



Security Analysis of Ripple Consensus

Ignacio Amores-Sesar 

University of Bern, Switzerland
ignacio.amores@inf.unibe.ch

Christian Cachin 

University of Bern, Switzerland
christian.cachin@inf.unibe.ch

Jovana Mičić 

University of Bern, Switzerland
jovana.micic@inf.unibe.ch

Abstract

The Ripple network is one of the most prominent blockchain platforms and its native XRP token currently has one of the highest cryptocurrency market capitalizations. The *Ripple consensus protocol* powers this network and is generally considered to a Byzantine fault-tolerant agreement protocol, which can reach consensus in the presence of faulty or malicious nodes. In contrast to traditional Byzantine agreement protocols, there is no global knowledge of all participating nodes in Ripple consensus; instead, each node declares a list of other nodes that it trusts and from which it considers votes.

Previous work has brought up concerns about the liveness and safety of the consensus protocol under the general assumptions stated initially by Ripple, and there is currently no appropriate understanding of its workings and its properties in the literature. This paper closes this gap and makes two contributions. It first provides a detailed, abstract description of the protocol, which has been derived from the source code. Second, the paper points out that the abstract protocol may violate safety and liveness in several simple executions under relatively benign network assumptions.

2012 ACM Subject Classification Theory of computation → Cryptographic protocols; Software and its engineering → Distributed systems organizing principles

Keywords and phrases Ripple, Blockchain, Quorums, Consensus

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2020.10

Funding This work has been funded by the Swiss National Science Foundation (SNSF) under grant agreement Nr. 200021_188443 (Advanced Consensus Protocols).

1 Introduction

Ripple is one of the oldest and most established blockchain networks; its XRP token is ranked fourth in market capitalization in October 2020. The Ripple network is primarily aimed at fast global payments, asset exchange, and settlement. Its distributed consensus protocol is implemented by a peer-to-peer network of validator nodes that maintain a history of all transactions on the network [24]. Unlike Nakamoto’s consensus protocol [21] in Bitcoin or Ethereum, the Ripple consensus protocol does not rely on “mining,” but uses a voting process based on the identities of its validator nodes to reach consensus. This makes Ripple much more efficient than Bitcoin for processing transactions (up to 1500 transactions per second) and lets it achieve very low transaction settlement times (4–5 seconds).

However, Ripple’s consensus protocol does not follow the established models and algorithms for *Byzantine agreement* [22, 15] or *Byzantine fault-tolerant (BFT) consensus* [8]. Those systems start from a common set of nodes that are communicating with each other to reach consensus and the corresponding protocols have been investigated for decades. Instead, the Ripple consensus protocol introduces the idea of subjective validators, such that every



© Ignacio Amores-Sesar, Christian Cachin, and Jovana Mičić;
licensed under Creative Commons License CC-BY

24th International Conference on Principles of Distributed Systems (OPODIS 2020).

Editors: Quentin Bramas, Rotem Oshman, and Paolo Romano; Article No. 10; pp. 10:1–10:16

Leibniz International Proceedings in Informatics



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

10:2 Security Analysis of Ripple Consensus


node declares some *trusted validators* and effectively communicates only with those nodes for reaching agreement on transactions. With this mechanism, the designers of Ripple aimed at opening up membership in the set of validator nodes compared to BFT consensus. The trusted validators of a node are defined by a *Unique Node List (UNL)*, which plays an important role in the formalization of the protocol. Every node maintains a static UNL in its configuration file and considers only the opinions of nodes in its UNL during consensus. Figure 1 shows an example network, where two UNLs are defined: $UNL_1 = \{1, 2, 3, 4\}$ and $UNL_2 = \{3, 4, 5, 6\}$; for instance, nodes 1, 2 and 3 may trust UNL_1 , and nodes 4, 5 and 6 may trust UNL_2 .

Consensus in Ripple aims at delivering the transactions submitted by clients to all participating nodes in a common global order, despite faulty or malicious (Byzantine) nodes [27]. This ensures that the sequence of transactions, which are grouped into so-called *ledgers* and then processed by each node, is the same for all nodes. Hence, the states of all correct nodes remain synchronized, according to the blueprint of state-machine replication [26].

Cachin and Vukolić [7] have earlier pointed out that it is important to formally assess the properties of blockchain consensus protocols. Unfortunately, many systems have been designed and were deployed without following the agreed-on principles on protocol analysis from the literature. Ripple is no exception to this, as we show in this work.

Specifically, we focus on two properties that every sound protocol must satisfy [1]: *safety* and *liveness*. Safety means that nothing “bad” will ever happen, and liveness means that something “good” eventually happens. Safety ensures that the network does not fork or double-spend a token, for instance. A violation of liveness would mean that the network stops making progress and halts processing transactions, which creates as much harm as forking.

This work first presents a complete, abstract description of the Ripple consensus protocol (Section 3). The model has been obtained directly from the source code. It is formulated in the language spoken by designers of consensus protocols, in order to facilitate a better understanding of the properties of Ripple consensus. No formal description of Ripple consensus with comparable technical depth has been available so far (apart from the source itself).



■ **Figure 1** Example of a Ripple network configuration with six nodes and two UNLs, $UNL_1 = \{1, 2, 3, 4\}$ and $UNL_2 = \{3, 4, 5, 6\}$. Nodes 1, 2, and 3 (white) trust UNL_1 , and nodes 4, 5, and 6 (black) trust UNL_2 . Notice that nodes 3 and 4 have more influence than the rest of nodes since they are in the intersection of both UNLs.

Second, we exhibit examples of how safety and liveness may be violated in executions of the Ripple consensus protocol (Sections 4 and 5). In particular, the *network may fork* under the standard condition on UNL overlap stated by Ripple and in the presence of a constant *fraction* of Byzantine nodes. The malicious nodes may simply send conflicting messages to correct nodes and delay the reception of other messages among correct nodes. Furthermore, the consensus protocol may *lose liveness* even if all nodes have the same UNL and there is only *one* Byzantine node. If this would occur, the system has to be restarted manually.

Given these findings, we conclude that the consensus protocol of the Ripple network is brittle and does not ensure consensus in the usual sense. It relies heavily on synchronized clocks, timely message delivery, the presence of a fault-free network, and an a-priori agreement on common trusted nodes. The role of the UNLs, their overlap, and the creation of global consensus from subjective trust choices remain unclear. If Ripple instead had adopted a standard BFT consensus protocol [5], as done by Tendermint [4], versions of Hyperledger Fabric [2], Libra [16] or Concord [13], then the Ripple network would resist a much wider range of corruptions, tolerate temporary loss of connectivity, and continue operating despite loss of synchronization.

2 Related work

Despite Ripple’s prominence and its relatively high age among blockchain protocols – the system was first released in 2012 – there are only few research papers investigating the Ripple consensus protocol compared to the large number of papers on Bitcoin. The original Ripple white paper of 2014 [27] describes the UNL model and illustrates some ideas behind the protocol. It claims that under the assumption of requiring an 80%-quorum for declaring consensus, the intersection between the UNLs of any two nodes u and v should be larger than 20% of the size of the larger of their UNLs, i.e.,

$$|UNL_u \cap UNL_v| \geq \frac{1}{5} \max\{|UNL_u|, |UNL_v|\}.$$

The only earlier protocol analysis in the scientific literature of which we are aware was authored by Armknecht et al. in 2015 [3]. This work analyzes the Ripple consensus protocol and outlines the security and privacy of the network compared to Bitcoin. The authors prove that a 20%-overlap, as claimed in the white paper, cannot be sufficient for reaching consensus and they increase the bound on the overlap to at least 40%, i.e.,

$$|UNL_u \cap UNL_v| > \frac{2}{5} \max\{|UNL_u|, |UNL_v|\}$$

In a preprint of 2018, Chase and MacBrough [10] further strengthen the required UNL overlap. They introduce a high-level model of the consensus protocol and describe some of its properties, but many details appear unclear or are left out. This work concludes that the overlap between UNLs should actually be larger than 90%. The paper also gives an example with 102 nodes that shows how liveness can be violated, even if the UNLs overlap almost completely (by 99%) and there are no faulty nodes. The authors conclude that manual intervention would be needed to resurrect the protocol after this.

Christodoulou et al. [11] further investigate the decentralization of a Ripple network by running simulations with different configurations. They observe how the convergence time and the so-called “Network Health Indicator,” a synthetic measure computed by a tool available from Ripple, change depending on the overlap between the nodes’ UNLs. Their experiments suggest that a large UNL overlap is required only with more than 20% of malicious nodes in the network. They also indicate a possibility for a dynamic determination of an optimal UNL overlap.

An analysis whose goal is similar to that of our work has been conducted by Mauri et al. [19]. Based on the source code, they give a verbal description of the consensus protocol, but do not analyze dynamic protocol properties. Our analysis, in contrast, provides a detailed, formal description with pseudocode and achieves a much better understanding of how the “preferred ledger” is chosen. Moreover, our work shows possible violations of safety and liveness, whereas Mauri et al. address only on the safety of the consensus protocol through sufficient conditions.

Other academic work mostly addresses network structure, transaction graph, and privacy aspects of payments on the Ripple blockchain [18, 20], which is orthogonal to our focus.

3 A description of the Ripple consensus protocol

The main part of our analysis consists of a detailed presentation of the Ripple consensus protocol in this section and formally in Algorithms 1–3. Before we describe this, we define the task that the protocol intends to solve.

3.1 Specification

Informally, the goal of the Ripple consensus protocol is “to ensure that the same transactions are processed and validated ledgers are consistent across the peer-to-peer XRP Ledger network” [25]. More precisely, this protocol implements the task of synchronizing the nodes so that they proceed through a common execution, by appending successive *ledgers* to an initially empty history and where each ledger consists of a number of *transactions*. This is the problem of replicating a service in a distributed system, which goes back to Lamport et al.’s pioneering work on Byzantine agreement [22, 15]. The problem has a long history and a good summary can be found in the book “30-year perspective on replication” [9].

For replicating an abstract service among a set of nodes, the service is formulated as a deterministic state machine that executes *transactions* submitted by clients or, for simplicity, by the nodes themselves. The *consensus protocol* disseminates the transactions among the nodes, such that each node locally executes the *same sequence of transactions* on its copy of the state. The task provided by this protocol is also called *atomic broadcast*, indicating that the nodes actually disseminate the transactions. When each node locally executes the same sequence of transactions, as directed by the protocol, and since each transaction is deterministic, all nodes will maintain the same copy of the state [26].

More formally, *atomic broadcast* is characterized by two events dealing with transactions: *submission* and *execution*, which may each occur multiple times. Every node may submit a transaction tx by invoking $submit(tx)$ and atomic broadcast applies tx to the application state on the node through $execute(tx)$. A protocol for atomic broadcast then ensures these properties [14, 5]:

Validity: If a correct node p submits a transaction tx , then p eventually executes tx .

Agreement: If a transaction tx is executed by some correct node, then tx is eventually executed by every correct node.

Integrity: No correct node executes a transaction more than once; moreover, if a correct node executes a transaction tx and the submitter p of tx is correct, then tx was previously submitted by p .

Total order: For transactions tx and tx' , suppose p and q are two correct nodes that both execute tx and tx' . Then p executes tx before tx' if and only if q executes tx before tx' .

Our specification does not refer to the heterogeneous trust structure defined by the UNLs and simply assumes all nodes should execute the same transactions. This corresponds to the implicit assumption in Ripple’s code and documentation. We note that the question of establishing global consistency in a distributed system with subjective trust structures is a topic of current research, as addressed by asymmetric quorum systems [6] or in the context of Stellar’s protocol [17], for example.

3.2 Overview

The following description was obtained directly from the source code. Its overall structure retains many elements and function names found in the code, so that it may serve as a guide to the source for others and to explain its working. If the goal had been to compare Ripple consensus to the existing literature on synchronous Byzantine agreement protocols, the formalization would differ considerably.

The protocol is highly *synchronous* and relies on a common notion of time. It is structured into successive *rounds of consensus*, whereby each round agrees on a *ledger* (a set of transactions to execute). Each round roughly takes a predefined amount of time and is driven by a heartbeat timer, which triggers a state update once per second. This contrasts with the Byzantine consensus protocols with partial synchrony [12], such as PBFT [8], which can tolerate arbitrarily long periods of asynchrony and rely on clocks or timeouts only for liveness. The Ripple protocol aims to agree on a transaction set within each synchronized round. The round ends when all nodes collectively declare to have reached consensus on a *proposal* for the round. The protocol is then said to *close* and later *validate* a ledger containing the agreed-on transaction set. However, the transactions in the ledger are executed only after another protocol step, once the ledger has become *fully validated*; this occurs in an asynchronous process in the background. Transaction execution is only logically synchronized with the consensus round.

A *ledger* consists of a batch of transactions that result from a consensus round and contains a hash of the logically preceding ledger. Ledgers are stored persistently and roughly play the role of blocks in other blockchain protocols. Each node locally maintains three different ledgers: the *current ledger*, which is in the process of building during a consensus round, the *previous ledger*, representing the most recently closed ledger and the *valid ledger*, which is the last fully validated ledger in the network.

In more detail, a consensus round has three *phases*: *open*, *establish*, and *accepted*. The usual phase transition goes from *open* to *establish* to *accepted* and then proceeds to the next consensus round, which starts again from *open*. However, it is also possible that the phase changes from *establish* to *open*, if a node detects that it has been forked from the others to a wrong ledger and resumes processing after switching to the ledger agreed by the network.

Nodes may submit transactions at any time, concurrently to executing the consensus rounds. They are disseminated among the nodes through a *gossip layer* that ensures only weak consistency. All transactions that have been received from gossip are placed into a buffer. Apparently, the original design assumed that the gossip layer ensures a notion of consistency that prevents Byzantine nodes from equivocating, in the sense of correct nodes never receive different messages from them. This assumption has been dropped later [10].

The protocol rounds and their phases are implemented by a state machine, which is invoked every second, when the global *heartbeat* timer ticks. Messages from other nodes are received asynchronously in the background and processed during the next timer interrupt.

The timeout handler (L56) first checks if the local *previous ledger* is the same as the *preferred ledger* of a sufficient majority of the nodes in the network. If not, the node has been forked or lost synchronization with the rest of the network and must bring itself back to the state agreed by the network. In this case, it starts a new consensus round from scratch.

When the node enters a new round of consensus, it sets the phase to *open*, resets round-specific data structures, and simply waits for the buffer to fill up with submitted transactions. Once the node has been in the *open* phase for more than half of the duration of the previous consensus round, the node moves to the *establish* phase (L63–L64; function *closeLedger*). It locally closes the ledger, which means to initialize its proposal for the consensus round and to send this to the other nodes in its UNL.

During the *establish* phase, the nodes exchange their proposals for the transactions to decide in this consensus round (using PROPOSAL messages). Obviously, these proposals may contain different transaction sets. All transactions on which the proposals from other nodes differ become *disputed*. Every node keeps track of how many other nodes in its UNL have proposed a disputed transaction and represents this information as *votes* by the other nodes. The node may remove a disputed transaction from its own proposal, or add one to its proposal, based on the votes of the others and based on the time that has passed. Specifically, the node increases the necessary threshold of votes for changing its own vote on a disputed transaction depending on the duration of the *establish* phase with respect to the time taken by the previous consensus round.

The node leaves the *establish* phase when it has found that there is a consensus on its proposal (L69–L71; functions *haveConsensus* and *onAccept*). The node constructs the next ledger (the “last closed ledger”) by “applying” the decided transactions. This ledger is signed and broadcast to the other nodes in a VALIDATION message.

The node then moves to the *accepted* phase and immediately initializes a new consensus round. Concurrently, the node receives VALIDATION messages from the nodes in its UNL. It verifies them and counts how many other nodes in its UNL have issued the same validation. When this number reaches 80% of the nodes in its UNL, the ledger becomes fully validated and the node executes the transactions contained in it.

3.3 Details

Functions. For simplicity, there are some functions that are not fully explained in the pseudocode. These functions are:

- *startTimer(timer, duration)* starts *timer*, which expires after the time passed as *duration*.
- *clock.now()* returns the current time.
- *Hash()* creates a unique identifier (often denoted *ID*) of a data structure by converting the data to a canonical representation and applying a cryptographic hash function.
- $A \triangle B$ denotes the symmetric set difference.
- *boolToInt(b)* converts a logical value *b* to an integer and returns *b*? 0:1.
- *sign_i(L)* creates a cryptographic digital signature for ledger *L* by node *i*.
- *verify_i(L, σ)* checks if the digital signature on *L* from node *i* is valid.
- *siblings(M)* returns the set of nodes, different from *M*, that have the same parent as *M*.

Remarks on the pseudocode. Next to every function name, a comment points to a specific file and line in the source code which contains its implementation. The Ripple source contains a large number of files and most of the consensus protocol implementation is actually spread

over multiple header (.h) files, which complicates the analysis of the code. The references in this work are based on version 1.4.0¹ of `rippled` [23].

Phase open. Function `beginConsensus` starts a consensus round for the next ledger (L50). Each ledger (L11) contains a hash (`ID`) that serves as its identifier, a sequence number (`seq`), a hash of the parent ledger (`parentID`), and a transaction set (`txns`), denoting the transactions applied by the ledger.

The node records the time when the `open` phase started (`openTime`, L54), so that it can later calculate how long the `open` phase has taken. This is important because the duration of the `open` phase determines when to close the ledger locally. If the time that has passed since `openTime` is longer or equal to half of the previous round time (`prevRoundTime`), consensus moves to phase `establish` by calling the function `closeLedger` (L64). Meanwhile all nodes submit transactions with the gossip layer (L46) and each node stores the transaction received via gossip messages in its transaction set S (L48). We model transactions as bit strings. In some places, and as in the source code, we use a short, unique *transaction identifier* (of type `int`) for each transaction $tx \in \{0, 1\}^*$, computed by a function `TxID(tx)`. A transaction set is a set of binary strings here, but the source code maintains a transaction set using a hash map, containing the transaction data indexed by their identifiers.

Phase establish. When the node moves from `open` to `establish`, it calls `closeLedger` that creates an initial proposal (stored in `result.proposal`), containing all transactions received from the gossip layer (L79) that have not been executed yet. A proposal structure (L16) contains the hash of the previous ledger (`prevLedgerID`), a sequence number (`seq`), the actual set (`txns`) of proposed transactions (in the source code named `position`), an identifier of the node (`node`) that created this proposal, and a timestamp (`time`) when this proposal is created (L79).

The node then broadcasts the new proposal as a `PROPOSAL` message (L81) to all nodes in its UNL. When they receive it, they will store its contents in their `currPeerProposals` collection of proposals (L85), if the message originates from a node their respective UNL. The `closeLedger` function also sets `result.roundTime` to the current time (L80). This serves to measure the duration of the `establish` phase and will be used later to determine how far the consensus process has converged.

Based on the proposals from other nodes, each node computes a set of disputed transactions (L88). A disputed transaction (`DisputedTx`, L5) contains the transaction itself (`tx`), a binary vote (`ourVote`) by the node on whether this transaction should be included in the ledger, the number of “yes” and “no” votes from other nodes on the transaction (`yays` and `nays`), taken from their `PROPOSAL` messages, and the list of votes on this transaction from the other nodes (`votes`).

A transaction becomes disputed when it is proposed by the node itself and some other node does not propose it, or vice versa. The node determines these by comparing its own transaction set with the transaction sets of all other nodes (L89). Every disputed transaction is recorded (as a `DisputedTx` structure) in the collection `result.disputes` (L90–L97).

During the `establish` phase, the node constantly updates its votes on all disputed transactions (L68; L99; L117) for responding to further `PROPOSAL` messages that have been received. A vote may change based on the number of nodes in favor of the transaction, the *convergence*

¹ The latest release (16. November 2020) is version 1.6.0. Compared to version 1.4.0, the current release has no significant changes concerning the consensus protocol.

10:8 Security Analysis of Ripple Consensus

ratio (*converge*) and a *threshold*. Convergence measures the expected progress in one single consensus round and is computed from the duration of the *establish* phase, the duration of the previous round, and an assumed maximal consensus-round time (L67). The value for the threshold is predefined. The further the consensus converges, the higher is the threshold that the number of opposing votes needs to reach so that the node changes its own vote (L126). Whenever the node's proposal is updated, the node broadcasts its new proposal to the other nodes (L113) and the disputed transactions are recomputed (L115).

Afterwards, the node checks if consensus on its proposed transaction set *result.txns* is reached, by calling the function *haveConsensus* (L69). The node counts agreements (L130) and disagreements (L131) with *result.txns*. If the fraction of agreeing nodes is at least 80% with respect to the UNL (L132), then consensus is reached. The node proceeds to the *accepted* phase by calling the function *onAccept* (L71).

Phase accepted. The function *onAccept* (L133) “applies” the agreed-on transaction set and thereby creates the next ledger (called the “last closed ledger” in the source code; L134). This ledger is then signed (L136) and broadcast to the other nodes as a VALIDATION message (L137). This marks the end of the *accepted* phase and a new consensus round is initiated by the node (L140).

Meanwhile, in the background, the node receives VALIDATION messages from other nodes in its UNL and tries to verify them (L141). This verification checks the signature and if the sequence number of the received ledger is the same as the sequence number of the own ledger. All validations that satisfy both conditions and contain the node's own agreed-on ledger are counted (L145); this comparison uses the cryptographic hash of the ledger structure in the source code. Again, if 80% of nodes have validated the same ledger and if the sequence number of that ledger is larger than that of the last fully validated ledger (L146), the ledger becomes fully validated (L147). The node then executes the transactions in the ledger (L150). In other words, the consensus decision has become final.

Preferred ledger. A node participating in consensus regularly computes the *preferred ledger*, which denotes the current ledger on which the network has decided. Due to possible faults and network delays, the node's *prevLedger* may have diverged from the preferred ledger, which is determined by calling the function *getPreferred(validLedger)* (L151). Should the network have adopted a different ledger than the *prevLedger* of the node, the node switches to this ledger and restarts the consensus round with the new ledger.

Notice that the validated ledgers from all correct nodes form a tree, rooted in the initial ledger (*genesisLedger*). Each node stores all valid ledgers that it receives in a tree-structured variable *tree*. Whenever the node receives a VALIDATION message containing a ledger L' , it adds L' to *tree* (L143). In order to compute the preferred ledger, we define the following functions, which are derived from the ledgers in *tree* and in the received VALIDATION messages:

- *tip-support*(L) for a ledger L is the number of validators in the UNL that have validated L . In other words,

$$\text{tip-support}(L) = |\{ j \in \text{UNL} \mid \text{validations}[j] = L \}|.$$

- *support*(L) for a given ledger L is the sum of the tip support of L and all its descendants in *tree*, i.e.,

$$\text{support}(L) = \text{tip-support}(L) + \sum_{L' \text{ is a child of } L \text{ in } \text{tree}} \text{tip-support}(L').$$

Algorithm 1 Ripple consensus protocol for node i (continues on next pages).

```

1: Type
2:   Enum Phase = {open, establish, accepted}
3:   Tx = {0, 1}* // a transaction
4:   TxSet = 2Tx
5:   DisputedTx( // DisputedTx.h:50
6:     Tx tx, // disputed transaction
7:     bool ourVote, // binary vote on whether transaction should be included
8:     int yays, // number of yes votes from others
9:     int nays, // number of no votes from others
10:    HashMap[int → bool] votes) // collection of votes indexed by node
11:   Ledger( // Ledger.h:77
12:     Hash ID, // identifier
13:     int seq, // sequence number of this ledger
14:     Hash parentID, // identifier of ledger's parent
15:     TxSet txns) // set of transactions applied by ledger
16:   Proposal( // ConsensusProposal.h:52
17:     Hash prevLedgerID, // hash of the previous ledger, on which this proposal builds
18:     int seq, // sequence number
19:     TxSet txns, // proposed transaction set, called position at ConsensusProposal.h:73
20:     int node, // node that proposes this
21:     milliseconds time) // time when proposal is created
22:   ConsensusResult( // ConsensusTypes.h:201
23:     TxSet txns, // set of transactions consensus agrees on
24:     Proposal proposal, // proposal containing transaction set
25:     HashMap[int → DisputedTx] disputes, // collection of disputed transactions
26:     milliseconds roundTime) // duration of the establish phase

27: State
28:   Phase phase // phase of the consensus round for agreeing on one ledger
29:   Tree tree // tree representation of received valid ledgers
30:   Ledger L // current working ledger
31:   Ledger prevLedger // last agreed-on ("closed") ledger according to the network
32:   Ledger validLedger // ledger that was most recently fully validated by the node
33:   TxSet S // transactions submitted by clients that have not yet been executed
34:   ConsensusResult result // data relevant for the outcome of consensus on a single ledger
35:   HashMap[int → Proposal] currPeerProposals // collection of proposals indexed by node
36:   HashMap[int → Ledger] validations // collection of validations indexed by node
37:   milliseconds prevRoundTime // time taken by the previous consensus round, initialized to 15s
38:   float converge ∈ [0, 1] // ratio of round time to prevRoundTime
39:   UNL ⊆ {1, ..., M} // validator nodes trusted by node  $i$ , taken from the configuration file
40:   milliseconds openTime // time when the last open phase started

41: function initialization()
42:   prevLedger ← genesisLedger // genesisLedger is the first ledger in the history of the network
43:   S ← {}
44:   beginConsensus() // start the first round of consensus
45:   startTimer(heartbeat, 1s) // NetworkOPs.cpp:673

46: upon submission of a transaction  $tx$  do
47:   send message [SUBMIT,  $tx$ ] with the gossip layer

48: upon receiving a message [SUBMIT,  $tx$ ] from the gossip layer do
49:   S ← S ∪ { $tx$ }

50: function beginConsensus() // start a new round of consensus, Consensus.h:663
51:   phase ← open // Consensus.h:669
52:   result ← ({}, ⊥, [], 0) // Consensus.h:674
53:   converge ← 0 // Consensus.h:675
54:   openTime ← clock.now() // remember the time when this consensus round started
55:   currPeerProposals ← [] // reset the proposals for this consensus round

```

Algorithm 2 Ripple consensus protocol for node i (continued).

```

56: upon timeout(heartbeat) do // Consensus.h:818
57:    $L' \leftarrow \text{getPreferred}(\text{validLedger})$ 
58:   if  $L' \neq \text{prevLedger}$  then
59:      $\text{prevLedger} \leftarrow L'$ 
60:      $\text{beginConsensus}(\text{prevLedger})$ 
61:   if  $\text{phase} = \text{open}$  then // wait until the closing ledger can be determined locally
62:     if  $(\text{clock.now}() - \text{openTime}) \geq \frac{\text{prevRoundTime}}{2}$  then // Consensus.cpp:75
63:        $\text{phase} \leftarrow \text{establish}$ 
64:        $\text{closeLedger}()$  // initialize consensus value in  $\text{result}$ 
65:     else if  $\text{phase} = \text{establish}$  then // agree on the contents of the ledger to close
66:        $\text{result.roundTime} \leftarrow \text{clock.now}() - \text{result.roundTime}$ 
67:        $\text{converge} \leftarrow \frac{\text{result.roundTime}}{\max\{\text{prevRoundTime}, 5s\}}$ 
68:        $\text{updateOurProposals}()$  // update consensus value in  $\text{result}$ 
69:       if  $\text{haveConsensus}()$  then
70:          $\text{phase} \leftarrow \text{accepted}$ 
71:          $\text{onAccept}()$  // note this immediately sets  $\text{phase} = \text{open}$  inside  $\text{beginConsensus}()$ 
72:       else if  $\text{phase} = \text{accepted}$  then // Consensus.h:821
73:         // do nothing
74:          $\text{startTimer}(\text{heartbeat}, 1s)$ 

75: // transition from  $\text{open}$  to  $\text{establish}$  phase
76: function  $\text{closeLedger}()$  // Consensus.h:1309
77:    $L \leftarrow (\perp, \text{prevLedger.seq} + 1, \perp, \{\})$ 
78:    $\text{result.txns} \leftarrow S$  // propose the current set of submitted transactions
79:    $\text{result.proposal} \leftarrow (\text{Hash}(\text{prevLedger}), 0, \text{result.txns}, i, \text{clock.now}())$ 
80:    $\text{result.roundTime} \leftarrow \text{clock.now}()$ 
81:   broadcast message [PROPOSAL,  $\text{result.proposal}$ ]
82:    $\text{result.disputes} \leftarrow []$  // disputes for transactions not proposed by all nodes in the UNL
83:   for  $j \in \text{UNL}$  such that  $\text{currPeerProposals}[j] \neq \perp$  do
84:      $\text{createDisputes}(\text{currPeerProposals}[j].\text{txns})$  // compared to  $\text{result.txns}$ , Consensus.h:1334

85: upon receiving a message [PROPOSAL,  $\text{prop}$ ] such that  $\text{prop} = (nl, \cdot, \cdot, j, \cdot)$  and
86:    $j \in \text{UNL}$  and  $nl = \text{Hash}(\text{prevLedger})$  do
87:    $\text{currPeerProposals}[j] \leftarrow \text{prop}$  // Consensus.h:781

88: function  $\text{createDisputes}(\text{TxSet set})$  // Consensus.h:1623
89:   for  $tx \in \text{result.txns} \Delta \text{set}$  do // all transactions that differ between  $\text{result.txns}$  and  $\text{set}$ 
90:      $dt \leftarrow (tx, (tx \in \text{result.txns}), 0, 0, [])$  //  $dt$  is a disputed transaction
91:     for  $k \in \text{UNL}$  such that  $\text{currPeerProposals}[k] \neq \perp$  do
92:       if  $tx \in \text{currPeerProposals}[k].\text{txns}$  then
93:          $dt.\text{votes}[k] \leftarrow 1$  // record node's vote for the disputed transaction
94:          $dt.\text{yays} \leftarrow dt.\text{yays} + 1$ 
95:       else
96:          $dt.\text{votes}[k] \leftarrow 0$  // record node's vote against the disputed transaction
97:          $dt.\text{nays} \leftarrow dt.\text{nays} + 1$ 
98:      $\text{result.disputes}[\text{TxID}(tx)] \leftarrow dt$  // phase  $\text{establish}$ 

99: function  $\text{updateOurProposals}()$  // Consensus.h:1361
100:   for  $j \in \text{UNL}$  such that  $(\text{clock.now}() - \text{currPeerProposals}[j].\text{time}) > 20s$  do
101:      $\text{currPeerProposals}[j] \leftarrow \perp$  // remove stale proposals
102:    $T \leftarrow \text{result.txns}$  // current set of transactions, to update from disputed ones
103:   for  $dt \in \text{result.disputes}$  do //  $dt$  is a disputed transaction
104:     if  $\text{updateVote}(dt)$  then // if vote on  $dt$  changes, update the dispute set
105:        $dt.\text{ourVote} \leftarrow \neg dt.\text{ourVote}$ 
106:       if  $dt.\text{ourVote}$  then // should the transaction be included? DisputedTx.h:77
107:          $T \leftarrow T \cup \{dt.tx\}$  //  $dt.\text{ourVote}$  is initially set in  $\text{createDisputes}(\text{TxSet set})$ 
108:       else
109:          $T \leftarrow T \setminus \{dt.tx\}$ 
110:   if  $T \neq \text{result.txns}$  then // if  $\text{txns}$  changed, then update  $\text{result}$  and tell the other nodes
111:      $\text{result.txns} \leftarrow T$ 
112:      $\text{result.proposal} \leftarrow (\text{Hash}(\text{prevLedger}), \text{result.proposal.seq} + 1, \text{result.txns}, i)$ 
113:     broadcast message [PROPOSAL,  $\text{result.proposal}$ ]
114:      $\text{result.disputes} \leftarrow []$  // recompute  $\text{disputes}$  after updating  $\text{result.txns}$ 
115:     for  $j \in \text{UNL}$  such that  $\text{currPeerProposals}[j] \neq \perp$  do // Consensus.h:1679
116:        $\text{createDisputes}(\text{currPeerProposals}[j].\text{txns})$ 

```

■ **Algorithm 3** Ripple consensus protocol for node i (continued).

```

117: function updateVote(DisputedTx dt) // DisputedTx.h:197
118:   if converge < 0.5 then // set threshold based on duration of the establish phase
119:     threshold ← 0.5
120:   else if converge < 0.85 then
121:     threshold ← 0.65
122:   else if converge < 2 then
123:     threshold ← 0.7
124:   else
125:     threshold ← 0.95
126:   newVote ←  $\left( \frac{dt.yays + \text{boolToInt}(dt.ourVote)}{dt.yays + dt.nays + 1} > \text{threshold} \right)$ 
127:   return (newVote ≠ dt.ourVote) // the vote changes

128: function haveConsensus() // Consensus.h:1545
129:   // count number of agreements and disagreements with our proposal
130:   agree ←  $|\{j | \text{currPeerProposals}[j] = \text{result.proposal}\}|$ 
131:   disagree ←  $|\{j | \text{currPeerProposals}[j] \neq \perp \wedge \text{currPeerProposals}[j] \neq \text{result.proposal}\}|$ 
132:   return  $\left( \frac{\text{agree} + 1}{\text{agree} + \text{disagree} + 1} \geq 0.8 \right)$  // 0.8 is defined in ConsensusParams.h, Consensus.cpp:104
   // phase accepted

133: function onAccept() // RCLConsensus.cpp:408
134:   L ← (prevLedger, result.txns) // L is the last closed ledger, RCLConsensus.cpp:708
135:   validations[i] ← L
136:   σ ← signi(L) // validate the ledger, RCLConsensus.cpp:743
137:   broadcast message [VALIDATION, i, σ, L]
138:   prevLedger ← L // store the last closed ledger
139:   prevRoundTime ← result.roundTime
140:   beginConsensus() // advance to the next round of consensus, NetworkOPs.cpp:1584

141: upon receiving a message [VALIDATION, j, σ, L'] such that // LedgerMaster.cpp:858
142:   L'.seq = L.seq and verifyj(L', σ) do
143:   add L' to tree
144:   validations[j] ← L' // store received validation
145:   valCount ←  $|\{k \in \text{UNL} | \text{validations}[k] = L'\}|$  // count the number of validations
146:   if valCount ≥ 0.8 · |UNL| and L.seq > validLedger.seq then
147:     validLedger ← L // ledger becomes fully validated
148:     S ← S ∪ {L.txns}
149:     for tx ∈ L.txns do // in some deterministic order
150:       execute(tx)

151: function getPreferred(Ledger L) // LedgerTrie.h:677
152:   if L is a leaf node in tree then
153:     return L
154:   else
155:     M ← arg max{support(N) | N is a child of L in the tree}
156:     if uncommitted(M) ≥ support(M) then
157:       return L
158:     else if max{support(N) | N ∈ siblings(M)} + uncommitted(M) < support(M) then
159:       return getPreferred(M)
160:     else
161:       return L

```

10:12 Security Analysis of Ripple Consensus

- $uncommitted(L)$ for a ledger L denotes the number of validators whose last validated ledger has a sequence number that is strictly smaller than the sequence number of L . More formally,


$$uncommitted(L) = |\{j \in UNL \mid validations[j].seq < L.seq\}|.$$

With these definitions, we now explain how $getPreferred(Ledger L)$ proceeds (L151–L161). If L has no children in $tree$, it returns L itself. Otherwise, the function considers the child of L that has the highest support among all children (M). If the support of M is still smaller than the number of validators that are yet uncommitted at this ledger-sequence number, then L is still the preferred ledger (L157). Otherwise, if the support of M is guaranteed to exceed the support of any of its siblings N , even when the uncommitted validators would also support N , then the function recursively calls $getPreferred$ on M , which outputs the preferred ledger for M and returns this as the preferred ledger for L . Otherwise, L itself is returned as the preferred ledger. Observe that in the case when M has no siblings conditions in L156 and L158 are equivalent. Then is enough to check if support of M is greater than uncommitted od M .

4 Violation of safety

In this section, we address the safety of the Ripple consensus protocol. We describe a simple scenario that violates consensus in an execution with seven nodes, of which one is Byzantine. Actually, one can generalize this to executions with more nodes.

To show that the Ripple consensus protocol violates safety and may let two correct nodes execute different transactions, we use the following scenario with seven nodes. Figure 2 gives a graphical representation of our scenario. Nodes are named by numbers. We let $UNL_1 = \{1, 2, 3, 4, 5\}$ and $UNL_2 = \{3, 4, 5, 6, 7\}$, as illustrated by the two hatched areas in the figure. Nodes 1, 2, and 3 (white) trust UNL_1 , nodes 5, 6, and 7 (black) trust UNL_2 , and they are all correct; node 4 (gray) is Byzantine. With this setup, we achieve 60% overlap between the UNLs of any two nodes.



■ **Figure 2** Example setup for showing a safety violation in the Ripple consensus protocol. The setup consists of seven nodes, one of them Byzantine, and two UNLs. Nodes 1, 2, and 3 (white) adopt UNL_1 , vertically hatched, and nodes 5, 6, and 7 adopt UNL_2 , horizontally hatched. Node 4 (gray) is Byzantine.

The key idea is that the Byzantine node (4) changes its behavior depending on the group of nodes to which it communicates. It will cause nodes 1, 2, and 3 (white) to propose some transaction tx and nodes 5, 6, and 7 (black) to propose a transaction tx' for the next ledger. No other transaction exists. The Byzantine node (4) follows the protocol as if it had proposed tx when interacting with the white nodes and behaves as if it had proposed tx' when interacting with the black nodes. Assuming that all nodes start the consensus roughly at the same time and they do not switch the preferred ledger, the protocol does the following:


- The Byzantine node 4 submits tx and tx' using gossip and causes $[\text{SUBMIT}, tx]$ to be received by nodes 1, 2, and 3 and $[\text{SUBMIT}, tx']$ to be received by nodes 5, 6, and 7 from the gossip layer. During the repeated *heartbeat* timer executions in the *open* phase, all correct nodes have the same value of *prevLedger* and send no further messages.
- Suppose at a common execution of the *heartbeat* timer execution (L56) all correct nodes proceed to the *establish* phase and call *closeLedger*. They broadcast the message $[\text{PROPOSAL}, S]$, with S containing tx or tx' , respectively (L81). Node 4 sends a PROPOSAL message containing tx to nodes 1, 2, and 3 and one containing tx' to nodes 5-7. Furthermore, every correct node executes *createDisputes* with the transaction set *txns* received in each PROPOSAL message, which creates *result.disputes* (L88). For nodes 1, 2, and 3, transaction tx' is disputed and for nodes 5, 6, and 7, transaction tx is disputed.
- During *establish* phase, all nodes update their vote for each disputed transaction (L117). Nodes 1, 2, and 3 consider tx' but do not change their *no* vote on tx' because only 20% of nodes in their UNL (namely, node 5) vote *yes* on tx' ; this is less than required *threshold* of 50% or more (L126). The same holds for nodes 5, 6, and 7 with respect to transaction tx . Hence, *result.txns* remains unchanged and no correct node sends another PROPOSAL message.
- Eventually, function *haveConsensus* returns TRUE for each correct node because the required $4/5 = 80\%$ of its UNL has issued the same proposal as the node itself (L128). Every correct node moves to the *accepted* phase.
- During *onAccept*, nodes 1, 2, and 3 send a VALIDATION message with ledger $L = (\text{prevLedger}, \{tx\})$, whereas nodes 5, 6, and 7 send a VALIDATION message containing $L' = (\text{prevLedger}, \{tx'\})$ (L137). Node 4 sends a VALIDATION message containing tx to nodes 1, 2, and 3 and a different one, containing tx' , to nodes 5, 6, and 7.
- Every correct node subsequently receives five validation messages, from all nodes in its UNL, and finds that 80% among them contain the same ledger (L141). Observe that no node changes its preferred ledger after calling *getPreferred*. This implies that nodes 1, 2, and 3 fully validate L and execute tx , whereas nodes 5, 6, and 7 fully validate L' and execute tx' . Hence, the agreement condition of consensus is violated.

5 Violation of liveness

Here we show how the Ripple consensus protocol may violate liveness even when all nodes have the same UNL and only one node is Byzantine. One can bring the protocol to a state, in which it cannot produce a correct ledger and where it stops making progress.

Consider a system with $2n$ correct nodes and one single Byzantine node. All nodes are assumed to trust each other, i.e., there is one common UNL containing all $2n + 1$ nodes. Observe that in this system, the fraction of Byzantines nodes can be made arbitrary small by increasing n .

As illustrated in Figure 3, node $n + 1$, which is Byzantine, exhibits a split-brain behavior and follows the protocol for an input transaction tx when interacting with nodes $1, \dots, n$, and operates with a different input transaction tx' when interacting with nodes $n + 2, \dots, 2n +$



■ **Figure 3** Setup in which liveness is violated in the Ripple network. The network consists of $2n + 1$ nodes with one single UNL and 1 Byzantine (black). The n first nodes propose transaction tx while the last n propose transaction tx' . The Byzantine proposes transaction tx to the n first nodes and transaction tx' to the last n .

1. This implies that the first half of the correct nodes, denoted $1, \dots, n$, will propose a transaction tx and the other half, nodes $n + 2, \dots, 2n + 1$, will propose transaction tx' . Similar to the execution shown in Section 4, the nodes start the consensus protocol roughly at the same time and they do not switch the preferred ledger, they proceed like this:

- Byzantine node $n + 1$ sends two messages, $[\text{SUBMIT}, tx]$ and $[\text{SUBMIT}, tx']$, using the gossip layer and causes tx to be received by the first n correct nodes and tx' to be received by the last n correct nodes.
- After some time has passed, the correct nodes start to close the ledger and move to the *establish* phase. Every correct node sends a PROPOSAL message, containing only the submitted transaction of which it knows (L81), namely tx for the first n correct nodes and tx' for the last n correct nodes.
- During *establish* phase, the correct nodes receive the PROPOSAL messages from all nodes (including the Byzantine node) and store them in *currPeerProposals* (L85). Since they all use the same UNL, all obtain the same PROPOSAL messages from the correct nodes.
- Each node creates disputes (L88) and updates them while more PROPOSAL messages arrive. Since the proposed transaction sets differ, each node creates a dispute for tx and for tx' .
- While the PROPOSAL messages are being processed, votes are counted in *updateVotes* (L117), using the *yays* and *nays* of each disputed transaction. For a correct node in $\{1, \dots, n\}$, notice that the first n nodes and the Byzantine node vote *no* for tx' and the last n nodes vote *yes*. Thus, the fraction of nodes voting *yes* for tx' is less than required threshold (50%), and so the first n nodes continue to vote *no* for tx' . Similarly, nodes $n + 2$ to $2n + 1$ never update their vote on tx and always vote *no* for tx .
- The *haveConsensus* function called periodically during the *establish* phase checks if at least 80% of the nodes in the UNL agree on the the proposal of the node itself (L128). From the perspective of each one of the first n correct nodes, n other nodes agree and n nodes disagree with its proposal, which contains tx . That is not enough support for achieving consensus and the function will return FALSE. The same holds from the perspective of the last n correct nodes, which also continuously return FALSE.
- Finally, the correct nodes will continue trying to update votes and get enough support, but without being able to generate a correct ledger. No correct node proceeds to validating the ledger. In other words, liveness of the protocol is not guaranteed.

6 Conclusion

Ripple is one of the oldest public blockchain platforms. For a long time, its native XRP token has been the third-most valuable in terms of its total market capitalization. The Ripple network is implemented as a peer-to-peer network of validator nodes, which should reach consensus even in the presence of faulty or malicious nodes. Its consensus protocol is generally considered to be a Byzantine fault-tolerant protocol, but without global knowledge of all participating nodes and where a node only communicates with other nodes it knows from its UNL. Previous work regarding the Ripple consensus protocol has already brought up some concerns about its liveness and safety. In order to better analyze the protocol, this work has presented an independent, abstract description derived directly from the implementation. Furthermore, this work has identified relatively simple cases, in which the protocol may violate safety and/or liveness and which have devastating effects on the health of the network. Our analysis illustrates the need for very close synchronization, tight interconnection, and fault-free operations among the participating validators in the Ripple network.

References

- 1 Bowen Alpern and Fred B. Schneider. Defining liveness. *Inf. Process. Lett.*, 21(4):181–185, 1985. doi:10.1016/0020-0190(85)90056-0.
- 2 Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolic, Sharon Weed Cocco, and Jason Yellick. Hyperledger fabric: a distributed operating system for permissioned blockchains. In Rui Oliveira, Pascal Felber, and Y. Charlie Hu, editors, *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*, pages 30:1–30:15. ACM, 2018. doi:10.1145/3190508.3190538.
- 3 Frederik Armknecht, Ghassan O. Karame, Avikarsha Mandal, Franck Youssef, and Erik Zenner. Ripple: Overview and outlook. In Mauro Conti, Matthias Schunter, and Ioannis G. Askoxylakis, editors, *Trust and Trustworthy Computing - 8th International Conference, TRUST 2015, Heraklion, Greece, August 24-26, 2015, Proceedings*, volume 9229 of *Lecture Notes in Computer Science*, pages 163–180. Springer, 2015. doi:10.1007/978-3-319-22846-4_10.
- 4 Ethan Buchman, Jae Kwon, and Zarko Milosevic. The latest gossip on BFT consensus. *CoRR*, abs/1807.04938, 2018. arXiv:1807.04938.
- 5 Christian Cachin, Rachid Guerraoui, and Luís E. T. Rodrigues. *Introduction to Reliable and Secure Distributed Programming (2. ed.)*. Springer, 2011. doi:10.1007/978-3-642-15260-3.
- 6 Christian Cachin and Björn Tackmann. Asymmetric distributed trust. In Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller, editors, *23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland*, volume 153 of *LIPICs*, pages 7:1–7:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.OPODIS.2019.7.
- 7 Christian Cachin and Marko Vukolic. Blockchain consensus protocols in the wild (keynote talk). In Andréa W. Richa, editor, *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria*, volume 91 of *LIPICs*, pages 1:1–1:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.DISC.2017.1.
- 8 Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance and proactive recovery. *ACM Trans. Comput. Syst.*, 20(4):398–461, 2002. doi:10.1145/571637.571640.
- 9 Bernadette Charron-Bost, Fernando Pedone, and André Schiper, editors. *Replication: Theory and Practice*, volume 5959 of *Lecture Notes in Computer Science*. Springer, 2010. doi:10.1007/978-3-642-11294-2.

- 10 Brad Chase and Ethan MacBrough. Analysis of the XRP ledger consensus protocol. *CoRR*, abs/1802.07242, 2018. [arXiv:1802.07242](https://arxiv.org/abs/1802.07242).
- 11 Klitias Christodoulou, Elias Iosif, Antonios Inglezakis, and Marinos Themistocleous. Consensus crash testing: Exploring ripple’s decentralization degree in adversarial environments. *Future Internet*, 12(3):53, 2020. doi:10.3390/fi12030053.
- 12 Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988. doi:10.1145/42282.42283.
- 13 Guy Golan-Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael K. Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. SBFT: A scalable and decentralized trust infrastructure. In *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2019, Portland, OR, USA, June 24-27, 2019*, pages 568–580. IEEE, 2019. doi:10.1109/DSN.2019.00063.
- 14 Vassos Hadzilacos and Sam Toueg. Fault-tolerant broadcasts and related problems. In Sape J. Mullender, editor, *Distributed Systems (2nd Ed.)*. ACM Press & Addison-Wesley, New York, 1993.
- 15 Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982. doi:10.1145/357172.357176.
- 16 LibraBFT Team. State machine replication in the Libra blockchain. Technical report, 2020. URL: <https://developers.libra.org/docs/state-machine-replication-paper>.
- 17 Giuliano Losa, Eli Gafni, and David Mazières. Stellar consensus by instantiation. In Jukka Suomela, editor, *33rd International Symposium on Distributed Computing, DISC 2019, October 14-18, 2019, Budapest, Hungary*, volume 146 of *LIPICs*, pages 27:1–27:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.DISC.2019.27.
- 18 Adriano Di Luzio, Alessandro Mei, and Julinda Stefa. Consensus robustness and transaction de-anonymization in the ripple currency exchange system. In Kisung Lee and Ling Liu, editors, *37th IEEE International Conference on Distributed Computing Systems, ICDCS 2017, Atlanta, GA, USA, June 5-8, 2017*, pages 140–150. IEEE Computer Society, 2017. doi:10.1109/ICDCS.2017.52.
- 19 Lara Mauri, Stelvio Cimato, and Ernesto Damiani. A formal approach for the analysis of the XRP ledger consensus protocol. In Steven Furnell, Paolo Mori, Edgar R. Weippl, and Olivier Camp, editors, *Proceedings of the 6th International Conference on Information Systems Security and Privacy, ICISSP 2020, Valletta, Malta, February 25-27, 2020*, pages 52–63. SCITEPRESS, 2020. doi:10.5220/0008954200520063.
- 20 Pedro Moreno-Sanchez, Navin Modi, Raghuvir Songhela, Aniket Kate, and Sonia Fahmy. Mind your credit: Assessing the health of the ripple credit network. In Pierre-Antoine Champin, Fabien L. Gandon, Mounia Lalmas, and Panagiotis G. Ipeirotis, editors, *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, pages 329–338. ACM, 2018. doi:10.1145/3178876.3186099.
- 21 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Whitepaper, 2009. URL: <http://bitcoin.org/bitcoin.pdf>.
- 22 Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980. doi:10.1145/322186.322188.
- 23 Ripple Labs. Ripple 1.4.0. <https://github.com/ripple/rippled/releases/tag/1.4.0>.
- 24 Ripple Labs. XRP Ledger Documentation > Concepts > Introduction > XRP Ledger Overview. Available online, <https://xrpl.org/xrp-ledger-overview.html>.
- 25 Ripple Labs. XRP Ledger Documentation > Concepts > Consensus Network > Consensus. Available online, <https://xrpl.org/consensus.html>, 2020.
- 26 Fred B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Comput. Surv.*, 22(4):299–319, 1990. doi:10.1145/98163.98167.
- 27 David Schwartz, Noah Youngs, and Arthur Britto. The Ripple protocol consensus algorithm. Ripple Labs Inc., available online, https://ripple.com/files/ripple_consensus_whitepaper.pdf, 2014.