Federated Byzantine Agreement Protocol Robustness to Targeted Network Attacks

Vytautas Tumas, Sean Rivera, Damien Magoni, Radu State

Federated Byzantine Agreement Protocol Robustness to Targeted Network Attacks

Vytautas Tumas
SEDAN - SnT
University of Luxembourg
Luxembourg, Luxembourg
vytautas.tumas@uni.lu

Sean Rivera
SEDAN - SnT
University of Luxembourg
Luxembourg, Luxembourg
sean.rivera@uni.lu

Damien Magoni
LaBRI - CNRS
University of Bordeaux
Talence, France
magoni@labri.fr

Radu State
SEDAN - SnT
University of Luxembourg
Luxembourg, Luxembourg
radu.state@uni.lu

***Abstract*—Federated Byzantine Agreement protocols applied in the XRP Ledger and Stellar use voting to reach a consensus. Participants of these protocols select whom to trust in the network and effectively communicate with the trustees to reach an agreement on transactions. Most trustees, for example 80% in the XRP Ledger, must agree on the same transactions for them to appear in the blockchain. However, disruptions to the communication between the trustees can prevent the trustees from reaching an agreement. Thus, halting the blockchain.**
In this paper, we propose a novel robustness metric to measure the Federated Byzantine Agreement protocol tolerance to node failures. We show that the XRP Ledger Consensus Protocol is vulnerable to targeted attacks. An attacker has to disconnect only 9% of the highest-degree nodes to halt the blockchain. We propose a mitigation strategy which maintains critical XRP Ledger network topology properties whilst increasing the robustness up to 45%.

***Index Terms*—XRP Ledger, Blockchain, Security, Robustness, Federated Byzantine Agreement.**

I. Introduction

Blockchain Consensus is a process by which independent participants agree to accept or refuse some transactions. The so-called *Nakamoto Consensus* [1] adopted by multiple blockchains, including Bitcoin, rely on a notion of *Proof-of-Work*. Proof-of-Work is purposefully a computationally intensive process. It requires participants to solve mathematical puzzles to advance blockchain history. The most popular alternative to Proof-of-Work, adopted by Ethereum 2.0 [2], is *Proof-of-Stake*. This protocol allows participants to affect which transactions to include based on the stake they possess in the blockchain.

At the opposite end of the spectrum lie communication-based consensus protocols. Participants achieve a consensus when a group of trusted participants (a quorum) agree on a set of transactions to include in the next ledger version. In classic consensus systems, all participants have to trust the quorum. However, this does not reflect the individual trust choices made in reality. *Federated Byzantine Agreement* (FBA) protocols improve the trust model by allowing each participant in the blockchain to select whom to trust and effectively work with the trustees to advance the ledger version [3]. Two prominent blockchains use the FBA protocol: the XRP Ledger [4] and Stellar [5].

XRP Ledger is one of the oldest, well-established cryptocurrencies. It uses the *XRP Ledger Consensus Protocol* to advance its blockchain history. It is the 6th largest cryptocurrency with a market capitalization of more than 18 billion USD [6]. It is the most valuable blockchain running an FBA consensus protocol.

To participate in the XRP Ledger Consensus Protocol, a participant must run a server called *validator*, capable of accepting and processing transactions. A validator maintains a static list of other validators it trusts called *Unique Node List* (UNL). At least 80% UNL members must propose the same transaction set to include in the next ledger version. The progress of the ledger stops if validators are unresponsive or cannot reach an agreement. The fact that nodes run a curated UNL [7] further compounds the problem as an attack has to disconnect only 20% or around seven validators to halt the ledger.

Consensus Protocols such as the one implemented in the XRP Ledger can tolerate two types of failures: *crash faults* and *Byzantine faults*. Under Byzantine failures, a participant may behave arbitrarily. They might respond correctly, not respond at all, or reply incorrectly. A significant corpus of research examines the Byzantine fault tolerance of the XRP Ledger [8]–[10].

In contrast, crash faults are less complex. A participant does not respond to and does not perform any operations. Percolation theory dictates that a sudden crash (absence) of a single node will have little to no impact on the network. However, a critical threshold of failures exists, after which the network fragments into isolated components [11]. Consider an attacker whose goal is to halt the XRP Ledger. Assuming that validators run the recommended secure configuration, an attacker is unlikely to target them directly. However, the attacker may disable other nodes until the validators "fall off" the network. Tumas *et al.* [12] showed that a small subset of around 40 nodes forms the connectivity backbone of the XRP Ledger. This backbone makes the XRP Ledger especially vulnerable to attacks directed at these nodes. In network science literature, *Robustness* is a network's ability to continue providing its services when its nodes are absent. More strictly, it is the percentage of nodes that have to disappear for the Largest Connected Component to have fewer than half

of the remaining nodes [11]. However, the FBA protocols depend on reliable message delivery by the underlying peer-to-peer network. Disruptions to the network would undermine the reliability of such protocols. Therefore, we propose a stricter alternative to Robustness, *Quorum Robustness*. It expresses the *Federated Byzantine Agreement* Protocol robustness to node failures, it is:

> *"The percentage of nodes to be removed, such that there are not enough trusted nodes to reach a consensus."*

To distinguish the two robustness metrics, throughout this paper, we will refer to the classic robustness metric as *Network Robustness* and our proposed metric as *Quorum Robustness*.

We summarize the remaining contributions as follows:

1) We conduct an empirical analysis of the Network and Quorum Robustness of the XRP Ledger Peer-to-Peer.
2) We show the conditions under which network fragmentation will halt the consensus protocol.
3) We provide an effective defence strategy to improve the robustness of the XRP Ledger.

We organize the remainder of this paper as follows. In Section II, we review related studies. We provided the necessary background information in Section III. In Sections V and VI, we describe the evaluation methodology and results. In Section VII, we showcase our mitigation strategy and discuss our findings in Section VIII. Finally, we conclude our work in Section IX.

II. Related Work

Cohen *et al.* [13], [14] provided theoretic thresholds for scale-free network resilience to random failures and targeted attacks. The authors demonstrated that scale-free networks are highly resistant to arbitrary failures but are susceptible to targeted attacks. Salah *et al.* [15] extended the model for directed networks and analysed the KAD structured P2P overlay. Recently, Balashov *et al.* [16] provided an optimal strategy for fragmenting scale-free networks and performance bounds on optimal attack strategies on scale-free networks.

Seres *et al.* [17] performed a topological analysis of Bitcoin's Lightning Network (LN). They found that LN had a critical threshold of 14% in the face of targeted attacks. However, the authors did not provide defence strategies. Lee *et al.* [18] continued the analysis of the Lightning Network. The authors evaluated LN's robustness to several attack types and proposed defence strategies against them. They found that default configuration settings in the LN client led to centralisation and recommended changes to fix them. Rohrer *et al.* [19] crafted a measure of the attacker's advantage based on network topology and performed a categorical analysis of potential attack vectors.

Zhao *et al.*. [20] conducted a temporal evolution analysis of Ethereum network interactions. The authors found that the network growth follows a preferential attachment model, and they get sparser as they mature over time. Furthermore, the studied blockchains are not resilient against partitioning and message propagation delay attacks. Gao *et al.* [21] scraped the P2P layer of the Ethereum network and conducted a graph-theoretic analysis of the derived topology. They showed that the Ethereum network is resilient against random failures and targeted attacks.

Research conducted on XRP Ledger predominantly focuses on the XRP Ledger Consensus Protocol. Chase *et al.* [9] described and analysed the XRP Ledger Consensus Protocol (LCP). The authors demonstrated that at least 90% overlap between UNLs is required to ensure network safety. In a later study, Christodoulou *et al.* [22] showed that when fewer than 20% of nodes are malicious, the overlap of UNLs can be relaxed. Otherwise, an overlap of 90-99% is required. In a similar study, Amores-Sesar *et al.* [8] demonstrated that, in the presence of Byzantine nodes, the ledger could fork under standard UNL overlap requirements. Furthermore, the authors showed that a single Byzantine node could cause consensus protocol to lose liveliness. Finally, Mauri *et al.* [10] formalised the XRP LCP and analysed the limitations of the security provisions. They demonstrated a fundamental trade-off between improving the network's responsiveness and security. To the best of our knowledge, there are no studies on the robustness of FBA protocols to network failures. With this work, we aim to close this gap.

III. Background

A. XRP Ledger Topology

The XRP Ledger consists of servers running the *rippled* [4] software. The interconnected *rippled* servers form the decentralized peer-to-peer overlay network. The node owners configure it to accept some inbound and outbound connections. Each outgoing connection corresponds to an incoming connection at another node.

There are two types of servers in the XRP Ledger: *tracking* and *validators*. The tracking servers accept client candidate transactions and broadcast them across the network. The validator servers agree on which transactions to include and work on progressing the ledger version.

A unique node ID identifies each node in the network. The identifier is the node's public key. Other participants in the network use the public key to verify the communication. In addition to the node ID, each validator has a unique validator ID. The ID verifies the Consensus Protocol messages sent by a validator.

B. XRP Ledger Consensus Protocol

An in-depth review of the XRP Ledger Consensus Protocol is outside the scope of this paper. We refer the reader to the works of Amores-Sesar *et al.* [8], Chase *et al.* [9] and Mauri *et.al.* [10] for an in-depth analysis.

Briefly, every participant of the XRP Ledger has the flexibility to select a set of validators they trust, called *Unique Node List* (UNL). There are two stages in the XRP Ledger Consensus Protocol: *consensus* and *validation*.

During the consensus phase, each validator proposes a set of transactions to include in the next ledger version. A validator includes transactions proposed by at least 80% of its UNL members. Similarly, it removes not proposed transactions. The consensus phase continues until at least 80% UNL validators agree on a set of transactions.

In the validation phase, validators calculate the new block version and broadcast it to the participants of the XRP Ledger. The participants declare that the new ledger version is valid only if at least 80% of the validators in their UNL submit the same ledger version. Otherwise, it downloads the block version from the network.

IV. Robustness

Failure of a single node has a limited impact on the integrity of the network. However, the network fragments into multiple isolated components as more nodes fail.

A. Network Robustness

Robustness, sometimes called resilience, is the ability of a network to maintain its functions when some of its nodes are missing. It's quantified as the percentage of nodes that have to fail until the *largest connected component* contains fewer than half of the remaining active nodes [11].

When measuring network robustness, nodes are removed using one of the two strategies: *random* and *targeted*. Depending on the network topology, they produce significantly different robustness results.

a) Random Strategy: The random strategy assumes that all nodes can fail with an equal likelihood. Random strategy models typical node behaviour, such as crashes or restarts. The Scale-Free networks, i.e. those whose degree distribution follows a Power-Law, are known to be remarkably resilient to random breakdowns [13].

b) Targeted Strategy: A targeted removal strategy is a model of a malicious attacker whose goal is to cause as much damage to the network as possible. The attacker can use the readily available network topology to identify authoritative nodes whose absence would cause the most damage to the network. Scale-free networks are especially vulnerable to targeted attacks [14]. Removal of only a few hubs causes the network to begin fragmenting. Continuing the attack breaks the network into small clusters rapidly.

B. Quorum robustness

The robustness metric previously discussed provides a threshold for complete network fragmentation, at which point it is considered non-functional. However, the core function of the XRP Ledger is to process user transactions using an FBA consensus protocol. Thus, an attacker may halt the XRP Ledger by preventing the FBA quorum from forming.

The version of the XRP Ledger advances when 80% of trusted validators agree on a set of transactions, and a participant receives the same new ledger version from 80% of its trusted validators. Participants in the XRP Ledger are recommended to use the UNL curated by the XRP Ledger Foundation [7]. If trusted validators become unavailable, the ledger may halt, or its resistance to Sybil Attacks [23] may weaken. In light of the previously mentioned limitations, we provide a novel, stricter definition of robustness called *Quorum Robustness*:

> *"The percentage of nodes to be removed, such that there are not enough trusted nodes to reach a consensus."*

In the remainder of this paper, we compare and contrast the two robustness metrics and propose a mitigation strategy to improve these metrics for the XRP Ledger.

V. Methodology

In this section, we outline the methodology for evaluating the robustness of XRP Ledger.

a) XRP Ledger Topology: The XRP Ledger community created the open-source XRP Ledger Crawler [4] to construct an accurate active network topology. We used this tool to crawl the XRP Ledger at one-hour intervals for two months between 05/01/2022 and 01/03/2022. We collected 1,290 network snapshots. We made this data openly available online for further research [24]. We performed the robustness analysis on a representative topology snapshot whose size, edge count, average degree, and other properties are closest to the mean of the whole dataset.

b) Synthetic Networks: We compare the robustness of the XRP Ledger overlay network to other network topologies. We generated three topologically different but equal graphs. Graphs are equal when they contain the same number of nodes and edges. The robustness of random graphs is a well-studied topic. However, we include a random graph generated using the Erdős–Rényi (ER) model for completeness. Real-world networks are not random [11], [25]: their degrees are prone to follow an exponential-like distribution. Therefore, we compare the robustness to a scale-free network generated with the Barabási–Albert model. The XRP Ledger topology is small-world, and the degree is power-law-like [12]. We used Klemm-Eguiluz (KE) model [26] to generate a small-world, scale-free network. We use this network to reveal the existence of some latent properties of the XRP Ledger, which make it less robust to targeted attacks.

c) Node Selection: We consider two node failure models: *random failures* and failures due to *targeted attacks*. We examined two metrics for target selection for targeted attacks: *node degree* and *betweenness centrality*.

The *betweenness centrality* of a node v is the fraction of shortest paths that pass through v. Both metrics produce similar robustness values. Therefore, due to space limitations, we only present the results of the degree metric.

After the highest-degree node fails, the overall network degree distribution changes. Thus we recalculate the priorities of each target after node removal. Furthermore, we assume the attacker cannot target validators directly.

d) Validator Selection: XRP Ledger developers recommend running a validator connected to a cluster of trusted tracking servers. The tracking servers, in turn, connect to the remaining network and relay incoming and outgoing validator messages. As a result, the validators are not present in the crawled topology. A set of validators is required to measure the robustness of the XRP Ledger Consensus Protocol. However, as they are not in the network crawl, we label a random set of 34[1] existing nodes in the network as validators.

e) Measuring Robustness: We compute the robustness with a Monte-Carlo simulation to ensure the randomly selected validators do not skew the robustness metric. The simulator works as follows: At each iteration, the simulator labels 34 randomly picked nodes as validators. It then computes the robustness metrics by removing nodes using one of the random or targeted node selection strategies. Once the network reaches the failure threshold, the simulator captures the percentage of nodes removed, resets the network state and begins the next iteration. It continues running until the standard deviation converges. The results presented in Section IV are an average of the simulator results (including standard deviation).

VI. Evaluation

We present the evaluation results in Table I. The entries are the percentage of nodes (including standard deviation) removed before reaching the robustness threshold.

Network Type	Strategy	Network (std)	Quorum (std)
XRP	Attack	20% (7%)	9% (3%)
XRP	Failure	94% (2%)	84% (12%)
Scale-Free	Attack	78% (2%)	78% (9%)
Scale-Free	Failure	95% (5%)	91% (4%)
Random	Attack	86% (0.07%)	88% (3%)
Random	Failure	95% (0.00%)	92% (7%)
Klemm-Eguiluz	Attack	68% (3%)	64% (8%)
Klemm-Eguiluz	Failure	94% (1%)	90% (11%)

TABLE I: The robustness of different network topologies.

a) Random Failures: Although random failures are well-studied, we include our measurements for completeness. We assume the validators are impervious to these failures. Otherwise, the simulator removes all the validators from the network resulting in skewed results. Our findings indicate that the networks are highly robust under both metrics. The fragmentation occurs only when over 80% of the nodes fail. All networks, excluding the Erdős-Rényi (ER) network, have a disproportionate number of small-degree nodes compared to high-degree nodes. Consequently, the likelihood of a failed node having a small degree is higher. The failure of small-degree nodes has a minimal impact on the overall robustness of the network, making fragmentation only possible after most nodes have failed.

b) Targeted Attacks: Our evaluation shows that the XRP Ledger is the least resilient among the studied networks. The network fully fragments when approx. 20% of highest degree nodes fail. Furthermore, we measured that the consensus protocol may halt when only 9% of the authoritative nodes fail.

Around 40 authoritative nodes form the backbone of the XRP Ledger [12]. The failure of a few of these nodes has a limited impact on the overall network connectivity, as multiple redundant paths connect these nodes. However, as the attack progresses, the overlay quickly begins to fragment.

In comparison, the Scale-Free and KE synthetic networks are much more resilient. We measured 78% robustness for both metrics in the Scale-Free network. In the KE network, we captured Network Robustness of 68% and Quorum Robustness of 64%. The Random network was the most resilient, with observed values of 86% and 88% of Network and Quorum robustness, respectively.

The values we observed in artificial networks are consistent with expectations. Generated networks tend to have more connections than expected for their size, making them more robust to targeted attacks.

To improve the liveliness of the ledger, XRP Ledger developers implemented a *Negative UNL* feature [4]. It allows the XRP Ledger to make forward progress in the event of a partial outage. The participants adjust their effective agreement threshold based on which validators from their UNL are operational. For example, if only 70% of UNL members are available, a validator will lower its consensus threshold to 70%. The lower bound for negative UNLs is 60%. Under these conditions, the quorum robustness increases to 11% with a standard deviation of 7%, but still vulnerable when compared to other topologies.

c) Standard Deviation: We observed varying standard deviation values for all networks irrelevant to the robustness metric. In our simulation, given some topology, the only changing variable is the set of *validator* nodes, which the simulator will ignore when selecting which node to remove. When considering *Quorum Robustness*, the existence of standard deviation suggests that the location of the validators in the network topology has a non-negligible effect on the robustness of the FBA protocol. Therefore, an optimal topological position for the validators which maximises the *Quorum Robustness* might exist for any given topology.

d) Effects of targeted attacks: The removal of authoritative nodes affects various properties of the network. We illustrate the largest connected component size (LCC) degradation in Figure 1a. The X-Axis indicates the fraction of nodes removed, while the Y-Axis refers to the relative LCC size. Initially, all networks are connected, and thus the LCC size is one. We observe that in comparison to other networks, XRP Ledger deteriorates rapidly. Scale-Free, KE and ER show a significantly slower fragmentation rate, but all three networks fragment immediately once they reach a critical failure threshold.

In Figure 1b, we provide the connected component size distribution after the network fragment. The X-axis displays the size of the connected component, and the Y-axis represents

[1]At the time of writing, the recommended UNL contains 34 validators.

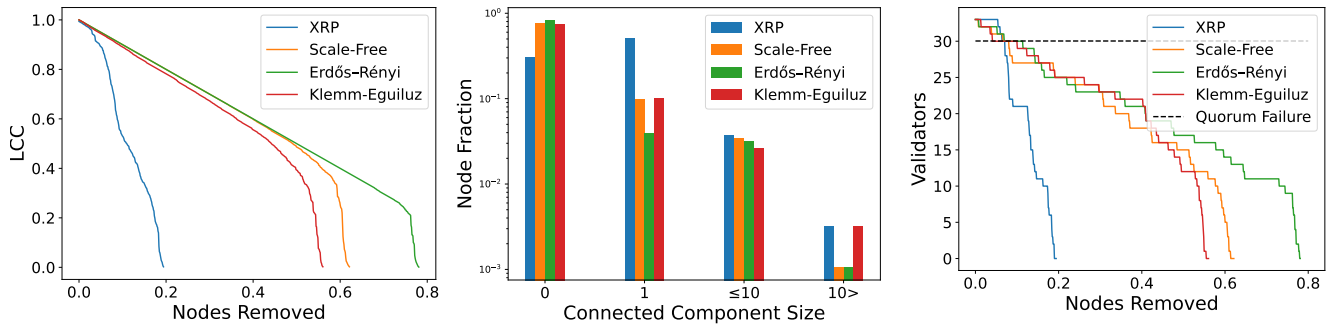

(a) Largest Connected Component Size. (b) Network post fragmentation. (c) Validator Count.

Fig. 1: Illustrative example of network property degradation under targeted attacks.

the fraction of nodes accounted for by each. The zero-size column accounts for the removed nodes.

The results indicate that the XRP Ledger network fragments earlier and to a greater extent than the other networks, as demonstrated by the 401 individual nodes. It is worth noting that the quorum is lost well before the network collapses, as depicted in Figure 1c. The shape of the figure is similar to that of the LCC size degradation, reflecting the relationship between the LCC size and the number of validators in it.

VII. Mitigation

We first outline critical XRP Ledger network properties for effective message dissemination.

A. Background

The XRP Ledger topology is small-world and highly disassortative [12].

a) Small-World: The small-world property pertains to the observation that the average shortest path length is small relative to the network size. At the same time, it exhibits a high degree of clustering. This combination of local clustering and global connectedness is what characterizes small-world networks [11].

b) Degree Correlation: Degree correlation, also called assortativity, refers to the tendency of nodes in a network to connect to other nodes with a similar degree. A network is assortative if high-degree nodes preferentially connect with other high-degree nodes, and low-degree nodes tend to connect with other low-degree nodes. Conversely, a network is disassortative if high-degree nodes connect with low-degree nodes and vice versa, forming a *hub-and-spoke* structure. As we showed in Section IV, disassortative networks are vulnerable to targeted attacks [11], as the small-degree nodes disconnect once high-degree nodes fail.

c) Trade-off: The XRP Ledger hub nodes create the high disassortativity and the small-world property. However, the hubs rapidly disseminate messages across the network. Thus, they are vital to the healthy operation of the blockchain.

This reliance on the hub nodes creates a trade-off. Most nodes depend on these hubs for access to the XRP Ledger. When they fail, the network fragments. In other words, while the hub nodes are essential for the efficient functioning of the XRP Ledger, they also represent a potential point of failure.

d) PeerFinder: The *PeerFinder* module of the rippled server provides the service for initial entry, establishing and accepting connections.

A server gains initial entry into the network by connecting to a set of hub nodes whose domain names are hardcoded into the rippled implementation or alternative servers provided by a list of IP addresses in the configuration file. Once a server has gained entry, it requests its peers for additional IP addresses of servers with available incoming connection slots. It repeats the process until it reaches the desired number of outgoing connections.

A node owner can configure it to accept incoming connections and advertises its open connections to its peers through a probabilistic broadcasting algorithm. This way, most servers can discover available slots. Once the node reaches its configured incoming connection limit, it directs new nodes to other servers with available incoming slots.

B. Mitigation Strategy

Algorithm 1 Selecting a new slot for node N.

```
1:  procedure SLOTS(S : slots[], desiredRatio, desiredPeers)
2:      S ← sortAscending(S)
3:      while True do
4:          smallPeers, highPeers ← N.peersByDegree()
5:          ratio ← |smallPeers| ÷ |highPeers|
6:          if |N.peers()| > desiredPeers then
7:              if ratio < desiredRatio then
8:                  N.replace(highPeers.first(), S.first())
9:              else
10:                 N.replace(smallPeers.first(), S.last())
11:         else
12:             if ratio < desiredRatio then
13:                 N.connect(S.first())
14:             else
15:                 N.connect(S.last())
```

a) Assumptions: We developed the mitigation strategy for improving the robustness of the XRP Ledger with the following assumptions in mind:

- The cost of adding a connection is higher than reestablishing an existing connection to a different node. Therefore, the strategy maintains the degree of each node and only rewires existing connections.
- The small-world property of the XRP Ledger is considered critical to its function. Thus the strategy must preserve it.
- We assume there are enough incoming connection slots to satisfy the demand.

We define a node as *high-degree* when it has exactly or more than *100* peers. Otherwise, the node is *small-degree*. This threshold covers the top 9% of nodes whose removal would cause the quorum failure.

b) Mitigation Approach: Based on the previously mentioned assumptions, we propose a mitigation strategy to maintain a ratio between low-degree and high-degree peers. To illustrate this ratio, consider a node with *ratio* of three. For each high-degree peer, the node will maintain three low-degree peers. The strategy increases the assortativity of the XRP Ledger, thereby improving its robustness while preserving its small-world property. In case of an attack, the failure of hub nodes will result in an increased average shortest path, but it will not halt the consensus process.

Our proposed mitigation strategy involves several modifications to the *PeerFinder* module of the *rippled* server. The first change is to extend the slot advertisements to include the current degree of the node. Although the degree information may become outdated, the short lifespan of the slot advertisements will result in a negligible impact.

We detail the new peer selection procedure in Algorithm 1. The algorithm calculates the ratio between the low-degree and high-degree peers and then makes connections accordingly. The algorithm replaces existing peers when it has the desired number of peers but not the ratio. Otherwise, it connects to new nodes. The node connects to a low-degree node if the computed ratio is lower than the desired one. Otherwise, the node connects to a high-degree slot.

C. Evaluation

a) Robustness: We illustrate the impact of the mitigation strategy in Figure 2. We computed these metrics following the methodology outlined in Section V. For brevity, we restrict our results to the XRP Ledger. The X-axis in the figure displays the balance between low-degree and high-degree peers, where a ratio of 3.0 implies three low-degree peers per each high-degree peer. The Y-axis shows the percentage of nodes removed through a targeted removal process before network fragmentation occurs. The shaded area indicates the standard deviation. The results offer insights into the efficacy of our proposed mitigation strategy in enhancing the stability and robustness of the XRP Ledger. The optimal ratio is 1:1, resulting in 52% network robustness and 45% quorum robustness. Although higher ratios offer a slight increase in robustness, the improvement rate slows.

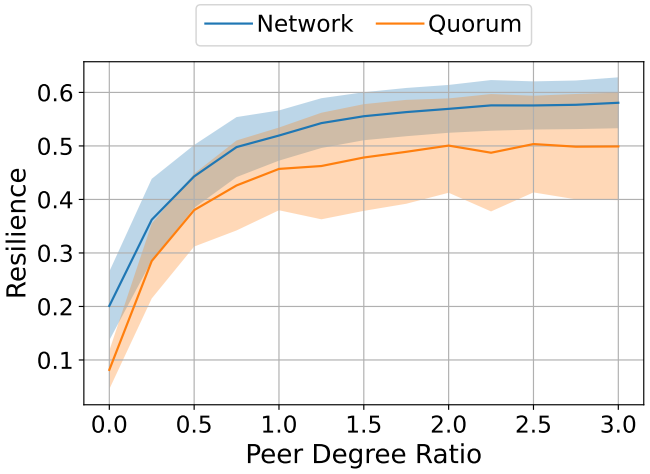

Fig. 2: Illustrative example of the mitigation strategy impact on XRP Ledger robustness.

	Original	Rewired (1.0)
Assortativity	-0.46	-0.21
Avg. Shortest Path	2.33	2.34
Avg. Clustering Coefficient	0.74	0.39
Diameter	5	4
Avg. Low/High Deg. Peer Ratio	0.16	0.91

TABLE II: Basic XRP Ledger properties.

b) Graph Properties: In Table II, we summarize the effects of the mitigation strategy on various network properties using a ratio of *1.0*. We reduced the assortativity by half whilst preserving the small-world property, as indicated by the Average Clustering Coefficient and the Average Shortest path. Interestingly, we also reduced the network diameter. In addition, the algorithm did not achieve the exact ratio across the whole network. Just under half of the nodes have an *odd* number of peers. Therefore, they cannot reach an exact 1:1 peer ratio.

VIII. Discussion

A. Network Robustness

The *Network Robustness* of the XRP Ledger is 20%. In contrast, the *Network Robustness* of Bitcoin and Ethereum are 6% and 4% respectively [27]. These blockchains do not use communication-based consensus. Thus, we don't calculate their *Quorum Robustness*.

However, note that these blockchains are much larger than the XRP Ledger. For instance, Bitcoin has 50,000 nodes, and Ethereum has 12,000 nodes [27]. In contrast, XRP Ledger has only approx. 950 nodes [12]. The other blockchains exhibit a lower robustness percentage. However, they are safe against network-based targeted attacks due to their sheer size.

B. Quorum Robustness

The stability of Federated Byzantine Agreement protocols is heavily dependent on the robustness of the underlying peer-to-peer network. In the XRP Ledger, small-degree nodes tend to connect to high-degree nodes [12]. The failure of high-degree nodes leads to the progressive disconnection of small-degree nodes from the network, causing a cascading network failure.

Validators in the XRP Ledger, when configured following recommended best practices, typically have a small degree and thus depend on hub nodes for access to the rest of the network. Our experiments have shown that the failure of approximately 9% of high-degree nodes will halt the consensus process. The standard deviation we observed also indicates that the topological position of validators has a considerable effect on the robustness of the blockchain, potentially making it more or less robust.

While the introduction of Negative UNLs results in a marginal improvement in robustness of 2%, this is insufficient as it does not address the underlying disassortative structure of the XRP Ledger.

C. Mitigation Strategy

A simple yet effective strategy to improve the robustness of the topology is to replace some of the existing connections to high-degree nodes with those to low-degree nodes. By maintaining a 1:1 ratio of connections to low and high degree nodes, we improve the *quorum robustness* by approx. 36%.

However, our solution is not without limitations. We assume that low-degree nodes will have available open slots to accept new incoming connections. It may not be the case in reality. Link analysis of XRP Ledger identified that most nodes establish outgoing but do not accept any or only accept a small number of incoming connections [12]. By design, there are no direct incentives to participate in the XRP Ledger [28]. However, running a node that accepts incoming connections requires significant investment.

IX. Conclusion

In this paper, we measured the classical *Network Robustness* of the XRP Ledger. We showed that the network fully fragments once $\tilde{2}0\%$ of highest-degree nodes fail. Furthermore, we introduce a novel, more strict *Quorum Robustness* metric for the *Federated Byzantine Agreement Protocols*. We measured that after $\tilde{9}\%$ of the highest-degree nodes fail, the XRP Ledger Consensus Protocol will halt. To improve these metrics, we proposed a mitigation strategy which increases the robustness by up to 45% whilst maintaining the critical network properties.

Acknowledgment

This work was supported by Ripple UBRI.

References

[1] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System." [Online]. Available: www.bitcoin.org

[2] "Proof-of-stake (PoS) | ethereum.org," Feb. 2023, [Online; accessed 3. Feb. 2023]. [Online]. Available: https://ethereum.org/en/developers/docs/consensus-mechanisms/pos

[3] G. A. F. Rebello, G. F. Camilo, L. C. B. Guimarães, L. A. C. de Souza, and O. C. M. B. Duarte, "Security and performance analysis of quorum-based blockchain consensus protocols," *2022 6th Cyber Security in Networking Conference (CSNet)*, pp. 1–7, 2022.

[4] "Home | XRPL.org," Feb. 2023, [Online; accessed 3. Feb. 2023]. [Online]. Available: https://xrpl.org

[5] "Stellar - an open network for money," Oct. 2022. [Online]. Available: https://www.stellar.org

[6] "Cryptocurrency Prices, Charts, and Crypto Market Cap | CoinGecko," Mar. 2023, [Online; accessed 7. Mar. 2023]. [Online]. Available: https://www.coingecko.com

[7] "UNL – XRP Ledger Foundation," Oct. 2022. [Online]. Available: https://foundation.xrpl.org/unl

[8] I. Amores-Sesar, C. Cachin, and J. Mićić, "Security analysis of ripple consensus," 2020. [Online]. Available: https://arxiv.org/abs/2011.14816

[9] B. Chase and E. MacBrough, "Analysis of the xrp ledger consensus protocol," *ArXiv*, vol. abs/1802.07242, 2018.

[10] L. Mauri, S. Cimato, and E. Damiani, "A formal approach for the analysis of the xrp ledger consensus protocol," in *International Conference on Information Systems Security and Privacy*, 2020.

[11] A.-L. Barabási and M. Pósfai, *Network science*. Cambridge University Press, 2016. [Online]. Available: http://barabasi.com/networksciencebook/

[12] V. Tumas, S. Rivera, D. Magoni, and R. State, "Topology analysis of the xrp network," in *38th ACM SIGAPP Symposium on Applied Computing (SAC'23)*, 2023.

[13] Cohen, Erez, ben Avraham, and Havlin, "Resilience of the internet to random breakdowns," *Physical review letters*, vol. 85 21, pp. 4626–8, 2000.

[14] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin, "Breakdown of the internet under intentional attack." *Physical review letters*, vol. 86 16, pp. 3682–5, 2000.

[15] H. Salah, S. Roos, and T. Strufe, "Characterizing graph-theoretic properties of a large-scale dht: Measurements vs. simulations," *2014 IEEE Symposium on Computers and Communications (ISCC)*, pp. 1–7, 2014.

[16] N. Balashov, R. Cohen, A. Haber, M. Krivelevich, and S. Haber, "Optimal shattering of complex networks," *Applied Network Science*, vol. 4, pp. 1–9, 2019.

[17] I. A. Seres, L. Gulyás, D. A. Nagy, and P. Burcsi, "Topological analysis of bitcoin's lightning network," in *Mathematical Research for Blockchain Economy*. Springer, 2020, pp. 1–12.

[18] S. Lee and H. Kim, "On the robustness of lightning network in bitcoin," *Pervasive and Mobile Computing*, vol. 61, p. 101108, 2020.

[19] E. Rohrer, J. Malliaris, and F. Tschorsch, "Discharged payment channels: Quantifying the lightning network's resilience to topology-based attacks," in *IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2019, pp. 347–356.

[20] L. Zhao, S. Sengupta, A. Khan, and R. Luo, "Temporal analysis of the entire ethereum blockchain network," *Proceedings of the Web Conference 2021*, 2021.

[21] Y. Gao, J. Shi, X. Wang, Q. Tan, C. Zhao, and Z. Yin, "Topology Measurement and Analysis on Ethereum P2P Network," in *IEEE Symposium on Computers and Communications*, 2019.

[22] K. Christodoulou, E. Iosif, A. Inglezakis, and M. Themistocleous, "Consensus crash testing: Exploring ripple's decentralization degree in adversarial environments," *Future Internet*, vol. 12, p. 53, 2020.

[23] J. R. Douceur, "The sybil attack," in *International Workshop on Peer-to-Peer Systems*, 2002.

[24] V. Tumas, "XRP Ledger Topology - 05/01/2022 - 01/03/2022," 2023. [Online]. Available: https://doi.org/10.7910/DVN/N44PJG

[25] A. D. Broido and A. Clauset, "Scale-free networks are rare," *Nature Communications*, vol. 10, 2018.

[26] K. Klemm and V. M. Eguíluz, "Growing scale-free networks with small-world behavior." *Physical review. E, Statistical, nonlinear, and soft matter physics*, vol. 65 5 Pt 2, p. 057102, 2001.

[27] A. Paphitis, N. Kourtellis, and M. Sirivianos, "A first look into the structural properties and resilience of blockchain overlays," *ArXiv*, vol. abs/2104.03044, 2021.

[28] "Running an XRP Ledger Validator," Sept 2022. [Online]. Available: https://xrpl.org/blog/2020/running-an-xrp-ledger-validator.html