

The Stellar Consensus Protocol: A Federated Model for Internet-level Consensus

David Mazières
Stellar Development Foundation

April 8, 2015 DRAFT

Abstract

This paper introduces federated Byzantine agreement (FBA), a new approach to consensus. FBA’s robustness stems from quorum slices, individual trust decisions made by each node that together determine system-level quorums. Slices bind the system together much the way individual networks’ peering and transit decisions now unify the Internet.

We also present the Stellar Consensus Protocol (SCP), a construction for FBA. Like all Byzantine agreement protocols, SCP makes no assumptions about the rational behavior of attackers. Unlike prior Byzantine agreement models, which presuppose a unanimously accepted membership list, SCP enjoys open membership that promotes organic network growth. Compared to decentralized proof of-work and proof-of-stake schemes, SCP has modest computing and financial requirements, lowering the barrier to entry and potentially opening up financial systems to new participants.

1 Introduction

Financial infrastructure is currently a mess of closed systems. Gaps between these systems mean that transaction costs are high [38] and money moves slowly across political and geographic boundaries [4, 12]. This friction has curtailed the growth of financial services, leaving billions of people underserved financially [41].

To solve these problems, we need financial infrastructure that supports the kind of organic growth and innovation we’ve seen from the Internet, yet still ensures the integrity of financial transactions. Historically, we have relied on high

barriers to entry to ensure integrity. We trust established financial institutions and do our best to regulate them. But this exclusivity conflicts with the goal of organic growth. Growth demands new, innovative participants, who may possess only modest financial and computing resources.

We need a worldwide financial network open to anyone, so that new organizations can join and extend financial access to unserved communities. The challenge for such a network is ensuring participants record transactions correctly. With a low barrier to entry, users won't trust providers to police themselves. With worldwide reach, providers won't all trust a single entity to operate the network. A compelling alternative is a decentralized system in which participants together ensure integrity by agreeing on the validity of one another's transactions. Such agreement hinges on a mechanism for worldwide consensus.

This paper presents federated Byzantine agreement (FBA), a model suitable for worldwide consensus. In FBA, each participant knows of others it considers important. It waits for the vast majority of those others to agree on any transaction before considering the transaction settled. In turn, those important participants do not agree to the transaction until the participants *they* consider important agree as well, and so on. Eventually, enough of the network accepts a transaction that it cannot be undone. Only then do any participants consider the transaction settled. FBA's consensus can ensure the integrity of a financial network. Its decentralized control can spur organic growth.

This paper further presents the Stellar consensus protocol (SCP), a construction for FBA. We prove that SCP's safety is optimal for an asynchronous protocol, in that it guarantees agreement under any node-failure scenario that admits such a guarantee. We also show that SCP is free from *blocked* states—in which consensus is no longer possible—unless participant failures make it impossible to satisfy trust dependencies. SCP is the first provably safe consensus mechanism to enjoy four key properties simultaneously:

- **Decentralized control.** Anyone is able to participate and no central authority dictates whose approval is required for consensus.
- **Low latency.** In practice, nodes can reach consensus at timescales humans expect for web or payment transactions—i.e., a few seconds at most.
- **Flexible trust.** Users have the freedom to trust any combination of parties they see fit. In particular, a small non-profit may play a key role in keeping much larger institutions honest.

- **Asymptotic security.** Safety rests on digital signatures and hash families whose parameters can realistically be tuned to protect against adversaries with unimaginably vast computing power.

SCP has applications beyond financial markets for ensuring organizations perform important functions honestly. A good example is certificate authorities (CAs), who literally hold the keys to the web. Experience shows that CAs sign incorrect certificates that get used in the wild [1, 28]. Several proposals address this problem through certificate transparency [21, 29, 5, 31]. SCP can strengthen certificate transparency: not only would anyone who so pleases be able to audit CA behavior, auditors would agree on the set of all issued certificates, making it difficult to backpedal or override previously issued certificates without detection.

The next section discusses previous approaches to consensus. Section 3 defines federated Byzantine agreement and lays out the notions of safety and liveness such a system hopes to achieve. Section 4 discusses optimal failure resilience in an FBA system, thereby setting out the notion of security we hope to achieve with SCP. Section 5 builds the tools needed to guarantee consensus in a federated setting. Section 6 presents SCP itself, and proves its safety and absence of blocked states for nodes with satisfiable dependencies. Section 7 discusses limitations of SCP. Finally, Section 8 summarizes results. For readers less familiar with mathematical notation, Appendix A defines some symbols used throughout the paper.

2 Related work

Figure 1 summarizes how SCP differs from previous consensus mechanisms. The most famous decentralized consensus mechanism is the proof-of-work scheme advanced by Bitcoin [32]. Bitcoin takes a two-pronged approach to consensus. First, it provides incentives for rational actors to behave well. Second, it settles transactions through a proof-of-work algorithm designed to protect against ill-behaved actors who do not possess the majority of the system’s computing power. Bitcoin has overwhelmingly demonstrated the appeal of decentralized consensus [7].

Proof of work has limitations, however. First, it wastes resources: by one estimate from 2014, Bitcoin consumed as much electric power as the entire country of Ireland [34]. Second, secure transaction settlement suffers from expected latencies in the minutes or tens of minutes [20]. Finally, in contrast to traditional cryptographic protocols, proof of work offers no asymptotic security. Given non-rational attackers—or ones with extrinsic incentives to sabotage consensus—small computational advantages can invalidate the security assumption, leading to so-called

mechanism	decentralized control	low latency	flexible trust	asymptotic security
proof of work [32]	✓			
proof of stake [22]	✓	maybe		maybe
Byzantine agreement [11]		✓	✓	✓
Tendermint [23]	✓	✓		✓
SCP (this work)	✓	✓	✓	✓

Figure 1: Properties of different consensus mechanisms

“51%” attacks, which in fact require less than 51% of computation [18]. As the leading digital currency backed by the most computing power, Bitcoin enjoys a measure of protection against 51% attacks. Smaller systems have fallen victim [14, 9], however, posing a problem for any proof-of-work system not built on the Bitcoin block chain.

An alternative to proof of work is proof of stake [22], in which consensus depends on parties that have posted collateral. Like proof of work, rewards encourage rational participants to obey the protocol; some systems additionally penalize bad behavior [15]. Proof of stake opens the possibility of so-called “nothing at stake” attacks, in which parties that previously posted collateral but later cashed it in and spent the money can go back and rewrite history from a point where they still had stake. To mitigate such attacks, systems effectively combine proof of stake with proof of work—scaling down the required work in proportion to stake—or delay refunding collateral long enough for some other (sometimes informal) consensus mechanism to establish an irreversible checkpoint.

Still another approach to consensus is Byzantine agreement [37, 27], the best known variant of which is PBFT [11]. Byzantine agreement ensures consensus despite arbitrary (including non-rational) behavior on the part of some fraction of participants. This approach has two appealing properties. First, consensus can be fast and efficient. Second, trust is entirely decoupled from resource ownership, which makes it possible for a small non-profit to help keep more powerful organizations such as banks or CAs honest. Complicating matters, however, all parties must agree on the the exact list of participants. Moreover, attackers must be prevented from joining multiple times and exceeding the system’s failure tolerance, a so-called Sybil attack [16]. BFT-CUP [2] accommodates unknown participants, but still presupposes a Sybil-proof centralized admission-control mechanism.

Generally, membership in Byzantine agreement systems is set by a central authority or closed negotiation. Prior attempts to decentralize admission have given up some of the benefits. One approach, taken by Ripple, is to publish a “starter” membership list that participants can edit for themselves, hoping people’s edits are either inconsequential or reproduced by an overwhelming fraction of participants. Unfortunately, because divergent lists invalidate safety guarantees [39], users are reluctant to edit the list in practice and a great deal of power ends up concentrated in the maintainer of the starter list. Another approach, taken by Tendermint [23], is to base membership on proof of stake. However, doing so once again ties trust to resource ownership. SCP is the first Byzantine agreement protocol to give each participant maximum freedom in choosing which combinations of other participants to trust.

3 Federated Byzantine agreement systems

This section introduces the federated Byzantine agreement model (FBA). Like non-federated agreement, FBA addresses the problem of updating replicated state such as a transaction ledger or certificate tree. By agreeing on what updates to apply, nodes avoid contradictory, irreconcilable states. We identify each update by a unique *slot* from which inter-update dependencies can be inferred. For instance, slots may be consecutively numbered positions in a sequentially applied log.

An FBA system runs a *consensus protocol* that ensures nodes agree on slot contents. A node v can safely apply update x in slot i when it has safely applied updates in all slots upon which i depends and, additionally, it believes all correctly functioning nodes will eventually agree on x for slot i . At this point, we say v has *externalized* x for slot i . The outside world may react to externalized values in irreversible ways, so a node cannot later change its mind about them.

A challenge for FBA is that malicious parties can join many times and outnumber honest nodes. Hence, traditional majority-based quorums do not work. Instead, FBA determines quorums in a decentralized way, by each node selecting what we call quorum slices, described in the next subsection. The following subsection discusses examples. Finally, we define the key properties of safety and liveness that a consensus protocol should hope to achieve.




Figure 2: v_1 's quorum slice is not a quorum.

3.1 Quorum slices

In a consensus protocol, nodes exchange messages asserting statements about slots. We assume such assertions cannot be forged, for instance because nodes are named by public key and they digitally sign messages. When a node hears a sufficient set of nodes assert a statement, it assumes no functioning node will ever contradict that statement. We call such a sufficient set a *quorum slice*, or, more concisely, just a *slice*. To permit progress in the face of node failures, a node may have multiple slices, any one of which is sufficient to convince it of a statement. At a high level, then, an FBA system consists of a loose confederation of nodes each of which has chosen one or more slices. More formally:

Definition (FBAS). A federated Byzantine agreement system, or *FBAS*, is a pair $\langle \mathbf{V}, \mathbf{Q} \rangle$ comprising a set of nodes \mathbf{V} and a quorum function $\mathbf{Q} : \mathbf{V} \rightarrow 2^{2^{\mathbf{V}}} \setminus \emptyset$ specifying one or more quorum slices for each node, where a node belongs to all of its own quorum slices, i.e., $\forall v \in \mathbf{V}, \forall q \in \mathbf{Q}(v), v \in q$.

Definition (quorum). A set of nodes $U \subseteq \mathbf{V}$ in FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ is a *quorum* iff $U \neq \emptyset$ and U contains a slice for each member, i.e., $\forall v \in U, \exists q \in \mathbf{Q}(v)$ such that $q \subseteq U$.

A quorum is a set of nodes sufficient to reach agreement. A quorum slice is the subset of a quorum convincing one particular node of agreement. A quorum slice may be smaller than a quorum. Consider the four-node system in Figure 2, where each node has a single slice and arrows point to the other members of that slice. Node v_1 's slice $\{v_1, v_2, v_3\}$ is sufficient to convince v_1 of a statement. But v_2 's and v_3 's slices include v_4 , meaning neither v_2 nor v_3 can assert a statement without v_4 's agreement. Hence, no agreement is possible without v_4 's participation, and the only quorum including v_1 is the set of all nodes $\{v_1, v_2, v_3, v_4\}$.

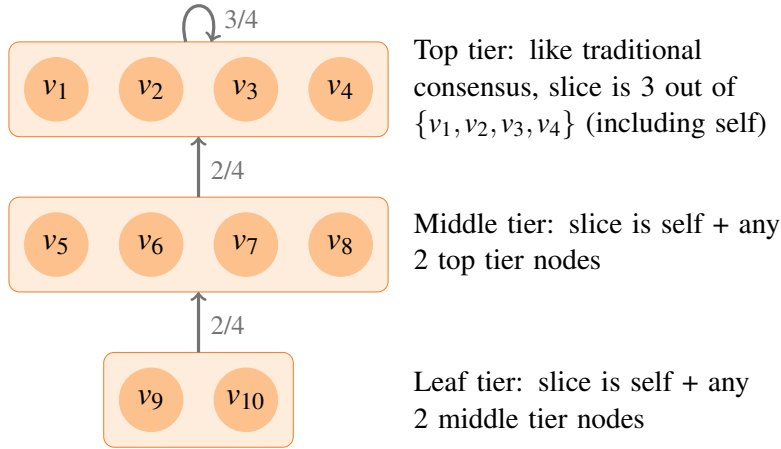


Figure 3: Tiered quorum structure example

Traditional non-federated consensus requires all nodes to accept the same slices, meaning $\forall v_1, v_2, \mathbf{Q}(v_1) = \mathbf{Q}(v_2)$. Because any slice suffices to convince all nodes of an assertion, centralized systems do not distinguish between slices and quorums. The downside is that membership and quorums must somehow be pre-ordained, precluding open membership and decentralized control. A typical traditional system such as PBFT [11] has $3f + 1$ nodes, any $2f + 1$ of which constitute a quorum. Here f is the maximum number of Byzantine failures—meaning nodes acting arbitrarily—the system can survive.

FBA, introduced by this paper, generalizes centralized consensus to accommodate a greater range of settings. FBA’s key difference is that each node v chooses its own quorum slice set $\mathbf{Q}(v)$. System-wide quorums thus arise from individual decisions made by each node. Nodes may select slices based on arbitrary criteria such as reputation or a financial arrangements. In some settings, it may be impractical for any single node to track the complete set \mathbf{V} of all nodes in the system, yet consensus should still be possible.

3.2 Examples and discussion

Figure 3 shows an example of a tiered system in which different nodes have very different slice sets, something possible only with FBA. A top tier, comprising v_1, \dots, v_4 , is structured like a PBFT system with $f = 1$, meaning it can tolerate one Byzantine failure so long as the other three nodes are reachable and well-behaved. Nodes v_5, \dots, v_8 constitute a middle tier and depend not on each other,

but rather on the top tier. Only two top tier nodes are required to form a slice for a middle tier node. (The top tier assumes at most one Byzantine failure, so two top tier nodes cannot both fail unless the whole system has failed.) Nodes v_9 and v_{10} are in a leaf tier for which a slice consists of any two middle tier nodes. Note that v_9 and v_{10} may pick disjoint slices such as $\{v_5, v_6\}$ and $\{v_7, v_8\}$; nonetheless, both will indirectly depend on the top tier.

In practice, the top tier could consist of anywhere from four to dozens of widely known and trusted financial institutions. As the size of the top tier grows, there may not be exact agreement on its membership, but there will be significant overlap between most parties' notions of top tier. Additionally, one can imagine multiple middle tiers, for instance one for each country or geographic region.

This tiered structure resembles inter-domain network routing. The Internet today is held together by individual peering and transit relationships between pairs of networks. No central authority dictates or arbitrates these arrangements. Yet these pairwise relationships have sufficed to create a notion of *de facto* tier one ISPs [33]. Though Internet reachability does suffer from firewalls, *transitive* reachability is nearly complete—e.g., a firewall might block the New York Times, but if it allows Google and Google can reach the New York Times, then the New York Times is transitively reachable. Transitive reachability may be of limited utility for web sites, but it is crucial for consensus; the equivalent example would be Google accepting statements only if the New York Times does.

If we think of quorum slices as analogous to network reachability and quorums as analogous to transitive reachability, then the Internet's near complete transitive reachability suggests we can likewise ensure worldwide consensus with FBA. In many ways, consensus is an easier problem than inter-domain routing. While transit consumes resources and costs money, slice inclusion merely requires checking digital signatures. Hence, FBA nodes can err on the side of inclusiveness, constructing conservative slices with greater interdependence and redundancy than typically seen in peering and transit arrangements.

Another example not possible with centralized consensus is cyclic dependency structures, such as the one depicted in Figure 4. Such a cycle is unlikely to arise intentionally, but when individual nodes choose their own slices, it is possible for the overall system to end up embedding dependency cycles. The bigger point is that, compared to traditional Byzantine agreement, an FBA protocol must cope with a far wider variety of quorum structures.




Figure 4: Cyclic quorum structure example

3.3 Safety and liveness

We categorize nodes as either *well-behaved* or *ill-behaved*. A well-behaved node chooses sensible quorum slices (discussed further in Section 4.1) and obeys the protocol, including eventually responding to all requests. An ill-behaved node does not. Ill-behaved nodes suffer Byzantine failure, meaning they behave arbitrarily. For instance, an ill-behaved node may be compromised, its owner may have maliciously modified the software, or it may have crashed.

The goal of Byzantine agreement is to ensure that well-behaved nodes externalize the same values despite the presence of such ill-behaved nodes. There are two parts to this goal. First, we would like to prevent nodes from diverging and externalizing different values for the same slot. Second, we would like to ensure nodes can actually externalize values, as opposed to getting blocked in some dead-end state from which consensus is no longer possible. We introduce the following two terms for these properties:

Definition (safety). A set of nodes in an FBAS enjoy *safety* if no two of them ever externalize different values for the same slot.

Definition (liveness). A node in an FBAS enjoys *liveness* if, given appropriate message delivery and timing, it can externalize new values.

We call well-behaved nodes that enjoy both safety and liveness *correct*. Nodes that are not correct have *failed*. All ill-behaved nodes have failed, but a well-behaved node can fail, too, by waiting indefinitely for messages from ill-behaved nodes, or, worse, by having its state poisoned by incorrect messages from ill-behaved nodes. Figure 5 illustrates the possible kinds of node failure. To the left are Byzantine failures, the ill-behaved nodes. To the right are two kinds of




Figure 5: Venn diagram of node failures

well-behaved but failed nodes. Nodes that lack liveness are termed *blocked*, while those that lack safety are termed *divergent*. An attack violating safety is strictly more powerful than one violating only liveness, so we classify divergent nodes as a subset of blocked ones.

Our definition of liveness is weak in that it admits *perpetual preemption*, a state in which consensus remains forever possible, yet the network continually thwarts it by delaying or reordering critical messages in just the wrong way. Perpetual preemption is inevitable in a purely asynchronous, deterministic system that survives node failure [19]. Fortunately, preemption is transient. It does not indicate node failure, because the system can recover at any time. Protocols can mitigate the problem through randomness (if nodes keep choosing random candidate values until enough happen to pick the same one [6, 8]) or through realistic assumptions about message latency [17]. The latter is more practical when one needs to bound execution time. Of course, only termination and not safety should depend on message timing.

4 Optimal resilience

Whether or not nodes enjoy safety and liveness depends on several factors: what quorum slices they have chosen, which nodes are ill-behaved, and of course the concrete consensus protocol and network behavior. As is common for asynchronous systems, we assume the network eventually delivers messages between well-behaved nodes, but can otherwise arbitrarily delay or reorder messages.

This section answers the following question: given a specific $\langle \mathbf{V}, \mathbf{Q} \rangle$ and particular subset of \mathbf{V} that is ill-behaved, what are the best safety and liveness that any federated Byzantine agreement protocol can guarantee regardless of the network? We first discuss quorum intersection, a property without which safety is impossible to guarantee. We then introduce a notion of dispensable sets—sets of failed nodes in spite of which it is possible to guarantee both safety and liveness.




Figure 6: FBAS lacking quorum intersection

4.1 Quorum intersection

A protocol can guarantee agreement only if the quorum slices represented by function \mathbf{Q} satisfy a validity property we call quorum intersection.

Definition (quorum intersection). An FBAS enjoys *quorum intersection* iff any two of its quorums share a node, i.e., for all quorums U_1 and U_2 , $U_1 \cap U_2 \neq \emptyset$.

Figure 6 illustrates a system lacking this property, where \mathbf{Q} permits two quorums, $\{v_1, v_2, v_3\}$ and $\{v_4, v_5, v_6\}$, that do not intersect. Disjoint quorums can independently agree on contradictory statements, undermining system-wide agreement. When many quorums exist, quorum intersection fails if any two do not intersect. For example, the set of all nodes $\{v_1, \dots, v_6\}$ in Figure 6 is a quorum that intersects the other two, but the system still lacks quorum intersection.

No protocol can guarantee safety in the absence of quorum intersection, since such a configuration can operate as two different FBAS systems that do not know anything about each other. However, even with quorum intersection, safety may be impossible to guarantee in the presence of ill-behaved nodes. Compare Figure 6, in which there are two disjoint quorums, to Figure 7, in which two quorums intersect at a single node v_7 and v_7 is ill-behaved. If v_7 makes inconsistent statements to the left and right quorums, the effect is equivalent to disjoint quorums.

In fact, since ill-behaved nodes contribute nothing to safety, no protocol can guarantee safety without the well-behaved nodes enjoying quorum intersection on their own. After all, in a worst-case scenario for safety, all ill-behaved nodes can constantly make inconsistent statements in an attempt to make quorums diverge. Two quorums overlapping only at ill-behaved nodes will again be able to operate like two different FBAS systems thanks to the duplicity of the ill-behaved nodes. In short, FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ can survive Byzantine failure by a set of nodes $B \subseteq \mathbf{V}$ iff $\langle \mathbf{V}, \mathbf{Q} \rangle$ enjoys quorum intersection after deleting the nodes in B from \mathbf{V} and from all slices in \mathbf{Q} . More formally:




Figure 7: Ill-behaved node v_7 can undermine quorum intersection.

Definition (delete). If $\langle \mathbf{V}, \mathbf{Q} \rangle$ is an FBAS and $B \subseteq \mathbf{V}$ is a set of nodes, then to *delete* B from $\langle \mathbf{V}, \mathbf{Q} \rangle$, written $\langle \mathbf{V}, \mathbf{Q} \rangle^B$, means to compute the modified FBAS $\langle \mathbf{V} \setminus B, \mathbf{Q}^B \rangle$ where $\mathbf{Q}^B(v) = \{q \setminus B \mid q \in \mathbf{Q}(v)\}$.

It is the responsibility of each node v to ensure $\mathbf{Q}(v)$ does not violate quorum intersection. One way to do so is to pick conservative slices that lead to large quorums. Of course, a malicious v may intentionally pick $\mathbf{Q}(v)$ to violate quorum intersection. But a malicious v can also lie about the value of $\mathbf{Q}(v)$ or ignore $\mathbf{Q}(v)$ to make arbitrary assertions. In short, $\mathbf{Q}(v)$'s value is not meaningful when v is ill-behaved. This is why the necessary property for safety—quorum intersection of well-behaved nodes after deleting ill-behaved nodes—is unaffected by the slices of ill-behaved nodes.

Suppose Figure 6 evolved from a three-node FBAS v_1, v_2, v_3 with quorum intersection to a six-node FBAS without. When v_4, v_5, v_6 join, they maliciously choose slices that violate quorum intersection and no protocol can guarantee safety for \mathbf{V} . Fortunately, deleting the bad nodes to yield $\langle \mathbf{V}, \mathbf{Q} \rangle^{\{v_4, v_5, v_6\}}$ restores quorum intersection, meaning at least $\{v_1, v_2, v_3\}$ can enjoy safety. Note that deletion is conceptual, for the sake of describing optimal safety. A protocol should guarantee safety for v_1, v_2, v_3 without their needing to know that v_4, v_5, v_6 are ill-behaved.

4.2 Dispensable sets (DSets)

We capture the fault-tolerance of nodes' slice selections through a notion of *dispensable set* or *DSet*. Informally, the safety and liveness of nodes outside a DSet can be guaranteed regardless of the behavior of nodes inside the DSet. Put another way, in an optimally resilient FBAS, if a single DSet encompasses every ill-behaved node, it also contains every failed node, and conversely all nodes outside the DSet are correct. As an example, in a centralized PBFT system with

$3f + 1$ nodes, any f or fewer nodes constitute a DSet. Since PBFT in fact survives up to f Byzantine failures, its robustness is optimal.

In the less regular example of Figure 3, $\{v_1\}$ is a DSet, since one top tier node can fail without affecting the rest of the system. $\{v_9\}$ is also a DSet because no other node depends on v_9 for correctness. $\{v_6, \dots, v_{10}\}$ is a DSet, because neither v_5 nor the top tier depend on any of those five nodes. $\{v_5, v_6\}$ is *not* a DSet, as it is a slice for v_9 and v_{10} and hence, if entirely malicious, can lie to v_9 and v_{10} and convince them of assertions inconsistent with each other or the rest of the system.

To prevent a misbehaving DSet from affecting the correctness of other nodes, two properties must hold. For safety, deleting the DSet cannot undermine quorum intersection. For liveness, the DSet cannot deny other nodes a functioning quorum. This leads to the following definition:

Definition (DSet). Let $\langle \mathbf{V}, \mathbf{Q} \rangle$ be an FBAS and $B \subseteq \mathbf{V}$ be a set of nodes. We say B is a dispensible set, or *DSet*, iff:

1. (*quorum intersection despite B*) $\langle \mathbf{V}, \mathbf{Q} \rangle^B$ enjoys quorum intersection, and
2. (*quorum availability despite B*) Either $\mathbf{V} \setminus B$ is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle$ or $B = \mathbf{V}$.

Quorum availability despite B protects against nodes in B refusing to answer requests and blocking other nodes' progress. Quorum intersection despite B protects against the opposite—nodes in B making contradictory assertions that enable other nodes to externalize inconsistent values for the same slot. Nodes must balance the two threats in slice selection. All else equal, bigger slices lead to bigger quorums with greater overlap, meaning fewer failed node sets B will undermine quorum intersection when deleted. On the other hand, bigger slices are more likely to contain failed nodes, endangering quorum availability.

The smallest DSet containing all ill-behaved nodes may encompass well-behaved nodes as well; this reflects the fact that a sufficiently large set of ill-behaved nodes can cause well-behaved nodes to fail. For instance, in Figure 3, the smallest DSet containing v_5 and v_6 is $\{v_5, v_6, v_9, v_{10}\}$. As a special case, the set of all nodes, \mathbf{V} , is a DSet. This is because if all nodes fail, the remaining (zero) nodes are vacuously correct. Given sufficiently many ill-behaved nodes, the set \mathbf{V} of all nodes may be the smallest DSet to contain all ill-behaved ones, which means no protocol can guarantee anything better than complete system failure.

The set of DSets in an FBAS is determined *a priori* by the quorum function \mathbf{Q} . Which nodes are well- and ill-behaved depends on runtime behavior, such as machines getting compromised. The DSets we care about are those that encompass

well-behaved / ill-behaved	Local property of nodes, independent of other nodes (except for the validity of slice selection).
intact / befouled	Property of nodes given their quorum slices and a particular set of ill-behaved nodes. Befouled nodes are ill-behaved or depend, possibly indirectly, on too many ill-behaved nodes.
correct / failed	Property of nodes given their quorum slices, a concrete protocol, and actual network behavior. The goal of a consensus protocol is to guarantee node correctness whenever possible.

Figure 8: Key properties of FBAS nodes

all ill-behaved nodes, as they help us distinguish nodes that should be guaranteed correct from ones for which such a guarantee is impossible. To this end, we introduce the following terms:

Definition (intact). A node v in an FBAS is *intact* iff there exists a DSet B containing all ill-behaved nodes and such that $v \notin B$.

Definition (befouled). A node v in an FBAS is *befouled* iff it is not intact.

A befouled node v is surrounded by enough failed nodes to block its progress or poison its state, even if v itself is well-behaved. No FBAS can guarantee the correctness of a befouled node. However, an optimal FBAS guarantees that every intact node remains correct. Figure 8 summarizes the key properties of nodes. The following theorems facilitate analysis by showing that the set of befouled nodes is always a DSet in an FBAS with quorum intersection.

Theorem 1. *Let U be a quorum in FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$, let $B \subseteq \mathbf{V}$ be a set of nodes, and let $U' = U \setminus B$. If $U' \neq \emptyset$ then U' is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$.*

Proof. Because U is a quorum, every node $v \in U$ has a $q \in \mathbf{Q}(v)$ such that $q \subseteq U$. Since $U' \subseteq U$, it follows that every $v \in U'$ has a $q \in \mathbf{Q}(v)$ such that $q \setminus B \subseteq U'$. Rewriting with deletion notation yields $\forall v \in U', \exists q \in \mathbf{Q}^B(v)$ such that $q \subseteq U'$, which, because $U' \subseteq \mathbf{V} \setminus B$, means that U' is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. \square

Theorem 2. *If B_1 and B_2 are DSets in an FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ enjoying quorum intersection, then $B = B_1 \cap B_2$ is a DSet, too.*

Proof. Let $U_1 = \mathbf{V} \setminus B_1$ and $U_2 = \mathbf{V} \setminus B_2$. If $U_1 = \emptyset$, then $B_1 = \mathbf{V}$ and $B = B_2$ (a DSet), so we are done. Similarly, if $U_2 = \emptyset$ then $B = B_1$ and we are done. Otherwise, note that by quorum availability despite DSets B_1 and B_2 , U_1 and U_2 are quorums in $\langle \mathbf{V}, \mathbf{Q} \rangle$. By definition, the union of two quorums is also a quorum. Hence $\mathbf{V} \setminus B = U_1 \cup U_2$ is a quorum and we have quorum availability despite B .

We must now show quorum intersection despite B . Let U_a and U_b be any two quorums in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. Let $U = U_1 \cap U_2 = U_2 \setminus B_1$. By quorum intersection of $\langle \mathbf{V}, \mathbf{Q} \rangle$, $U = U_1 \cap U_2 \neq \emptyset$. But then by Theorem 1, $U = U_2 \setminus B_1$ must be a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_1}$. Now consider that $U_a \setminus B_1$ and $U_a \setminus B_2$ cannot both be empty, or else $U_a \setminus B = U_a$ would be. Hence, by Theorem 1, either $U_a \setminus B_1$ is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_1}$, or $U_a \setminus B_2$ is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_2}$, or both. In the former case, note that if $U_a \setminus B_1$ is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_1}$, then by quorum intersection of $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_1}$, $(U_a \setminus B_1) \cap U \neq \emptyset$; since $(U_a \setminus B_1) \cap U = (U_a \setminus B_1) \setminus B_2$, it follows that $U_a \setminus B_2 \neq \emptyset$, making $U_a \setminus B_2$ a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_2}$. By a similar argument, $U_b \setminus B_2$ must be a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^{B_2}$. But then by quorum intersection despite B_2 tells us that $(U_a \setminus B_2) \cap (U_b \setminus B_2) \neq \emptyset$, which is only possible if $U_a \cap U_b \neq \emptyset$. \square

Theorem 3. *In an FBAS with quorum intersection, the set of befouled nodes is a DSet.*

Proof. Let B_{\min} be the intersection of every DSet that contains all ill-behaved nodes. It follows from the definition of intact that a node v is intact iff $v \notin B_{\min}$. Thus, B_{\min} is precisely the set of befouled nodes. By Theorem 2, DSets are closed under intersection, so B_{\min} is also a DSet. \square

5 Federated voting

This section develops a federated voting technique that FBAS nodes can use to agree on a statement. At a high level, the process for agreeing on some statement a involves nodes exchanging two sets of messages. First, nodes *vote* for a . Then, if the vote was successful, nodes *confirm* a , effectively holding a second vote on the fact that the first vote succeeded.

From each node's perspective, the two rounds of messages divide agreement on a statement a into three phases: unknown, accepted, and confirmed. Initially, a 's status is completely *unknown* to a node v — a could end up true, false, or even *stuck* in a permanently indeterminate state. If the first vote succeeds, v may come to *accept* a . No two intact nodes ever accept contradictory statements, so if v is intact and accepts a , then a cannot be false.

For two reasons, however, v accepting a does not make a true. First, the fact that v accepted a does not mean all intact nodes can; a could be stuck for other nodes. Second, if v is befouled, then accepting a means nothing— a may be false at intact nodes. Yet even if v is befouled—which v does not know—the system may still enjoy quorum intersection of well-behaved nodes, in which case for optimal safety v needs greater assurance of a . Holding a second vote addresses both problems. If the second vote succeeds, v moves to the *confirmed* phase in which it can finally deem a true and act on it.

The next few subsections detail the federated voting process. Because voting does not rule out the possibility of stuck statements, Section 5.6 discusses how to cope with them. Section 6 will turn federated voting into a consensus protocol that avoids the possibility of stuck slots for intact nodes.

5.1 Voting with open membership

A correct node in a Byzantine agreement system accepts a statement a only when it knows that other correct nodes will never agree to statements contradicting a . Most protocols employ voting for this purpose. Well-behaved nodes vote for a statement a only if it is valid. Well-behaved nodes also never change their votes. Hence, in centralized Byzantine agreement, it is safe to accept a if a majority of well-behaved nodes—a quorum—has voted for it. We say a statement is *ratified* once it has received the necessary votes.

In a federated setting, we must adapt voting to accommodate open membership. One difference is that a quorum no longer corresponds to a majority of well-behaved nodes. However, the majority requirement primarily serves to ensure quorum intersection of well-behaved nodes, which Section 4.1 already adapted to FBA. Another implication of open membership is that nodes must discover what constitutes a quorum as part of the voting process. To implement quorum discovery, a protocol can include $\mathbf{Q}(v)$ in all messages from v or provide some other means of querying v for $\mathbf{Q}(v)$.

Definition (vote). A node v *votes* for an (abstract) statement a iff

1. v asserts a is valid and consistent with all statements v has accepted, and
2. v asserts it has never *voted against* a —i.e., voted for a statement that contradicts a —and v promises never to vote against a in the future.

Definition (ratify). A quorum U_a *ratifies* a statement a iff every member of U_a votes for a . A node v *ratifies* a iff v is a member of a quorum U_a that ratifies a .

Theorem 4. *Two contradictory statements a and \bar{a} cannot both be ratified in an FBAS that enjoys quorum intersection and contains no ill-behaved nodes.*

Proof. By contradiction. Suppose quorum U_1 ratifies a and quorum U_2 ratifies \bar{a} . By quorum intersection, $\exists v \in U_1 \cap U_2$. Such a v must have illegally voted for both a and \bar{a} , violating the assumption of no ill-behaved nodes. \square

Theorem 5. *Let $\langle \mathbf{V}, \mathbf{Q} \rangle$ be an FBAS enjoying quorum intersection despite B , and suppose B contains all ill-behaved nodes. Let v_1 and v_2 be two nodes not in B . Let a and \bar{a} be contradictory statements. If v_1 ratifies a then v_2 cannot ratify \bar{a} .*

Proof. By contradiction. Suppose v_1 ratifies a and v_2 ratifies \bar{a} . By definition, there must exist a quorum U_1 containing v_1 that ratified a and quorum U_2 containing v_2 that ratified \bar{a} . By Theorem 1, since $U_1 \setminus B \neq \emptyset$ and $U_2 \setminus B \neq \emptyset$, both must be quorums in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$, meaning they ratified a and \bar{a} respectively in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. But $\langle \mathbf{V}, \mathbf{Q} \rangle^B$ enjoys quorum intersection and has no ill-behaved nodes, so Theorem 4 tell us a and \bar{a} cannot both be ratified. \square

Theorem 6. *Two intact nodes in an FBAS with quorum intersection cannot ratify contradictory statements.*

Proof. Let B be the set of befouled nodes. By Theorem 3, B is a DSet. By the definition of DSet, $\langle \mathbf{V}, \mathbf{Q} \rangle$ enjoys quorum intersection despite B . By Theorem 5, two nodes not in B cannot ratify contradictory statements. \square

5.2 Blocking sets

In centralized consensus, liveness is an all-or-nothing property of the system. Either a unanimously well-behaved quorum exists, or else ill-behaved nodes can prevent the rest of the system from accepting new statements. In FBA, by contrast, liveness may differ across nodes. For instance, in the tiered quorum example of Figure 3, if middle tier nodes v_6, v_7, v_8 crash, the leaf tier will be blocked while the top tier and node v_5 will continue to enjoy liveness.

An FBA protocol can guarantee liveness to a node v only if $\mathbf{Q}(v)$ contains at least one quorum slice comprising only correct nodes. A set B of failed nodes can violate this property if B contains at least one member of each of v 's slices. We term such a set B v -blocking, because it has the power to block progress by v .

Definition (v -blocking). Let $v \in \mathbf{V}$ be a node in FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$. A set $B \subseteq \mathbf{V}$ is v -blocking iff it overlaps every one of v 's slices, i.e., $\forall q \in \mathbf{Q}(v), q \cap B \neq \emptyset$.




Figure 9: v_4 voted for \bar{a} , which contradicts ratified statement a .

Theorem 7. Let $B \subseteq \mathbf{V}$ be a set of nodes in FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$. $\langle \mathbf{V}, \mathbf{Q} \rangle$ enjoys quorum availability despite B iff B is not v -blocking for any $v \in \mathbf{V} \setminus B$.

Proof. “ $\forall v \in \mathbf{V} \setminus B, B$ is not v -blocking” is equivalent to “ $\forall v \in \mathbf{V} \setminus B, \exists q \in \mathbf{Q}(v)$ such that $q \subseteq \mathbf{V} \setminus B$.” By the definition of quorum, the latter holds iff $\mathbf{V} \setminus B$ is a quorum or $B = \mathbf{V}$, the exact definition of quorum availability despite B . \square

As a corollary, the DSet of befouled nodes is not v -blocking for any intact v .

5.3 Accepting statements

When an intact node v learns that it has ratified a statement, Theorem 6 tells v that other intact nodes will not ratify contradictory statements. This condition is sufficient for v to accept a , but we cannot make it necessary. Ratifying a statement requires voting for it, and some nodes may have voted for contradictory statements. In Figure 9, for example, v_4 votes for \bar{a} before learning that the other three nodes ratified the contradictory statement a . Though v_4 cannot now vote for a , we would still like it to accept a to be consistent with the other nodes.

A key insight is that if a node v is intact, then no v -blocking set B can consist entirely of befouled nodes. Now suppose B is a v -blocking set and every member of B claims to accept statement a . If v is intact, at least one member of B must be, too. The intact member will not lie about accepting a ; hence, a is true and v can accept it. Of course, if v is befouled, then a might not be true. But a befouled node can accept anything and vacuously not affect the safety of intact nodes.

Definition (accept). An FBAS node v *accepts* a statement a iff it has never accepted a statement contradicting a and it determines that either

1. There exists a quorum U such that $v \in U$ and each member of U either voted for a or claims to accept a , or

2. Each member of a v -blocking set claims to accept a .

Though a well-behaved node cannot vote for contradictory statements, condition 2 above allows a node to *vote* for one statement and later *accept* a contradictory one.

Theorem 8. *Two intact nodes in an FBAS that enjoys quorum intersection cannot accept contradictory statements.*

Proof. Let $\langle \mathbf{V}, \mathbf{Q} \rangle$ be an FBAS with quorum intersection and let B be its DSet of befouled nodes (which exists by Theorem 3). Suppose an intact node accepts statement a . Let v be the first intact node to accept a . At the point v accepts a , only befouled nodes in B can claim to accept it. Since B cannot be v -blocking for an intact v , there must be a slice $q \in \mathbf{Q}(v)$ such that each member of $q \setminus B$ voted for a . In other words, v ratified a in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. Generalizing, any statement accepted by an intact node in $\langle \mathbf{V}, \mathbf{Q} \rangle$ must be ratified in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. Because B is a DSet, $\langle \mathbf{V}, \mathbf{Q} \rangle^B$ enjoys quorum intersection. Because additionally B contains all ill-behaved nodes, Theorem 4 rules out ratification of contradictory statements. \square

5.4 Accepting is not enough

Unfortunately, for nodes to assume the truth of accepted statements would yield suboptimal safety and liveness guarantees in a federated consensus protocol. We discuss the issues with safety and liveness in turn. For intuition, we then explain why these issues are thornier in FBA than in centralized Byzantine agreement.

5.4.1 Safety

Consider an FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ in which the only quorum is unanimous consent, i.e., $\forall v, \mathbf{Q}(v) = \{\mathbf{V}\}$. This ought to be a conservative choice for safety—don't do anything unless everyone agrees. Yet since every node is v -blocking for every v , any node can single-handedly convince any other node to accept arbitrary statements.

The problem is that accepted statements are only safe among intact nodes. But as discussed in Section 4.1, the only condition necessary to guarantee safety is quorum intersection of well-behaved nodes, which might hold even in the case that some well-behaved nodes are befouled. In particular, when $\mathbf{Q}(v) = \{\mathbf{V}\}$, the only DSets are \emptyset and \mathbf{V} , meaning any node failure befouls the whole system. By contrast, quorum intersection holds despite every $B \subseteq \mathbf{V}$.

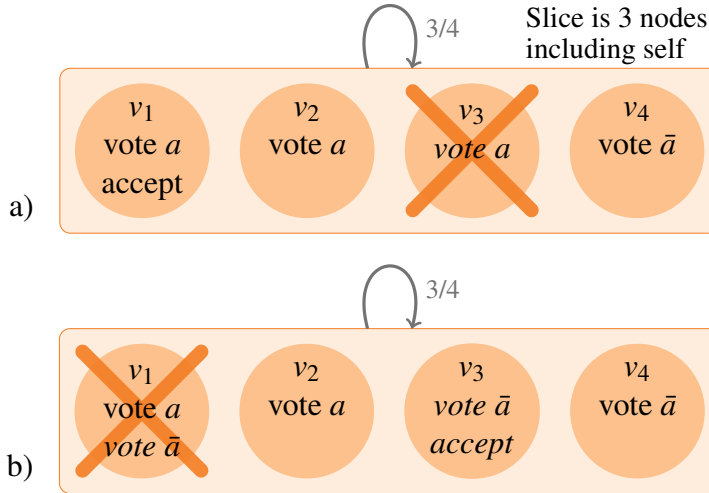


Figure 10: Scenarios indistinguishable to v_2 when v_2 does not see italic messages

5.4.2 Liveness

Another limitation of accepted statements is that other intact nodes may be unable to accept them. This possibility makes reliance on accepted statements problematic for liveness. If a node proceeds to act on a statement because it accepted the statement, other nodes could be unable to proceed in a similar fashion.

Consider Figure 10a, in which node v_3 crashes after helping v_1 ratify and accept statement a . Though v_1 accepts a , v_2 and v_4 cannot. In particular, from v_2 's perspective, the situation depicted is indistinguishable from Figure 10b, in which v_3 voted for \bar{a} and is well-behaved but slow to respond, while v_1 is ill-behaved and sent v_3 a vote for \bar{a} (thereby causing v_3 to accept \bar{a}) while illegally also sending v_2 a vote for a .

To support a protocol-level notion of liveness in cases like Figure 10a, v_1 needs a way to ensure every other intact node can eventually accept a before v_1 acts on a . Once this is the case, it makes sense to say the system agrees on a .

Definition (agree). An FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ *agrees* on a statement a iff, regardless of what subsequently transpires, once sufficient messages are delivered and processed, every intact node will accept a .

5.4.3 Comparison to centralized voting

To understand why the above issues arise in federated voting, consider a centralized Byzantine agreement system of N nodes with quorum size T . Such a system enjoys quorum availability with $f_L = N - T$ or fewer node failures. Since any two quorums share at least $2T - N$ nodes, quorum intersection of well-behaved nodes holds up to $f_S = 2T - N - 1$ Byzantine failures.

Centralized Byzantine agreement systems typically set $N = 3f + 1$ and $T = 2f + 1$ to yield $f_L = f_S = f$, the equilibrium point at which safety and liveness have the same fault tolerance. If safety is more important than liveness, some protocols increase T so that $f_S > f_L$ [30]. In FBA, because quorums arise organically, systems are unlikely to find themselves at equilibrium, making it far more important to protect safety in the absence of liveness.

Now consider a node v that votes against a ratified statement a in a centralized system. If v hears $f_S + 1$ nodes claim a was ratified, v knows that either one of them is well-behaved or all safety guarantees have collapsed. Either way, v can immediately act on a with no loss of safety. The FBA equivalent would be to hear from a set B where B , if deleted, undermines quorum intersection of well-behaved nodes. Identifying such a B is a tricky proposition, first because quorums are discovered dynamically, second because ill-behaved nodes may lie about slices, and third because v does not know which nodes are well-behaved. Instead, we defined federated voting to accept a when a v -blocking set does. The v -blocking property has the advantage of being easily checkable, but is equivalent to hearing from $f_L + 1$ nodes in a centralized system when we really want $f_S + 1$.

To guarantee agreement among all well-behaved nodes in a centralized system, one merely needs $f_L + f_S + 1$ nodes to acknowledge that a statement was ratified. If more than f_L of them fail, we do not expect liveness anyway. If f_L or fewer fail, then we know $f_S + 1$ nodes remain willing to attest to ratification, which will in turn convince all other well-behaved nodes. Again, the reliance on f_S has no easy analogue in the FBA model.

Put another way, at some point nodes need to believe a statement strongly enough to depend on its truth for safety. A centralized system offers two ways to reach this point for a statement a : ratify a first-hand, or reason backwards from $f_S + 1$ nodes claiming a was ratified, figuring safety is hopeless if they have all lied. FBA lacks the latter approach; the only tool it has for safety among well-behaved nodes is first-hand ratification. Since nodes still need a way to overcome votes against ratified statements, we introduced a notion of accepting, but it provides a weaker consistency guarantee limited to intact nodes.

5.5 Statement confirmation

Both limitations of accepted statements stem from the fact that an intact node v may end up voting against a statement a that is nonetheless ratified. After voting against a , v cannot vote for it, preventing v from ever ratifying a . To provide v a means of accepting a after voting against it, the definition of accept has a second criterion based on v -blocking sets. But the second criterion is weaker than ratification, offering no guarantees to befouled nodes that enjoy quorum intersection.

Now if a statement a has the property that no intact node ever votes against it, then we have no need to accept it. Intact nodes can all simply ratify a , and we can require them to do so before acting on a . We call such statements irrefutable.

Definition (irrefutable). A statement a is *irrefutable* in an FBAS if no intact node can ever vote against it.

Theorem 8 tells us that two intact nodes cannot accept contradictory statements. Thus, while some nodes may vote against a statement a that was accepted by an intact node, the *fact* that an intact node accepted a is irrefutable. This suggests holding a second vote to ratify the fact that an intact node accepted a .

Definition (confirm). A quorum U_a in an FBAS *confirms* a statement a iff $\forall v \in U_a$, v claims to accept a . A node *confirms* a iff it is in such a quorum.

Nodes express that they have accepted statement a by stating “accept(a),” an abbreviation of the statement, “An intact node accepted a .” To confirm a means to ratify accept(a). A well-behaved node v can vote for accept(a) only after accepting a , as v cannot assume any other nodes are intact. If v itself is befouled, accept(a) might be false, in which case voting for it may cost v liveness, but a befouled node has no guarantee of liveness anyway.

The next theorem shows that nodes can rely on confirmed statements without losing optimal safety. Theorem 10 then shows that confirmed statements meet the definition of agreement from Section 5.4.2, meaning nodes can rely on confirmed statements without endangering the liveness of intact nodes.

Theorem 9. *Let $\langle \mathbf{V}, \mathbf{Q} \rangle$ be an FBAS enjoying quorum intersection despite B , and suppose B contains all ill-behaved nodes. Let v_1 and v_2 be two nodes not in B . Let a and \bar{a} be contradictory statements. If v_1 confirms a then v_2 cannot confirm \bar{a} .*

Proof. First note that accept(a) contradicts accept(\bar{a})—no well-behaved node can vote for both. Note further that v_1 must ratify accept(a) to confirm a . By Theorem 5, v_2 cannot ratify accept(\bar{a}) and hence cannot confirm \bar{a} . \square

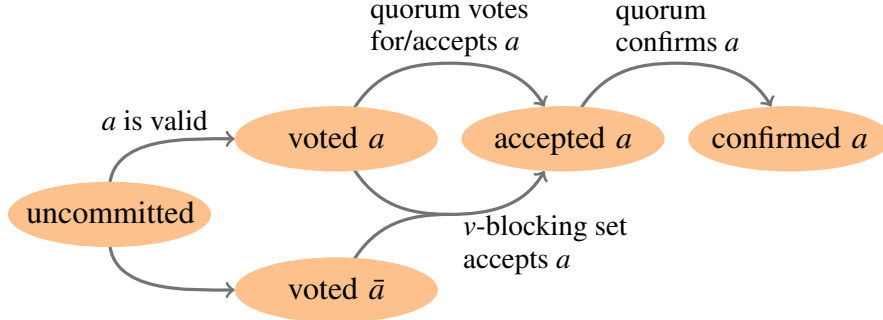


Figure 11: Possible states of an accepted statement a at a single node v

Theorem 10. *If an intact node in an FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ with quorum intersection confirms a statement a , then, whatever subsequently transpires, once sufficient messages are delivered and processed, every intact node will accept and confirm a .*

Proof. Let B be the DSet of befouled nodes and let $U_a \not\subseteq B$ be the quorum through which an intact node confirmed a . Let nodes in $U_a \setminus B$ broadcast $\text{accept}(a)$. By definition, any node v , regardless of how it has voted, accepts a after receiving $\text{accept}(a)$ from a v -blocking set. Hence, these messages may convince additional nodes to accept a . Let these additional nodes in turn broadcast $\text{accept}(a)$ until a point is reached at which, regardless of future communication, no further intact nodes can ever accept a . Once this process is complete, let V^+ be the set of intact nodes that have accepted a , and let $V^- = (\mathbf{V} \setminus B) \setminus V^+$ be the intact nodes that cannot accept a . To show every intact node accepted a , we must show $V^- = \emptyset$.

By Theorem 1, $U_a \setminus B$ is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. Because V^+ is not v -blocking for any $v \in V^-$, then by Theorem 7 either $V^- = \emptyset$ or V^- is a quorum in $\langle \mathbf{V}, \mathbf{Q} \rangle^B$. The latter case leads to a contradiction: Since V^- consists of only intact (and thus well-behaved) nodes, none of them would state $\text{accept}(a)$ without first actually accepting a , meaning $(U_a \setminus B) \cap V^- = \emptyset$. But this is impossible, as quorum intersection despite B (a DSet) tells us the opposite, namely $(U_a \setminus B) \cap V^- \neq \emptyset$. The remaining possibility is that $V^- = \emptyset$. Once every intact accepts a , all can vote to confirm a , which will succeed since intact nodes form a quorum. \square

Figure 11 summarizes the paths an intact node v can take to confirm a . Given no knowledge, v might vote for either a or the contradictory \bar{a} . If v votes for \bar{a} , it cannot later vote for a , but can nonetheless *accept* a if a v -blocking set accepts it. A subsequent quorum of confirmation messages allows v to *confirm* a , which by Theorem 10 means the system agrees on a .




Figure 12: Possible system-wide status of a statement a

5.6 Liveness and neutralization

The central challenge of distributed consensus, whether centralized or not, is that a statement can get stuck in a permanently indeterminate state before the system reaches agreement on it. Hence, a protocol must not attempt to ratify externalized values directly. Should statement “The value of slot i is x ” get stuck, the system will be forever unable to agree on slot i , losing liveness. The solution is to craft the statements in votes carefully. It must be possible to break a stuck statement’s hold on the question we really care about, namely slot contents. We call the process of obsoleting a stuck statement *neutralization*.

More concretely, Figure 12 depicts the potential status a statement a can have system-wide. Initially, the system is *bivalent*, by which we mean there is one sequence of possible events through which all intact nodes will accept a , and another sequence through which all intact nodes will *reject* a (i.e., accept a statement \bar{a} contradicting a). At some point, one of these two outcomes may cease to be possible. If no intact node can ever reject a , we say the system is *a-valent*; conversely if no intact node can ever accept a , we say the system is *\bar{a} -valent*.

At the time an FBAS transitions from bivalent to *a-valent*, there is a possible outcome in which all intact nodes accept a . However, this might not remain the case. Consider a PBFT-like four-node system $\{v_1, \dots, v_4\}$ in which any three nodes constitute a quorum. If v_1 and v_2 vote for a , the system becomes *a-valent*; no three nodes can ratify a contradictory statement. However, if v_3 and v_4 subsequently vote for \bar{a} contradicting a , it also becomes impossible to ratify a . In this case, a ’s state is permanently indeterminate, or *stuck*.

As seen in Figure 10a, even once an intact node accepts a , the system may still fail to reach system-wide agreement on a . However, by Theorem 10, once an intact node confirms a , all intact nodes can eventually come to accept it; hence the

Local state	System-wide status of a
uncommitted	unknown (any)
voted a	unknown (any)
voted \bar{a}	unknown (any)
accepted a	stuck, a -valent, or a agreed
confirmed a	a agreed

Figure 13: What an intact node knows about the status of statement a

system has agreed upon a . Figure 13 summarizes what intact nodes know about the global state of a statement from their own local state.

To preserve the possibility of consensus, a protocol must ensure that every statement is either irrefutable—and hence cannot get stuck—or neutralizable—and hence cannot block progress if stuck. There are two popular approaches to crafting neutralizable statements, the *view-based* approach, pioneered by view-stamped replication [35] and favored by PBFT [11], and the *ballot-based* approach, invented by Paxos [24]. The ballot-based approach may be harder to understand [36]. Compounding confusion, people often call viewstamped replication “Paxos” or assert the two algorithms are the same when they are not [40].

View-based protocols associate the *slots* in votes with monotonically increasing view numbers. Should consensus get stuck on the i th slot in view n , nodes recover by agreeing that view n had fewer than i meaningful slots and moving to a higher view number. Ballot-based protocols associate the *values* in votes with monotonically increasing ballot numbers. Should a ballot get stuck, nodes retry with a higher ballot number, taking care to select values that do not contradict any previous indeterminate ballot. This work takes a ballot-based approach, as doing so makes it easier to do away with the notion of a distinguished primary node or leader. (E.g., leader behavior can be emulated [26].)

6 SCP: A federated Byzantine agreement protocol

This section presents the Stellar Consensus Protocol, SCP. At a high level, the protocol revolves around ballots associated with values. Nodes employ federated voting to commit and abort these ballots. If nodes agree to commit a ballot, consensus terminates and the value associated with the ballot can be externalized for the slot in question. If nodes agree to abort a ballot, its value becomes irrelevant.

If a ballot gets stuck in a state where one or more intact nodes cannot commit or abort it, then nodes try again with a higher ballot, but must associate the new ballot with the same value as the stuck one in case any node believes the stuck ballot was committed.

We begin by describing SCP, then show its correctness, and finally discuss termination. Intuitively, safety results from ensuring that all stuck and committed ballots are associated with the same value. Liveness follows from the fact that a stuck ballot can be neutralized by moving to a higher ballot.

6.1 Description

SCP achieves consensus on slots, but treats each slot independently. It can be viewed as many separate instances of a single-slot consensus protocol (analogous to the more whimsically named “single-decree synod” in Paxos [24]). Unless otherwise stated, all discussion should be interpreted within the context of a single slot. In particular, each ballot belongs to one slot, even if the slot is implicit.

SCP revolves around committing and aborting ballots associated with values. A ballot b is a pair of the form $b = \langle n, x \rangle$, where x is its associated value. The ballot should therefore be viewed as a referendum on externalizing x , not a choice between multiple candidate values. The value n in ballot $b = \langle n, x \rangle$ is a counter to ensure higher ballot numbers are always available. (Acceptable n s can be bounded by the current time to avoid exhausting the counter.) Ballots are totally ordered with n more significant than x . For convenience, a special invalid null ballot $\mathbf{0}$ is less than all other ballots.

We speak of committing and aborting a ballot b as a shorthand for using federated voting to agree on the statements *commit* b and *abort* b , respectively. *commit* and *abort* are contradictory, so a well-behaved node may vote for at most one of them. Section 5 would use “ $\overline{\text{commit}}$ ” as the opposite of *commit*; *abort* is just a more intuitive notation for the same statement. Note that *commit* b and *abort* b are conceptual statements about ballot b . SCP messages concisely encode sets of such statements over multiple ballots.

Because at most one value can be chosen for a given slot, all committed and stuck ballots must contain the same value. Roughly speaking, this means *commit* statements are invalid if they conflict with lower-numbered unaborted ballots.

Definition (compatible). Two ballots $b_1 = \langle n_1, x_1 \rangle$ and $b_2 = \langle n_2, x_2 \rangle$ are *compatible*, written $b_1 \sim b_2$, iff $x_1 = x_2$ and *incompatible*, written $b_1 \not\sim b_2$, iff $x_1 \neq x_2$. We also write $b_1 \lesssim b_2$ or $b_2 \gtrsim b_1$ iff $b_1 < b_2$ (or equivalently $b_2 > b_1$) and $b_1 \approx b_2$.

Variable	Meaning
b	Current ballot v is attempting to prepare and commit (Invariant: if $c \neq \mathbf{0}$ then $b \sim c$)
p	Highest ballot accepted as prepared, if any, otherwise $\mathbf{0}$
p'	Highest ballot accepted as prepared such that $p' \lesssim p$, otherwise $\mathbf{0}$
c	Highest ballot for which v has voted <i>commit</i> c without later accepting <i>abort</i> c , or $\mathbf{0}$ if there is no such ballot
M	Set containing the latest message seen from each node
φ	current phase: one of PREPARE, FINISH, or EXTERNALIZE

Figure 14: State maintained by each node v for each slot i

Definition (prepared). A ballot b is *prepared* iff every statement in the following set is true: $\{ \text{abort } b_{\text{old}} \mid b_{\text{old}} \lesssim b \}$.

More precisely, then, *commit* b is valid—and hence can legitimately appear in votes—only if b is prepared. Nodes can ensure b is prepared through federated voting on the corresponding *abort* statements. It is convenient to vote on these statements en masse, so wherever we write “ b is prepared,” the surrounding context applies to the whole set of *abort* statements. In particular, a node votes, accepts, or confirms that b is prepared iff it votes for, accepts, or confirms, respectively, all of these *aborts*.

Informally, to commit a ballot b and externalize its value, SCP nodes must first agree b is prepared, then agree on *commit* b . Each of the prepare and commit steps employs federated voting. Since federated voting itself involves two messages, the whole SCP protocol can commit a ballot in four rounds. However, for simplicity, the protocol has only two messages. The first three rounds use the first message, and the final round the second.

Figure 14 illustrates the state a node v must store for slot i . It stores its current ballot b , the highest ballot c (if any) it has voted to *commit* and for which it has not subsequently accepted an *abort*, the two most recent incompatible ballots it has prepared (p, p'), the latest message received from each node, and a phase φ . Ballots b, p , and p' must be non-decreasing. In addition, if $c \neq \mathbf{0}$ —meaning v may have participated in ratifying *commit* c —code must ensure $b \sim c$. This invariant guarantees PREPARE messages are always valid for the current ballot b .

Figure 15 lists the two concrete messages that constitute SCP. A node v ini-

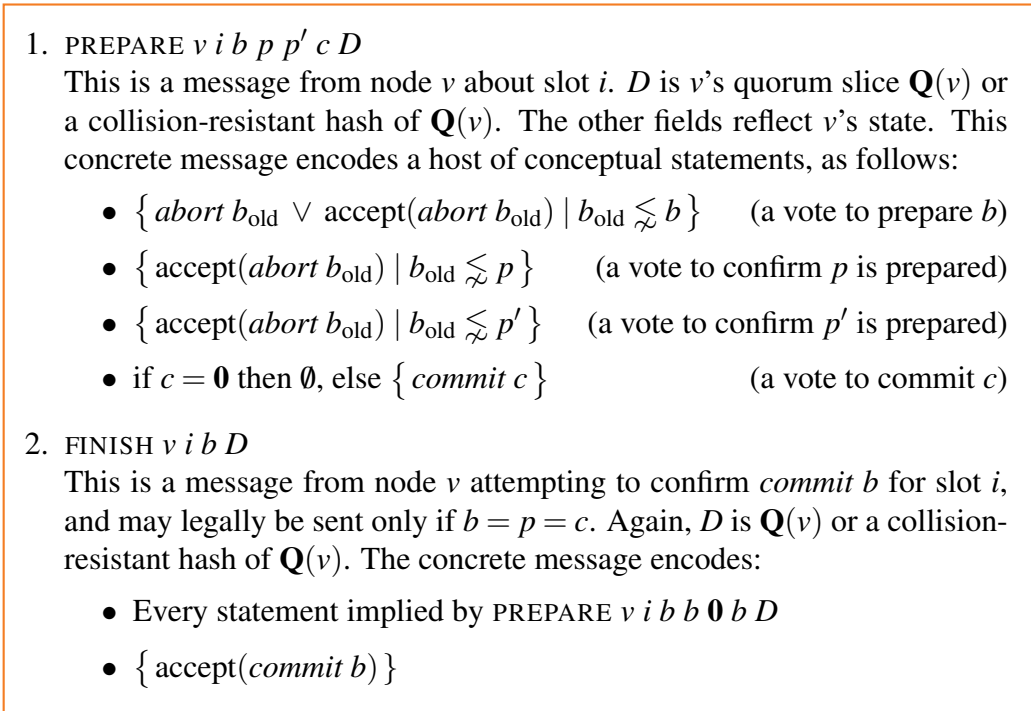


Figure 15: Messages in SCP

tializes its state by setting b to a ballot with a value it hopes to externalize, as well as $p \leftarrow \mathbf{0}$, $p' \leftarrow \mathbf{0}$, $c \leftarrow \mathbf{0}$, $M \leftarrow \emptyset$, and $\varphi \leftarrow \text{PREPARE}$. It then repeatedly exchanges messages with peers, sending whichever message is indicated by φ (initially PREPARE). Upon adding a message to M , v updates its state as follows:

1. If $\varphi = \text{PREPARE}$ and a message allows v to accept that new ballots are prepared by either of accept's two criteria, then update p and p' accordingly. If afterwards $c \neq \mathbf{0}$ but $p \gtrsim c$ or $p' \gtrsim c$, then set $c \leftarrow \mathbf{0}$.
2. If $b \neq c$, $b = p$, $\varphi = \text{PREPARE}$, and v confirms b is prepared, then set $c \leftarrow b$.
3. If $b = p = c$, $\varphi = \text{PREPARE}$, and v accepts *commit* b , set $\varphi \leftarrow \text{FINISH}$.
4. If $b = p = c$ and $\varphi = \text{FINISH}$ and v confirms *commit* b , then set $\varphi \leftarrow \text{EXTERNALIZE}$ and externalize the value in b .

While $b \neq c$, the above protocol implements federated voting to confirm b is prepared; once $b = p = c$, it implements federated voting on *commit* b . Note all

messages from a node are totally ordered. All PREPARE messages precede all FINISH messages. PREPARE messages are ordered by $\langle b, p, p', c \rangle$, with b the most significant and c the least significant. FINISH messages are ordered by the ballot b . This ordering makes it possible to keep only the latest ballot from each node in M without relying on message delivery order to know which is latest.

A few details of the protocol merit explanation. First, note that in PREPARE, the *abort* messages derived from b do not specify whether v is voting for *abort* b_{old} or *accept*(*abort* b_{old}). The distinction between voting for and confirming is unimportant for the definition of *accept*. Glossing over the distinction allows v to forget about old ballots it voted to *commit* (and hence cannot vote to *abort*), so long as it accepted an *abort* message for them. Though including four ballots in every PREPARE may seem excessive, p and p' help unblock nodes that may have a stale c that prevents them from updating b . It follows from the definition of prepare that the two highest incompatible ballots a node has accepted as prepared subsume all ballots the node has accepted as prepared.

6.2 Correctness

Theorem 11. *Let v_1 and v_2 be nodes in an FBAS enjoying quorum intersection despite B , where $v_1 \notin B$ and $v_2 \notin B$. If, after running SCP, v_1 externalizes value x_1 for slot i and v_2 externalizes value x_2 for slot i , then $x_1 = x_2$.*

Proof. In order to externalize these values, v_1 must confirm *commit* b_1 where $b_1 = \langle n_1, x_1 \rangle$, and similarly v_2 for $b_2 = \langle n_2, x_2 \rangle$. If $x_1 \neq x_2$, then $b_1 \approx b_2$. In that case, assume without loss of generality that $b_1 \lesssim b_2$. By virtue of confirming *commit* b_2 , v_2 must have set $c \leftarrow b_2$ at some point, which could only happen in condition 2 on page 28, where b_2 is confirmed prepared. But since $b_1 \lesssim b_2$, that implies b_2 confirmed *abort* b_1 . Theorem 9 tells us this is impossible since b_1 confirmed *commit* b_1 . Hence it must be that $x_1 = x_2$. \square

Theorem 12. *Given a slot in an SCP system with quorum intersection, there is always a sequence of events through which intact nodes can reach agreement on and commit a value. (I.e., there are no blocked states for intact nodes.)*

Proof. Because befouled nodes form a DSet, quorum availability despite that DSet means intact nodes can ratify statements without the cooperation of befouled nodes. Suppose no intact node has voted to *commit* any ballot. Intact nodes can then take any ballot higher than any they have voted to *abort* and work through the protocol to commit that ballot. Lack of response from befouled nodes will

not block intact nodes from assembling quorums and making progress. (We are only trying to prove some sequence of events exists, so the fact that in a different sequence the network could delay intact nodes and/or prioritize messages from befouled ones need not concern us.)

Otherwise, let $c = \langle n, x \rangle$ be the highest ballot in any *commit* vote by an intact node. An intact node cannot vote *commit* c before confirming *abort* b_{old} for all $b_{\text{old}} \prec c$. By Theorem 10, once these *aborts* are confirmed, every intact node can come to confirm them and hence confirm c is prepared. Let $b = \langle n', x \rangle \geq c$ be a ballot compatible with c that no intact node has voted to *abort*. Once nodes confirm c is prepared, because no intact node has voted to *commit* any ballot over c , intact nodes can also vote and confirm *abort* statements for all ballots between c and b , at which point they have confirmed b is prepared. From this point, again intact nodes can complete the protocol regardless of lack of response from befouled nodes. \square

6.3 Termination

An important implementation issue is making sure that SCP terminates in practice, even if it is not possible to guarantee deterministic termination for a general asynchronous system [19].

First, note that some applications of SCP are not the general case. Because slots need not be totally ordered, some applications will typically admit only one plausible ballot per slot, making protocol messages effectively irrefutable. For example, in a payment network, slots might be $\langle \text{account}, \text{serial number} \rangle$ pairs, where serial numbers are consecutive integers and ballots must be signed by the account holder. Such a scheme totally orders all payments from a particular account, but not across accounts. As long as each account holder signs only one ballot per serial number, there is no possibility of preemption by conflicting values. Should an account holder lose state and accidentally duplicate a serial number, the person can easily recover by signing a higher ballot number that subsumes one of the existing ballots. Such an approach also works for some certificate transparency applications, when each namespace authority has an independent history.

In other applications, there may be multiple candidate ballots for each slot. For example, the Stellar financial network, because of order books, cannot treat account histories independently and instead agrees on batches of transactions for each slot. One strategy is to cycle between a synchronous phase, that assumes all intact nodes will respond within a certain time, and the asynchronous SCP protocol [25]. If the synchronous phase succeeds, the asynchronous one can guarantee

safety. If the synchronous phase fails, the asynchronous phase will likely fail as well, but then one can simply restart the synchronous phase with a longer timeout.

A synchronous protocol that always terminates but is not always safe is sometimes known as a *conciliator* [3]. Here we suggest one strategy for a federated conciliator. Let H be a cryptographic hash function. For a given slot i and ballot number n , define the *permuted id* of a node v to be $H(i, n, v)$. Define the *slice weight* of node v_2 at node v_1 to be the fraction of v_1 's slices in which v_2 is the node with the highest permuted id. To attempt synchronous agreement, each node v starts with a candidate value x . Then, for a fixed period of time it overwrites x with the candidate value of the node with the highest slice weight it hears from, further propagating the new value if necessary. Slice weight ties can be broken by taking the node with the higher permuted id. Of course, if v has the highest permuted id, it may not need to overwrite x .

With finite probability, the slice weight technique will propagate the value x of the node with the highest permuted id everywhere. At that point, if x does not conflict with unaborted past commit messages, nodes can use SCP on ballot $\langle n, x \rangle$. Because of the federated structure of quorums, however, the slice weight technique may not work, even given timely message delivery and well-behaved nodes. In that case, SCP may fail, but nodes can restart the conciliator process with a higher ballot number n . There is a finite probability that an intact and central node will end up with the highest permuted id, meaning if timeouts are adjusted appropriately each cycle, consensus can be reached with probability 1.

One advantage of the slice weight technique is that the node creating the final value of x can insert a timestamp or other non-deterministic data that other nodes would not have been likely to derive independently. Moreover, if the system is attempting consensus on a batch of transactions, other nodes may be able to send transactions directly to the node with the highest slice weight, rather than broadcasting them everywhere. In this case, SCP would likely be used on a cryptographic hash of a batch of transactions, and the node aggregating transactions can use peer-to-peer multicast [10] to disseminate them. Of course, if the node with the highest permuted id has failed, the ballot is also likely to fail. But permuted ids change with each new value of n .

7 Limitations

SCP can only guarantee safety when nodes choose adequate quorum slices. Section 3.2 discusses why we can reasonably expect them to do so. Nonetheless,

when security depends upon a user-configurable parameter, there is always the possibility people will set it wrong.

Even guaranteed safety, when people set quorum slices right, does not rule out other security issues that may arise in a federated system. For example, Byzantine nodes may attempt to filter transactions on the input side of SCP while otherwise producing the correct output, thereby delaying a victim transaction. In a financial market, widely-trusted nodes could leverage their position in the network to gain information with which to engage in front running or other unethical activities.

Though SCP's safety is optimal, its performance and communication latency are not. In the common case that nodes have not previously voted to *commit* ballots incompatible with the current one, it is possible to reduce the number of communication rounds from four to three. An earlier version of SCP did so, but the protocol description was more complex. First, it required nodes to cache and retransmit signed messages previously sent by failed nodes. Second, it was no longer possible to gloss over the distinction between votes and confirmations of *abort* statements in PREPARE messages, so nodes had to send around potentially unbounded lists of exceptions to their *abort* votes.

SCP guarantees safety to nodes with quorum intersection, even when those nodes are befouled. However, such befouled nodes may be blocked. The protocol has no way to unblock them. Obviously an important step is to update \mathbf{Q} to replace failed peers. Unfortunately, updating \mathbf{Q} may be insufficient to unblock previously befouled nodes that were tricked into voting for conflicting *commit* messages. In such situations, a node v may need to disavow past votes, which it can only do by rejoining the system under a new node id v' . There may exist a way to automate such recovery when \mathbf{Q} is updated.

The FBA model requires continuity of participants over time. Should all nodes simultaneously and permanently leave, restarting consensus would require central coordination or human-level agreement. By contrast, a proof-of-work system such as Bitcoin could undergo sudden complete turnover yet continue to operate with little human intervention. On the other hand, if nodes do return, an FBAS can recover from an arbitrarily long outage, while a proof-of-work scheme would face the possibility of an attacker working on a fork during the outage.

An intriguing possibility for the FBA model is to leverage federated voting to mediate tussles [13]. For instance, could nodes vote to upgrade an application-level protocol running on top of SCP? What about periodically voting on parameters such as transaction fees in response to spam levels? SCP currently has no such capabilities. There are significant challenges to synchronizing policy votes (which may fail or get stuck) with the main transaction stream without jeopardiz-

ing liveness. But the appeal of decentralized policy configuration is tantalizing.

8 Summary

Byzantine agreement has long enabled distributed systems to achieve consensus with efficiency, standard cryptographic security, and flexibility in designating trusted participants. More recently, Bitcoin introduced the revolutionary notion of decentralized consensus, leading to many new systems and research challenges. This paper introduces the notion of federated Byzantine agreement (FBA), which achieves decentralized consensus while preserving the traditional benefits of Byzantine agreement. The key distinction between FBA and prior Byzantine agreement systems is that FBA forms quorums from participants' individual trust decisions, allowing an organic growth model similar to that of the Internet. The Stellar Consensus Protocol (SCP) is a construction for FBA that achieves optimal resilience against ill-behaved participants.

Acknowledgments

Jed McCaleb inspired this work and provided feedback, terminology suggestions, and help thinking through numerous conjectures. Jessica Collier helped write the paper. Stan Polu created the first implementation of SCP and provided invaluable corrections, suggestions, simplifications, and feedback in the process. Jelle van den Hooff provided the key idea to restructure the paper around quorum intersection and federated voting, as well as other crucial suggestions for terminology, organization, and presentation. Nicolas Barry, Bekki Bolthouse, Mike Hamburg, Graydon Hoare, Joyce Kim, Robert Morris, Lucas Ryan, and Katherine Tom slogged through drafts of the paper, identifying errors and sources of confusion as well as providing helpful suggestions. Eva Gantz provided helpful motivation and references. Winnie Lim provided guidance on figures. Finally, the author would like to thank the whole Stellar team for their support, feedback, and encouragement.

Disclaimer

Professor Mazières's contribution to this publication was as a paid consultant, and was not part of his Stanford University duties or responsibilities.

References

- [1] Fraudulent digital certificates could allow spoofing. Microsoft Security Advisory 2798897, January 2013. <https://technet.microsoft.com/en-us/library/security/2798897.aspx>.
- [2] Eduardo A. Alchieri, Alysson Neves Bessani, Joni Silva Fraga, and Fabíola Greve. Byzantine consensus with unknown participants. In *Proceedings of the 12th International Conference on Principles of Distributed Systems*, pages 22–40, December 2008.
- [3] James Aspnes. A modular approach to shared-memory consensus, with applications to the probabilistic-write model. In *Proceedings of the 29th Symposium on Principles of Distributed Computing*, pages 460–467, July 2010.
- [4] Rachel Banning-Lover. Boatfuls of cash: how do you get money into fragile states? February 2015. <http://www.theguardian.com/global-development-professionals-network/2015/feb/19/boatfuls-of-cash-how-do-you-get-money-into-fragile-states>.
- [5] David Basin, Cas Cremers, Tiffany Hyun-Jin Kim, Adrian Perrig, Ralf Sasse, and Pawel Szalachowski. ARPKI: Attack resilient public-key infrastructure. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 382–393, November 2014.
- [6] Michael Ben-Or. Another advantage of free choice (extended abstract): Completely asynchronous agreement protocols. In *Proceedings of the 2nd Symposium on Principles of Distributed Computing*, pages 27–30, June 1983.
- [7] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. Research perspectives and challenges for bitcoin and cryptocurrencies. In *Proceedings of the 36th IEEE Symposium on Security and Privacy*, May 2015.
- [8] Gabriel Bracha and Sam Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM*, 32(4):824–840, October 1985.
- [9] Danny Bradbury. Feathercoin hit by massive attack, June 2013. <http://www.coindesk.com/feathercoin-hit-by-massive-attack/>.
- [10] Miguel Castro, Peter Druschel, Anne-Marie Kermarrec, Animesh Nandi, Antony Rowstron, and Atul Singh. Splitstream: High-bandwidth multicast in cooperative environments. In *Proceedings of the 19th Symposium on Operating Systems Principles*, pages 298–313, October 2003.

- [11] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*, pages 173–186, February 1999.
- [12] CGAP. Making money transfers work for microfinance institutions, March 2008. <http://www.cgap.org/sites/default/files/CGAP-Technical-Guide-Making-Money-Transfers-Work-for-Microfinance-Institutions-A-Technical-Guide-to-Developing-and-Delivering-Money-Transfers-Mar-2008.pdf>.
- [13] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in cyberspace: Defining tomorrow’s internet. *IEEE/ACM Transactions on Networking*, 13(3):462–475, June 2005.
- [14] crazyearner. Terracoin attack over 1.2TH attack confirmd [sic], July 2013. <https://bitcointalk.org/index.php?topic=261986.0>.
- [15] Kourosh Davarpanah, Dan Kaufman, and Ophelie Pubellier. NeuCoin: the first secure, cost-efficient and decentralized cryptocurrency, March 2015. <http://www.neucoin.org/en/whitepaper/download>.
- [16] John R. Douceur. The Sybil attack. In *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, pages 251–260, March 2002.
- [17] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, April 1988.
- [18] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable, November 2013. <http://arxiv.org/abs/1311.0243>.
- [19] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985.
- [20] Ghassan O. Karame, Elli Androulaki, and Srdjan Capkun. Double-spending fast payments in bitcoin. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 906–917, October 2012.
- [21] Tiffany Hyun-Jin Kim, Lin-Shung Huang, Adrian Perring, Collin Jackson, and Virgil Gligor. Accountable key infrastructure (AKI): A proposal for a public-key validation infrastructure. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 679–690, May 2013.

- [22] Sunny King and Scott Nadal. PPCoin: Peer-to-peer crypto-currency with proof-of-stake, August 2012.
<http://peercoin.net/assets/paper/peercoin-paper.pdf>.
- [23] Jae Kwon. Tendermint: Consensus without mining, 2014.
<http://tendermint.com/docs/tendermint.pdf>.
- [24] Leslie Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, May 1998.
- [25] Leslie Lamport. Brief announcement: Leaderless byzantine paxos. In *Proceedings of the 25th International Conference on Distributed Computing*, pages 141–142, September 2011.
- [26] Leslie Lamport. Byzantizing paxos by refinement. In *Proceedings of the 25th International Conference on Distributed Computing*, pages 211–224, September 2011.
- [27] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.
- [28] Adam Langley. Maintaining digital certificate security, March 2015.
<http://googleonlinesecurity.blogspot.com/2015/03/maintaining-digital-certificate-security.html>.
- [29] Ben Laurie, Adam Langley, and Emilia Kasper. Certificate transparency. RFC 6962, Internet Engineering Task Force (IETF), June 2013.
<http://tools.ietf.org/html/rfc6962>.
- [30] Jinyuan Li and David Mazières. Beyond one-third faulty replicas in Byzantine fault tolerant systems. In *Proceedings of the 4th Symposium on Networked Systems Design and Implementation*, pages 131–144, April 2007.
- [31] Marcela S. Melara, Aaron Blankstein, Joseph Bonneau, Michael J. Freedman, and Edward W. Felten. CONIKS: A privacy-preserving consistent key service for secure end-to-end communication. Cryptology ePrint Archive, Report 2014/1004, December 2014. <http://eprint.iacr.org/2014/1004>.
- [32] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
<http://bitcoin.org/bitcoin.pdf>.
- [33] William B. Norton. The art of peering: The peering playbook, August 2010.
<http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html>.

- [34] Karl J. O’Dwyer and David Malone. Bitcoin mining and its energy footprint. In *Irish Signals and Systems Conference*, pages 280–285, Limerick, Ireland, June 2014.
- [35] Brian M. Oki and Barbara H. Liskov. Viewstamped replication: A new primary copy method to support highly-available distributed systems. In *Proceedings of the 7th Symposium on Principles of Distributed Computing*, pages 8–17, August 1988.
- [36] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference*, pages 305–319, June 2014.
- [37] Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, April 1980.
- [38] Claire Provost. Why do africans pay the most to send money home? January 2013. <http://www.theguardian.com/global-development/2013/jan/30/africans-pay-most-send-money>.
- [39] David Schwartz, Noah Youngs, and Arthur Britto. The Ripple protocol consensus algorithm, 2014. https://ripple.com/files/ripple_consensus_whitepaper.pdf.
- [40] Robbert van Renesse, Nicolas Schiper, and Fred B. Schneider. Vive la différence: Paxos vs. Viewstamped Replication vs. Zab. *IEEE Transactions on Dependable and Secure Computing*, September 2014.
- [41] World Bank. Global financial development report 2014: Financial inclusion, 2014. http://siteresources.worldbank.org/EXTGLOBALFINREPORT/Resources/8816096-1361888425203/9062080-1364927957721/GFDR-2014_Complete_Report.pdf.

A Glossary of notation

Notation	Name	Definition
iff		An abbreviation of “if and only if.”
$f(x)$	function	The result of calculating function f on argument x .
\bar{a}	complement	An overbar connotes the opposite, i.e., \bar{a} is the opposite of a .
$\langle a_1, \dots, a_n \rangle$	tuple	A structure (compound value) with field values a_1, \dots, a_n .
$A \wedge B$	logical and	Both A and B are true.
$A \vee B$	logical or	At least one, possibly both, of A and B are true.
$\exists e, C(e)$	there exists	There is at least one value e for which condition $C(e)$ is true.
$\forall e, C(e)$	for all	$C(e)$ is true of every value e .
$\{a, b, \dots\}$	set	A set containing the listed elements (a, b, \dots) .
$\{e \mid C(e)\}$	set-builder	The set of all elements e for which $C(e)$ is true.
\emptyset	empty set	The set containing no elements.
$e \in S$	element of	Element e is a member of set S .
$A \subseteq B$	subset	Every member of set A is also a member of set B .
$A \subsetneq B$	strict subset	$A \subseteq B$ and $A \neq B$.
2^A	powerset	The set of sets containing every possible combination of members of A , i.e., $2^A = \{B \mid B \subseteq A\}$.
$A \cup B$	union	The set containing all elements that are members of A or members of B , i.e., $A \cup B = \{e \mid e \in A \vee e \in B\}$.
$A \cap B$	intersection	The set containing all elements that are members of both A and B , i.e., $A \cap B = \{e \mid e \in A \wedge e \in B\}$.
$A \setminus B$	set difference	The set containing every element of A that is not a member of B , i.e., $A \setminus B = \{e \mid e \in A \wedge e \notin B\}$.
/	not	Negates a symbol’s meaning. E.g., $e \notin A$ means $e \in A$ is false, while $\nexists e, C(e)$ means no e exists such that $C(e)$ is true.