

Decentralization in Open Quorum Systems: Limitative Results for Ripple and Stellar

Andrea Bracciali 

Department of Computer Science, University of Stirling, UK
<http://www.cs.stir.ac.uk/~abb/>
abbracciali@gmail.com

Davide Grossi 

Bernoulli Institute for Maths, CS and AI, University of Groningen, The Netherlands
Amsterdam Center for Law and Economics, University of Amsterdam, The Netherlands
Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands
<http://www.davidegrossi.me>
d.grossi@rug.nl

Ronald de Haan 

Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands
<https://staff.science.uva.nl/r.dehaan/>
me@ronalddhaan.eu

Abstract

Decentralisation is one of the promises introduced by blockchain technologies: fair and secure interaction amongst peers with no dominant positions, single points of failure or censorship. Decentralisation, however, appears difficult to be formally defined, possibly a continuum property of systems that can be more or less decentralised, or can tend to decentralisation in their lifetime. In this paper we focus on decentralisation in quorum-based approaches to open (permissionless) consensus as illustrated in influential protocols such as the Ripple and Stellar protocols. Drawing from game theory and computational complexity, we establish limiting results concerning the decentralisation vs. safety trade-off in Ripple and Stellar, and we propose a novel methodology to formalise and quantitatively analyse decentralisation in this type of blockchains.

2012 ACM Subject Classification Security and privacy → Distributed systems security

Keywords and phrases Blockchain, decentralization, game theory, computational complexity

Digital Object Identifier 10.4230/OASICS.Tokenomics.2020.5

Related Version Earlier versions of the paper are available at <https://export.arxiv.org/abs/1911.08182>.

Acknowledgements We are indebted to the anonymous reviewers of Tokenomics 2020 for several punctual and helpful suggestions on an earlier version of the paper.

1 Introduction

To allow “any two willing parties to transact directly with each other without the need for a trusted third party” [29] was one of the main motivations for the introduction of the Bitcoin blockchain, and several earlier attempts at digital currencies. A blockchain is a distributed state machine in charge of guaranteeing the correctness and trustability of data,¹ e.g., monetary transactions in the case of Bitcoin. State updates are recorded in a chain of data blocks. Data are protected by replication of the state, i.e., of the chain of

¹ Blockchains can also make the computation trustable, e.g., guaranteeing the fair and untamperable execution of agreements among peers encoded as programs.



© Andrea Bracciali, Davide Grossi, and Ronald de Haan;
licensed under Creative Commons License CC-BY

2nd International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2020).

Editors: Emmanuelle Anceaume, Christophe Bisière, Matthieu Bouvard, Quentin Bramas, and Catherine Casamatta; Article No. 5; pp. 5:1–5:20



OpenAccess Series in Informatics

OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

blocks, within a network of peers. The blockchain protocol must guarantee some form of distributed consensus allowing peers to agree on the information contained in the blockchain, e.g., who has been paid, and that no double spending of virtual coins has occurred, without the supervision of a centralised authority – a currency without a central bank.

Context: the quest for decentralisation

In this paper we address decentralisation in distributed consensus. Decentralisation is a key concern in *permissionless* blockchains, where participation is allowed in a generally unrestricted way. Permissionless blockchains are clearly exposed to the presence of Byzantine peers, i.e., dishonest peers trying to exploit the network and not bound to the blockchain protocol. Byzantine distributed consensus is a long-standing problem, from Lamport’s characterisation [25] and the FLP impossibility result [12], to the subsequent research on data replication and consistency based on Byzantine Fault-tolerant consensus (BFT), [27, 36]. Several proposals are currently competing in a multi-billion market, addressing the so-called *blockchain trilemma*, i.e., achieving *security*, *scalability* and *decentralisation* together.

One of the breakthroughs of Bitcoin was the introduction of the Proof-of-Work (PoW) [9] as a mechanism to enable a probabilistic Byzantine distributed consensus. Informally speaking, by solving a computationally hard problem one of the peers is entitled to create the next block, cryptographically linked to the previous ones. Under the assumption that Byzantine computational power is suitably limited within the network, the probability that enough work can be channeled to alter block history decreases with the ageing of the blocks [15]. Bitcoin reaches finality with an acceptable probability in about one hour (6 blocks), with limited transactions per second.²

In *Proof-Of-Stake* (PoS) blockchains peers contribute to the definition of the next block with a probability proportional to the stake (coins), rather than computational power, they detain in the system. Safety follows from the honest peers detaining the majority of stake. Scalability improves in Proof-of-Stake, but the management of security typically results in being more complex.

The *BFT paradigm* has also been proposed for blockchain consensus, providing scalability in transaction per second thanks to low transaction latency and high throughput. BFT, however, is more constrained in terms of the scalability in the number of peers [40], since the number and identity of peers needs to be known and in some cases fixed [5]. This kind of blockchain has been proposed, for instance, for financial services, where a limited number of known and certified peers need to exchange fast and numerous transactions. It is worth remarking that if consensus requires control on peers, a centralised authority might be required, with implications also on identity, privacy and censorship.

Research question

In this paper we focus on BFT blockchains based on quorum systems [39]. In such systems consensus emerges from neighbourhoods of peers. The properties of such neighborhoods, together with assumptions on Byzantine failure thresholds, are then essential to guarantee the liveness and safety of the consensus protocol, that is, whether honest peers are able to eventually reach consensus on a correct next state. At the heart of these protocols is a notion of trust between peers: nodes select which other nodes to trust, and listen to, in the network. The paper focuses on the following research question:

To what extent can consensus be decentralized, when based on peer-to-peer trust relations?

² More technical and comprehensive introductions to blockchains can be found in [2, 30, 41].

We take an analytical approach and establish theorems that point to the existence of inherent tradeoffs between safety and decentralization for approaches based on such trust networks. We use tools from game theory (specifically the theory of command games [20, 19]) and computational complexity theory. The interface of methods from theoretical economics and computational complexity have proven extremely prolific in other areas of computer science and artificial intelligence, such as computational social choice theory [3, 18]. Our paper showcases these methods for general investigations on blockchain consensus and decentralisation.

We will specifically consider Ripple [6, 37] and Stellar [28], two quorum-based blockchains attempting to extend the applicability of the BFT paradigm from a permissioned to a permissionless setting, aiming at improving decentralisation. Ripple provides frictionless global payments and corporate-oriented efficient transactions. It currently relies on a list of “authorised” validators³ in charge of the correctness of transactions. Access is permissioned and each peer will need to have in their neighbourhood of trust a number of validators from the list. While the list was originally entirely composed by Ripple validators, today third-party validators, e.g., private companies and universities, have been included. Stellar provides payments and asset management to corporate and individuals, and aims to push decentralisation further by offering open membership and allowing peers to autonomously define their trust networks, i.e., the set of validators that they trust. However, strong constraints hold on the topology of such trust networks.

Both Ripple and Stellar have been object of criticism with respect to the level of decentralisation of their current implementations, and the need for further research on protocols like Ripple and Stellar is emphasised, for instance, in [4].

Related work

Even though at the time of writing Ripple and Stellar are, respectively, the third and twelfth blockchain systems in terms of market capitalization,⁴ little foundational work exists on their protocols. Correctness analyses of Ripple have been proposed in [6], and of Stellar in [28, 17]. A specific study on the issue of decentralisation in Stellar has also very recently been presented in [23]. The paper investigates the network of Stellar’s peer-to-peer trust relations by means of an extended version of PageRank used to evaluate nodes’ influence. Findings about the current status show centralisation on two critical validators, which are controlled by the Stellar Foundation. Our paper contributes further general results on the level of decentralization that could reasonably be achieved in systems like Ripple and Stellar.

Paper contribution and outline

The paper makes three contributions. *First* it develops a novel theoretical framework, rooted in economic theory (command games [20, 19], power indices [33, 1]), to ascertain the influence that peers can exert on each other in quorum systems based on trust networks. This contributes a novel methodology for a much needed quantitative evaluation of decentralisation in blockchain with respect to its consensus layer. The proposed methods are applied to Ripple and Stellar (Theorem 25). *Second*, it establishes an “impossibility of decentralisation” result for a class of consensus protocols of the Ripple type (Theorem 16), which are based on

³ At the time of writing the list consists of about 30 validators, available at <https://xrpscharts.ripple.com>

⁴ Source: <https://coinmarketcap.com/all/views/all/>. Retrieved on 11th May 2020.

trust networks with a fixed threshold of tolerable Byzantine peers. This results shows that, in such systems, a necessary condition for safety is the existence of validators that must be trusted by every peer in the network, hindering the possibility of full decentralisation. *Third*, it develops an appraisal of computational barriers to decentralization in protocols like Stellar, that are based on so-called federated Byzantine agreement systems. Specifically, we show that constraints that are necessary to guarantee the safety of the network require peers to be able to solve problems that are computationally intractable in principle (Theorems 21 and 22). This result identifies computational difficulties for the construction of safe peer-to-peer trust networks, thereby pointing to computational barriers for the full decentralisation of Stellar.

A Byzantine model of trust network is introduced in Section 2, impossibility and intractability results are presented in Section 3, and decentralisation measures in Section 4. Section 5 concludes. All proofs are provided in a technical appendix.

2 Preliminaries

The paper is concerned with systems based on the following high-level blueprint: nodes hold opinions on a value (e.g, whether a transaction should be recorded or not); they validate an opinion when they observe enough nodes, among those they trust, that hold the same opinion (agreement); some nodes may be Byzantine; and no two honest nodes should be able to validate opinions with different values (safety). Crucially, nodes are able to autonomously decide which other nodes to trust. This section defines the above set-up formally.

2.1 Opinions and Opinion Profiles

Let N be the set of nodes, with $n = |N|$, and let $H \subseteq N$ the set of honest nodes and $B = N \setminus H$ the set of Byzantine nodes. The opinion of a honest node $i \in N$ is a value in $\{0, 1\}$. At any given time, the collection of each node's opinions defines an opinion profile that associates a “genuine” opinion from $\{0, 1\}$ to every honest node (the opinion that the node reveals to the network). And to each Byzantine node in $N \setminus H$ it associates a function from honest nodes to opinions. This function represents the value that each Byzantine node would reveal to each honest node.

► **Definition 1.** An opinion profile $\mathbf{o} : N \rightarrow \{0, 1\} \cup \{0, 1\}^H$ such that $\mathbf{o}(i) \in \{0, 1\}$ if $i \in H$ and $\mathbf{o}(i) \in \{0, 1\}^H$ if $i \in B$. Given $x \in \{0, 1\}$, \bar{x} represents the element of singleton $\{0, 1\} \setminus \{x\}$.

The definition makes some simplifications, which are worth flagging. Even from the perspective of binary valued opinions, it would be more realistic to work with a ternary set of opinions containing 1, 0 plus an undefined value representing undecided nodes. Also, the definition rules out the possibility for a Byzantine node to reveal inconsistent values (that is, 1 and 0 to a same node). However, such simplifications are not fundamental, and the results we present would carry over to these more general settings.

We define then what it means for a node to observe an agreement among other nodes, in a given opinion profile.

► **Definition 2.** Let an opinion profile \mathbf{o} be given. A set of nodes $C \subseteq N$ is said to agree (in \mathbf{o}) from the perspective of $i \in N$ if for all $j, k \in H \cap C$ and $j', k' \in B \cap C$, $\mathbf{o}(j) = \mathbf{o}(k) = \mathbf{o}(j')(i) = \mathbf{o}(k')(i)$.

That is, a node i observes agreement in a set of nodes C whenever the honest nodes in that set hold the same opinion, and that opinion is also the opinion revealed to i by the Byzantine nodes in C .

2.2 Byzantine Trust Networks

We now shift to the definition of the structures of peer-to-peer trust relations underpinning consensus in the systems we are focusing on.

► **Definition 3.** A Byzantine trust network (BTN) is a tuple $\mathcal{T} = \langle N, H, T_i, \mathcal{C}_i \rangle$ where:

- $N = \{1, \dots, n\}$ is a finite set of nodes.
- $H \subseteq N$ is the set of honest nodes. $B = N \setminus H$ is the set of Byzantine nodes. We denote with $b_i = \frac{|B \cap T_i|}{|T_i|}$ the ratio of Byzantine nodes in T_i , and with $b = \max \{b_i\}_{i \in H}$ the largest such ratio.
- $T_i \subseteq N$, for each node $i \in H$, is the non-empty set of nodes that i trusts, i.e., its trust set.⁵
- $\mathcal{C}_i \subset 2^{T_i}$, with $i \in H$. We refer to \mathcal{C}_i as the set of winning coalitions for node i . For ease of presentation we will refer to winning coalitions also via a function $\mathcal{C} : N \rightarrow 2^{2^N}$ assigning the set \mathcal{C}_i of subsets of N to each $i \in H$.

We will sometimes assume that, for all i , $i \in T_i$. We will also sometimes assume that for all $i \in H$, $\{i\} \in \mathcal{C}_i$. In such a case the BTN is said to be vetoed.

Intuitively, a winning coalition $C \in \mathcal{C}_i$ for i is a set of nodes such that, if all its members agree from the perspective of i , then their opinion is validated by i (cf. [32]). We define the notion of validation formally below (Definition 8). When $\{i\}$ belongs to \mathcal{C}_i , i cannot validate an opinion unless it also holds such opinion (it holds a veto for its own validation).

► **Remark 4.** In the Stellar white paper [28] BTN are referred to as *federated Byzantine agreement systems* (FBAS), or as *federated Byzantine quorum systems* in [17], and the winning coalitions of a node are referred to as *quorum slices*. BTNs are also known structures in the economic theory literature, where they are referred to as *command games* [20, 19], or as *simple game structures* [22].

A natural class of BTNs is obtained by associating a quota, or threshold, $q_i \in (0.5, 1]$ to each honest node i .⁶

► **Definition 5.** A Quota Byzantine Trust Network (QBTN) is a BTN such that for all $i \in H$ there exists a quota $q_i \in (0.5, 1]$ such that:

$$\mathcal{C}_i = \{C \subseteq T_i \mid |C| \geq q_i \cdot |T_i|\}.$$

A QBTN is therefore denoted by a tuple $\mathcal{T} = \langle N, H, T_i, q_i \rangle$. A QBTN is said to be uniform whenever $q_i = q_j$ for any $i, j \in H$. It is said to be effective whenever $q_i \in (0.5 + b, 1 - b]$ for any $i \in H$.

Intuitively, QBTNs are BTNs where the winning coalitions of a node are determined by a numerical quota: i 's validation is determined whenever at least a fraction q_i of nodes in T_i hold that opinion. This quota should be set in such a way that: *i*) the quota is met whenever the honest nodes in T_i agree, and *ii*) if the quota is met for an opinion x , then there is at least

⁵ Ripple and Stellar refer to trust sets as unique node lists (UNLs).

⁶ Cf. [16].

an honest majority of nodes with opinion x in the trust set of i . That is, $q_i \in (0.5 + b, 1 - b]$, which is the case for what we called effective QBTNs. For this constraint to be met the fraction of Byzantine nodes in each T_i should therefore fall in the interval $[0, 0.25)$.

► **Remark 6.** It is worth observing that in this context the largest possible ration of Byzantine nodes b in a BTN plays a slightly different role than in the standard BFT failure models where, classically, safety would require $b < \frac{1}{3}|T_i|$ [31, 8, 25]. In the standard BFT failure models one is interested in making sure that any two simple majorities intersect in some honest node in order for the system not to fork. In a BTN, the role of b is slightly different: it is to guarantee that an honest node can always safely validate an opinion because a honest majority exists in its trust set, that agrees on that opinion.

The Ripple consensus protocol [6, 37] is based on uniform QBTNs with quotas $q_i = 0.8$ and $b = 0.2$. The Stellar consensus protocol as described in [28] does not rely on QBTNs but requires the more general class of vetoed BTNs.

► **Remark 7.** In Definition 3 we associate winning coalitions only to honest nodes. We do this for simplicity but it should be clear that trivial collections of winning coalitions can be associate also to Byzantine nodes. Since the validation by a Byzantine node i is not influenced by any other node its trivial collection of winning coalitions is the set $\{C \subseteq N \mid \{i\} \subseteq C\}$, that is, the set of all coalitions containing i . Intuitively, this amounts to stating that i is the only node influencing its own opinion.

2.3 Validation and Safety

Intuitively, a node i validates an opinion whenever a winning coalition of nodes trusted by i agrees on that opinion, from the perspective of i . Given an opinion profile \mathbf{o} we denote by

$$T_j^{\mathbf{o}}(x) = \{i \in T_j \mid \mathbf{o}(i) = x \text{ if } i \in H, \text{ and } \mathbf{o}(i)(j) = x \text{ if } i \in B\} \quad (1)$$

the set of nodes trusted by i who hold opinion x (if they are honest), or reveal opinion x to i (if they are Byzantine).

► **Definition 8.** Let a BTN and an opinion profile \mathbf{o} be given. We say that $i \in N$ validates $x \in \{0, 1\}$ (in \mathbf{o}) if $T_i^{\mathbf{o}}(x) \in \mathcal{C}_i$.

In a QBTN, an honest node i validates an opinion whenever there are at least $q_i \cdot |T_i|$ nodes with the same opinion among the nodes it trusts.

► **Remark 9.** BTNs are a generalization of so-called Byzantine quorum systems [27]. Each BTN naturally induces the set of quora $\{Q \subseteq N \mid \forall i \in Q, \exists C \in \mathcal{C}_i : C \subseteq Q\}$. In words, a *quorum* is a set of nodes that contains a winning coalition for each node in the set (cf. [28]). Intuitively, it is a set of nodes that have the means to stably validate an opinion. For a set of quora to be a Byzantine quorum system, quora also need to pairwise intersect. We will come back to this in Section 3.2.

To introduce safety formally, we need some auxiliary notions. Let $s : N \rightarrow 2^N$ be a function picking, for any agent i , one coalition out of \mathcal{C}_i . Each function s induces an operator $F_s : 2^N \rightarrow 2^N$ such that $F_s(C) = \bigcup_{i \in C} s(i)$, collecting, for each i in C , the winning coalition $s(i)$ picked by function s . We can then recursively construct sets of nodes by joining winning coalitions of nodes in earlier sets. Such a construction reaches a fixpoint where for each node in the set a winning coalition is already contained in the set. Formally, for $C \subseteq N$: $F_s^0(C) = C$ and $F_s^{n+1}(C) = F_s(\bigcup_{0 \leq k \leq n} F_s^k(C))$. Define then $F_s^*(C) = \bigcup_{0 \leq n} F_s^n(C)$. As N

is finite, there exists a non-negative integer n such that $F_s^n(C) = F_s^*(C)$. Observe that for any $C \subseteq N$, $F_s^*(C)$ is such that for all $i \in F_s^*(C)$ there is a $D \in \mathcal{C}_i$ such that $D \subseteq F_s^*(C)$. That is, $F_s^*(C)$ is a quorum (cf. Remark 9).

► **Definition 10.** Let a BTN and an opinion profile \mathbf{o} be given. We then say that an opinion profile \mathbf{o} is:

- forked (or, is a fork) if there are two honest nodes $i, j \in H$ such that i validates x and j validates \bar{x} in \mathbf{o} .
- strongly forked (or, is a strong fork) if there are two honest nodes $i, j \in H$ and a function s such that all nodes in $F_s^*(\{i\})$ agree on x and all nodes in $F_s^*(\{j\})$ agree on \bar{x} .

So, a fork is a profile where two honest nodes i and j have validated two different opinions. A strong fork is a fork where, in addition, there is a winning coalition for i agreeing on x and a winning coalition for j agreeing on \bar{x} , and all nodes in that winning coalition for i also have a winning coalition agreeing on x and all nodes in that winning coalition for j have a winning coalition agreeing on \bar{x} , and so on. In short, a strong fork is a “stable” fork.

► **Definition 11.** A BTN is safe if there exists no forked profile for it. It is weakly safe if there exists no strongly forked profile for it.

Safety rules out the possibility that two honest nodes may settle on different opinions, and therefore the possibility that the run of any consensus protocol on the BTN would generate a stream of opinion profiles that contains a profile where two nodes have validated different values. Weak safety allows for forks of only a limited kind. It rules out the possibility that forks are of a “deep” kind in the sense that they involve all winning coalitions upon which the diverging opinions are rooted. Clearly safety implies weak safety but not vice versa.

It is finally worth stressing that the above notions of safety and weak safety are protocol independent: they concern all consensus protocols where validation depends locally on the values relayed by trusted nodes on a given BTN, like the Ripple or Stellar protocols.⁷

3 (De)centralisation and (In)tractability

This section explores inherent limitations present in the above notion of safety for BTNs, establishing general limitative results for the class of consensus protocols based on them, such as Ripple and Stellar. First, it focuses on uniform QBTNs (Definition 5), as exemplified by the Ripple consensus protocol, showing that safety drastically limits the freedom of nodes in selecting trustees. Second, it focuses on safety for general BTNs (Definition 3), as exemplified by the Stellar consensus protocol, showing that, even though safety in such settings allows for more freedom on the part of nodes, it does require single nodes to solve decision problems that are, in principle, computationally intractable.

3.1 Safety Implies Centralization in Uniform QBTNs

We show that the safety of uniform QBTNs implies that nodes cannot be fully free to choose their trust set. The result builds on ideas and techniques from [6].

We first define some notation:

$$\beta_{ij} = \min \{|T_i \cap T_j|, b \cdot |T_i|, b \cdot |T_j|\}. \quad (2)$$

⁷ However, it is worth stressing we are not considering protocols where validation may depend on information richer than just nodes’ opinions.

Intuitively, β_{ij} denotes the maximum possible number of Byzantine nodes present in the intersection of the trust sets of i and j . Such a number equals the maximum amount of Byzantine nodes assumed by the node, either i or j , that tolerates fewer Byzantine nodes. Such a number cannot obviously exceed the size of the intersection itself.

► **Lemma 12** ([6]). *Let $\langle N, H, T_i, q_i \rangle$ be an effective QBTN. For any profile \mathbf{o} and node $i \in H$, if $|T_i^{\mathbf{o}}(x)| > 0$ then for any $j \in H$,*

$$|T_j^{\mathbf{o}}(x) \cap H| \geq |T_i \cap T_j| + |T_i^{\mathbf{o}}(x)| - |T_i| - \beta_{ij} \quad (3)$$

$$|T_j^{\mathbf{o}}(\bar{x})| \leq |T_j| - |T_i \cap T_j| - |T_i^{\mathbf{o}}(x)| + |T_i| + \beta_{ij} \quad (4)$$

This lemma establishes a lower bound on the number of honest nodes with opinion x that a honest node j can observe in its trust set, given the number of nodes (not necessarily honest) that another honest node i observes. It is used in the proof of Lemma 14. Notice that the bound in (3) and (4) are not necessarily strict, as illustrated in the following example.

► **Example 13.** Let $\langle N, H, T_i, q_i \rangle$ be such that: $N = \{1, \dots, 9\}$, $B = \{5\}$ (recall $N = H \cup B$), $T_1 = T_2 = T_3 = T_4 = \{1, 2, 3, 4, 5\}$ and $T_6 = T_7 = T_8 = T_9 = \{5, 6, 7, 8, 9\}$, $q_1 = q_2 = q_4 = q_5 = 1 - \frac{1}{9}$. Let then \mathbf{o} be such that $\mathbf{o}(1) = \mathbf{o}(2) = \mathbf{o}(3) = \mathbf{o}(4) = 1$, $\mathbf{o}(6) = \mathbf{o}(7) = \mathbf{o}(8) = \mathbf{o}(9) = 0$, finally $\mathbf{o}(5)$ be such that $\mathbf{o}(5)(1) = \mathbf{o}(5)(2) = \mathbf{o}(5)(3) = \mathbf{o}(5)(4) = 0$ and $\mathbf{o}(5)(6) = \mathbf{o}(5)(7) = \mathbf{o}(5)(8) = \mathbf{o}(5)(9) = 1$. So no honest node in $\{1, 2, 3, 4, 5\}$ can see a honest node with opinion 0 and, vice versa, no honest node in $\{5, 6, 7, 8, 9\}$ can see a honest node with opinion 1. But honest nodes in both set can see a (Byzantine) node with opposite opinion. Let now $j = 1$ and $i = 2$. We have, $|T_j^1 \cap H| = 5 + 5 - 5 - 1 = 4$ and $|T_j^0| = 5 - 5 - 5 + 5 + 1 = 1$.

► **Lemma 14.** *Let $\mathcal{T} = \langle N, H, T_i, q_i \rangle$ be a safe, uniform and effective QBTN. Then for all $i, j \in H$: $|T_i \cap T_j| > \frac{b}{1-b} \cdot (|T_i| + |T_j|)$.*

Notice that the maximum size of the intersection of two trust sets is obviously $\frac{1}{2}(|T_i| + |T_j|)$.⁸ Lemma 14 establishes a lower bound on the size of the intersection of trust sets required by safety. The intuition behind the lemma is the following one. In order to make it impossible for two honest nodes to validate opposite values, their trust sets should intersect to the extent that any two winning coalitions for the two nodes would also have to intersect *and* contain at least one honest node.

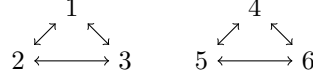
► **Lemma 15.** *Let $\langle N, H, T_i, q_i \rangle$ be a uniform BTN. If for all $i, j \in H$, $|T_i \cap T_j| > 0.25 \cdot (|T_i| + |T_j|)$, then $\bigcap_{i \in H} T_i \neq \emptyset$.*

Intuitively, if any two trust sets have a large enough intersection (larger than a quarter of their combined size or, equivalently, larger than half their average size) then they must all intersect.

► **Theorem 16.** *In uniform QBTNs with effective quotas, if $b \geq 0.2$ then safety implies the existence of nodes that are trusted by all honest nodes.*

Proof. The result follows directly from Lemmas 14 and 15 and the observations that: in effective QBTNs where $b \geq 0.2$ it follows that $b \leq 0.25$; and for $b \in [0.2, 0.25]$ we have that $\frac{b}{1-b} \geq 0.25$ as desired. ◀

⁸ This is the case, for instance, in QBTNs where $|T_i| = |T_j|$ for all $i, j \in H$.



■ **Figure 1** Example from [28] of a vetoed BTN lacking QI. Arrows denote which nodes each node trusts (reflexive arrows omitted). $\mathcal{C}(1) = \mathcal{C}(2) = \mathcal{C}(3) = \{\{1, 2, 3\}\}$ and $\mathcal{C}(4) = \mathcal{C}(5) = \mathcal{C}(6) = \{\{4, 5, 6\}\}$.

If we understand decentralisation as the property of trust networks in which nodes have full freedom on whom to trust in the network, then the theorem can be interpreted as a general impossibility result for consensus based on QBTNs: if quotas are uniform and set appropriately w.r.t. the postulated maximum share of Byzantine nodes in trust sets, and if such share is large enough, then the existence of nodes that are trusted by everyone is a necessary condition for the safety of consensus. Furthermore, beyond limiting the choice of nodes, a (limited) set of such nodes clearly represents a dominant position and a risk factor.

In general, Theorem 16 is relevant for any consensus protocol that could be run on uniform QBTNs. In particular, it applies to the Ripple consensus protocol when the maximum fraction b of Byzantine nodes in trust sets is set to 0.2. In a way, Theorem 16 provides an ex-post analytical justification to the current design of the Ripple trust network where all trust sets are required to include a same subset of nodes (cf. [6]). Currently Ripple relies on a single UNL mostly controlled by Ripple, although plans for further decentralisation are under discussion.⁹

3.2 Safety and Quorum Intersection in BTNs

In this and the next section we consider the general case of (vetoed) BTNs, to which Theorem 16 does not apply. This more general setting applies instead to Stellar as described in its white paper [28], where the Stellar consensus protocol does not presuppose uniformity of quotas. Actually, Stellar aims to offer open membership and freedom in choosing its own trust networks, which, together with BFT good scalability, would yield a decentralised and efficient blockchain. In such a setting an intuitive necessary condition for safety is that trust networks are sufficiently “interconnected”, in the following sense.

► **Definition 17** ([28]). *A vetoed BTN enjoys quorum intersection (QI) whenever for any two sets $Q_1, Q_2 \subseteq H$, if Q_1 and Q_2 are quora, then $Q_1 \cap Q_2 \neq \emptyset$.*

► **Example 18.** In Figure 1 the quora are $\{1, 2, 3\}$, $\{4, 5, 6\}$, $\{1, 2, 3, 4, 5, 6\}$. This BTN does not enjoy QI, but both of its disjoint components (with support $\{1, 2, 3\}$ and $\{4, 5, 6\}$) do. Suppose instead that $\{1, 2, 3, 5\} \in \mathcal{C}(3)$, that is, 3 also trusts 5. Then the system would satisfy QI with quora: $\{4, 5, 6\}$, $\{1, 2, 3, 4, 5, 6\}$.

► **Example 19.** In a BTN \mathcal{T} where $\forall i \in N, \mathcal{C}(i) = \{N\}$, the set N of all nodes is the unique quorum, and \mathcal{T} trivially enjoys quorum intersection.

In fact, there is a close relationship between quorum intersection and the property of weak safety:

► **Theorem 20.** *A vetoed BTN is weakly safe iff any two quora intersect and such intersection contains at least one honest node.*

⁹ Cf. <https://xrpcharts.ripple.com/>.

Clearly, nodes in a BTN cannot know which nodes are Byzantine so their best effort in order to guarantee weak safety is to guarantee QI is not violated.¹⁰

3.3 The Intractability of Maintaining Quorum Intersection

Quorum intersection is in fact assumed by all the existing correctness analyses of Stellar [28, 17]. It is furthermore stressed in [28, p. 9] that: “[...] it is the responsibility of each node i to ensure \mathcal{C}_i [notation adapted] does not violate quorum intersection.” The key question, from a safety perspective, becomes therefore whether single nodes can reasonably be tasked with maintaining QI. Apart from incentive issues, which have also been flagged [23], we argue that this is a problematic requirement from a merely computational standpoint. This might not be an issue in the current, small-scale, Stellar configuration (although an instance of QI failure has been recently reported [26]), but it is something to be considered in a path towards full decentralisation with a full-scale number of nodes and validators. As our analysis below shows, maintaining QI is a computationally intractable problem.

We present two results. First we show that deciding whether a given BTN satisfies QI is intractable.¹¹ Second, we show that deciding whether adding a new trust set with winning coalitions preserves QI on a given BTN is also computationally intractable (again coNP-hard). This is arguably the decision problem that nodes need to solve when linking to the Stellar network. We argue that these results point to a possible computational bottleneck for the scalability of the consensus model of Stellar.

We first define the problem of deciding whether QI holds in a given BTN.

QUORUM-INTERSECTION

Input: A BTN $\mathcal{T} = \langle N, \mathcal{C} \rangle$ where the sets $\mathcal{C}(i)$ for $i \in N$ are listed explicitly.¹²

Question: Is it the case that for each two quora Q_1, Q_2 , $Q_1 \cap Q_2 \neq \emptyset$?

► **Theorem 21.** *QUORUM-INTERSECTION is coNP-complete.*

The intractability result of Theorem 21 says that it may be computationally hard, in practice, to check QI. Such a result is robust in the sense that the related problem of checking whether QI holds after the insertion of one new winning coalition by a node into a system that already satisfies QI, is also coNP-complete. (actually coNP-hard).

COALITION-ADDITION-QUORUM-INTERSECTION

Input: Two BTNs $\mathcal{T} = \langle N, \mathcal{C} \rangle$ and $\mathcal{T}' = \langle N, \mathcal{C}' \rangle$, such that \mathcal{T}' satisfies QI, and such that \mathcal{T} is obtained from \mathcal{T}' by adding one single coalition to $\mathcal{C}'(i)$ for some $i \in N$, where the sets $\mathcal{C}(i)$ and $\mathcal{C}'(i)$ for all $i \in N$ are listed explicitly.

Question: Is it the case that for each two quora Q_1, Q_2 of \mathcal{T} $Q_1 \cap Q_2 \neq \emptyset$?

► **Theorem 22.** *COALITION-ADDITION-QUORUM-INTERSECTION is coNP-complete.*

¹⁰ However, a BTN could be complemented by a more fine-grained failure model consisting of a set of sets of possible Byzantine nodes representing the possible failure scenarios that nodes may encounter (cf. [17]).

¹¹ An equivalent result has also been recently presented in [24]. That paper provides a proof of NP-completeness (via reduction from the Set Splitting Problem) of the complementary problem for which we prove coNP-completeness (via 3SAT).

¹² For the purpose of this and the following result we do not need to take into consideration the H and T_i elements of a BTN (Definition 3). We therefore omit them for conciseness.

4 Quantifying Influence on Consensus in BTNs

Theorem 16 showed that, in uniform QBTNs, safety implies the existence of nodes that are trusted by all honest nodes. While this can definitely be interpreted as a high level of centralisation required by safety, it is worth trying to precisely quantify the effect that the existence of all-trusted nodes has on consensus. In PoW and PoS protocols it is straightforward, at least by first approximation, to understand what the influence of each node is on the consensus process: each node will be able to determine a fraction of blocks corresponding to the node's share of total hashing power (PoW) or of total stakes (PoS). For consensus based on voting over trust structures, like in Ripple and Stellar, quantifying nodes' influence in a principled way is not obvious. This section proposes a methodology for such quantification that leverages the theory of voting games.

4.1 Influence Matrices

A BTN (Definition 3) associates to each honest node i a structure $\langle T_i, \mathcal{C}_i \rangle$. Such structures are known in game theory as simple games, that is a set of agents endowed with a set of winning coalitions. Such structures have been extensively studied to provide exact quantifications of power, for instance, in voting. In this section we show how techniques based on simple games provide a principled way to quantify influence in BTNs. We study the influence of j over i in a BTN as the power of j in the simple game that the BTN associates to i .

The *Penrose-Banzhaf index* [33, 1] of a node j in the simple game $\langle T_i, \mathcal{C}_i \rangle$ of node i is

$$B_i(j) = \frac{1}{2^{n-1}} \sum_{C \subseteq N \setminus \{j\}} v(C \cup \{j\}) - v(C) \quad (5)$$

where v is the characteristic function of \mathcal{C}_i .¹³ Essentially, the index counts the number of times in which j is decisive in turning a losing coalition into a winning one, that is, one that can determine the validation of an opinion by i . Equivalently, $B_i(j)$ can be interpreted as the probability that j determines whether i validates a specific opinion, assuming all other agents in T_i have opinions distributed uniformly at random.

The normalised version $NB_i(j)$ of the Penrose-Banzhaf index is:

$$NB_i(j) = \frac{B_i(j)}{\sum_{k \in N} B_i(k)} \quad (6)$$

For any agent $j \notin T_i$ we stipulate $NB(j) = 0$, as nodes that i does not trust cannot influence i 's opinion directly – in game-theoretic jargon they are “dummy agents” in i 's simple game. Byzantine nodes are assigned degenerate simple games containing a singleton winning coalition that has themselves as only member (cf. Remark 7 above). Byzantine agents cannot be influenced: for all $i \in N \setminus H$ from the degenerate simple game associated to i , $NB_i(i) = 1$, and $NB_i(j) = 0$ for each $j \neq i$.

Given a BTN, we associate to each honest node i a vector $[NB_i(1), \dots, NB_i(n)]$ of normalized Penrose-Banzhaf indices capturing the influence that each node has on i . Clearly $\sum_{j \in N} NB_i(j) = 1$ and $NB_i(j) > 0$ only if $j \in T_i$, for any $i \in N$. Notice that the vector of a Byzantine node i is therefore degenerate: $NB_i(i) = 1$ and $NB_i(j) = 0$ for each $j \neq i$. It

¹³That is, for any $C \subseteq N$, $v(C) = 1$ iff $C \in \mathcal{C}_i$.

follows that each BTN \mathcal{T} induces a stochastic $N \times N$ matrix

$$I(\mathcal{T}) = \begin{bmatrix} I_{11} & I_{12} & I_{13} & \dots & I_{1n} \\ I_{21} & I_{22} & I_{23} & \dots & I_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ I_{n1} & I_{n2} & I_{n3} & \dots & I_{nn} \end{bmatrix}$$

where I_{ij} denotes the normalized Penrose-Banzhaf index $\text{NB}_i(j)$ of node j in the simple game associated to i . We call such matrix $I(\mathcal{T}) = [I_{ij}]_{i,j \in N}$ the *influence matrix* (of \mathcal{T}). We will drop reference to \mathcal{T} when no confusion arises. The matrix encodes the direct influence that each node has on each other in the sense of being decisive for the validation of an opinion. Matrices of this type have a long history in the mathematical modeling of influence in economics and the social sciences dating back to [13, 7], and have recently received renewed attention [21].¹⁴ Similar matrices, but based on the Shapley-Shubik power index [38] instead of the Penrose-Banzhaf one, have been studied in [20, 19].¹⁵

► **Example 23.** Consider the following BTN with no Byzantine nodes and consisting of 6 agents all having a same set of 5 agents as trust set: $N = H = \{1, 2, 3, 4, 5, 6\}$, $T_i = \{1, 2, 3, 4, 5\}$ for all $i \in N$ and $q_i = 0.8$ for all $i \in N$. By (5) and (6) for each $i \in \{1, 2, 3, 4, 5, 6\}$ we have $B_j(i) = \frac{2}{8}$ and $\text{NB}_j(i) = \frac{1}{5}$, for each node $j \in \{1, 2, 3, 4, 5\}$. The influence matrix describing this BTN consists of 6 identical row vectors $[\frac{1}{5} \ \frac{1}{5} \ \frac{1}{5} \ \frac{1}{5} \ \frac{1}{5} \ 0.]$. Consider now a variant of the above BTN where node 5 is Byzantine. The influence matrix describing this variant consists of 5 identical vectors $[\frac{1}{5} \ \frac{1}{5} \ \frac{1}{5} \ \frac{1}{5} \ \frac{1}{5} \ 0.]$ for the rows corresponding to nodes 1 – 4 and 6, and the degenerate row vector $[0 \ 0 \ 0 \ 0 \ 1 \ 0.]$ for the row of node 5. That is, all the nodes in T_i have the same influence on honest nodes, but no honest node influences 5.

4.2 Limit Influence

Influence matrices describe the extent of direct influence between nodes. But influence is not only direct: by influencing nodes k that in turn influence node i , a node j can exert indirect influence on the opinions that i validates. This type of indirect influence is captured by the powers of the influence matrix:

$I_{ij}^1 = I_{ij}$ represents the probability that j can directly sway i to validate a value x . This is j 's direct influence on i .

$I_{ij}^2 = \sum_{k \in N} I_{ik} \cdot I_{kj}$ represents the probability that j can sway i 's validation in two steps, by swaying the validation of intermediate nodes k which in turn sway i 's validation directly.

This is j 's indirect (2-step) influence on i .

I_{ij}^k more generally represents j 's indirect (k -step) influence on i .

So the influence (direct or indirect) of j on i in a BTN is given by the total probability of all ways in which j can determine the value of i 's validation. Formally this amounts to $\lim_{t \rightarrow \infty} (I^t)_{ij}$, provided such limit exists. In yet other words, this denotes the likelihood that j is decisive for i to validate an opinion.

We are then in the position to quantify what the influence is of each node on every other node by taking the limit of the power of the influence matrix of the BTN \mathcal{T} , that is:

$$I(\mathcal{T})^\infty = \lim_{t \rightarrow \infty} I(\mathcal{T})^t \tag{7}$$

¹⁴ See also [34, 35] for an overview of such models.

¹⁵ For a comparison between these two power indices we refer the reader to [11].

If the limit matrix in (7) exists, we say that the influence matrix $I(\mathcal{T})$ is *regular*. We say that it is *fully regular* when its limit matrix exists and it is such that all rows are identical.¹⁶ Intuitively, regularity means that it is possible to precisely quantify the influence of each node on each other node; full regularity means that every node has the same influence on every other node.

► **Example 24.** Consider again the two BTNs introduced in Example 23. In the first case, where all nodes are honest, all nodes belonging to some trust set have positive and – given the symmetry built in the example – the same influence:

$$\begin{bmatrix} \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix}^2 = \lim_{t \rightarrow \infty} \begin{bmatrix} \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix}^t$$

In the second case, where node 5 is Byzantine, the only node having positive influence (total influence 1 in this example) is precisely 5:

$$\lim_{t \rightarrow \infty} \begin{bmatrix} \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix}^t = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

In other words, the only node having influence on which values will be validated by other nodes, and therefore on agreement, is the Byzantine node.

4.3 Limit Influence in Ripple and Stellar

Theorem 16 established that in uniform QBTNs, and therefore Ripple, safety requires centralisation in the sense of demanding the existence of a non-empty set of nodes trusted by all other nodes. While this does not apply in general to Stellar, recent studies have highlighted that Stellar enjoys a similar level of centralisation.¹⁷

Here we put the above methodology at work to study limit influence in centralised BTNs, that is BTNs where nodes exist that are trusted by all nodes. We show (Theorem 25) that: the existence of nodes trusted by all nodes makes it possible to establish limit influence (first claim); this limit influence is such that every node has the same limit influence on every other node (second claim) when at most one Byzantine node exists in the BTN; but if only just one all-trusted node trusts a Byzantine node, no honest node has limit influence on any other honest node (third claim). That is, in a centralised BTN the power of deciding whether an opinion can become consensus or not is, in principle, all in the hands of Byzantine nodes.

¹⁶The “regularity” and “full regularity” terminology are borrowed from [14] and [34].

¹⁷Data analysis of the current Stellar network has shown [23] that one of the three Stellar foundations validators is included in all trust sets. If we treat the Stellar foundation to be operating as one node, Stellar satisfies *de facto* the same level of centralisation that we have shown is analytically required for Ripple.

► **Theorem 25.** *Let \mathcal{T} be a BTN. If \mathcal{T} is such that $\bigcap_{i \in H} T_i \neq \emptyset$ then:*

- 1) $I(\mathcal{T})$ is regular;
- 2) $I(\mathcal{T})$ is fully regular if in addition \mathcal{T} is such that $|B| \leq 1$;
- 3) and, if there exists $j \in (\bigcap_{i \in H} T_i) \cap H$ such that $T_j \cap B \neq \emptyset$ then for all $j, k \in H$, $I(\mathcal{T})_{jk}^\infty = 0$.

Again, it is worth noticing that this is a general protocol-independent result: it concerns all protocols working on centralized BTNs where consensus is determined through validations locally dependent on trusted nodes. In particular, it applies to the setup of the Ripple trust network under the assumption of safety (by Theorem 16) and to the current setup of the Stellar trust network.

5 Conclusions

We addressed decentralisation in the specific context of BFT consensus based on open quorum systems, showcasing the relevance of tools from economic theory (command games, power indices) and computational complexity theory. In doing so we focused on a general class of consensus, linking decentralisation to a precise measure of the influence of each peer in the network (a theme extensively studied in economics), an analysis of the structural properties of the consensus network, and the computational complexity of problems related to safe consensus. The obtained limiting results on Ripple and Stellar are coherent with the current practice and the proposals that industry is putting forward to improve decentralisation. We argue that the obtained results show this is a promising approach to the formal analysis of decentralisation.

Our results point to several avenues of future research. We are planning to extend our analysis to other blockchains based on BFT consensus that are currently being developed, noticeably Cobalt [10] as an evolution of the Ripple/Stellar tradition. More generally, we also want to explore the applicability of the methodology beyond the framework of Byzantine trust networks, since measures of the relative influence of peers are of interest for other blockchain frameworks, e.g. PoS. At the same time, we also intend to build on such measures to address the relationships between influence, decentralisation and, crucially, revenue. Properly understanding such mechanisms will serve to the long-term goal of designing more reliable and robust blockchains. On the application side, the development of a prototype analysis toolkit and collection of relevant data is also an ongoing activity.

References

- 1 J. Banzhaf. Weighted voting doesn't work: A mathematical analysis. *Rutgeres Law Review*, 19:317–343, 1965.
- 2 J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. Kroll, and E. Felten. Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 104–121. IEEE, 2015.
- 3 F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- 4 C. Cachin and M. Vukolic. Blockchain consensus protocols in the wild. Technical report, CoRR abs/1707.01873, 2017.
- 5 M. Castro and B. Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, OSDI '99, pages 173–186, Berkeley, CA, USA, 1999. USENIX Association. URL: <http://dl.acm.org/citation.cfm?id=296806.296824>.

- 6 B. Chase and MacBrough E. Analysis of the XRP ledger consensus protocol. Technical report, Ripple Research, 2018.
- 7 Morris H. DeGroot. Reaching a Consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- 8 D. Dolev. Byzantine generals stike again. *Journal of Algorithms*, 3(1):14–30, 1982.
- 9 C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, *Advances in Cryptology — CRYPTO’ 92*, pages 139–147, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.
- 10 MacBrough E. Cobalt: BFT governance in open networks. Technical report, Ripple Research, 2018.
- 11 D. Felsenthal and M. Machover. *The Measurement of Voting Power*. Edward Elgar Publishing, 1998.
- 12 M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985. doi:10.1145/3149.214121.
- 13 J. French. A formal theory of social power. *Psychological Review*, 61:181–194, 1956.
- 14 F. Gantmacher. *The Theory of Matrices*. AMS Chealsea Publishing, 1959.
- 15 J. Garay, A. Kiayias, and Nikos L. The bitcoin backbone protocol: Analysis and applications. In E. Oswald and M. Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015*. Springer, 2015. doi:10.1007/978-3-662-46803-6_10.
- 16 H. Garcia-Molina and D. Barbara. How to assign votes in a distributed system. *J. ACM*, 32(4):841–860, October 1985. doi:10.1145/4221.4223.
- 17 C. García-Pérez and A. Gotsman. Federated Byzantine Quorum Systems. In J. Cao, F. Ellen, L. Rodrigues, and B. Ferreira, editors, *22nd International Conference on Principles of Distributed Systems (OPODIS 2018)*, volume 125 of *LIPIcs*, pages 17:1–17:16, 2018. doi:10.4230/LIPIcs.OPODIS.2018.17.
- 18 U. Grandi. Social choice on social networks. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 169–184. COST, 2018.
- 19 X. Hu and L. Shapley. On authority distributions in organizations: Controls. *Games and Economic Behavior*, 45:153–170, 2003.
- 20 X. Hu and L. Shapley. On authority distributions in organizations: Equilibrium. *Games and Economic Behavior*, 45:132–152, 2003.
- 21 M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA, 2008.
- 22 D. Karos and H. Peters. Indirect control and power in mutual control structures. *Games and Economic Behavior*, 92, 2015.
- 23 M. Kim, Y. Kwon, and Y. Kim. Is stellar as secure as you think? In *IEEE Security and Privacy on the Blockchain (IEEE S&B ’19)*. IEEE, 2019.
- 24 L. Lachowski. Complexity of the quorum intersection property of the federated byzantine agreement system, 2019. URL: <https://arxiv.org/abs/1902.06493>.
- 25 L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- 26 M. Lokhava, G. Losa, D. Mazières, G. Hoare, N. Barry, E. Gafni, J. Jove, and Jed McCaleb R. Malinowsky. Fast and secure global payments with Stellar. In *Proceedings of SOSP’19*. ACM, 2019.
- 27 D. Malkhi and M. Reiter. Byzantine quorum systems. *Distributed Computing*, 11:203–213, 1998.
- 28 D. Mazierès. The stellar consensus protocol: A federated model for internet-level consensus. Stellar Development Foundation, 2016.
- 29 S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Bitcoin project white paper*, 2009.
- 30 A. Narayanan, J. Bonneau, E. Felten, A. Miller, and S. Goldfeder. *Bitcoin and Cryptocurrency Technologies*. Princeton University Press, 2016.

- 31 M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *Journal of the Association for Computing Machinery*, 27(2):228–234, 1980.
- 32 D. Peleg. Local majorities, coalitions and monopolies in graphs: a review. *Theor. Comput. Sci.*, 282(2):231–257, 2002.
- 33 L. Penrose. The elementary statistics of majority voting. *Journal of the Royal Statistical Society*, 109(1):53–57, 1946.
- 34 A. Proskurnikov and R. Tempo. A tutorial on modeling and analysis of dynamic social networks. Part I. *Annual Reviews in Control*, 43:65–79, 2017.
- 35 A. Proskurnikov and R. Tempo. A tutorial on modeling and analysis of dynamic social networks. Part II. *Annual Reviews in Control*, 45:166–190, 2018.
- 36 F. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, 1990.
- 37 D. Schwartz, N. Youngs, and A. Britto. The ripple protocol consensus algorithm. Technical report, Ripple Labs, 2014.
- 38 L. Shapley and M. Shubik. A method for evaluating the distribution of power in a committee system. *American Political Science Review*, 48:787–792, 1954.
- 39 M. Vukolic. *Quorum Systems with Applications to Storage and Consensus*. Morgan & Claypool Publishers, 2012.
- 40 M. Vukolic. The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication. In *Proceedings of iNetSec’15*, volume 9591 of *LNCS*, pages 112–125, 2015.
- 41 R. Wattenhofer. *Distributed Ledger Technology: The Science of the Blockchain*. Createspace Independent Publishing Platform, 2017.

A Technical appendix

A.1 Proofs of Section 2

A.1.1 Lemma 12

Proof. (3) By assumption $|T_i^o(x)| > 0$. So j observes at least $|T_i^o(x) \cap T_j \cap H|$ nodes with opinion x in T_j . Those are the honest nodes among the nodes with opinion x that both i and j can observe. So

$$|T_j^o(x) \cap H| \geq |T_i^o(x) \cap T_j \cap H|.$$

Now among the nodes in $T_i^o(x) \cap T_j$ there are at most β_{ij} Byzantine nodes that could reveal the opposite opinion \bar{x} to j . So,

$$|T_i^o(x) \cap T_j \cap H| \geq |T_i^o(x) \cap T_j| - \beta_{ij}.$$

The claim is finally established by the following series of (in)equalities:

$$\begin{aligned} |T_i^o(x) \cap T_j \cap H| &\geq |T_i^o(x) \cap T_j| - \beta_{ij} \\ &\geq |T_i^o(x)| - |T_i \setminus T_j| - \beta_{ij} \\ &= |T_i^o(x)| - |T_i| + |T_i \cap T_j| - \beta_{ij} \end{aligned}$$

(4) By assumption $|T_i^o(x)| > 0$. So, by (3), $|T_j^o(x) \cap H| = |T_i \cap T_j| + |T_i^o(x)| - |T_i| - \beta_{ij}$ whenever only the honest nodes in T_j have opinion x . It follows that

$$\begin{aligned} |T_j^o(\bar{x})| &\leq |T_j| - (|T_i \cap T_j| + |T_i^o(x)| - |T_i| - \beta_{ij}) \\ &= |T_j| - |T_i \cap T_j| - |T_i^o(x)| + |T_i| + \beta_{ij}. \end{aligned}$$

This completes the proof. ◀

A.1.2 Lemma 14

Proof. The proof consists of two sub-arguments. First we show that safety implies that, for all $i, j \in H$:

$$|T_i \cap T_j| > b \cdot (|T_i| + |T_j|) + \beta_{ij} \quad (8)$$

By safety (Definition 11), if $|T_i^\circ(x)| \geq q_i |T_i|$ with $i \in H$, then for all $j \in H$ $T_j^\circ(\bar{x}) < q |T_j|$. Assume $T_i^\circ(x) \geq q |T_i|$ with $i \in H$. By Lemma 12 and safety we have:

$$\begin{aligned} T_j^\circ(\bar{x}) &\leq |T_j| - |T_i \cap T_j| - q |T_i| + |T_i| + \beta_{ij} \\ &< q |T_j|. \end{aligned}$$

From $|T_j| - |T_i \cap T_j| - q |T_i| + |T_i| + \beta_{ij} < q |T_j|$ and the assumption on the effectiveness of q (that is, $q \in (0.5 + b, 1 - b]$) we thus obtain

$$\begin{aligned} b \cdot (|T_i| + |T_j|) + \beta_{ij} &\leq (1 - q) \cdot (|T_i| + |T_j|) + \beta_{ij} \\ &= |T_j| - q |T_j| - q |T_i| + |T_i| + \beta_{ij} \\ &< |T_i \cap T_j| \end{aligned}$$

as desired.¹⁸

Second We show that safety also implies that, for all $i, j \in H$:

$$b \cdot (|T_i| + |T_j|) + \beta_{ij} > \frac{b}{1 - b} (|T_i| + |T_j|) \quad (9)$$

We have established that safety implies that for all $i, j \in H$, $|T_i \cap T_j| > b \cdot (|T_i| + |T_j|) + \beta_{ij}$ (8), that is, the size of the intersection of the trust sets of i and j should be larger than the maximum possible fraction of Byzantine nodes times the combined size of the trust sets, plus β_{ij} . Now recall the definition of β_{ij} (2). By (8) it cannot be the case that $\beta_{ij} = |T_i \cap T_j|$. So $\beta_{ij} = b |T_k|$ where T_k is the the smallest set between T_i and T_j . Now assume, w.l.o.g. that $|T_i| \geq |T_j|$ and so that $|T_j| = x \cdot |T_i|$ with $x \in (0, 1]$. By (8) we have:

$$\begin{aligned} |T_i \cap T_j| &> b(|T_i| + |T_j|) + \beta_{ij} \\ &= b(|T_i| + x|T_i|) + bx|T_i| \\ &= b|T_i|(1 + 2x) \end{aligned}$$

From this, and the fact that a set is always at least as large as its intersection with another we obtain a lower bound for x by the following series of inequalities:

$$\begin{aligned} x|T_i| &\geq |T_i \cap T_j| \\ x|T_i| &> b|T_i|(1 + 2x) \\ x &> b(1 + 2x) \\ x &> b + 2bx \\ x - 2bx &> b \\ x(1 - 2b) &> b \\ x &> \frac{b}{1 - 2b} \end{aligned}$$

¹⁸Cf. [6, Proposition 4].

By substituting $\frac{b}{1-2b}$ for x in (8) we thus obtain a lower bound for $|T_i \cap T_j|$ in b . We then reformulate (8) in terms of the combined size $\alpha = |T_i| + |T_j|$ of the two trust sets:

$$\begin{aligned}
|T_i \cap T_j| &> b(\underbrace{|T_i| + x|T_i|}_{\alpha}) + bx|T_i| \\
&> b(|T_i| + \frac{b}{1-2b}|T_i|) + b\frac{b}{1-2b}|T_i| \\
&= b\alpha + b\frac{b}{1-2b}\frac{1-2b}{1-b}\alpha & \text{as } |T_i| = \alpha\frac{1-2b}{1-b} \\
&= b\alpha(1 + \frac{b}{1-b}) \\
&= \frac{b}{1-b}\alpha
\end{aligned}$$

So safety implies that the size of the intersection of T_i and T_j must be larger than the fraction $\frac{b}{1-b}$ of the combined size of the two sets. \blacktriangleleft

A.1.3 Lemma 15

Proof. The proof is by induction on $|H|$. Base if $|H| = 1$ the claim holds trivially. Step Now assume the claim holds for $|H| = m$ (IH). We prove it holds for $|H| = m + 1$. So assume for all $i, j \in H$, $|T_i \cap T_j| > 0.25 \cdot (|T_i| + |T_j|)$, and let k be the $m + 1^{\text{th}}$ node in H . By IH we know that $\bigcap_{i \in H \setminus \{k\}} T_i \neq \emptyset$. Now take one of the smallest (w.r.t. size) T_i with $i \in H \setminus \{k\}$ and call it T_j . There are two cases. $|T_k| \leq |T_j|$. Then $|T_j \cap T_k| > 0.5 \cdot |T_k|$. Since T_j was smallest amongst the T_i , it also holds that $\forall i \in H \setminus \{k\} \ |T_i \cap T_k| > 0.5 \cdot |T_k|$. From this we conclude that $\bigcap_{i \in H} T_i \neq \emptyset$. $|T_k| > |T_j|$ Then $|T_j \cap T_k| > 0.5 \cdot |T_j|$. Since T_j was smallest amongst the T_i , it also holds that $\forall i \in H \setminus \{k\} \ |T_i \cap T_j| > 0.5 \cdot |T_j|$, from which we also conclude $\bigcap_{i \in H} T_i \neq \emptyset$. \blacktriangleleft

A.1.4 Theorem 20

Proof. Left to right Straightforwardly proven by contraposition. Right to left Proceed by contraposition and assume there is a profile \mathbf{o} , a function s and agents 1 and 2 such that all $k \in C_1 = F_s^*(\{1\})$ agree on x and all $k \in C_2 = F_s^*(\{2\})$ agree on \bar{x} . Observe that C_1 and C_2 are quora containing (since the BTN is vetoed) 1 and 2. There are two cases. Either $C_1 \cap C_2 = \emptyset$, or if that is not the case then $C_1 \cap C_2 \subseteq B$ as only Byzantine nodes can reveal different opinions to different nodes. Hence C_1 and C_2 are either disjoint or their intersection contains only Byzantine nodes. \blacktriangleleft

A.1.5 Theorem 21

Proof. To see that the problem is contained in coNP, we describe a nondeterministic polynomial-time algorithm to decide whether $\mathcal{T} = \langle N, \mathcal{C} \rangle$ does not have the quorum intersection property. The algorithm guesses two disjoint sets $Q_1, Q_2 \subseteq N$. Then, for each $u \in [2]$ and for each $i \in Q_u$, the algorithm checks if there is some $C \in \mathcal{C}(i)$ such that $C \subseteq Q_u$. That is, the algorithm verifies that Q_1 and Q_2 are quora (which is the case if and only if all checks succeed). Clearly, all checks can be performed in polynomial time. Thus, deciding whether \mathcal{T} has the quorum intersection property is in coNP.

To show coNP-hardness, we reduce from the coNP-complete propositional unsatisfiability problem (UNSAT). Let φ be a propositional formula containing the propositional

variables x_1, \dots, x_n . Without loss of generality, we may assume that φ is in 3CNF, i.e., that $\varphi = c_1 \wedge \dots \wedge c_m$ and that for each $j \in [m]$, $c_j = (T_{j,1} \vee T_{j,2} \vee T_{j,3})$, where $T_{j,1}, T_{j,2}, T_{j,3}$ are literals. We construct a command game $\mathcal{T} = \langle N, \mathcal{C} \rangle$ that has the quorum intersection property if and only if φ is unsatisfiable.

We let:

$$N = \{z_0, z_1\} \cup \{c_j \mid j \in [m]\} \cup \{y_i, p_i, n_i \mid i \in [n]\}.$$

That is, we have nodes z_0, z_1 , a node c_j for each clause of φ , and nodes y_i, p_i, n_i for each variable occurring in φ .

We define the sets of winning coalitions of the nodes in N as follows:

$$\begin{aligned} \mathcal{C}(z_0) &= \{\{z_0, y_1, \dots, y_n\}\}; \\ \mathcal{C}(z_1) &= \{\{z_1, c_1, \dots, c_m\}\}; \\ \mathcal{C}(y_i) &= \{\{y_i, p_i\}, \{y_i, n_i\}\} && \text{foreach } i \in [n]; \\ \mathcal{C}(c_j) &= \{\{c_j, \sigma(T_{j,1})\}, \{c_j, \sigma(T_{j,2})\}, \{c_j, \sigma(T_{j,3})\}\} && \text{foreach } j \in [m]; \\ \mathcal{C}(p_i) &= \{\{p_i, z_0\}, \{p_i, z_1\}\} && \text{foreach } i \in [n]; \text{ and} \\ \mathcal{C}(n_i) &= \{\{n_i, z_0\}, \{n_i, z_1\}\} && \text{foreach } i \in [n]; \end{aligned}$$

where for each positive literal x_i , we let $\sigma(x_i) = p_i$; and for each negative literal $\neg x_i$, we let $\sigma(\neg x_i) = n_i$.

We argue that $\mathcal{T} = \langle N, \mathcal{C} \rangle$ has the quorum intersection property if and only if φ is unsatisfiable. We show the equivalent statement that $\mathcal{T} = \langle N, \mathcal{C} \rangle$ does **not** have the quorum intersection property if and only if φ is **satisfiable**.

(\Rightarrow) Suppose that there exist $Q_1, Q_2 \in \mathbf{Q}^{\mathcal{T}}$ such that $Q_1 \cap Q_2 = \emptyset$. We may assume without loss of generality that Q_1 and Q_2 are core quora. We know that neither Q_1 nor Q_2 can contain both z_0 and z_1 , because each quorum of \mathcal{T} must contain either z_0 or z_1 (by the specific construction of \mathcal{T}). Thus, we may assume that $z_0 \in Q_2$ and $z_1 \in Q_1$.

Then also $\{y_1, \dots, y_n\} \subseteq Q_2$. Moreover, for each $i \in [n]$, we know that then either $p_i \in Q_2$ or $n_i \in Q_2$ (and not both). We also know that $\{c_1, \dots, c_m\} \subseteq Q_1$. Now define the truth assignment $\alpha : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ as follows. For each $i \in [n]$, we let $\alpha(x_i) = 1$ if $n_i \in Q_2$ and we let $\alpha(x_i) = 0$ if $p_i \in Q_2$.

We show that α satisfies φ . Take an arbitrary clause c_j of φ . Due to the construction of $\mathcal{C}(c_j)$, we know that Q_1 contains (at least) one of $\sigma(T_{j,1}), \sigma(T_{j,2}), \sigma(T_{j,3})$. Take some $u \in [3]$ such that $\sigma(T_{j,u}) \in Q_1$. We show that α satisfies $T_{j,u}$. To derive a contradiction, suppose the contrary, i.e., that α does not satisfy $T_{j,u}$. Then $\sigma(T_{j,u}) \in Q_2$ (by the construction of α), and thus $Q_1 \cap Q_2 \neq \emptyset$, which contradicts our initial assumption that $Q_1 \cap Q_2 = \emptyset$. Thus, we can conclude that α satisfies $T_{j,u}$. This concludes our proof that φ is satisfiable.

(\Leftarrow) Conversely, suppose that φ is satisfiable, i.e., that there is some truth assignment $\alpha : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ that satisfies all clauses of φ . For each clause c_j , define $\rho(c_j)$ to be some literal $T_{j,u}$ in c_j that is satisfied by α . Moreover, for each $i \in [n]$, let $\mu(x_i) = n_i$ if $\alpha(x_i) = 1$ and $\mu(x_i) = p_i$ if $\alpha(x_i) = 0$. Consider the following two sets $Q_1, Q_2 \subseteq N$:

$$\begin{aligned} Q_1 &= \{z_1, c_1, \dots, c_m\} \cup \{\sigma(\rho(c_j)) \mid j \in [m]\}; \text{ and} \\ Q_2 &= \{z_0, y_1, \dots, y_n\} \cup \{\mu(x_i) \mid i \in [n]\}; \end{aligned}$$

It is straightforward to verify that Q_1 and Q_2 are both quora, i.e., that $Q_1, Q_2 \in \mathbf{Q}^{\mathcal{T}}$. Moreover, since it holds that $Q_1 \cap Q_2 = \emptyset$, we know that \mathcal{T} does not satisfy the quorum intersection property. \blacktriangleleft

A.1.6 Theorem 22

Proof sketch. Membership in coNP follows directly from Proposition 21. We show coNP-hardness by modifying the reduction given in the proof of Proposition 21. We describe a reduction from UNSAT. Let φ be a propositional formula containing the propositional variables x_1, \dots, x_n . Without loss of generality, we may assume that φ is in 3CNF. Moreover, without loss of generality, we may assume that $\varphi[x_1 \mapsto 1]$ is unsatisfiable. We construct the command game $\mathcal{T} = \langle N, \mathcal{C} \rangle$ as in the proof of Proposition 21. Moreover, we transform \mathcal{T} into \mathcal{T}' by removing the coalition $\{y_1, n_1\}$ from $\mathcal{C}(y_1)$. By a similar argument as the one used in the proof of Proposition 21, we know that the command game \mathcal{T}' satisfies the quorum intersection property, because $\varphi[x_1 \mapsto 1]$ is unsatisfiable. Moreover, \mathcal{T} satisfies the quorum intersection property if and only if φ is unsatisfiable. \blacktriangleleft

A.2 Proofs of Section 4

A.2.1 Theorem 25

Proof. We first need to introduce some auxiliary notation. Given an influence matrix I , $\mathcal{G}(I) = \langle N, E \rangle$ denotes the (directed) graph of I , where $E = \{ij \mid I_{ji} > 0\}$. Intuitively $ji \in E$ whenever j influences i (i.e., has a positive Banzhaf-Penrose index in i 's simple game). By assumption there exist nodes that influence all other honest nodes (themselves included). Let now $E(\bigcap_{i \in H} T_i) = \{i \in N \mid \exists j \in \bigcap_{i \in H} T_i, ij \in E\}$, that is, the set of nodes that influence some node that influences all honest nodes. We distinguish three cases.

$|B| = 0$ So there are no Byzantine nodes in \mathcal{T} , and $N = H = E(\bigcap_{i \in H} T_i)$. It follows that $\mathcal{G}(I(\mathcal{T}))$ is strongly connected (there exists a path from every node to every node) and aperiodic (there are no two cycles in the graph whose length is divided by an integer larger than 1). Trivially, it is also closed (there exists no node outside $\mathcal{G}(I(\mathcal{T}))$ that influences nodes in $\mathcal{G}(I(\mathcal{T}))$). The full regularity of I then follows from known results on influence matrices (cf. [34, Lemma 11]): if $\mathcal{G}(I)$ contains only one strongly connected component and is aperiodic, then I is fully regular.

$|B| = 1$ and $E(\bigcap_{i \in H} T_i) \cap B \neq \emptyset$ So there exists exactly one Byzantine node in N , which furthermore belongs to $E(\bigcap_{i \in H} T_i)$ (that is, it influences at least one honest node influencing in turn all honest nodes). Call i such Byzantine agent. Recall that, by construction, $ii \in E$. So the subgraph consisting of i and the self-loop ii is the only closed, aperiodic, strongly connected component of $\mathcal{G}(I(\mathcal{T}))$. Full regularity therefore follows by known results as in the previous case.

$|B| \geq 1$ or $E(\bigcap_{i \in H} T_i) \cap B = \emptyset$ So there exist several Byzantine nodes in N or there are Byzantine nodes which do not influence nodes in $\bigcap_{i \in H} T_i$. Both such cases determine the existence, by arguments analogous to those provided for the previous two cases, of several closed, aperiodic strongly connected components in $\mathcal{G}(I(\mathcal{T}))$. The regularity of I then follows again from known results on influence matrices (cf. [34, Theorem 12], [21, Theorem 8.1]: if all closed strongly connected components of $\mathcal{G}(I)$ are aperiodic, then I is regular.

In all three cases I is regular, proving claim 1). In the first two cases ($|B| \leq 1$) I is furthermore fully regular, establishing claim 2). Finally, to prove claim 3) we reason as follows. If there exists a honest agent in $\bigcap_{i \in H} T_i$ trusting a Byzantine agent, then the only closed strongly connected components of $\mathcal{G}(I)$ are the Byzantine nodes. In the limit, such nodes will therefore be the only ones having positive influence. \blacktriangleleft