



HAL
open science

Topology Analysis of the XRP Ledger

Vytautas Tumas, Sean Rivera, Damien Magoni, Radu State

► **To cite this version:**

Vytautas Tumas, Sean Rivera, Damien Magoni, Radu State. Topology Analysis of the XRP Ledger. 38th ACM SIGAPP Symposium on Applied Computing, Mar 2023, Tallinn, Estonia. 10.1145/3555776.3577611 . hal-03927624

HAL Id: hal-03927624

<https://hal.science/hal-03927624>

Submitted on 6 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Topology Analysis of the XRP Ledger

Vytautas Tumas*
University of Luxembourg
Luxembourg, Luxembourg
vytautas.tumas@uni.lu

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

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

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

ABSTRACT

XRP Ledger is one of the oldest, well-established blockchains. Despite the popularity of the XRP Ledger, little is known about its underlying peer-to-peer network. The structural properties of a network impact its efficiency, security and robustness. We aim to close the knowledge gap by providing a detailed analysis of the XRP overlay network.

In this paper we examine the graph-theoretic properties of the XRP Ledger peer-to-peer network and its temporal characteristics. We crawl the XRP Ledger over two months and collect 1,290 unique network snapshots. We uncover a small group of nodes that act as a networking backbone. In addition, we observe a high network churn, with a third of the nodes changing every five days. Our findings have strong implications for the resilience and safety of the XRP Ledger.

CCS CONCEPTS

• **Networks** → **Logical / virtual topologies**; *Network dynamics*;

KEYWORDS

Blockchain, XRP Ledger, Topology Analysis, Measurement

ACM Reference Format:

Vytautas Tumas, Sean Rivera, Damien Magoni, and Radu State. 2023. Topology Analysis of the XRP Ledger. In *The 38th ACM/SIGAPP Symposium on Applied Computing (SAC '23)*, March 27 - March 31, 2023, Tallinn, Estonia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3555776.3577611>

1 INTRODUCTION

XRP Ledger is one of the oldest, well-established cryptocurrencies. In 2022 it ranked seventh by market capitalization. The XRP Ledger aims to provide high transaction throughput whilst maintaining security against Byzantine failures. The XRP Ledger Consensus Protocol is a Federated Byzantine Agreement protocol [28], in which each participant selects a Unique Node List (UNL) of validators.

*Corresponding author.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SAC '23, March 27 - March 31, 2023, Tallinn, Estonia
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9517-5/23/03.
<https://doi.org/10.1145/3555776.3577611>

These validators are not necessarily trusted individually but are believed not to collude as a collective.

Servers running the *rippled* software join into a single peer-to-peer network. The peer-to-peer network's topological structure affects the blockchain's security, resilience, and efficiency. By design, there are no direct incentives to run *rippled* software [4]. Those who participate do so because they are interested in the long-term health of or are participants on the XRP Ledger. A corpus of research focuses on the study of structural properties of Bitcoin [14, 23] and Ethereum [17, 26, 31]. To the best of our knowledge, there are no works examining the overlay network of the XRP Ledger.

The overlay network is uniquely suited for study. Unlike other blockchains that focus on hiding their topology, XRP Ledger natively supports network crawling [3]. The public availability of data enables researchers to determine the accurate topology of the network.

In this paper, we provide an in-depth analysis of the graph-theoretic properties of the XRP Ledger overlay network. Our main contributions are as follows:

- We measure the structural properties of the network, as well as their evolution over time. We discover a central component of the network.
- We examine the stability of the nodes and their uptime. We show that less than 50% of nodes maintained their presence during the measurement period.
- Finally, we show that the network may be vulnerable to Autonomous System failures.

The remainder of this paper is structured as follows. We discuss related work in Section 2. In Section 3, we introduce the relevant aspects of the XRP network. We describe our findings in Section 4 and in Section 5 we discuss network changes over time. Finally, we conclude our work in Section 6.

2 RELATED WORK

We discovered a significant corpus studying cryptocurrency networks, predominantly Bitcoin and Ethereum. We provide a summary of these works in this Section.

Miller *et al.* [23] were the first to determine the topology of the Bitcoin network. The authors discovered "*extremely high-degree nodes*", which persist in the network over time. Furthermore, the Bitcoin network is not purely random. Delgado-Segura *et al.* [14] inferred the topology of Bitcoin using orphaned transactions. Due to the limitations of their method, they performed measurements

only in the Bitcoin *testnet*. Their results indicate that the testnet is not a random graph.

Paphitis et.al. [26] conducted a graph-theoretic analysis of several different blockchain overlay networks. The results indicate that blockchain overlays have varying network properties and degree distributions. Despite the significant variance, there is a strong correlation between the node's session length and the degree. In addition, the networks have small average shortest paths, but they are not small-world. Finally, the overlay networks are resilient to random node failures, but targeted attacks can considerably affect their connectivity.

Similar studies focus on the Ethereum block-chain. Zhao *et al.* [31] performed a temporal, evolutionary analysis of the Ethereum blockchain interaction networks. The authors found a link between anomalies in structural properties and real-life events. Furthermore, they discovered that the network expansion follows a preferential attachment model.

In a later study, Gao *et al.* [17] conducted a graph-theoretic analysis of the peer-to-peer layer of the Ethereum network. They discovered an abundance of nodes that do not contribute to the Ethereum network. Furthermore, they showed that the degree distribution does not follow a power-law. In contradiction to the work of Paphitis *et al.*, the authors found evidence of small-world property.

The research conducted in the XRP Ledger context is predominantly on the Consensus Protocol. Chase *et al.* [10] provide a detailed description and analysis of the Consensus Protocol. They demonstrate that at least a 90% overlap of the UNLs is required to ensure network safety. In a later study, Christodoulou *et al.* [11] show 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.* [6] demonstrate that, in the presence of Byzantine nodes, the ledger may fork under standard UNL overlap requirements. Furthermore, the authors show that a single Byzantine node may cause consensus protocol to lose liveliness.

In a different line of research, Roma *et al.* [29] studied the energy efficiency of an XRP validator. They found that the annual validator running cost is significantly lower than that of a miner.

Aoyama[7] provides a unique view of the XRP network from the perspective of its transactions. They found a clear divide between groups accepting transactions and groups receiving transactions.

To the best of our knowledge, there are no previous studies on the topological properties of the XRP network. With this work, we aim to close this gap.

3 BACKGROUND

The XRP Ledger consists of nodes running the *rippled* [16] software. The interconnected *rippled* servers form the decentralized peer-to-peer overlay network.

The node owners configure it to accept some number of inbound and outbound connections. Each outgoing connection corresponds to an incoming connection at another node. When nodes connect, the communication over the link is bidirectional. We, therefore, represent the overlay network as a directed graph. The direction of an edge identifies the node that initiated the connection.

A node gets initial entry into the overlay network by connecting to several hardcoded bootstrapping hubs. These hubs share the

addresses of other nodes with available inbound connections. The node continues to establish links to others until it reaches the desired limit of outgoing connections. When a node has reached its maximum number of inbound links, it rejects further connection attempts. At the time of data collection, the default number of outgoing connections was 10.

A node periodically advertises to its peers when it has open inbound connection slots. Nodes store this information and communicate it with their peers. As a result, available incoming connections propagate throughout the network.

4 NETWORK ANALYSIS

Network Topology. We used the XRP Ledger Crawler [15] to discover the nodes in the overlay network. The crawler starts by querying the peers of a single *rippled* server¹. It adds the new nodes to a list and calls every node with a known IP address. The crawler repeats this process until it no longer discovers new nodes. We enriched the snapshot data with Autonomous System information. We collected the XRP Ledger network snapshot for two months, between 05/01/2022 and 01/03/2022. We crawled the network topology at one-hour intervals. In two months, we collected 1,290 snapshots. We made the datasets available online for further research [5].

Table 1: Basic XRP network properties.


	Mean	STD
Nodes	948.53	18.54
Edges	15010.26	508.92
In-Degree	15.82	45.62
Out-Degree	15.82	19.94
Connected Component	1	0.00
Assortativity	-0.48	0.02
Global Clustering Coefficient	0.76	0.02
Density	0.03	0.00
Avg. Shortest Path	2.31	0.03
Diameter	5.1	0.33

We summarize the basic properties of the XRP Ledger Network in Table 1.


Size. The network is relatively small. We observed 948 nodes and 15,010 links on average. In comparison, Bitcoin has 50,000 nodes and Ethereum 12,000 nodes [26]. We measured a fluctuation of 2% in the total node count and 1% in the edge count.

In&Out Degrees. Each outgoing connection corresponds to an incoming one, and the nodes report only the active links (not the potential ones). Therefore, the means of incoming and outgoing degrees are equal. The standard deviation of incoming connections is 45.62. This is more than two times greater than that of outgoing ones. The difference in deviations suggests that some nodes in the network accept significantly more incoming connections than others.

¹We used *r.ripple.com* as the starting node



(a) Total connection distribution.



(b) Outgoing (top) and Incoming (bottom) degree distribution.

Figure 1: Illustrative node degree distribution.

Connected Components. A network is consistently connected when, at any point in time, there is a path between two nodes in the network. XRP Ledger is consistently connected, as indicated by the single connected component and zero-value standard deviation. However, this may also be due to the nature of the crawler. The crawler can only discover the nodes that are members of the same connected component as the initial entry node. However, any node not connected to the core could not participate in the Ledger.

Network Density. Network Density is the proportion of the possible and actual connections in the network. Higher values indicate a denser network. In dense networks, messages have a lower propagation delay but at the cost of increased redundancy [8]. The XRP Ledger network has a density of 0.03. In comparison, the density of Bitcoin and Ethereum are 0.002 and 0.0006, respectively [26].

Clustering Coefficient. The Clustering Coefficient quantifies the node’s tendency to form tightly knit groups with high-density ties. The Global Clustering coefficient is 0.78. The low average shortest path and high clustering coefficient of XRP Ledger Network suggest that it may exhibit the small-world property.

4.1 Single Network

We conducted an in-depth analysis of a single XRP Ledger network snapshot. We selected the sample whose node and edge counts are the closest to the mean of the dataset. In the remainder of this Section, we discuss our findings.

Degree Distribution. The network degree distribution impacts many of its properties, such as message propagation delay and the resilience of the network [8]. Random networks have binomial degree distributions, whereas real-world networks contain a small number of highly connected nodes that cannot be accounted for by random models [9].

In Figure 1a, we illustrate the percentage of nodes (y-axis) with a given number of combined incoming and outgoing connections (x-axis). The distribution’s shape is similar to a gamma distribution

with a long right tail. The majority of nodes, 32.5%, have between 10 and 15 connections. At the tail end, nodes have over 325 peers, six times more than nodes at the beginning of the tail.


In Figure 1b, we show separated incoming and outgoing distributions. The upper plot depicts the outgoing connections. The majority of nodes establish between 1 and 22 connections. The largest bin holds 50% of all the nodes, with ten peers. The spike reflects the default *rippled* configuration. At the time of the data collection, the default number of outgoing connections was 10. This value was since updated to 21 [2]. Other nodes connect to between 22 and 90 peers. We also found three outliers; Two nodes with well over 150 and one with just under 100 connections.

In the lower plot, we depict the incoming connection distribution. The first bin contains 60% of nodes without incoming connections. There is no incentive to accept connections, but there is a server maintenance cost. Therefore, the majority of servers only establish outgoing connections. In addition, the first bin may also include validators. By default, for security reasons, they do not accept incoming connections.


The second largest group represents 19% of nodes with between 9 and 11 connections. The remaining bins contain 11% of nodes. These account for the vast majority of the incoming connections in the network. Nodes with around 150 incoming connections are the hubs.

In Figure 2a, we illustrate the cumulative sum of incoming and outgoing connections. There are two outgoing connection groups. The first group, indicated by the exponential portion of the curve, holds nodes whose out-degree is above the mean. It contains 15% of the nodes that account for 50% of all connections. The second group, indicated by the linear portion of the curve, holds the remaining 85% of the nodes. Finally, a deeper inspection revealed two outlier nodes with over 150 outgoing connections.

We similarly grouped the incoming connections. The first group, indicated by the sharp spike of the curve, dominates the overall network connectivity. It contains 11% of nodes with an in-degree



(a) Connection Distribution amongst the nodes.



(b) In/Out Degree Ratio.

Figure 2: Representative network properties.

above the mean. These nodes account for 85% of incoming connections. The second group, depicted by the short linear portion of the curve, contains 27% of nodes. They account for approx. 15% of the incoming links. The final group, reflected by the plateau, holds nodes without incoming connections and accounts for the remaining 62% of nodes.

The incoming and outgoing connection distributions are heavy-tailed. However, they seem to have different shapes. We discuss which model best describes these distributions in Section 4.1. We also observe that a small subset of nodes holds the majority of connections. Our findings suggest that the network has a group of authoritative nodes.

Scale-Free Property. Across scientific domains, it is often claimed that real-world networks are scale-free. Details vary, but in general, a network is scale-free, when nodes with degree k follow a power-law distribution $k^{-\alpha}$, where α is the scaling criterion $\alpha > 1$. However, other versions of this hypothesis have stronger restrictions, e.g. $2 < \alpha < 3$ [8]. Cohen *et al.* show that scale-free networks are highly resilient to random attacks but are vulnerable to targeted attacks [12]. Therefore, it is important to understand the type of degree distribution.

We used the *fitter* [1] Python library to find the most accurate model. We used the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) to determine the quality of a fit. A lower AIC or BIC value indicates a better fit. The analysis in Section 4.1 revealed that the in and out degrees are likely to have different distributions. Therefore we modelled the in, out, and combined distributions separately. We found that the in-degree distribution was best captured by the *power-law* distribution, with α 1.2. On the other hand, the out-degree was best described by the *generalized normal distribution*, with a heavy, long tail. Likewise, *generalized normal distribution* fits the overall degree distribution the best.

The ubiquity of scale-free networks in the real world has been questioned [9]. Therefore, we avoid claiming that the XRP network

is scale-free, as a deeper analysis is required. However, our findings indicate that the XRP network is not random. Furthermore, the power-law distribution fit of the in-degree offers further evidence that the network relies on a subset of nodes for its connectivity.

Small-World Property. The well-studied small-world property indicates that a short path connects any two nodes in the network [8]. An average shortest path l is short when $l \approx \frac{\ln N}{\ln \langle k \rangle}$, where N is the size of the network, and $\langle k \rangle$ is the average degree.

Manfred Kochen and Ithiel de Sola Pool [13] formalized the effect. Which was later popularized by the well-known Milgram experiment that inspired the *six degrees of separation* phrase.

Network G is said to be small-world if it has a similar average shortest path length but a greater clustering coefficient than an equivalent random graph. Two graphs are equivalent when they have an equal number of nodes. Let L_g be the average shortest path length of G and C_g its clustering coefficient. Equivalent properties for a random graph are L_{rand} and C_{rand} . Network G is said to be small-world if $L_g \geq L_{rand}$ and $C_g \gg C_{rand}$.

A quantitative measure of small-worldness is expressed as follows: $\gamma_g = \frac{C_g}{C_{rand}}$ and $\lambda_g = \frac{L_g}{L_{rand}}$, where γ_g is the clustering coefficient ratio and λ_g is the average shortest path ratio of network G and an equivalent random graph. Then measure of small-worldness is expressed as $S = \frac{\gamma_g}{\lambda_g}$. A network is considered small-world when $S > 1$ [19].

We used the Erdős-Rényi (ER) model to generate random graphs. To ensure the robustness of the small-worldness calculation, we used Monte Carlo sampling of 1000 equivalent ER graphs. We measured $S = 8.3$ for the XRP Network. We, therefore, conclude that the XRP network has the small-world property.

In/Out Degree Analysis. Link analysis is a method to identify authoritative nodes in a network [21, 24]. We use it to identify selfish nodes that do not reciprocate the connections they establish by accepting incoming links.




Figure 3: Distribution of the average shortest path length.

We express the link ratio as $\lambda = \frac{In+1}{Out+1}$, a ratio between incoming and outgoing number of connections. All degrees are incremented by 1 to account for no incoming or outgoing connections. A high ratio $\lambda > 1$ suggests that a node is altruistic - it establishes more incoming connections than outgoing ones. Conversely, $\lambda < 1$ indicates nodes that consume more connectivity than they provide.

We illustrate the ratio distribution in Figure 2b. We observe that 15% of nodes have $\lambda < 1$. Interestingly, we find that a significant percentage of nodes have a $\lambda = 0.09$. These nodes use the default *rippled* configuration, with ten outgoing and zero incoming connections. In contrast, only about 10% of nodes have more incoming than outgoing connections, and only 3% $\lambda > 1$.

There are no direct incentives to participate in the XRP network. However, running a node that accepts incoming connections requires significant investment. Such a server has to be reliable and available. Therefore, we see that most nodes connect to the network as *consumers*, and only relatively few behave altruistically. In the next section, we discuss the preference of nodes to connect to other similar nodes.

Degree Correlation. The *degree correlation* captures the node's preference to form connections with others that are similar in some way [8]. In the context of this study, we consider similarity in terms of node degree. A network is *assortative* when nodes tend to connect to others with a similar degree. In a *disassortative* network, small-degree nodes prefer to link with high-degree nodes, and hubs tend to avoid each other. Finally, a network is considered *neutral* when the wiring between the nodes is random.

The *degree correlation* impacts the robustness of a network [30]. In an assortative network, node removal causes little fragmentation, as high-degree nodes form a core group and are redundant. In contrast, disassortative networks are easier to fragment [8]. High-degree nodes connect to many small-degree nodes, forming a hub-and-spoke structure. The small-degree nodes become disconnected once a high-degree node fails.

Degree correlation coefficient r characterizes degree correlation using a single number r [25]. In general it varies between $-1 \leq r \leq 1$ [8]. For $r > 0$ the network is assortative, for $r = 0$ the network is neutral, and for $r < 0$ the network is disassortative. The degree correlation coefficient for XRP Ledger is -0.48 . In comparison, the degree correlation of an equivalent ER network is zero. We conclude




Figure 4: Node uptime CDF.

that the XRP Ledger network is disassortative. It has a hub-and-spoke network structure and may be vulnerable to targeted attacks.

Average Shortest Path Distribution. The average shortest path is the mean of all shortest paths from a node to every other node in the network. We show the distribution of these distances in Figure 3.

The X-Axis is the average path length (in hops), rounded to the tenth. The Y-Axis is the percentage of nodes with the given path length. We observe that around 50% of the nodes have an average distance between 2.2 and 2.3 hops. The distribution has a bell-like shape with a long tail towards longer distances. IP networks have a similar distribution, although the average path length is around nine hops [22].

XRP Ledger uses broadcasting to propagate messages in the network. The shortest path distribution suggests that the network's topology assists in timely message delivery.


Node Distribution over Autonomous Systems. The Autonomous System (AS) number is a 32-bit unique identifier. It represents a collection of IP networks administered by a single entity. We used the AS number to compute the node distribution per AS. In the remainder of this section, we discuss our findings.

Most systems contain only a handful of nodes, while a few AS have a large number of nodes. Around 18% of discovered nodes did not reveal their IP addresses. Therefore, we do not know their AS details.


We split AS details across two tables to illustrate the heavy-tailed nature of node distribution between the AS. Table 2 shows the top ten AS by the number of nodes. These systems, owned by the largest cloud service providers (Amazon, Google, Microsoft), hold nearly 62% of the XRP Ledger nodes. Routing failures in one or more of the dominant Autonomous Systems may disrupt XRP Ledger operations. Therefore, a concentration of nodes may represent a weakness for the ledger.

The remaining 38% of the nodes are evenly distributed over 117 AS. Table 3 shows the number of AS possessing a given number of nodes. We observe that 84 AS only host one XRP node. As the node count per AS increases, the number of autonomous systems rapidly approaches 1.

Node Uptime Distribution. The *rippled* software reports the number of seconds (uptime) it has been running. We plot the cumulative



(a) Incoming connections.



(b) Outgoing connections.

Figure 5: Temporal complimentary cumulative degree distribution.

Table 2: AS with the highest number of nodes.

Rank	AS number	AS name	XRP nodes
1	16509	Amazon.com	177
2	24940	Hetzner Online	115
3	14618	Amazon.com	71
4	8987	Amazon DS Ireland	70
5	396982	Google	52
6	8075	Microsoft Corporation	25
7	16276	OVH	23
8	14061	DigitalOcean	18
9	38895	Amazon.com	18
10	134963	Alibaba.com Singapore	17

Table 3: Number of AS with a given number of nodes.

Nodes per AS	Total number of AS
1	84
2	13
3	8
4	6
5	3
6	1
8	2

distribution function (CDF) of the reported uptime in Figure 4. The X-Axis depicts the uptime in days, rounded to the closest hour.

The average uptime is 9.7 days, with a standard deviation of 18.4 days. Just under 18% of nodes reported uptime of fewer than 12 hours, whereas the oldest node was running for 259 days. Approximately 18% of nodes reported an uptime between 20 and 60 days, and 2% were running for up to 80 days. We observed only a handful of nodes older than 100 days.

5 TEMPORAL ANALYSIS

In this section, we discuss the evolution of the network over time. We begin with a summary of the temporally stable properties. We observe that the preference for small-degree nodes to connect to high-degree nodes remains constant over time. Likewise, the global clustering coefficient and average shortest path are stable. Furthermore, all network snapshots have the small-world property. These findings suggest that there were no significant disruptions in the network during the observation period.

The relatively small change in the network's size had a non-negligible effect on the average incoming and outgoing degree, as indicated by the standard deviation. We dedicate the rest of this section to discussing these changes.

Degree Distribution. We illustrate the complementary cumulative distribution function (CCDF) of the node degrees in Figure 5. Overall, both distributions have long tails, and their shapes remain stable. However, we see some variance over time in both figures, as indicated by the changing thickness of the plots.

We plot the CCDF of incoming connections in Figure 5a. Our first observation is that consistently 60% of nodes do not accept incoming connections. Likewise, we see little variance at the tail-end of the spectrum. We see slightly more variance in nodes close to the mean and nodes with a degree between 250 and 300. We observe the largest variance in nodes with in-degree between 50-150. Our observations suggest that the nodes at the ends of the distribution are saturated. They cannot accept new peers. Therefore, nodes in the middle of the distribution handle the new connections to the network. Furthermore, the majority of new nodes do not accept incoming links.

We depicted the CCDF of the outgoing connections in Figure 5b. The long, thin tail of the distribution suggests the existence of a few stable nodes with a high number of outgoing connections. We see a much higher variance in the group of nodes with an out-degree between 50 and 100. Finally, the majority of new nodes had an out-degree under the mean.




Figure 6: Temporal in/out degree ratio.

Two versions of the *rippled* software came out during the data collection. Some observed variances may be explained by nodes leaving the network to update their version. However, overall the network has a stable member group.

In/Out Degree Analysis. In Figure 6, we display the degree ratio plot for all captured snapshots. We observe little change in the overall degree ratio. The majority of nodes establish more outgoing than incoming connections. Only 10% of nodes establish more incoming than outgoing connections.

The lack of change in the shape of the curve confirms our initial observation that nodes do not reciprocate the connections they consume.

Membership Stability. Over the collection period, we discovered 3,000 unique nodes. In Figure 7, we outline the lifespan of these nodes. The green, striped bar indicates nodes with the shortest lifespan. These nodes were present in around 5% of all network snapshots. On the other side of the plot, the blue crossed bar represents the most stable nodes. They were present in at least 95% of all the snapshots. The remaining 1/5th of the nodes have a gradually decreasing lifespan.

The fully present nodes have an average in-degree of 26.1. In comparison, the nodes we observed in the 5% of snapshots have an average in-degree of 16. The difference between the values suggests that the fully present nodes are the ones that form the network backbone, which we discovered in Section 4.1.

We further analyzed the presence of the top 10% of the highest in-degree nodes in the network over time. The group of the first network snapshot contains 95 nodes. The last network snapshot holds 98 nodes. However, 23 nodes or 24% from the first group are not present in the second group. Four nodes changed their IDs but had the old IP addresses and similar degree profiles. However, we did not find the other 19.

Node Uptime Over Time.

We measured the uptime of the 410 nodes present during every network crawl. Figure 8 presents two illustrative examples of the




Figure 7: Number of times a node was observed in a network.




Figure 8: Illustrative uptime of representative nodes.

observed uptime. We limit our selection to the most representative nodes, in which we can observe clear patterns.

The top graph depicts the uptime graph of the 259 days old node we discussed in Section 4.1. There are two distinct features in the figure. The top line indicates that a server was running smoothly for the observation period. In contrast, the bottom feature suggests a server started and failed multiple times. These features suggest two separate instances of *rippled* running behind a single IP address.

The bottom plot offers a better illustration of two servers behind one IP address. There are two parallel lines of similar length. When we query the server uptime, we receive a response from one of the two servers. The fragmentation in the lines is information gaps caused by a different server handling the uptime request.

In all plots, we notice inexplicable uptime values. These could suggest that a new *rippled* instance started or, more worryingly, uptime reporting issues. However, we leave the study of these observations for future work.

6 CONCLUSION & FUTURE WORK

A decentralized peer-to-peer overlay network forms the backbone of the XRP Ledger. In this paper, we provided an in-depth analysis of the graph-theoretic properties of this overlay.

We use a publicly available crawler to capture 1,290 snapshots of the underlying overlay network over two months. We find that it is significantly smaller than other blockchain overlay networks. The nodes are connected via short paths and are tightly clustered. Furthermore, the clusters tend to have a hub-and-spoke structure, as shown by the high assortativity of the network. Unlike other blockchain overlay networks, XRP has a small-world topology.

XRP does share some similarities with other blockchains. The network degree distribution has an exponential-like shape. We did not find conclusive evidence that it is scale-free. However, like other blockchains, the topology is not random.

Overall, the size of the network is consistent over time. However, we captured a significant amount of churn. Given these observations, we suspect that many nodes join the network to conduct their business and leave shortly.

The XRP overlay network may be vulnerable to targeted attacks. We discovered the existence of a small subset of influential nodes that provide the backbone of the network connectivity. Furthermore, a malicious actor can use the publicly available topology to identify these nodes.

We revealed a vast disparity between nodes that accept incoming connections and nodes that do not. Furthermore, link analysis showed that many nodes do not accept incoming connections. These nodes increase the dependence on the influential nodes. Thus, contributing to the network centralization. We suspect that a lack of financial incentive contributes to this behavior, as there are significant costs associated with running a reliable node. Natural centralization is a common problem in decentralized peer-to-peer networks[20][18]. A common solution is to introduce communal incentives or mandatory behavior.

Our results raise further questions about the security and vulnerability of the XRP Ledger. Research works [12] [27] show that networks with a long-tail degree distribution are susceptible to targeted attacks. For future work, we intend to evaluate the resilience of the XRP network to random and targeted attacks, and to identify mitigation strategies.

ACKNOWLEDGMENT

This work was supported by Ripple UBRI².

REFERENCES

- [1] 2022. FITTER Documentation. <https://fitter.readthedocs.io/en/latest> [Online; accessed 9. Sep. 2022].
- [2] 2022. Maximum Number of Peers. <https://xrpl.org/set-max-number-of-peers.html> [Online; accessed 9. Sep. 2022].
- [3] 2022. Peer Crawler. <https://xrpl.org/peer-crawler.html> [Online; accessed 9. Sep. 2022].
- [4] 2022. Running an XRP Ledger Validator. <https://xrpl.org/blog/2020/running-an-xrp-ledger-validator.html> [Online; accessed 9. Sep. 2022].
- [5] 2022. XRP Network Snapshots. <https://drive.google.com/drive/folders/1SY4IemcQsr0FCiagL0dSyNIOYmQ6SLdg> [Online; accessed 9. Sep. 2022].
- [6] Ignacio Amores-Sesar, Christian Cachin, and Jovana Mičić. 2020. Security Analysis of Ripple Consensus. *arXiv:cs.DC/2011.14816*

- [7] Hideaki Aoyama. 2021. *XRP Network and Proposal of Flow Index*. <https://doi.org/10.7566/JPSCP.36.011003>
- [8] Albert-László Barabási and Márton Pósfai. 2016. *Network science*. Cambridge University Press, Cambridge. <http://barabasi.com/networksciencebook/>
- [9] Anna D. Broido and Aaron Clauset. 2018. Scale-free networks are rare. *Nature Communications* 10 (2018). <https://doi.org/10.1038/s41467-019-08746-5>
- [10] Brad Chase and Ethan MacBrough. 2018. Analysis of the XRP Ledger Consensus Protocol. *arXiv:cs.DC/1802.07242*
- [11] Klitos Christodoulou, Elias Iosif, Antonios Inglezakis, and Marinos Themistocleous. 2020. Consensus Crash Testing: Exploring Ripple's Decentralization Degree in Adversarial Environments. *Future Internet* 12 (2020), 53. <https://doi.org/10.3390/fi12030053>
- [12] Reuven Cohen, Keren Erez, Daniel ben Avraham, and Shlomo Havlin. 2000. Breakdown of the internet under intentional attack. *Physical review letters* 86 16 (2000), 3682–5.
- [13] Ithiel de Sola Pool and Manfred Kochen. 1978. Contacts and influence. *Social Networks* 1, 1 (1978), 5–51. [https://doi.org/10.1016/0378-8733\(78\)90011-4](https://doi.org/10.1016/0378-8733(78)90011-4)
- [14] Sergi Delgado-Segura, Surya Bakshi, Cristina Pérez-Solà, James Litton, Andrew Pachulski, Andrew Miller, and Bobby Bhattacharjee. 2018. TxProbe: Discovering Bitcoin's Network Topology Using Orphan Transactions. <https://doi.org/10.48550/ARXIV.1812.00942>
- [15] RippleX Engineering. 2022. XRPL Network Crawler. <https://github.com/xpring-eng/rippled-network-crawler> [Online; accessed 9. Sep. 2022].
- [16] XRP Ledger Foundation. 2022. rippled. <https://github.com/XRPLF/rippled> [Online; accessed 9. Sep. 2022].
- [17] Yue Gao, Jinqiao Shi, Xuebin Wang, Qingfeng Tan, Can Zhao, and Zelin Yin. 2019. Topology Measurement and Analysis on Ethereum P2P Network. *Proceedings - IEEE Symposium on Computers and Communications* 2019-June (2019). <https://doi.org/10.1109/ISCC47284.2019.8969695>
- [18] D. Hughes, G. Coulson, and J. Walkerdine. 2005. Free riding on Gnutella revisited: the bell tolls? *IEEE Distributed Systems Online* 6, 6 (2005). <https://doi.org/10.1109/MDSO.2005.31>
- [19] Mark D. Humphries and Kevin Gurney. 2008. Network 'small-world-ness': A quantitative method for determining canonical network equivalence. *PLoS ONE* 3, 4 (apr 2008). <https://doi.org/10.1371/JOURNAL.PONE.0002051>
- [20] Murat Karakaya, Ibrahim Korpeoglu, and Özgür Ulusoy. 2009. Free Riding in Peer-to-Peer Networks. *IEEE Internet Computing* 13, 2 (2009), 92–98. <https://doi.org/10.1109/MIC.2009.33>
- [21] Jon Kleinberg and Steve Lawrence. 2001. The Structure of the Web. *Science* 294, 5548 (2001), 1849–1850. <https://doi.org/10.1126/science.1067014>
- [22] D. Magoni and J.-J. Pansiot. 2002. Evaluation of Internet topology generators by power law and distance indicators. In *10th IEEE International Conference on Networks*. 401–406. <https://doi.org/10.1109/ICON.2002.1033345>
- [23] Andrew K. Miller, James Litton, Andrew Pachulski, Neal Gupta, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2015. *Discovering Bitcoin's Public Topology and Influential Nodes*. Technical Report. University of Maryland, College Park.
- [24] Alan Mislove, Massimiliano Marcon, Krishna P. Gummadi, Peter Druschel, and Bobby Bhattacharjee. 2007. Measurement and Analysis of Online Social Networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC '07)*. Association for Computing Machinery, New York, NY, USA, 29–42. <https://doi.org/10.1145/1298306.1298311>
- [25] M.E.J. Newman. 2003. Mixing patterns in networks. *Physical review. E, Statistical, nonlinear, and soft matter physics* 67 (03 2003), 026126. <https://doi.org/10.1103/PhysRevE.67.026126>
- [26] Aristodemos Paphitis, Nicolas Kourtellis, and Michael Sirivianos. 2021. A First Look into the Structural Properties and Resilience of Blockchain Overlays. *arXiv preprint arXiv:2104.03044* (2021).
- [27] Guan-Sheng Peng, Suo-Yi Tan, Jun Wu, and Petter Holme. 2016. Trade-offs between robustness and small-world effect in complex networks OPEN. (2016). <https://doi.org/10.1038/srep37317>
- [28] Gabriel Antonio F. Rebello, Gustavo Franco Camilo, Lucas C. B. Guimarães, Lucas Aíram C. de Souza, and Otto Carlos Muniz Bandeira Duarte. 2022. Security and Performance Analysis of Quorum-based Blockchain Consensus Protocols. *2022 6th Cyber Security in Networking Conference (CSNet)* (2022), 1–7. <https://doi.org/10.1109/CSNet56116.2022.9955597>
- [29] Crystal Andrea Roma and M Anwar Hasan. 2020. Energy consumption analysis of XRP validator. In *2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 1–3.
- [30] Alexei Vazquez and Yamir Moreno. 2003. Resilience to damage of graphs with degree correlations. *Physical review. E, Statistical, nonlinear, and soft matter physics* 67 (02 2003), 015101. <https://doi.org/10.1103/PhysRevE.67.015101>
- [31] Lin Zhao, Sourav Sen Gupta, Arijit Khan, and Robby Luo. [n. d.]. Temporal Analysis of the Entire Ethereum Blockchain Network; Temporal Analysis of the Entire Ethereum Blockchain Network. ([n. d.]), 12. <https://doi.org/10.1145/3442381.3449916>

²<https://ubri.ripple.com/>