO platforms

In the last years, several DAO platforms have emerged to facilitate DAO adoption [27]. They provide tools and templates to create a DAO, and a front-end to interact with it.⁴ These platforms typically implement solutions to tackle the governance scalability problem.

For example, *DAOHaus*⁵ provides a simple decision-making system that requires no minimum quorum to approve proposals. It just counts cast votes, and if there are more than 50% up-votes, the proposal passes [28]. However, if any member does not agree with the result, they can make *rage quitting*, exiting with their portion of resources.

On the contrary, the DAO platform *Colony*⁶ addresses the scalability problem with two measures. First, splitting DAOs into domains which are potentially independent, and can work without interactions among them in a stigmergic manner. Second, by avoiding voting, as much as they can. All decisions are approved by default unless someone has an objection, in which case it is discussed and solved by voting [29].

The DAO platform *Aragon*⁷ requires reaching a static quorum to approve proposals. However, Aragon DAOs are highly customizable, as they can install and configure 'apps' to change multiple features, including the voting system. For example, it is possible to use the *Conviction Voting* system [30], which prevents several typical attacks of other voting systems, and which was tested using agent-based simulation [31]. This system aggregates the preferences of users voters and the longer they keep their preference for the same proposal, the

⁴It is necessary to have a special browser or an interface to use these front-ends. For instance, MetaMask <https://metamask.io/> is a browser extension used to interact with Ethereum's Dapps.

⁵<https://daohaus.club>

⁶<https://colony.io>

⁷<https://aragon.org>

stronger their conviction gets. Moreover, in Aragon it is even possible to use the *DAOstack* voting system, which is discussed next.

3. Holographic Consensus in DAOstack

As mentioned above, absolute majority as a community decision-making mechanism is not scalable. To solve that, *DAOstack*⁸ has implemented the *Holographic Consensus* that aims to allow scalability without harming resilience [19].

Holographic Consensus requires to introduce some concepts in order to understand how it works. First, the *DAO's global opinion* is defined as the decision made by an absolute majority of voters (>50%). An approximation of the *DAO's opinion* involving a sufficient number of agents on the decision could be used instead (e.g. relative majority). However, in that case, a small number of agents could hijack the *DAO's global opinion*. Attackers could spam a group of proposals in a small frame time, and this would require the attention of most voters to prevent the attack. Therefore, relative majority increases the scale factor of a DAO, but also increases the potential attacks.

Because of that, *Holographic Consensus* sets an absolute majority to approve a given proposal, and only if some conditions are met, a relative majority suffices. The process that turns an absolute majority into a relative one is called *boosting*. Boosting a proposal consists of people spending tokens to promote it, until either a threshold is reached or the boosting period ends. Anyone can boost a proposal, not only the members of the DAO, i.e. foreign agents are also allowed to boost proposals. Due to that, we may consider that two "economies" arise in this context. In the first one, only DAO members have a portion of its tokens (reputation holders), and these reputational tokens cannot be bought or transferred, and are used to vote. The other economy is bound to the boosting process, using another kind of transferable tokens (akin to a cryptocurrency), which may be bought, and thus anyone, not just organization members, may use them to boost a proposal. The boosting process is explained in detail below.

When a proposal is created, a boosting period starts and anyone can "upstake" or "downstake" her tokens to that proposal, i.e. bet a certain number of tokens that the proposal will pass (upstake) or not pass (downstake) after being voted. If the *upstaked* tokens of a proposal reach a dynamic threshold, which grows exponentially with the number of active boosted proposals, it is considered that the proposal is relevant enough, and then it is boosted. Boosting proposals is also a way

⁸<https://daostack.io>

of fast-track them. Finally, when the proposal has an outcome (the result voted by the DAO members), stakers will lose or gain tokens, depending on their bet. The ones who were aligned with the *DAO's opinion* will gain tokens, while the ones who were not, will lose them.

Matan Field [19] argues that stakers are also predictors because if they do not want to lose tokens, they need to anticipate the *DAO's global opinion*. Therefore, if stakers are rational, they will boost 'good proposals' and filter out 'bad proposals'. Consequently, proposals boosted need less votes to be approved, which saves community attention and voting effort, while those not boosted will require a majority vote. Note that a non-boosted proposal is a proposal that has not received enough upstakes due to downstakes or inaction.

As a result, DAO members need stakers for the sake of scalability, and stakers need to anticipate *DAO's opinion* in order to earn a reward. Note that the boosting process can be thought as a poll to anticipate the *DAO's opinion*, or more accurately as a prediction market [32]. We will validate its predictive power in Section 5.

4. DAOstack in numbers

Before analyzing the use and predictive power of *Holographic Consensus* in *DAOstack*, we will analyze its DAOs to better understand the ecosystem.⁹ Data and source code are open-licensed and publicly available.¹⁰

Since its start in April 2019, *DAOstack* has deployed 22 DAOs on Ethereum main network (*mainnet*), which means that they are fully deployed in a public blockchain. At the beginning, it was not possible to directly create a DAO into *DAOstack*, and it was required to ask for it to its administrators, but now a template facilitates the creation process.

There are 6,083 users in all the DAOs within the platform. Each user is identified by an Ethereum address, but the same address can be used in more than one DAO. If we take this into account, *DAOstack* has 5,952 unique users, which means that around 2% of the addresses are used in more than one DAO. It is also worth to mention that, as in many online communities, a user may have different addresses (i.e. user ids), but there is no reliable way to detect them. Another important fact is that *Kyber DAO Exp#2* is a DAO with 4,946 users. If we excluded this large DAO, *DAOstack* would have 1,137 users operating in 21 DAOs.

⁹The historical data was collected on June 24th, 2020

¹⁰Standalone scripts can be found in <https://github.com/Grasia/daostack-census> and the DAO-Analyzer software in: <https://github.com/Grasia/dao-analyzer>. All the data used in this paper has been gathered from *DAOstack GraphQL* in <https://thegraph.com/explorer/subgraph/daostack/master>.

4.1. Users

Figure 1 shows the distribution of *DAOstack* DAOs according to its number of users. The y-axis represents the number of DAOs that have a given number of users (in bins of 10 users), with the red trace showing that one DAO has 0 users. To better display the histogram, we have removed from the plot the outlier *Kyber DAO Exp#2*, which has 4,946 users.

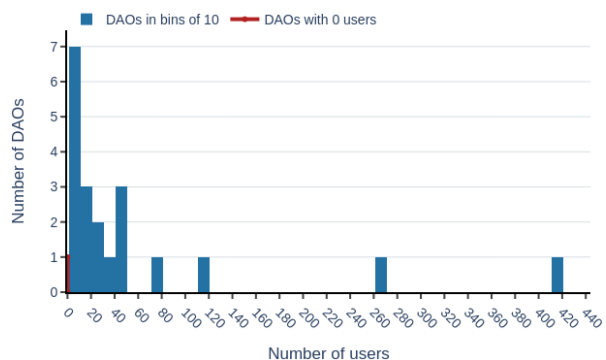


Figure 1. Community size distribution in *DAOstack*. *Kyber DAO Exp#2* (4,946 users) was filtered out.

The community size distribution shows that most DAOs have few users, while the majority of users are concentrated in a few big DAOs. More precisely, 50% of the DAOs have less or equal to 21 users, 75% of the DAOs has less than 47 users, and 95% of them has less than 409 users. This kind of heavy-tail user distribution is not different from that of other online communities, like users by *Wikipedia* language [33].

Figure 2 shows the monthly time series of new users.¹¹ It can be seen that new users join *DAOstack* each month, even if in the last few months the number has decreased towards a mean of 25 per month. *DAOstack* is still in early Beta and joining the platform currently requires a technical knowledge, these facts may be a barrier to new users.

The highest value was registered in June 2019 with 5,397 new users, which implies the 89% of all ecosystem users.¹² However, 4,944 of them joined to *Kyber DAO Exp#2*, the biggest DAO already mentioned. Still, even if we remove this DAO's users, June 2019 still registered the biggest amount of new users (453). This is almost three times more than the second month with more newcomers, April 2019, which was the first month of *DAOstack*. Besides these peaks, the number of new users is between 16 and 80 in the rest of months.

¹¹Note the series last month, June 2020, is drawn lighter because the month was not finished at the time of the data collection (June 24th), and thus its final value may be higher.

¹²The y-axis in the figure ignores the very high June 2019 value to favor plot legibility.




Figure 2. Time series of new users

4.2. Activity

The definition of activity depends on the purpose of the online community, e.g. wiki communities usually consider "edits", while free/open source development considers "commits". DAOs are mainly governance tool to submit proposals, and vote for or against them. *DAOstack* also includes stakes, which are a kind of vote not restricted to members of the DAO (see Section 3). Thus, we will consider the registration of a new proposal, voting, and staking as activity indicators.

According to that definition, there were a total of 7,719 actions in *DAOstack* since its creation. Registering proposals represents a 14% (1,085) of the actions performed, voting represents a 60% (4,617), and staking represents 26% (2,017). Unsurprisingly, the number of proposals is the smallest, but it is quite high considering the number of communities. Interestingly, the number of stakes is much smaller than the number of votes, which could be explained because staking implies spending both money (GEN token) and time for adequately pondering a proposal. Furthermore, staking is a mechanism designed for big DAOs, and we have seen that half of the DAOs have 21 or less members.

Figures 3, 4, and 5 show the distribution of the stated kinds of activity, i.e. proposals, votes, and stakes, across DAOs. The proposal distribution in Figure 3 shows that all DAOs but two have at most 90 proposals. The exceptions are *Genesis Alpha*, the DAO that governs the *DAOstack* project, with 382 proposals, and *dxDAO*, with 179 proposals. However, 77% of DAOs have made 50 or less proposals, and there are two DAOs that never made a proposal. It is clear that some DAOs are extremely active, while most of them exhibit a more modest activity profile. Again, this uneven behavior is similar to that in online communities such as wikis [34].

In the distribution of votes and stakes we see again the pattern of few communities capitalizing most of the activity, while most of them having moderate activity

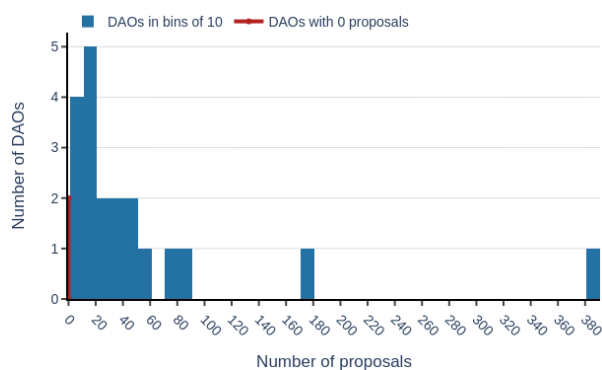


Figure 3. Distribution of proposals across DAOs.

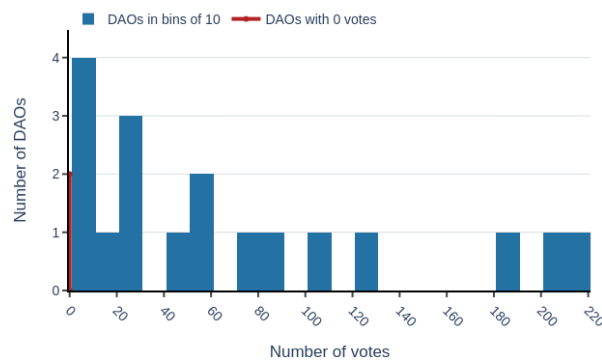


Figure 4. Distribution of votes across DAOs (excluding two large cases)




Figure 5. Distribution of stakes across DAOs (excluding one large case)

levels. We omitted the outliers for the sake of legibility: *Genesis Alpha* with 2,201 and *dxDAO* with 1,167 votes were excluded from Figure 4, while *Genesis Alpha* with 1,078 stakes was excluded from Figure 5.

The 75% of DAOs have at most 121 votes and 47 stakes, while 95% of them have at most 1,119 votes and 313 stakes. The mode in both cases is between 0 and 10: six DAOs in the case of votes and nine in the case of stakes. These values are very low, specially for the

case of votes, which probably means that these numbers correspond to abandoned DAOs.

It is worth noting that six DAOs registered no stakes (red trace), with five of them having less or equals to six users. The other is *Fortmatic DAO*, which has 22 users, but it has not registered any actions. It is remarkable that there is an eccentric case, *dOrg* with 41 users, 212 votes but just two stakes. This DAO has an engaged community, where proposals are usually approved by absolute majority. However, DAOs with that number of users usually use stakes to boost proposals.

All the activity distributions are highly skewed. We have not found a clear correlation between the number of users and activity, because there are mostly inactive DAOs with a lot of users and some small DAOs that generate a lot of activity. Still, there is some evidence of a positive correlation between users and activity. We have omitted the scatter plots and the correlation values, given the small sample and the presence of outliers.

4.3. Activity over time

In this section we focus in community activity over time. As above, our definition of activity includes creating a proposal, voting or staking. We consider that a DAO is ‘alive’ or active in a given month if at least one action took place in that month.




Figure 6. Monthly time series of active DAOs

Figure 6 shows the time series of active DAOs in each month (that is, excluding DAOs with no activity). The time series has an average of 7.8, which means that only around one third of the DAOs are active each month. As before, June is marked in light color as data collection finished before that month ended.

Due to the aggregation of all the DAOs in the mentioned chart, it is not easy to see if the active or inactive ones are always the same. Figure 7 complements such information. In Figure 7 we represent each DAO with a bar with the number of months in which the DAO was active. DAOs which have

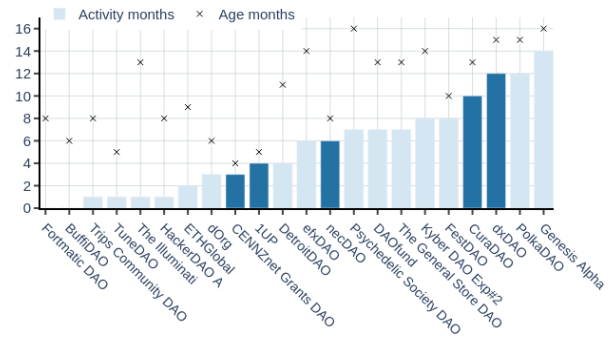


Figure 7. DAOs active months since birth date. DAOs in dark blue have registered activity in the last two months (May-June 2020).

performed at least one action on the last two months (May-June) are in dark blue, while others are in baby blue. Furthermore, a cross over each bar represents the age in months of the DAO, that is, the maximum number of months that it could have been active. We can see that only two of them have always been active, the young *CENNZnet Grants DAO*, and *IUP*. In the other side, we can see two DAOs that never registered an action: (*Fortmatic DAO*, and *BuffiDAO*). In general, we can see that the activity level varies, while some DAOs are usually active, some others seem to be less active, and a few of them seem to be abandoned. The abandonment of an online project is typical in other fields such as wikis or free/open source software communities [34, 35]. It is worth noting that, in DAOs, the level of activity may depend on the purpose of the DAO and the size of the community, since some of them may not need to vote proposals that frequently.




Figure 8. Monthly time series of active users

Finally, Figure 8 shows the number of active users over time. There is a mean of 98 users who have performed at least one action, which represents a 1.6% of total users. However, we remember that *Kyber DAO Exp#2* has 4,946 users, and this DAO is inactive. If we exclude them, the registered activity on *DAOstack*

is maintained by a small fraction of users (around 8%). Again, this kind of unequal participation is typical from other online communities [36, 37].

5. Validation of boosting as a predictor

As explained in Section 3, boosting through staking can be seen as a prediction system; more precisely, a prediction market. We want to assess whether its predictions succeed or not. The boosting process has two possible outcomes: the proposal can be boosted or not. A boosted proposal is a proposal that stakers think that DAO's members will approve. On the contrary, a non-boosted proposal, theoretically, is a proposal that stakers consider that will not be approved by the DAO community. Bear in mind that if no one stakes, the proposal, then it is non-boosted. As a result we define the following prediction outcomes for a proposal:

- True Positive (TP) = boosted and approved
- True Negative (TN) = non-boosted and rejected
- False Positive (FP) = boosted and rejected
- False Negative (FN) = non-boosted and approved

Given these outcomes, we will measure how "good" the prediction is using:

- Accuracy = $(TP + TN) / (TP + TN + FP + FN)$
- Precision (boosted) = $TP / (TP + FP)$
- Precision (non-boosted) = $TN / (TN + FN)$

5.1. Splitting DAOs in groups

We are going to assess the prediction ability of the boosting process across *DAOstack*, but since DAOs vastly differ in terms of users or proposals presented, we will segregate DAOs into various groups. For the segregation, we will analyze two hypotheses:

- i The use of boosting in a community is positively correlated with the size of the community, i.e. the number of users.
- ii The use of boosting in a community is positively correlated with the number of proposals presented in the community.

We will consider two possible definitions for characterizing the use of boosting in a community. First, using the number of proposals which have at least one stake. Second, the percentage of proposals of a given DAO that reached the boosted state. We use the

	Staked proposals	Boosted proposals
Users	0.71*	0.64*
Proposals	0.55	0.53

Table 1. Spearman's correlation, where (*) means significant at $p=0.01$

percentage to homogenize the differences in the number of proposals found in the DAOs.

To determine which definition of staking is more adequate, Table 1 shows their correlation with the numbers of users and proposals. We use Spearman's rank-order correlation to analyze the presence of a monotonic relationship, but not necessarily a linear one, as with Pearson's correlation. We can see that the number of users has positive correlation with both definitions of the use of boosting. By comparison, the number of proposals has a weaker correlation. Due to that, we find more evidence of *i* than of *ii*.




Figure 9. DAOs plotted by number of users vs the percentage of their proposals that were staked.

Figure 9 shows the relation between the number of proposals and the percentage of proposals which has received at least one stake. *Kyber DAO Exp#2* (4,946 users) was filtered out to improve the visualization. We see a monotonic direct relationship, but far from linear.

Given these results, we will split DAOs according to the number of users. We could consider user variability over time. However, we have found that all DAOs in *DAOstack* have registered their vast majority of users in their first months of life.

In our analysis, we will filter out two inactive DAOs, but decided to keep the rest, which at least have performed one action, to analyze the whole ecosystems. We will use the median to divide our population into two

groups with the same number of DAOs (10). First group includes all DAOs which have 23 or less members, while the second group includes DAOs which have more than 23 members. We will study the use of the boosting mechanism in each of these groups.

	Users	Proposals
Group A	95	178
Group B	6013	907

Table 2. Total users and proposals in Group A and B

Table 2 shows the total number of users and proposals of each group. It is remarkable that *Group A* has just a 1% of all users and a 16% of all proposals.

5.2. Results of the validation

Table 3 shows the accuracy (Acc.) and the precision for boosted (Prec. Boost) and non-boosted (Prec. Not Boost) proposals for each DAO group.

	Acc.	Prec. boost.	Prec. non-boost.
Group A	0.67	0.93	0.5
Group B	0.93	0.96	0.85

Table 3. Accuracy and precision of the boosted and non-boosted proposals for each group

The accuracy of both groups clearly differ: *Group A* has a moderate global accuracy while *Group B* shows a strong accuracy. A closer look at the precision values makes possible to determine where the difference lies. While the precision of the boosted proposals is quite high in both groups (0.93 in A and 0.96 in B), non-boosted proposals are not as good for predicting rejection. More precisely, the precision of Group A is just 0.5 which means that for small DAOs a non-boosted proposal is not a good indicator on whether it will be rejected by the community. Since in Group A DAOs are small, a potential explanation would be that members of the DAO may not feel inclined to staking because majority voting and the number of concurrent proposals are manageable. Moreover, proposals from smaller DAOs may not receive the same attention (and stakes) than those for the bigger ones. Interestingly, the precision of non-boosted proposals is much higher in Group B (0.85), which means that for bigger DAOs non-boosting is a good proxy for rejection. Thus, boosting and holographic consensus seem to serve their purpose for large DAOs, as it is the original intention.

To see how prediction works over time, Figures 10 and 11 show the evolution of the precision for Group A and Group B, respectively. According to Figure 10, *Group A* presents a stable ratio when it comes to

predicting approvals, a mean of 0.92, and a standard deviation of 0.18 over the time series. Whilst the behaviour of the prediction of rejection has a lower mean (0.63) and is more irregular (std. dev. of 0.35). The low precision for non-boosted proposals is due to the persistent presence of non-boosted and approved proposals that disappear in February 2020. Moreover, since that month the precision for boosted proposals increase due to new DAOs, *IUP* and *CENNZnet Grants DAO*, that use boosting in an effective way.

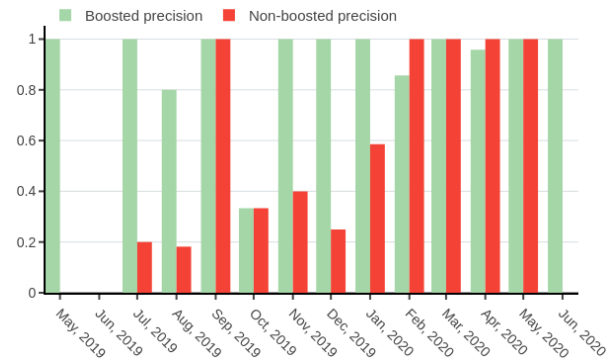


Figure 10. Boosted precision of Group A

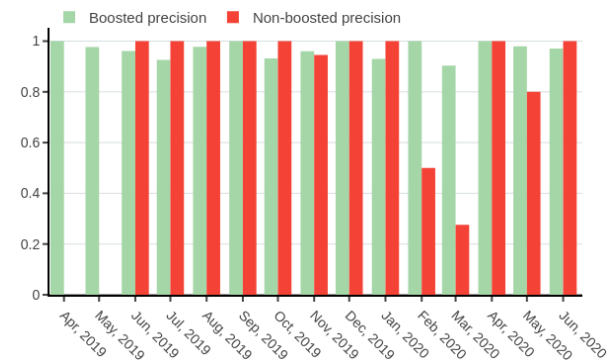


Figure 11. Boosted precision of Group B

Figure 11 shows an almost perfect precision for boosted proposals (mean of 0.97 and standard deviation of 0.03). Precision of non-boosted proposals registers a mean of 0.89 and a standard deviation of 0.22, which are good values and much better than those from *Group A*. The decrease in the non-boosted precision in February and March 2020 is due to the false negatives (non-boosted and approved) by *dOrg*, which has 41 users, and, as seen in Section 4.2, it uses no staking and always approves proposals by absolute majority.

Figure 12 shows the accuracy of each DAO, as defined in Section 5. DAOs are sorted by the number of users, and the color of the dot represents the number of proposals of the DAO. A vertical line separates DAOs from *Group A* and *Group B*.

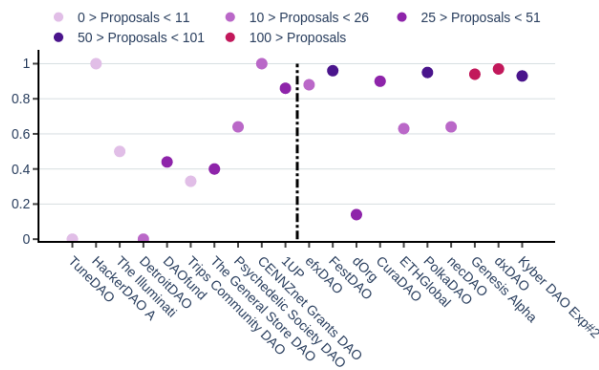


Figure 12. Accuracy for each DAO. A vertical line separates Group A (left) and Group B (right).

In *Group A*, accuracy is relatively low, which means that staking is not very useful, but three DAOs have an accuracy over 0.8. One was *HackerDAO A* that registered just one proposal, thus the value is not significant. The other two are the biggest ones in terms of users, *IUP* with 20 users and *CENNZnet Grants DAO* with 15 users. This suggests that small DAOs do not use effectively the boosting process, because probably they can approve everything by majority voting.

In *Group B*, besides *dOrg* discussed above, the other nine DAOs have an accuracy over 0.6, and seven of them over 0.88. Most of these DAOs have over 25 proposals. This suggests that when DAOs reach a certain number of members they tend to rely on boosting and that it works quite well in terms of anticipating the relevance of a proposal for the community.

6. Discussion

Our work contributes to validate to which extent DAOs facilitate scalable governance. In particular, we focus on the *DAOstack* platform and its decision-making mechanism, *Holographic Consensus*, which combines a voting system with a prediction market. This method facilitates scalability by highlighting relevant proposals and reducing the quorum required to approve them.

We analyze the 22 DAOs of *DAOstack* and their 6,000 users, which are early-adopters. Overall, *Holographic Consensus* seems to follow the expected behaviour: its prediction effectiveness is around ($\approx 90\%$) in large DAOs. In fact, the data hints that the larger a DAO is, the more *Holographic Consensus* (staking) is used. On the contrary, within small DAOs the results are mixed, but not unsatisfactory.

We acknowledge the limitations of this analysis. Even if we analyzed all *DAOstack* DAOs so far, its number is still small, as it is the size of most of the communities. While our results validate the idea that

Holographic Consensus is an appropriate mechanism to facilitate large-scale decentralized cooperation, some aspects may have influenced the results. For example, in DAOs voting has a transaction cost (i.e. a micro-payment), and thus the decisions must be relevant enough for decision-makers to pay such costs. There are novel methods to avoid users from incurring in such costs (e.g. the DAO absorbing the cost), although they had not yet been implemented in the platform at the time of our study. For the same reasons, in bigger DAOs, multiple trivial decisions, such as the need for paying workers, must be voted, and then are usually boosted and approved with just a few votes. This effect may artificially inflate the results.

Still, ours are promising results in a field which is not yet mature to make strong assertions. DAO governance is complex and it is affected by different aspects such as the unequal weight of votes that depends on the reputation of each individual (an issue that we did not explore). Thus, more research is clearly needed, e.g. including other DAO platforms (e.g. Aragon, DAOHaus) or other blockchain-enabled decision making mechanisms. Besides, triangulating quantitative analysis with qualitative data would provide further insights on the results.

Despite of its limitations, our analysis shows how online communities are achieving governance through decision-making mechanisms on the blockchain. The implications are diverse: first, e-voting systems may extract relevant lessons from DAOs, and the use of community-based filtering mechanisms. In addition, the success of this novel mechanisms for online communities may open the door to further experimentation, specifically combining blockchain-based tokens and its multiple applications in online communities and, more clearly, in DAO communities. Overall, this is also a good signal for DAO communities, which may support the claims that these may become the organizations of the future.

Acknowledgements

This work was supported by the project P2P Models (<https://p2pmodels.eu>) funded by the European Research Council ERC-2017-STG 625 [grant no.: 759207], and by the Spanish Ministry of Science, Innovation and Universities [grant no.: RTI2018-096820-A-100]. We also thank Ivan Sucharski and the reviewers for their constructive feedback.

References

- [1] S. Underwood, "Blockchain beyond bitcoin," *Communications of the ACM*, vol. 59, no. 11, p. 15–17,

- 2016.
- [2] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352–375, 2018.
 - [3] P. D. Filippi and S. Hassan, "Blockchain technology as a regulatory technology: From code is law to law is code," *First Monday*, vol. 21, no. 12, 2016.
 - [4] G. Wood et al., "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, no. 2014, pp. 1–32, 2014.
 - [5] F. Casino, T. K. Dasaklis, and C. Patsakis, "A systematic literature review of blockchain-based applications: current status, classification and open issues," *Telematics and Informatics*, vol. 36, pp. 55–81, 2019.
 - [6] L. Cocco, A. Pinna, and M. Marchesi, "Banking on Blockchain: Costs Savings Thanks to the Blockchain Technology," *Future Internet*, vol. 9, p. 25, 2017.
 - [7] D. Cawrey, "37Coins Plans Worldwide Bitcoin Access with SMS-Based Wallet." <https://www.coindesk.com/37coins-plans-worldwide-bitcoin-access-sms-based-wallet/>, 2014.
 - [8] K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292–2303, 2016.
 - [9] G. Miscione, "Blockchain as organizational technology," in *University of Zurich, Department of Informatics IFI Colloquium, Zurich, Switzerland, 28th February 2019*, University College Dublin, 2019.
 - [10] P. D. Filippi and A. Wright, *Blockchain and the Law: The Rule of Code*. Harvard University Press, 2018.
 - [11] M. Swan, *Blockchain: Blueprint for a new economy*. O'Reilly Media, Inc., 2015.
 - [12] N. Radziwill, "Blockchain revolution: How the technology behind bitcoin is changing money, business, and the world," *The Quality Management Journal*, vol. 25, no. 1, pp. 64–65, 2018.
 - [13] M. Atzori, "Blockchain technology and decentralized governance: Is the state still necessary?," *Available at SSRN 2709713*, 2015.
 - [14] Q. DuPont, "Experiments in algorithmic governance: A history and ethnography of "The DAO," a failed decentralized autonomous organization," *Bitcoin and beyond*, pp. 157–177, 2017.
 - [15] Y. Benkler, *The Wealth of Networks: How social production transforms markets and freedom*. Yale University Press, 2006.
 - [16] A. Shaw and B. M. Hill, "Laboratories of oligarchy? how the iron law extends to peer production," *Journal of Communication*, vol. 64, no. 2, pp. 215–238, 2014.
 - [17] D. Rozas, A. Tenorio-Fornés, S. Díaz-Molina, and S. Hassan, "When Ostrom meets blockchain: Exploring the potentials of blockchain for commons governance," *Available at SSRN 3272329*, 2018.
 - [18] "Decentralized autonomous organizations. the next frontier of human coordination." <https://aragon.org/dao>.
 - [19] M. Field, "Holographic consensus—part 1." <https://medium.com/daostack/holographic-consensus-part-1-116a73ba1e1c>, 2018.
 - [20] M. R. Clarkson, S. Chong, and A. C. Myers, "Civitas: Toward a secure voting system," in *2008 IEEE Symposium on Security and Privacy*, pp. 354–368, IEEE, 2008.
 - [21] N. Kshetri and J. Voas, "Blockchain-enabled e-voting," *IEEE Software*, vol. 35, no. 4, pp. 95–99, 2018.
 - [22] J. Xie, F. R. Yu, T. Huang, R. Xie, J. Liu, and Y. Liu, "A survey on the scalability of blockchain systems," *IEEE Network*, vol. 33, no. 5, pp. 166–173, 2019.
 - [23] N. Urbinati, *Representative democracy: principles and genealogy*. University of Chicago Press, 2006.
 - [24] J. Haskell, *Direct democracy or representative government? Dispelling the populist myth*. Routledge, 2018.
 - [25] M. Hilbert, "The maturing concept of e-democracy: From e-voting and online consultations to democratic value out of jumbled online chatter," *Journal of Information Technology & Politics*, vol. 6, no. 2, pp. 87–110, 2009.
 - [26] M. Field, "Decentralized Governance Matters." <https://medium.com/daostack/decentralized-governance-first-principles-1fc6ea492ed>, 2018.
 - [27] Y. El Faqir, J. Arroyo, and S. Hassan, "An overview of Decentralized Autonomous Organizations on the blockchain," in *Proceedings of the 16th International Symposium on Open Collaboration 11:1-11:8*. ACM, Opensym 2020, 2020.
 - [28] A. Soleimani, A. Bhuptani, J. Young, L. Haber, and R. Sethuram, "The Moloch DAO: Beating the tragedy of the commons using Decentralized Autonomous Organizations." <https://github.com/MolochVentures/Whitepaper/blob/master/Whitepaper.pdf>, 2019.
 - [29] M. Mannan, "Fostering Worker Cooperatives with Blockchain Technology: Lessons from the Colony Project," *Erasmus Law Review*, vol. 3, pp. 190–203, 2018.
 - [30] J. Emmett, "Conviction Voting: A Novel Continuous Decision Making Alternative to Governance." <https://link.medium.com/tJJ49Vdcb>, 2019.
 - [31] M. Zargham, "Sensor Fusion and Social Choice." <https://github.com/BlockScience/conviction/blob/master/social-sensorfusion.pdf>, 2018.
 - [32] J. Wolfers and E. Zitzewitz, "Prediction markets," *Journal of Economic Perspectives*, vol. 18, no. 2, pp. 107–126, 2004.
 - [33] F. Ortega, *Wikipedia: A quantitative analysis*. PhD thesis, Universidad Rey Juan Carlos, Madrid, 2009.
 - [34] G. Jimenez-Diaz, A. Serrano, and J. Arroyo, "A Wikia census: motives, tools and insights," in *Proceedings of the 14th International Symposium on Open Collaboration*, pp. 2:1–2:6, ACM, 2018.
 - [35] S. Krishnamurthy, "Cave or community?: An empirical examination of 100 mature open source projects," *First Monday*, vol. 7, no. 6, 2002.
 - [36] T. Chelkowski, P. Gloor, and D. Jemielniak, "Inequalities in open source software development: Analysis of contributor's commits in apache software foundation projects," *PLoS One*, vol. 11, no. 4, p. e0152976, 2016.
 - [37] A. Serrano, J. Arroyo, and S. Hassan, "Participation inequality in wikis: A temporal analysis using WikiChron," in *Proceedings of the 14th International Symposium on Open Collaboration*, pp. 12:1–12:7, ACM, 2018.