

Federated Blockchain Systems: A better trade-off between sustainability and decentralization?

Florian, Martin

Erstveröffentlichung / Primary Publication

Arbeitspapier / working paper

This work has been funded by the Federal Ministry of Education and Research of Germany (BMBF) (grant no.: 16DII121, 16DII122, 16DII123, 16DII124, 16DII125, 16DII126, 16DII127, 16DII128 - "Deutsches Internet-Institut").

Empfohlene Zitierung / Suggested Citation:

Florian, M. (2022). *Federated Blockchain Systems: A better trade-off between sustainability and decentralization?* (Weizenbaum Series, 26). Berlin: Weizenbaum Institute for the Networked Society - The German Internet Institute. <https://doi.org/10.34669/WI.WS/26>

Nutzungsbedingungen:

Dieser Text wird unter einer CC BY Lizenz (Namensnennung) zur Verfügung gestellt. Nähere Auskünfte zu den CC-Lizenzen finden Sie hier: <https://creativecommons.org/licenses/by/4.0/deed.de>

Terms of use:

This document is made available under a CC BY Licence (Attribution). For more information see: <https://creativecommons.org/licenses/by/4.0>

Martin Florian

Federated Blockchain Systems

A better trade-off between sustainability and decentralization?

Federated Blockchain Systems*

Martin Florian \ Weizenbaum Institute for the Networked Society and Humboldt-Universität zu Berlin \
martin.florian@hu-berlin.de

ISSN 2748-5587 \ DOI [10.34669/WI.WS/26](https://doi.org/10.34669/WI.WS/26)

EDITORS: The Managing Board members of the Weizenbaum-Institut e.V.
Prof. Dr. Christoph Neuberger
Prof. Dr. Sascha Friesike
Prof. Dr. Martin Krzywdzinski
Dr. Karin-Irene Eiermann

Hardenbergstraße 32 \ 10623 Berlin \ Tel.: +49 30 700141-001
info@weizenbaum-institut.de \ www.weizenbaum-institut.de

TYPESETTING: Roland Toth, M.A.

COPYRIGHT: This series is available open access and is licensed under Creative Commons Attribution 4.0 (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>

WEIZENBAUM INSTITUTE: The Weizenbaum Institute for the Networked Society – The German Internet Institute is a joint project funded by the Federal Ministry of Education and Research (BMBF). It conducts interdisciplinary and basic research on the changes in society caused by digitalisation and develops options for shaping politics, business and civil society.

This work has been funded by the Federal Ministry of Education and Research of Germany (BMBF) (grant no.: 16DII121, 16DII122, 16DII123, 16DII124, 16DII125, 16DII126, 16DII127, 16DII128 – “Deutsches Internet-Institut”).

Abstract

Blockchain-based systems are enjoying unbroken popularity. Different economic and social actors are investigating their application for fostering decentralization and separation of power. Whether a blockchain-based system can live up to such goals is heavily determined by the choice of a consensus protocol – the rules by which participants agree on what gets added to the blockchain. Bitcoin’s consensus protocol is inherently decentralization-enabling, at a notoriously high ecological cost. So-called permissioned protocols, while incomparably more efficient,

are dismissed as being closed-off and “centralized”. Federated blockchain systems represent a middle ground between these two extremes and promise to offer openness and security without sacrificing ecological sustainability. As a rough approximation, their approach can be described as bootstrapping consensus from a web of trust. In this overview article, after a short review of the Bitcoin approach and possible alternatives to it, we introduce the ideas behind federated blockchain systems and discuss their impact on future blockchain systems.

* Disclaimer: This article is intended for a technology-interested but essentially non-technical audience. For aiding comprehension, some concepts will be described in more high-level terms. Please consult the cited works for details and formal write-ups.

Table of Contents

1	Introduction	4
2	Non-Permissioned and yet Non-Wasteful Consensus	5
3	From FBASs to Federated Blockchain Systems	7
4	Are Federated Blockchain Systems “Centralized”?	9
5	Can Permissioned Systems be Improved?	10
6	Conclusion	11
7	References	11

1 Introduction

Blockchain-based systems [4] are enjoying unbroken popularity. Different economic and social actors are investigating their application for fostering decentralization and separation of power. Whether a blockchain-based system can live up to such goals is heavily determined by the choice of a *consensus protocol* – the rules by which participants agree¹ on what gets added to the blockchain², and in what order. Different consensus protocols build upon different assumptions about the identities of participants and the relationships between them. On one end of the spectrum, in what is commonly called *permissioned systems*, a static group of so-called *validators* is selected in an a priori, “top-down” manner. The permissioned approach thereby shifts much of the burden of achieving “decentralization” outside of the technical system, and is hence viewed with skepticism by proponents of radically open systems such as Bitcoin [21, 17]. Bitcoin-like systems (usually cryptocurrencies) occupy the other end of the spectrum and are commonly called *permissionless* – no knowledge about other participants is assumed and the ability to influence consensus is determined by the continuous investment of computing resources (*mining*). While this approach has so far proven effective at securing valuable global networks such as the Bitcoin network, it is also vastly resource-inefficient and thus unsustainable from an ecological standpoint. Various attempts have been made to develop more energy-efficient permissionless consensus approaches, for example so-called Proof-of-Stake [19] protocols. Despite significant research and development, critical obstacles to realizing fully permissionless, secure and yet non-wasteful consensus systems remain (we discuss some of these challenges below).

Federated blockchain systems are located on the middle ground of the permissioned–permissionless spectrum. They represent a promising new approach for reconciling aspirations to decentrality and openness with the necessity of ecological sustainability. We loosely define a federated blockchain system as a blockchain system that is built upon the *Federated Byzantine Agreement System* (FBAS) model of consensus [18] or a closely related model such as the *Asymmetric Quorum System* (AQS) or *Personal Byzantine Quorum System* (PBQS) [11] models. In very rough terms, all of these models describe structures reminiscent of a “web of trust” or social network, with some extra complexity. Each participant, modeled as a network *node*, defines its own rules about which groups of other nodes it requires agreement from. Protocols like the *Stellar Consensus Protocol* (SCP) [18] leverage the resulting structure for establishing a consensus system [6, 15, 10]. A federated blockchain system uses SCP, or a similar protocol, for agreeing on the state of the blockchain. A visual overview is given in Figure 1. Examples for federated blockchain systems include the popular *Stellar*³ [10] and *MobileCoin*⁴ [16] networks. In terms of ecological costs, federated blockchain systems are incomparably more sustainable than, for example, Bitcoin.

Not every web-like structure can be used for bootstrapping a useful federated blockchain system. Federated blockchain systems are only *live* (i.e., capable of producing new blocks) and *safe* (i.e., protected against inconsistencies) if the interplay between individual node configurations results in

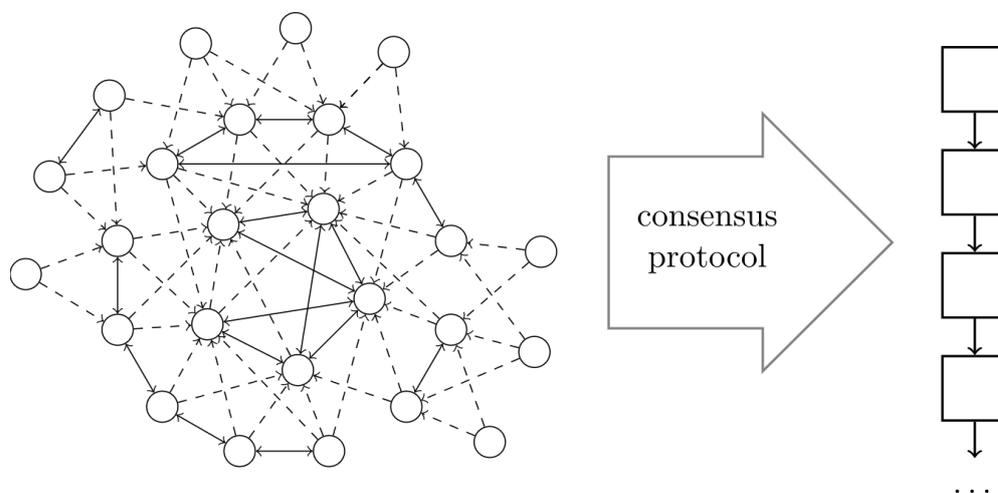
¹ Despite their name, consensus protocols do not typically require *all* participants to agree but are rather based on different forms of majority voting.

² *The blockchain* is essentially a shared register, often also called a *distributed ledger*. Full copies of the blockchain are stored by multiple, if not all, interested parties.

³ <https://stellar.org>

⁴ <https://mobilecoin.com/>

Figure 1: In federated blockchain systems, the blockchain (on the right) is secured by leveraging preexisting relationships between nodes (on the left)



certain global properties. In our work [2], we focus on the conditions and dynamics under which these properties emerge, or fail to emerge. One of our key results is that an often-small group of nodes is exclusively relevant for determining liveness buffers, i.e., how many nodes can become faulty before liveness is lost. We also find that it can be quite hard for nodes to become part of this *top tier*, raising the

question to what extent federated blockchain systems are actually better than permissioned ones.

In this article, we give an overview over possible alternatives to the energy-hungry “Bitcoin approach”, introduce federated blockchain systems and FBASs⁵ in greater detail and finally discuss the potential impact of these ideas on future blockchain systems.

2 Non-Permissioned and yet Non-Wasteful Consensus

Bitcoin [21, 17] is well known as the first and biggest (i.a., by market capitalization) decentralized cryptocurrency. Bitcoin is also well known, however, for its immense energy usage⁶, which is an inherent result of its underlying consensus protocol. Bitcoin consensus is based on the concept of *Proof-of-Work* (PoW): nodes are allowed to participate in the forming of new blockchain blocks in proportion to their investment of computing resources. The energy consumption of Bitcoin is (game-theoretically) designed

to rise with the price of the underlying cryptocurrency and may therefore potentially reach even higher levels in the future, further threatening global sustainability goals.

One of the main functions of PoW in Bitcoin is the protection against *Sybil attacks* [22], which is easily one of the greatest challenges when designing non-permissioned systems. In a Sybil attack, an adversary creates a large number of fake identities

⁵ We focus on the FBAS model in favor of related models due to its currently higher practical relevance. At the level of abstraction of this article, the differences to the AQS or PBQS models are negligible.

⁶ Up-to date estimates are available, for example, via the Cambridge Bitcoin Electricity Consumption Index (<https://cbeci.org/>). At the time of writing, the index estimates that the energy consumption of Bitcoin is higher than that of the Netherlands.

and uses them to influence voting-based processes and distort the “majority view”. Aiming at decentralization and openness, the Bitcoin network is comprised of peers that can join anonymously and without asking for permission in any way (hence the term *permissionless*). However, these are the exact preconditions for adversaries to create an unbounded number of fake identities and thereby undermine protocol mechanisms, i.e., to mount a Sybil attack. PoW-based consensus is resistant to such attacks by granting consensus “voting rights” not based on identities, but based on the amounts of continuously invested energy. Permissionless alternatives to PoW must identify a similar resource that cannot easily be scaled by an adversary.

A popular idea for realizing permissionless and yet energy-efficient consensus is to replace PoW with *Proof-of-Stake* (PoS) [19, 17, 7]. In a PoS-based consensus protocol, nodes participate in consensus in proportion to the amount of cryptocurrency units they possess and are willing to “invest” (the cryptocurrency units are locked away for a certain period in exchange for voting rights). PoS is often cited as a viable alternative to PoW. However, fundamental flaws in the PoS approach exist that make its merit for powering secure and decentralized systems questionable. Among other things, PoS systems are more vulnerable to history rewriting attacks (in what is also known as a *long-range attack* or *costless simulation* attack). If an adversary compromises enough cryptocurrency addresses to control the majority of funds at any point in the “past”, he can create a fork of the blockchain starting from that point [7]. The faultiness of this fork is only detectable under strong assumptions about network coherence⁷ or through the use of trusted third parties. Other arguments against PoS include the fact that the allocation of voting rights based on monetary resources can yield systems that are plutocratic, i.e., ultimately not quite egalitarian.

PoS is notably not the only contender in the field of permissionless and yet energy-efficient consensus. It is, however, the by far best-researched and -tested one. The amount of drawbacks and challenges that the PoS approach still faces is symptomatic of the more general difficulty of enabling Sybil-resistant consensus without resorting to any form of existing knowledge about nodes. Permissioned systems such as the *Diem*⁸ cryptocurrency (prominently pushed by Facebook/Meta) or blockchain networks based on the *Hyperledger Fabric* [13] platform incorporate such knowledge. They require that a group of nodes is agreed upon a priori to serve as *validators* and participate in consensus (exclusively). How this group is selected remains outside of the technical scope. For example, questions of decentralization can be offloaded to external governance structures such as foundations. The lack of trust in non-technical decision-making organs is possibly one of the reasons for the genesis of decentralized cryptocurrencies like Bitcoin, however. Can decentralization and separation of power still be realized through some clever technological design, creating a compelling alternative to energy-hungry systems such as Bitcoin?

In the more general context of designing Sybil-resistant systems, a well-researched middle way between non-technical access control and permissionless proof-of-X-type approaches is to leverage existing one-to-one relationships between participants, most notably their social network (see, e.g., [20] for an overview). Social network-based Sybil defense mechanisms leverage the fact that while an adversary might be able to create Sybil identities and simulate relationships between them, resulting Sybil clusters will be connected to the remaining network only weakly. The underlying assumption is that the forming of connections to honest nodes is hard and not arbitrarily scalable for an adversary, i.e., that he cannot fool an arbitrary number of honest participants an arbitrary number of times. To the

⁷ Which can be difficult to satisfy in practice, as we also demonstrate in our works such as [8].

⁸ <https://diem.com/>

best of our knowledge, the first system that builds upon this type of assumption for bootstrapping a secure consensus system is the *Ripple* network [14]. Here, each node operator decides by himself which subset of the node population, or *unique node list* (UNL), his nodes will consider as relevant. For example, he might select nodes operated by trustworthy organizations, or nodes operated by friends. In essence, the sum of all individual UNLs defines

a “web of trust” between the nodes. Given sufficient overlap between all UNLs, the resulting system becomes capable of safe and live consensus. However, the specific criteria for achieving safety are quite strict, resulting in a high risk of systems evolving to become de-facto permissioned [14, 10]. Ripple’s UNL model is, in essence, a strictly less expressive predecessor to the FBAS model, which we will discuss in the following.

3 From FBASs to Federated Blockchain Systems

Federated blockchain systems are blockchain systems bootstrapped from a *Federated Byzantine Agreement System* (FBAS) [18] or a related formalization (for example [1], [11]; we use the FBAS model in the following). The security and performance of any federated blockchain system is heavily influenced by the structure of its underlying FBAS. In the following, we give an informal overview of what an FBAS actually is, introducing key concepts and terminology. We defer to works such as [2] for a more in-depth exploration of the topic.

An FBAS constitutes of a set of nodes. Each node is associated with an individual description of which groups of other nodes it requires agreement from. In practice, nodes are *configured* in this way by the individuals or organizations that operate them (typically, their “owners”). Configurations can be based on arbitrary considerations. For example, an operator might choose to include nodes of operators that he subjectively “trusts”, or nodes belonging to organizations that he strongly wishes to remain “in sync” with. A *quorum set* is a handy format for configuring

an FBAS node. In their most simple form, quorum sets define a group of nodes along with a threshold value that denotes how many of the constituent nodes must be in agreement. Figure 2 illustrates this abstract definition with an example FBAS. Depicted are the quorum sets of five nodes⁹ (on the right) and a heuristic graph representation of the resulting FBAS¹⁰.

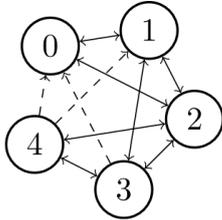
In the example, node 0 requires agreement from both nodes 1 and 2, whereas node 1 requires agreement from at least two nodes out of the set $\{0,2,3\}$. At least one of the combinations $\{0,2\}$, $\{0,3\}$, $\{2,3\}$ must agree on the same (valid) value in order for node 1 to accept and eventually confirm that value. In a federated blockchain system, the *value* typically corresponds to the contents of the next blockchain block.

Informally, an FBAS enables *liveness* for a consensus protocol like SCP if, when honoring the configurations of all nodes, the consensus protocol is able to make progress, e.g., extend a blockchain with new blocks. FBAS nodes agree on new values only

⁹ We use numbers here for naming nodes, for simplicity and conciseness. Still, the reader is free to imagine his own names based on entities that he considers particularly trustworthy. Node 0 could be the Weizenbaum Institute, for example, node 1 ...

¹⁰ A directed edge in the graph implies that the edge’s head node is included in the tail node’s quorum set. Note that any representation of an FBAS as a regular graph is necessarily incomplete, as constraints such as “ m out of n nodes must agree” cannot be modelled.

Figure 2: An example FBAS (V, Q) with nodes $V = \{0, 1, 2, 3, 4\}$ and quorum sets Q (on the right) defined in an informal manner. Also depicted (on the left) is a heuristic graph representation of the FBAS.



Quorum sets (informal):

$Q(0)$: all nodes in $\{1, 2\}$ must agree

$Q(1)$: 2 nodes out of $\{0, 2, 3\}$ must agree

$Q(2)$: 3 nodes out of $\{0, 1, 3, 4\}$ must agree

$Q(3)$: 3 nodes out of $\{0, 1, 2, 4\}$ must agree

$Q(4)$: 3 nodes out of $\{0, 1, 2, 3\}$ must agree

when sufficient nodes in their quorum sets agree on the same value. In the above example, node 0 will only accept a new block if it is certain that nodes 1 and 2 will accept it as well. A group of FBAS nodes that can by itself agree on new values is called a *quorum*. In the above example, nodes $\{0, 1, 2, 3\}$ form a quorum, but nodes $\{0, 1, 2\}$ do not (node 2 wouldn't be satisfied). Sets of nodes that intersect every quorum in an FBAS are known as *blocking sets*. If all nodes in a blocking set crash or become uncooperative, no quorums can be formed and liveness is necessarily lost. In the above example, $\{2\}$ is already a blocking set.

In order to guarantee *safety*, i.e., that no two sets of nodes agree on conflicting values, an FBAS must enjoy *quorum intersection*. This means that each two quorums should share at least one common node. Lack of quorum intersection in an FBAS can, for example, lead to forks and double spends. A *splitting set* is a set of nodes that can, if malicious, compromise quorum intersection, by lying about node configurations and accepted values. In the above example, $\{0, 1, 2\}$ form a splitting set and could pull off such an attack, causing nodes 3 and 4 to diverge from each other.

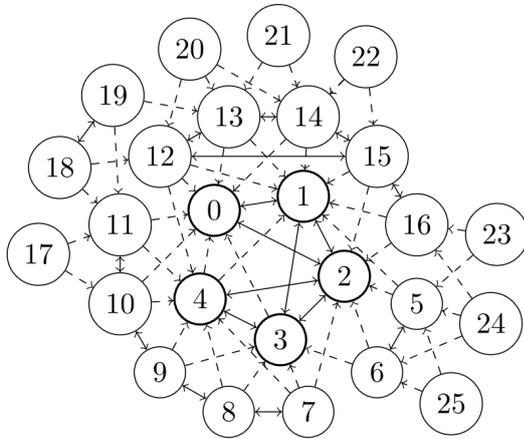
An important part of FBAS analysis consists in determining minimal blocking sets and minimal splitting sets, thereby giving a lower bound for the number of nodes that need to be compromised for liveness or safety to become threatened. Identifying

the relevant sets is a computationally intensive task. We developed algorithms and a comprehensive analysis framework¹¹ to be able to perform analyses efficiently [2].

Perhaps more interestingly, however, we also uncovered that minimal blocking sets and some of the most impactful minimal splitting sets are formed out of a group of nodes that is usually significantly smaller than the whole node population of an FBAS. We call this group of nodes the *top tier* of an FBAS – the group of nodes that are exclusively relevant for determining FBAS-wide liveness buffers. In the FBAS shown in Figure 3, the top tier T is formed by the nodes $T := \{0, 1, 2, 3, 4\}$. In terms of its most important properties, the complex FBAS from Figure 3 is identical to the much simpler one from Figure 2. For example, a failure of node 2 would still halt the whole system ($\{2\}$ being a blocking set) and a group of nodes as small as $\{0, 1, 2\}$ can still, by themselves, split the system and cause a blockchain fork.

¹¹ https://github.com/wiberlin/fbas_analyzer

Figure 3: A more complex example FBAS (V,Q) with nodes $V = \{0,1,\dots,24,25\}$ and the values of Q (on the right) defined in an informal manner. Also depicted (on the left) is a heuristic graph representation of the FBAS. Through analysis, it can be shown that nodes 0–4 form the top tier of (V,Q) .



Quorum configurations (informal):

- Q(0): all nodes in $\{1, 2\}$ must agree
- Q(1): 2 nodes out of $\{0, 2, 3\}$ must agree
- Q(2): 3 nodes out of $\{0, 1, 3, 4\}$ must agree
- Q(3): 3 nodes out of $\{0, 1, 2, 4\}$ must agree
- Q(4): 3 nodes out of $\{0, 1, 2, 3\}$ must agree
- Q(5): 2 nodes out of $\{1, 2, 6\}$ must agree
- Q(6): node 5 must agree, and 1 node out of $\{2, 3\}$
- ...

4 Are Federated Blockchain Systems “Centralized”?

Previous works such as [9] have noted that the Stellar network, the largest federated blockchain system deployed today, exhibits strong signs of centralization. Based on our own analysis of the Stellar network and others¹², we conclude that federated blockchain systems naturally develop small top tiers that become defining for the systems’ core safety and liveness properties [2]. To a large extent, top tier nodes are the only ones playing a relevant role in consensus – just like validator nodes in a permissioned system. We argue, however, that it is not only the existence and size of a top tier that determines “centralization”, but also the dynamism of that group’s membership, i.e., whether it is feasible for lower-tier nodes to rise to top-tier status and whether underperforming or misbehaving top-tier nodes can be effectively ejected from their privileged position. Our formal analysis in [2] concludes that changes to the top tier are a significant challenge if the residing top tier is opposed

to them and lower-tier node operators prioritize safety when configuring their nodes.

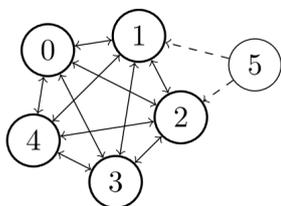
And yet, the composition of a federated system’s top tier can be significantly more dynamic than the set of validators of a typical permissioned system. The reason is that top tier nodes in a federated blockchain system can make *individual decisions* about who should be able to influence consensus, while validators in a permissioned system are bound to a global configuration that must be the same for all.

Consider the example in Figure 4. Nodes 0–4 form a top tier. Any one of these nodes’ operators can, if he believes it’s a good idea, change his node’s configuration to elevate node 5 to top tier status as well. For example, node 0 could be reconfigured so that it needs agreement from 4 nodes out of $\{1,2,3,4,5\}$ ¹³.

¹² See also our recent report on the MobileCoin network [3].

¹³ As a consequence, $\{3, 5\}$ and $\{4, 5\}$ would become minimal blocking sets and 2 of the 3 minimal quorums of the FBAS would include node 5.

Figure 4: An example FBAS (V, Q) with nodes $V = \{0, 1, 2, 3, 4, 5\}$ and the values of Q (on the right) defined in an informal manner. It is sufficient for any of the top tier nodes 0–4 to include 5 in its quorum configuration in order for 5 to become a top tier node and thereby relevant for consensus.



The operator of a validator node in a classical permissioned system¹⁴ doesn't have this level of autonomy. For illustration, note that the federated system implied by the original configurations of nodes 0–4 from Figure 4 is identical to a classical permissioned system with 5 validators (nodes 0–4) and the threshold rule “4 nodes out of $\{0, 1, 2, 3, 4\}$ must agree”. In order for node 5 to become a validator in such a permissioned system, nodes 0–4 must accept a new global configuration such as “5 nodes out of $\{0, 1, 2, 3, 4, 5\}$

Quorum configurations (informal):

Q(0): 3 nodes out of $\{1, 2, 3, 4\}$ must agree

Q(1): 3 nodes out of $\{0, 2, 3, 4\}$ must agree

Q(2): 3 nodes out of $\{0, 1, 3, 4\}$ must agree

Q(3): 3 nodes out of $\{0, 1, 2, 4\}$ must agree

Q(4): 3 nodes out of $\{0, 1, 2, 3\}$ must agree

Q(5): all nodes in $\{1, 2\}$ must agree

must agree”. The standard threshold model we reproduce here does not offer much more flexibility than that for fine-tuning a validator's importance. Existing validators would need to be persuaded (or coerced) by a non-technical governance body to accept the new global configuration, or they might vote on it “on the blockchain”. In any case, gathering support for such a change might be hard if validators (respectively the individuals or organizations behind them) have different opinions about node 5.

5 Can Permissioned Systems be Improved?

Permissioned blockchain systems are significantly easier to reconcile with sustainability goals than their permissionless counterparts, but are often written off as being “too centralized”. In summary of the preceding discussion, we can raise two specific criticisms about permissioned blockchain systems that are resolved by moving to a similarly resource-efficient federated design:

1. It's not possible to formulate agreement rules that are more complex than a single flat “ m out of n ”, for example for giving different validators a different weight in consensus.
2. All changes to the agreement rules require a coordinated, non-technical intervention or at least a majority vote.

¹⁴ More specifically, we are referring here to systems that use a threshold-based consensus protocol with finality, such as *Practical Byzantine Fault Tolerance* (PBFT) [23] (an early proposal that has inspired many modern protocols) and *HotStuff* [12] (a state-of-the-art protocol optimized for use in blockchain systems).

Arguably, these limitations are central to the image of permissioned blockchain systems as closed-up, static systems. Federated blockchain systems solve these limitations. However, they also come with their own drawbacks – they are hard to grasp and explain to users, hard to monitor [2], and require the use of consensus protocols that are, in general, harder to design and implement. Also, in cases where a permissioned system is already deployed, changing key components (such as the consensus protocol) for switching to a federated design might not be an easy option.

So is it possible to use ideas from federated blockchain systems to make permissioned systems more open and attractive? It seems that this is the case. With respect to criticism 1, Alpos and Cachin [5] have recently shown that popular permissioned consensus protocols can be adapted to support arbitrary quorum rules. With such adaptations, a permissioned

system could support all configurations that can be expressed as an FBAS.

What about criticism 2? A possible update to a permissioned blockchain system is to allow validators to change the global configuration in any way they wish, without requiring any votes from other nodes, as long as the changes impact only the validator's *own role* in consensus. If a validator gets 2 votes in consensus, for example, he should be able to unilaterally change the global consensus rules so that some new node gets one of its votes from now on. Essentially, we suggest that permissioned systems can be designed in which each validator is free to make all changes to the global consensus rules that it would also have been able to make if the system was a federated one. Whether a change is permissible in this respect can be determined locally by each validator, for example using an adaptation of our analysis methodology from [2].

6 Conclusion

This article is an attempt to introduce the intriguing idea of federated blockchain systems to a wider audience. Federated blockchain systems represent a middle way approach that can enable high levels of decentralization, openness and security at a negligible ecological cost. We gave an informal overview of how they work – how each participant can choose simple agreement rules for himself that can, when taken together, be used for forming system-wide consensus on the state of a blockchain. We also

discussed how federated blockchain systems differ from permissioned systems and how the ideas behind federated blockchain systems can be used for making permissioned systems more open. Openness to change and new consensus participants can improve the public trust in a deployed blockchain system. This might not only increase the acceptance and adoption of the specific system but may also, more generally, help drive interest away from ecologically questionable approaches that are popular today.

7 References

- [1] C. Cachin, “Asymmetric distributed trust”, in *International Conference on Distributed Computing and Networking 2021*, ser. ICDCN '21, Nara, Japan: ACM, 2021, p. 3.
- [2] M. Florian, S. Henningsen, C. Ndolo, and B. Scheuermann, *The sum of its parts: Analysis of federated byzantine agreement systems*, Under journal review, 2021. arXiv: [2002.08101](https://arxiv.org/abs/2002.08101) [cs.DC].
- [3] C. Ndolo, S. Henningsen, and M. Florian, *Crawling the MobileCoin quorum system*, 2021. arXiv: [2111.12364](https://arxiv.org/abs/2111.12364) [cs.DC].
- [4] M.-C. Valiente and F. Tschorsch, “Blockchain-based technologies”, *Internet Policy Review*, vol. 10, no. 2, 2021.
- [5] O. Alpos and C. Cachin, “Consensus beyond thresholds: Generalized Byzantine quorums made live”, in *2020 International Symposium on Reliable Distributed Systems (SRDS)*, IEEE, 2020, pp. 21–30.
- [6] Á. García-Pérez and M. A. Schett, “Deconstructing Stellar consensus”, in *23rd International Conference on Principles of Distributed Systems (OPODIS 2019)*, vol. 153, Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2020, 5:1–5:16.
- [7] P. Daian, R. Pass, and E. Shi, “Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake”, in *International Conference on Financial Cryptography and Data Security*, Springer, 2019, pp. 23–41.
- [8] S. A. Henningsen, D. Teunis, M. Florian, and B. Scheuermann, “Eclipsing Ethereum peers with false friends”, in *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, IEEE, 2019, pp. 300–309.
- [9] M. Kim, Y. Kwon, and Y. Kim, “Is Stellar as secure as you think?”, in *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, Stockholm, Sweden: IEEE, 2019, pp. 377–385.
- [10] M. Lokhava et al., “Fast and secure global payments with Stellar”, in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, ser. SOSP '19, Huntsville, Ontario, Canada: ACM, 2019, pp. 80–96.
- [11] G. Losa, E. Gafni, and D. Mazières, “Stellar consensus by instantiation”, in *33rd International Symposium on Distributed Computing (DISC 2019)*, vol. 146, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019, 27:1–27:15.
- [12] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, “HotStuff: BFT consensus with linearity and responsiveness”, in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, ser. PODC '19, Toronto ON, Canada: ACM, 2019, pp. 347–356.
- [13] E. Androulaki et al., “Hyperledger Fabric: A distributed operating system for permissioned blockchains”, in *Proceedings of the thirteenth EuroSys conference*, 2018, pp. 1–15.
- [14] B. Chase and E. MacBrough, *Analysis of the XRP ledger consensus protocol*, 2018. arXiv: [1802.07242](https://arxiv.org/abs/1802.07242).
- [15] Á. García-Pérez and A. Gotsman, “Federated Byzantine Quorum Systems”, in *22nd International Conference on Principles of Distributed Systems (OPODIS 2018)*, vol. 125, Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 17:1–17:16.
- [16] *MobileCoin*, Nov. 2017. [Online]. Available: https://mobilecoin.foundation/pdf/MobileCoin_White_Paper.pdf.
- [17] F. Tschorsch and B. Scheuermann, “Bitcoin and beyond: A technical survey on decentralized digital currencies”, *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 2084–2123, 2016.
- [18] D. Mazières, *The Stellar consensus protocol: A federated model for Internet-level consensus*, 2015. [Online]. Available: <https://www.stellar.org/papers/stellar-consensus-protocol.pdf>.
- [19] S. King and S. Nadal, *PPCoin: Peer-to-peer crypto-currency with proof-of-stake*, 2012. [Online]. Available: <https://archive.org/details/PPCoinPaper>.
- [20] B. Viswanath, A. Post, K. P. Gummadi, and A. Mislove, “An analysis of social network-based sybil defenses”, *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4, pp. 363–374, 2010.
- [21] S. Nakamoto, *Bitcoin: A peer-to-peer electronic cash system*, 2008. [Online]. Available: <http://nakamotoinstitute.org/bitcoin/>.
- [22] J. R. Douceur, “The Sybil attack”, in *Peer-to-peer Systems*, Berlin, Heidelberg: Springer, 2002, pp. 251–260.
- [23] M. Castro and B. Liskov, “Practical byzantine fault tolerance”, in *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)*, New Orleans, Louisiana, USA: USENIX, 1999, pp. 173–186.