DOI:
10.1007/978-3-030-79318-0_9

Document Version
Early version, also known as pre-print

Self-Governing Public Decentralised Systems*
Work in Progress

Moritz Platt[0000−0003−2137−4370] and Peter McBurney

Department of Informatics, King's College London, UK
{moritz.platt,peter.mcburney}@kcl.ac.uk

Abstract. The selection of members responsible for data replication is a challenge in decentralised record-keeping systems. In 'permissioned' systems, this crucial task is performed by a central authority or consortium. In 'permissionless' systems, however, the selection process is not trivial and comes with risks. Malicious actors, in a privileged position, can tamper with data, threatening the integrity of the system as a whole. Permissionless membership selection protocols, popularised with the dissemination of distributed ledger technology, have the objective of limiting the influence of a single entity on the wider network. They do so by approximating a participant's legitimacy to participate in record maintenance. These approximations come with downsides, in terms of attackability, system performance, supported use-cases and resource requirements. In this paper, we propose a prototypical membership selection protocol that uses the measure of personhood as an approximation of legitimacy. Interpreting a decentralised system as a political system, we frame the membership selection problem as one of political representation. We propose a protocol that democratically attributes a personhood score to members, thus creating a self-governing public decentralised system. This work in progress lays out a roadmap for the formal evaluation of self-governing public decentralised systems and describes the anticipated challenges in their implementation. Our proposals provide a means to evolve the membership selection protocol when a closed, permissioned system evolves to an open, permissionless system, as several commercial platforms intend to do.

Keywords: Socio-Technical Systems · Self-Governance · Distributed Ledger Technology · Blockchain · Decentralisation · Consensus · Permissionless Networks · Membership Selection.

1 Introduction

A challenge that any decentralised record-keeping system, that operates in a potentially-distrusting environment, faces is how to select members to perform

*We are grateful for financial support from the UK EPSRC VOLT Project, grant number EP/P031811/1. We thank the anonymous reviewers whose comments helped to improve this manuscript.

validation of records and record evolvement. Building on earlier work in distributed system design, the peer-to-peer electronic cash system 'Bitcoin' [36] solved this membership selection problem, bringing with it a new paradigm of decentralised record-keeping based on 'proof-of-work' (cf. section 2.1). It is distinguished from earlier approaches by being truly permissionless, in the sense that 'any network participant has the ability to create a candidate record' [43, p. 61]. This paradigm laid the foundation for a variety of similar protocols. Addressing the challenge of preventing illegitimate updates to shared data, by nefarious actors, is the main contribution underlying Nakamoto's work. Proof-of-work based protocols have been widely criticised for their environmental impact [35, 44] and for poor throughput characteristics when compared to centralised systems [11].

Recognising these shortcomings, alternative approaches for selecting participants in public decentralised systems have been proposed. These are commonly differentiated by their trust assumption, i.e. whether the protocol gives certain entities, on a network, extended permissions in the membership selection process [37, p. 7]. Given that these selection processes often resemble votes, parallels to the political realm are obvious. Following this line of thought, it can be speculated that membership selection in decentralised systems follows mechanisms similar to those present in archetypes of self-governance. This paper evaluates whether self-governance could provide inspiration for a novel membership selection protocol, that combines the advantages of different degrees of openness in the membership selection process of a decentralised system.

2 State of the Art[1]

Fault-tolerance of distributed systems has been part of the research agenda in computer science for a long time. Among the early findings, most relevant to self-governing public decentralised systems, is the work of Lamport *et al.* [27]. They show how a decentralised system behaves when actors spread incorrect or conflicting information, or withhold information. They describe how a system tolerates a limited fraction of these actors, often referred to as 'byzantine' actors. Douceur [14] makes another instrumental finding, showing how a 'single faulty entity', often referred to as a 'sybil' actor, can gain control of a redundant network by 'presenting multiple identities'.

In the absence of an overseeing authority with special privileges, public decentralised systems must be able to rely on the adherence of the majority of their members to the system's protocol. While redundancy allows systems to tolerate a proportion of 'byzantine' actors, 'sybil' actors can never be tolerated. Unsurprisingly, most of the research in distributed ledger technology revolves around these two concepts.

Hellwig *et al.* [21] define the terms 'anonymity' and 'pseudonymity' in the context of cryptocurrencies, describing anonymous transactions as those that 'do

[1] This section only discusses approaches that are relevant to self-governing systems. Bach *et al.* [1] and Natoli *et al.* [37] provide more comprehensive surveys of prevailing patterns.

not require a name' and pseudonymous transactions as transactions in which 'a false name is used'. This is reflected by the use of 'addresses', self-created pseudonymous identifiers, that are easy to generate within a very large address space. Pseudonymity on cryptocurrency networks is infamously heavily abused [47, 3, 24].

2.1 Proof-of-Work

In his work on the Bitcoin protocol, Nakamoto [36] formalised the need for consensus on two dimensions; the validity of blocks (i.e. the comprehensive validity of all transactions in a block) and the ordering of blocks. The predominant goal of the consensus algorithm, in this protocol, is to counteract the 'double-spend problem', a well-known problem in the realm of electronic payments, that allows dishonest actors (in the absence of a control mechanism) to over-spend their funds. Solutions to this problem were proposed much earlier but encompassed a pre-defined actor (or set of actors) to assume the role of trusted third party [9], or required tamper-proof hardware [52]. Thus, pre-Proof-of-Work payment systems relied on a permissioned approach, in which certain participants (e.g. banks or payment system operators) had special privileges.

Nakamoto's contribution was to introduce a permissionless approach to membership selection. The goal of membership selection in Bitcoin is to select a 'miner' to validate the transactional data and act as an ordering authority, immutably linking the current set of transactions with all previous sets. This problem can't be solved by selecting actors randomly, since they are not uniquely identifiable on a network, in fact, they can generate arbitrary numbers of pseudonyms, making a random selection highly susceptible to 'Sybil Attacks'. These are attacks in which a single malicious entity presents multiple identities [14], thus improving their chances of being selected. The main purpose of proof-of-work in membership selection is to create an environment in which a participant ('miner') is incentivised to act honestly. The Bitcoin protocol, similar to other proof-of-work protocols, is designed to be incentive-compatible, i.e. it should accomplish its goal of evolving the decentralised data correctly, if all participants follow the rules and are capable of handling the informational requirements [15]. This is based on the assumption that those who expended computing resources, by participating in proof-of-work, have an incentive not to introduce incorrect data because they would otherwise threaten the value of their reward.

Along with Bitcoin, numerous other decentralised record-keeping systems, specifically cryptocurrencies [29, 53], utilise proof-of-work protocols. Proof-of-work systems have been found to be 'dreaded with various attacks' with 'robust and practical security solutions' to those being absent [10]. Attacks can roughly be divided into two classes; 'Goldfinger' attacks [26], where adversaries seek to disrupt the validation of records and record evolvement on a public distributed system (e.g. double-spending [23] or brute-force attacks [22]), and attacks that do not threaten the system as a whole but lead to financial gain or loss for individual participants in the protocol (e.g. refund attacks [34], transaction malleability [54] or denial-of-service attacks against participants [51]).

2.2 Proof-of-Stake

Saleh [44] names avoiding the expenditure of an 'exorbitant amount of energy' as the main motivation for the blockchain community to research alternatives to proof-of-work. King and Nadal [25] proposed 'ppcoin', a peer-to-peer cryptocurrency, with proof-of-stake, taking inspiration from earlier discussions of the concept among Bitcoin circles [42]. In their design, being able to prove ownership of currency, along with proving how long it has been held, will determine the difficulty of creating a new block, thus making those participants who have held larger quantities of currency for longer more influential in record evolvement. Compared to proof-of-work, this shifts the responsibility of maintaining and evolving decentralised data from those who invest computing resources to those who hold the most currency.

This fundamentally changes the incentives for behaving dishonestly. Gui *et al.* [19] show that proof-of-stake is less vulnerable to both double-spending attacks and sabotage attacks than proof-of-work. Li *et al.* [30] discuss how this new paradigm has brought with it new forms of attacks; specifically, 'nothing at stake' attacks, where malicious validators generate conflicting blocks to slow down consensus time and 'long range' attacks in which malicious actors create forks from historic blocks, allowing them to form longer chains, based on an out-dated view of stake. A notable proof-of-stake protocol is 'Algorand', which 'assigns weights to users proportionally to the monetary value they have in the system' [17]. Thin *et al.* [49] verified that 'Tendermint', an exemplary proof-of-stake protocol, can reach consensus when at least $\frac{2}{3}$ of a proof-of-stake network are in agreement. The 'Ethereum' blockchain, utilising a proof-of-work approach from inception, is planned to migrate to a 'proof of stake-based finality system, which overlays an existing proof-of-work blockchain' [7] with the next version of the platform.

2.3 Delegated Proof-of-Stake

Larimer [28] proposed a variation to proof-of-stake, introducing a delegation scheme, in which 'shareholders may delegate their voting power to a representative'. This delegation is implemented via proxy signatures [5]. Delegating the right to validate and evolve records to other participants is useful in an environment where there is a majority of participants that are not interested in, or capable of, providing validation. Delegated proof-of-stake has been considered a suitable building block for election-based protocols, such as 'Snow White' [12], a protocol that supports committee reconfiguration and remains robust in the presence of sporadic participation.

2.4 Proof-of-Personhood

Borge *et al.* [6] show how, by conducting 'pseudonym parties', 'sybil attacks' (cf. section 2) can be prevented. They use this as the foundation of the 'PoPCoin' protocol in which proving 'personhood', i.e. the existence as a human individual, grants membership on a network.

2.5 Proof-of-Authority

Permissionless membership selection (cf. sections 2.1–2.4) is not appropriate in all contexts. An environment in which anyone can create a candidate record is undesirable in scenarios where there is a need for limiting the audience or participation. Drivers for rejection of an open approach can be privacy concerns (i.e. who can access the data to be validated), or regulatory concerns (e.g. where regulatory requirements exist that govern who can partake in a certain activity). In such contexts, membership selection can be achieved by policy, i.e. through a pre-defined list of privileged members. The practice of employing a central party, or consortium, to decide who is allowed to perform record validation activities is known as 'proof-of-authority'. An example of an implementation of this approach is the 'Ripple Protocol Consensus Algorithm' [45], that employs a pre-defined 'Unique Node List' of trusted servers, with no facility for altering this list via the standard protocol. The transaction ordering and timestamping services described in the 'Corda' protocol operate similarly. Here, it is the responsibility of the network governing body to establish and maintain a list of notaries [20]. Facebook's 'Libra' payment system plans to adopt a similar approach, with a set of pre-approved validators [31].

2.6 Voting 'On-Ledger'

In addition to concerns of membership selection, voting can be conducted on-ledger, i.e. on top of an already-established system. It is important to note that e-voting can be conducted on-ledger, irrespective of the membership selection paradigm. That means that entities that engage in on-ledger voting do not necessarily have to be 'members', in the sense of the outlined membership selection paradigms.

Dhillon *et al.* [13] point out that large-scale decentralised online voting poses challenges around the governance of voting networks and delegations of votes. On a smaller scale, McCorry *et al.* [33] show how a 'self-tallying internet voting protocol with maximum voter privacy' can be implemented on top of the Ethereum blockchain. The voting functionality here is implemented 'top-of-Stack', i.e. the voters do not have to participate in proof-of-work (cf. section 2.1) but can cast their votes via Ethereum transactions.

3 Problem Motivation

Some parallels can be drawn between membership selection methodologies, in decentralised systems, and processes of political representation that can be observed in the analogue world. The virtual constituency of those who can create candidate records, in a decentralised system (cf. figure 1), can be compared to a constituency in the political sense, albeit not a well-defined, stable state-level constituency, but a transnational fast-evolving one. This group can rely on the legal framework of the governing system to gain a high degree of certainty that the codified rules of the system will be enforced by the executive branch.

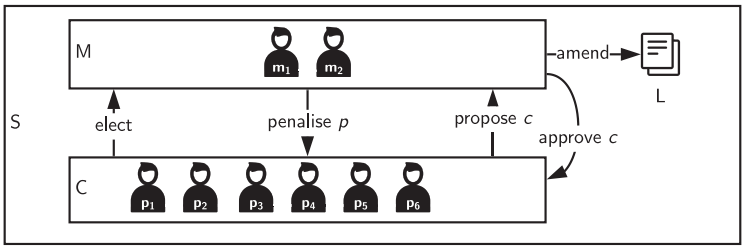

Fig. 1. A decentralised system S, comprised of regular participants ($p_{1..n}$) and participants with additional duties ('miners' $m_{1..n}$) who are appointed or elected to fulfil these duties. Participants propose candidate records, c, to be included in the entirety of public records. Miners decide, based on a legislative framework, L, whether a candidate record is permissible and, based on this evaluation, either approve it or, penalise the responsible participant for proposing an impermissible record.

The group of participants who validate candidate records, approve and evolve them, can be compared to a government, particularly the legislative and executive branches of government. Depending on the underlying membership protocol, miners are appointed (cf. section 2.5) or elected (cf. sections 2.1–2.4) via probabilistic methods that roughly resemble a majority vote (e.g. the majority computing power in the case of proof-of-work, or the majority of funds in the case of proof-of-stake). Participants who create candidate records make these available to miners, who in turn validate that these are compliant with the legal framework. Should a candidate record be found to be compliant, miners will endorse it publicly. Should the candidate record found to be in breach of the framework, participants may be penalised. It can be assumed that the legislative framework is deterministic, meaning all honest miners will come to an identical conclusion on compliancy when evaluating a record[2].

4 Solution

The previous section shows how democratic political representation can be considered an archetype for member selection in decentralised systems. While numerous approaches have been developed to approximate democratic member selection (cf. section 2), no scalable protocol to implement the democratic ideal of 'One Person/One Vote' in membership selection currently exists. This section shows how the reason for this is the difficulty of determining who should be considered eligible to vote, rather than the implementation of the voting protocol itself. It will employ the concept of 'personhood' as an approximation of eligibility, building on the definition of Borge *et al.* [6], who specify 'personhood' as the property of having a unique identity in the real world. We define personhood as

[2]Deterministic, automatically evaluable, 'smart contracts' [48] are a concept at the intersection of law and computer science. While natural language contracts require interpretation, computer language contracts are designed not to [8].

a binary property but recognise that there is no canonical definition of it [16], thereby allowing elected authorities entrusted with admitting participants to the network to apply their own definition. Extending this, a probabilistic 'personhood score' $ph \in [0, 1]$ can be calculated for participants. Here, $ph = 0$ indicates that there is no confidence that the identity presented is a unique identity (i.e. it is almost certainly a 'sybil' identity), whereas $ph = 1$ indicates maximum confidence in it being a unique identity.

4.1 'One person/One Vote' in Delegated Proof-of-Stake

In a delegated proof-of-stake system S (cf. section 2.3) members $p_{1..n}$ form a constituency C. They can delegate participation in the consensus protocol to other parties on the network, miners ($m_{1..n}$), thus approximating a vote. Since delegation privileges in this protocol are aligned with the currency holdings of the delegating party, this pattern corresponds to the 'One Share/One Vote' paradigm, well-known in the realm of corporate securities [18]. Given that delegated proof-of-stake effectively already implements a 'One *Share*/One Vote' paradigm, it can be easily restructured to support a 'One *Person*/One Vote' paradigm by introducing additional constraints to limit the number of shares and how they can circulate. These constraints could be introduced into 'system contracts'[3].

i. Delegated proof-of-stake is performed using personhood tokens as stake.
ii. Every person with voting rights on the network receives a fixed number of personhood tokens once they enter the network.
iii. There is no other source of personhood tokens.
iv. Personhood tokens cannot be traded and are not given out as a reward.

4.2 Establishing Personhood

The requirements *i*, *iii* and *iv* of the previous section can easily be satisfied through minor modifications of existing protocols. An implementation could be forking existing delegated proof-of-stake protocol implementations (cf. section 2.3) and changing validation code so that token movement becomes impermissible. The second requirement, however, raises a more significant problem. Here, it needs to be considered who should be admitted to the decentralised system, by whom and on what basis.

Gatekeeping Authority A trivial solution to this problem is the introduction of a central 'gatekeeping authority' A. They could assign stake in the form of personhood tokens. Conceivably they would assign all participants $ph = 1$, having confidence in their onboarding process. A hypothetical protocol that implements this paradigm could require prospective participants to generate an address (cf.

[3]The term has been popularised by the 'EOS.IO' blockchain [4], where it is used to describe fundamental functionality of the core protocol that is not modifiable by individual users, as compared to user defined smart contracts.

section 2) and to subsequently submit it to the authority along with proof of their identity. Should the identity information provided satisfy the requirements of the central authority, they would allocate voting right tokens and fund the address of the applicant.

This approach would effectively recreate a proof-of-authority protocol (cf. section 2.5) with one centralised admitting entity. The shortcomings of this approach coincide with the shortcomings of a proof-of-authority protocol, namely, the fact that the governance structure needs to be determined at the inception of the network and can only be amended by the initiative of the admitting entity itself. This approach requires a permanently high degree of trust in the admitting entity.

Self-Governed Evolving Constituencies While a central gatekeeping approach is well-suited for cases in which few changes to the constituency of a system are to be anticipated, and in which the central authority is irrefutably trusted, it is inappropriate in a scenario in which the constituency is evolving. An evolving constituency is conceivable in self-governance scenarios, e.g. when a group of constituents choose to administer common resources through a joint governance process [50, 39], or when a group decides they require self-governance, for resolving grievances or making political decisions outside of a wider context governed by an external authority [40].

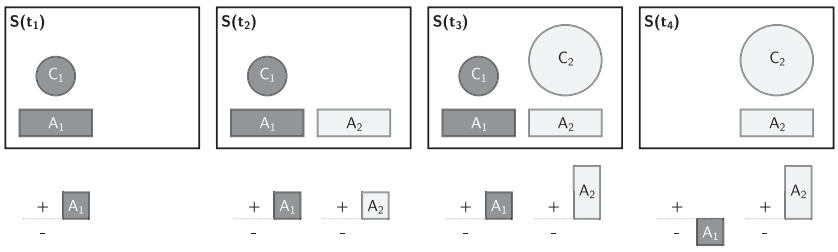

Fig. 2. A decentralised system S with a constituency that evolves over periods $t_{1..4}$. Initially, the first constituency C_1 is the only one to participate in delegated consensus on the network. They record their endorsement for identity authority A_1. In t_2 they also endorse authority A_2. Following this, in t_3, members endorsed by A_2 are on-boarded to the system. They endorse A_2. In the final period (t_4), a large number of participants discourage A_1, following which the value of the personhood tokens issued by them drops below the reputational threshold and C_1 loses their membership.

To replicate a self-governing structure in a decentralised system, the previously described 'gatekeeping' approach is a suitable point of departure. Conceptually, this means allowing multiple authorities $A_{1..n}$ in parallel. Delegated proof-of-stake protocols already bring all necessary technological primitives for voting, since delegation is a voting process in itself. Building on that, a second voting

layer can be proposed. This voting layer would be responsible for the election of gatekeeping authorities into and out of the decentralised system.

As shown in figure 2, a decentralised system making use of evolving constituencies would be initialised with a 'genesis gatekeeping authority'[4]. This authority would act in the same way as described above, i.e. they would assign voting right tokens to members of the genesis constituency, following their principles of identity validation. The self-governing aspect would manifest itself by introducing new capabilities to the system:

1. Members of a constituency can endorse a gatekeeping authority on the network;
2. Members of a constituency can discourage a gatekeeping authority on the network.

The personhood score is relevant as it will determine the stake of a participant, i.e. how much voting power they can delegate to a miner of their choice. It also constitutes a weighting factor on an endorsement/discouragement. In a system utilising personhood scoring, participants would likely delegate their continuous participation in voting processes to agents acting on their behalf. Therefore, these systems would effectively constitute open agent societies that are likely to implement deliberative democracies, in which participants are both consumers of political information (i.e. the personhood scores of others) and producers of political information (by endorsing or discouraging authorities) [32].

Identity Claim To show that participants are a member of a constituency, they can publicly broadcast a simple identity claim. This public record would consist of a signature, provided by the identity authority, that a participant claims identity from. Assuming that participants are identified by their public key, or public key hash, as is common in blockchain protocols [38, p. 25], an identity claim could simply comprise a signature by an identity authority over the participant's address. A valid claim would grant the participant a voting right token by the relevant authority.

Endorsement/Discouragement As shown in the previous section, members can endorse or discourage gatekeeping authorities via a broadcast message. These actions directly impact the reputation of the authority and thus the personhood score the authority can grant. Per authority $A_{1..n}$ a vector of endorsement scores $\boldsymbol{e}_{A_{1..n}}$ and a vector of discouragement scores $\boldsymbol{d}_{A_{1..n}}$ are kept publicly. Participants add to either of the vectors via a message they broadcast. They can add at most one message to each vector. The value they add to the vector represents their personhood score as determined by the reputation of their gatekeeping authority. This means that the influence a participant can exert on the reputation of another authority is proportional to their reputation. This allows for basic arithmetic on the personhood score values of participants, depending on which

[4]The naming is inspired by the term 'genesis block', the genesis of the Bitcoin Blockchain.

constituency they belong to. A dampened personhood score (dph) depends on the reputation of the issuing authority: $dph = ||\boldsymbol{e}||_1 / (||\boldsymbol{e}||_1 + ||\boldsymbol{d}||_1)$ if $||\boldsymbol{e}||_1 + ||\boldsymbol{d}||_1 > 0$, otherwise $dph = 0$. Here $||\cdot||_1$ denotes the L^1-norm. Consider an example of a network, with three identity authorities, $A_{1..3}$ with reputation of $A_1 = 0.8$ and $A_2 = 0.3$. Assuming that a member of A_1 has endorsed A_3 and a member of has A_2 discouraged A_3, the reputation value of A_3 would be $\frac{0.8}{0.8+0.3} \approx 0.7$. The reputational value of A_3 would be updated in the next iteration of the protocol.

Temporal Normalisation A single malevolent authority can flood the network with sybil actors[5], who can disrupt any record-keeping and record-evolving activity on the network, permanently. Temporal normalisation can mitigate sybil attacks that go along with a sudden influx of bogus identities. Such attacks are likely to be preceded by an event in which a previously trustworthy private key is used to generate bogus identities on a network. This allows for participant's scores to increase over time. A time-normalised personhood score can be calculated using any normalisation function $f : [0, 1] \times \mathbb{N} \mapsto [0, 1]$, for example $f(dph, t) = dph \cdot (1 - \frac{1}{1+t})$, where dph is the personhood score of the participant at the current time and t is a measure of elapsed time, such as 'block height' added, since admission of the participant.

Constituency Size Ceilings Assuming that attackers have perfect knowledge of the protocol and can therefore conduct attacks that take the maturing of personhood scores into account, the temporal normalisation approach alone cannot be effective. Under a reasonable temporal normalisation function, a large number of bogus identities can achieve a large cumulative time-normalised personhood score, even for small t. This makes the effectiveness of temporal normalisation dependent only on the rate at which attackers can create bogus identities. To counteract this phenomenon, an overall constituency size ceiling that limits the total number of identities, created by one authority, can be introduced. This can be a fixed value of maximum permissible identities issued per authority or a function that limits the permissible growth of the number of identities.

Diversity Dimension In case of such a well-planned attack, the temporal and reputational safeguards would have limited effect. For these scenarios, a quantitative safeguard enforcing diversity should be introduced. This should give reputational signals from diverse sources more weight, thus potentially alleviating scenarios in which a large constituency of malicious actors is built up. Attackers can simulate diversity by creating multiple malicious sybil authorities. These newly established authorities would, however, be subject to the measures previously discussed, thereby limiting their influence.

Ousting The result of removing an authority from a network has a notable side-effect: Any previous endorsements that any of the constituents of the ousted

[5]This is common in real-life, for example in the 2020 branch-stacking scandal in the Victorian branch of the Australian Labor Party, which led to resignations of senior ministers in the Victorian state government [41].

authority have made, become void once the authority is removed. Consequently, in a scenario where a malicious authority has constructed a tree-like endorsement structure to circumvent any of the safeguards discussed previously, child authorities will instantaneously lose the endorsement of the ousted parent's constituents.

Reputational Threshold Once authorities are disproportionally discouraged on the network, they are likely to pose a threat to the integrity of the system. For this reason, deteriorating identity authorities should eventually lose all influence on the network, even if, in arithmetical terms, personhood tokens issued by them still hold value. To implement this, a lower bound for personhood scores can be introduced. Once an authority falls under that threshold, the personhood tokens issued by them would be devalued.

5 Conclusion and Future Work

In this paper, we have shown how a delegated proof-of-stake protocol can be evolved into a protocol supporting a self-governing public decentralised system. This supports the strategy of transitioning systems operated by a single entity to permissionless systems over time, which can be observed in major commercial blockchain initiatives [2, 46]. While, as a work in progress, the protocol proposed lacks formalisation, intuition suggests that the concept of evolving constituencies, backed by identity authorities, that can be added to and removed from a network dynamically, has merit. We have described how such an evolving system is more flexible than a proof-of-authority system, when constituencies change. We also outline how the protocol proposed might come closer to the ideal of 'one person/one vote' than other common approximations. We introduce a numeric 'personhood score', that allows for probabilistic calculations, taking into account the likelihood of a presented public key being the only identity of a given natural person. We anticipate attacks on such networks and show how damping of 'personhood scores' can mitigate those. Future work must focus on formalising the protocol to evaluate its robustness. A formal approach will ultimately prove or disprove its advantages over existing membership selection protocols, in the context of attacks.

References

1. Bach, L.M., Mihaljevic, B., Zagar, M.: Comparative analysis of blockchain consensus algorithms. In: 2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO). pp. 1545–1550 (2018). https://doi.org/10.23919/MIPRO.2018.8400278
2. Baird, L., Harmon, M., Madsen, P.: Hedera: A public hashgraph network & governing council. Whitepaper 2.0, Hedera Hashgraph (September 2019), https://www.hedera.com/hh-whitepaper-v2.0-17Sep19.pdf

3. Barone, R., Masciandaro, D.: Cryptocurrency or usury? crime and alternative money laundering techniques. European Journal of Law and Economics **47**(2), 233–254 (February 2019). https://doi.org/10.1007/s10657-019-09609-6
4. block.one: EOS.IO technical white paper. Whitepaper v2 (2018), https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md
5. Boldyreva, A., Palacio, A., Warinschi, B.: Secure proxy signature schemes for delegation of signing rights. Journal of Cryptology **25**(1), 57–115 (October 2010). https://doi.org/10.1007/s00145-010-9082-x
6. Borge, M., Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Ford, B.: Proof-of-personhood: Redemocratizing permissionless cryptocurrencies. In: 2017 IEEE European Symposium on Security and Privacy Workshops (EuroS PW). pp. 23–26 (2017). https://doi.org/10.1109/EuroSPW.2017.46
7. Buterin, V., Griffith, V.: Casper the friendly finality gadget (2017)
8. Cannarsa, M.: Interpretation of contracts and smart contracts: Smart interpretation or interpretation of smart contracts? European Review of Private Law pp. 773–785 (2018)
9. Chaum, D., Fiat, A., Naor, M.: Untraceable electronic cash. In: Advances in Cryptology — CRYPTO' 88, pp. 319–327. Springer New York (1990). https://doi.org/10.1007/0-387-34799-2_25
10. Conti, M., Sandeep Kumar, E., Lal, C., Ruj, S.: A survey on security and privacy issues of Bitcoin. IEEE Communications Surveys & Tutorials **20**(4), 3416–3452 (2018). https://doi.org/10.1109/COMST.2018.2842460
11. Croman, K., Decker, C., Eyal, I., Gencer, A.E., Juels, A., Kosba, A., Miller, A., Saxena, P., Shi, E., Gün Sirer, E., Song, D., Wattenhofer, R.: On scaling decentralized blockchains. In: Clark, J., Meiklejohn, S., Ryan, P.Y., Wallach, D., Brenner, M., Rohloff, K. (eds.) Financial Cryptography and Data Security. pp. 106–125. Springer Berlin Heidelberg, Berlin, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53357-4_8
12. Daian, P., Pass, R., Shi, E.: Snow White: Robustly reconfigurable consensus and applications to provably secure proof of stake. In: Goldberg, I., Moore, T. (eds.) Financial Cryptography and Data Security. pp. 23–41. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-32101-7_2
13. Dhillon, A., Kotsialou, G., McBurney, P., Riley, L.: Introduction to voting and the blockchain: some open questions for economists. CAGE Online Working Paper Series 416, Competitive Advantage in the Global Economy (CAGE) (2019), https://ideas.repec.org/p/cge/wacage/416.html
14. Douceur, J.R.: The Sybil attack. In: Revised Papers from the First International Workshop on Peer-to-Peer Systems. pp. 251–260. IPTPS '01, Springer-Verlag, Berlin, Heidelberg (2002). https://doi.org/10.5555/646334.687813
15. Durlauf, S.N., Blume, L.E.: Incentive compatibility. In: Durlauf, S.N., Blume, L.E. (eds.) Game Theory. pp. 158–168. Palgrave Macmillan UK, London (2010). https://doi.org/10.1057/9780230280847_16
16. Foster, C., Herring, J.: Theories of Personhood, pp. 21–34. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-53459-6_2
17. Gilad, Y., Hemo, R., Micali, S., Vlachos, G., Zeldovich, N.: Algorand. In: Proceedings of the 26th Symposium on Operating Systems Principles. ACM (October 2017). https://doi.org/10.1145/3132747.3132757
18. Grossman, S.J., Hart, O.D.: One share/one vote and the market for corporate control. Working Paper 2347, National Bureau of Economic Research (August 1987). https://doi.org/10.3386/w2347, http://www.nber.org/papers/w2347

19. Gui, G., Hortacsu, A., Tudon, J.: A memo on the proof-of-stake mechanism (July 2018)
20. Hearn, M., Brown, R.G.: Corda: A distributed ledger. Whitepaper Version 1.0, R3 (August 2019), https://www.r3.com/wp-content/uploads/2019/08/corda-technical-whitepaper-August-29-2019.pdf
21. Hellwig, D., Karlic, G., Huchzermeier, A.: Privacy and Anonymity, pp. 99–121. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-40142-9_5
22. Heusser, J.: SAT solving–an alternative to brute force bitcoin mining (February 2013), https://jheusser.github.io/2013/02/03/satcoin.html
23. Karame, G.O., Androulaki, E., Capkun, S.: Double-spending fast payments in Bitcoin. In: Proceedings of the 2012 ACM Conference on Computer and Communications Security. pp. 906–917. CCS '12, Association for Computing Machinery, New York, NY, USA (2012). https://doi.org/10.1145/2382196.2382292
24. Kethineni, S., Cao, Y.: The rise in popularity of cryptocurrency and associated criminal activity. International Criminal Justice Review pp. 325–344 (February 2019). https://doi.org/10.1177/1057567719827051
25. King, S., Nadal, S.: PPCoin: Peer-to-peer crypto-currency with proof-of-stake. Self-published (August 2012), https://decred.org/research/king2012.pdf
26. Kroll, J.A., Davey, I.C., Felten, E.W.: The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In: Proceedings of WEIS. vol. 2013, p. 11 (2013)
27. Lamport, L., Shostak, R., Pease, M.: The byzantine generals problem. ACM Transactions on Programming Languages and Systems **4**(3), 382–401 (July 1982). https://doi.org/10.1145/357172.357176
28. Larimer, D.: Delegated proof-of-stake (DPOS) (April 2014), http://107.170.30.182/security/delegated-proof-of-stake.php
29. Lee, C.: Litecoin (2011)
30. Li, W., Andreina, S., Bohli, J.M., Karame, G.: Securing proof-of-stake blockchain protocols. In: Garcia-Alfaro, J., Navarro-Arribas, G., Hartenstein, H., Herrera-Joancomartí, J. (eds.) Data Privacy Management, Cryptocurrencies and Blockchain Technology. pp. 297–315. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-67816-0_17
31. Libra Association Members: The Libra payment system. Whitepaper 2.0, Libra Association, Geneva, Switzerland (April 2020), https://libra.org/en-US/wp-content/uploads/sites/23/2020/04/Libra_WhitePaperV2_April2020.pdf
32. McBurney, P., Parsons, S.: Engineering democracy in open agent systems. In: Engineering Societies in the Agents World IV, pp. 66–80. Springer Berlin Heidelberg (2004). https://doi.org/10.1007/978-3-540-25946-6_4
33. McCorry, P., Shahandashti, S.F., Hao, F.: A smart contract for boardroom voting with maximum voter privacy. In: Kiayias, A. (ed.) Financial Cryptography and Data Security. pp. 357–375. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-70972-7_20
34. Miller, A.: Feather-forks: enforcing a blacklist with sub-50% hash power. Bitcoin Forum Post (October 2013), https://bitcointalk.org/index.php?topic=312668.0
35. Mora, C., Rollins, R.L., Taladay, K., Kantar, M.B., Chock, M.K., Shimada, M., Franklin, E.C.: Bitcoin emissions alone could push global warming above 2°c. Nature Climate Change **8**(11), 931–933 (October 2018). https://doi.org/10.1038/s41558-018-0321-8
36. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008)
37. Natoli, C., Yu, J., Gramoli, V., Esteves-Verissimo, P.: Deconstructing blockchains: A comprehensive survey on consensus, membership and structure (2019)

38. Orman, H.: Blockchain: the Emperors New PKI? IEEE Internet Computing **22**(2), 23–28 (March 2018). https://doi.org/10.1109/mic.2018.022021659
39. Ostrom, E.: Self-governance and forest resources. Occasional Paper 20, Center for International Forestry Research (February 1999)
40. Ostrom, E., Walker, J., Gardner, R.: Covenants with and without a sword: Self-governance is possible. American Political Science Review **86**(2), 404–417 (June 1992). https://doi.org/10.2307/1964229
41. Patrick, A., Marin-Guzman, D.: 'Everyone knew what was going on'. The Australian Financial Review (June 2020)
42. QuantumMechanic: Proof of stake instead of proof of work. Bitcoin Forum Post (July 2011), https://bitcointalk.org/index.php?topic=27787.0
43. Rauchs, M., Glidden, A., Gordon, B., Pieters, G., Recanatini, M., Rostand, F., Vagneur, K., Zhang, B.: Distributed ledger technology systems. a conceptual framework. Tech. rep., Cambridge Centre for Alternative Finance, Cambridge, UK (August 2018)
44. Saleh, F.: Blockchain without waste: Proof-of-stake. The Review of Financial Studies (July 2020). https://doi.org/10.1093/rfs/hhaa075
45. Schwartz, D., Youngs, N., Britto, A.: The Ripple protocol consensus algorithm. Whitepaper, Ripple Labs Inc (2014)
46. Shehar, B., Catalini, C., Danezis, G., Doudchenko, N., Maurer, B., Sonnino, A., Wernerfelt, N.: Moving toward permissionless consensus. Tech. rep., Libra Association, Geneva, Switzerland (June 2019), https://libra.org/wp-content/uploads/2019/06/MovingTowardPermissionlessConsensus_en_US.pdf
47. Stroukal, D., Nedvědová, B.: Bitcoin and other cryptocurrency as an instrument of crime in cyberspace. Proceedings of Business and Management Conferences 4407036, International Institute of Social and Economic Sciences (November 2016), https://ideas.repec.org/p/sek/ibmpro/4407036.html
48. Szabo, N.: Formalizing and securing relationships on public networks. First Monday **2**(9) (September 1997). https://doi.org/10.5210/fm.v2i9.548
49. Thin, W.Y.M.M., Dong, N., Bai, G., Dong, J.S.: Formal analysis of a proof-of-stake blockchain. In: 2018 23rd International Conference on Engineering of Complex Computer Systems (ICECCS). IEEE (December 2018). https://doi.org/10.1109/iceccs2018.2018.00031
50. Townsend, R.E.: Fisheries self-governance: corporate or cooperative structures? Marine Policy **19**(1), 39–45 (January 1995). https://doi.org/10.1016/0308-597X(95)92571-N
51. Vasek, M., Thornton, M., Moore, T.: Empirical analysis of denial-of-service attacks in the Bitcoin ecosystem. In: Böhme, R., Brenner, M., Moore, T., Smith, M. (eds.) Financial Cryptography and Data Security. pp. 57–71. Springer Berlin Heidelberg, Berlin, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44774-1_5
52. Waidner, M., Pfitzmann, B.: Loss-tolerance for electronic wallets. In: Digest of Papers. Fault-Tolerant Computing: 20th International Symposium. pp. 140–147 (1990). https://doi.org/10.1109/FTCS.1990.89349
53. Wood, G.: Ethereum: A secure decentralised generalised transaction ledger. Yellow paper, Stiftung Ethereum (2017), https://ethereum.github.io/yellowpaper/paper.pdf
54. Wuille, P.: Dealing with malleability. Bitcoin Improvement Proposal (March 2014), https://github.com/bitcoin/bips/blob/master/bip-0062.mediawiki