# Mind Your Credit: Assessing the Health of the Ripple Credit Network

Pedro Moreno-Sanchez, Navin Modi, Raghuvir Songhela, Aniket Kate, Sonia Fahmy
Purdue University
{pmorenos, modin, rsonghel, aniket, fahmy}@purdue.edu

## ABSTRACT

Initially proposed as part of the cypherpunk-esque rebellion started by Bitcoin, the Ripple credit network is dramatically reshaping the financial and remittance industries. At its core, the Ripple network is a weighted directed graph of IOweYou credit links, where a transaction between two nodes is only allowed if there is sufficient max-flow between them. These path-based settlements set Ripple apart from the cryptocurrency space (such as Bitcoin and altcoins) both conceptually and in terms of applicability. This work takes an in-depth look at the Ripple network since its inception in order to understand its structure, evolution, and vulnerability to attacks.

We observe that gateway nodes determine not only the graph structure but also the existence of clear geographical communities. From the security point of view, although the core of the network is well-connected and provides high liquidity to users, there is a user base of around $50,000$ wallets prone to disruption by as few as 10 highly connected wallets. Finally, we find that redistribution of credit (i.e., rippling) among credit links as a result of a transaction puts at risk around 30M USD in the current Ripple network. Our study should serve as a wake-up call to inattentive users to improve their connectivity to the network, and ensure that their rippling flags and credit limits on the links are correctly set.

## 1. INTRODUCTION

The Ripple network [11,18,42] stands out among the plethora of flourishing cryptocurrencies and payment networks by allowing transactions in traditional fiat currencies, cryptocurrencies and user-defined currencies simultaneously. Its inherent capability to perform cross-currency transactions over credit paths in a matter of seconds for a small fee in a publicly verifiable manner paves the way for reducing financial institution costs by billions of dollars [15,16,35].

In this state of affairs, early adopters of Ripple [32,45] have been recently followed by a wave of financial institutions worldwide [10,17,36,46,47,49], including 12 of the world's top 50 banks [38], remittance institutions [6,33] and online exchange services for cryptocurrencies [5,31]. Ripple's market capitalization is currently second only to Bitcoin, and at least twice as large as the next competitor Ethereum [2].

Among early academic efforts, Armknecht et al. [18] present basic statistics of the Ripple network usage such as transaction volume, and consider the safety of the Ripple consensus process. Moreno-Sanchez et al. [41,42] focus on deanonymization attacks and privacy enhancing solutions for network users. Nevertheless, the Ripple network is yet to attract attention similar to Bitcoin [19,23,29,37,48] from the academic community. This is critical because Ripple's IOweYou credit network [18,25–27,30,42] and path-based transactions over those credit links clearly set it apart (structurally and functionally) from cryptocurrencies.

Against this background, this paper presents, to the best of our knowledge, the first extensive study of the Ripple network and transactions since its inception, shedding light on the current deployment, evolution, and security of the Ripple system. By analyzing the collected Ripple network data from December '16 with a total of $99,413$ wallets and $246,672$ credit links, as well as $27,406,877$ transactions during the period Jan '13 – Dec '16, we make two key contributions.

First, we characterize the structure of the Ripple network, the motifs (i.e., subgraphs) that best describe the graph, the evolution of its communities and its mixing time (Section 4). We observe that the Ripple network is formed with "gateways" as key players. Gateways are highly connected bootstrapping wallets trusted to set up links to new users. Further, the Ripple network is clearly structured into communities adhering to geographical regions, where each community is defined by (on average) two gateway wallets. We find that initially the Ripple network included communities mostly in Europe and Asia whereas currently Ripple is rapidly expanding to new communities such as Israel. Finally, we observe that the number of new credit links in the Ripple network grows linearly with the number of wallets, and hence the network density is decreasing. The network is slow-mixing, unclustered and disassortative.

Second, we assess the "health" of the network by investigating the network liquidity, the undesired redistribution of credit, and the resilience to disruptive wallets that, for instance, do not allow transactions through them (Section 5). We observe that the core of the Ripple network, composed of around $10,000$ wallets, provides high liquidity and is resilient to disruptive wallets. However, a large user base (around $50,000$) is prone to disruption by as few as 10 highly connected wallets, and their credit (currently about 14M USD) is at risk of being no longer connected to the main component of the Ripple network. Finally, we observe that 22% of the wallets in the Ripple network are prone to possibly undesired redistribution of credit (i.e., rippling) among their credit links associated with a total credit of about 30M USD. The Ripple community can enhance the health of the network by educating users on improving their connectivity and setting the upper limits of their credit links well below the default value.

## 2. RIPPLE OVERVIEW

**The Ripple Network.** With its roots in IOweYou credit networks [18, 27, 30, 42], the Ripple network essentially is a weighted, directed graph where nodes represent wallets and edges represent credit links between wallets. The non-negative weight on an edge ($u_1$, $u_2$) represents the amount that $u_1$ owes to $u_2$. The credit on a link is upper-bound by $\infty$ by default, although the wallet owner ($u_2$ in our example) can customize it. Figure 1 depicts an example.

**Ripple Transactions.** Ripple allows two types of transactions: direct XRP payments and path-based settlement transactions. XRP is the native currency in Ripple incorporated for users to pay a small fee per transaction towards curbing denial of service attacks and unbounded wallet creation. A direct XRP payment exchanges XRP between two wallets, even if they are not connected via a network path. XRP payments resemble debit payments between users rather than path-based credit settlements, which is the focus of this paper. Therefore, we omit XRP payments in our analysis and refer the readers to [42] for further details.

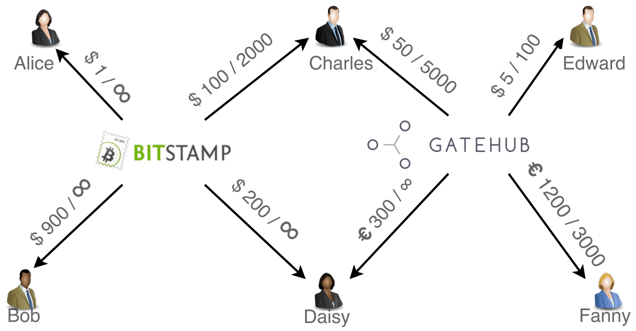

Figure 1: A Ripple credit network example. An edge weight shows two values $a$ / $b$, where $a$ denotes the current credit in the credit link and $b$ denotes the upper bound. The edge lower bound is always zero.

A path-based settlement transaction (or simply transaction hereby) uses a path of credit links between sender and receiver to settle credit between them. In the example shown in Figure 1, assume that Alice wants to pay $1 to Edward. At first, credit links are considered undirected to find a path from the sender to the receiver. The transaction can be routed using the path Alice $\leftarrow$ Bitstamp $\rightarrow$ Charles $\leftarrow$ GateHub $\rightarrow$ Edward. The transaction is carried out by updating the credit value on each credit link depending on its direction as follows: credit links in the direction from sender to receiver are increased by $1, while reverse credit links are decreased by $1. In the running example, Alice $\leftarrow$ Bitstamp and Charles $\leftarrow$ GateHub are decreased to $0 and $49, respectively, whereas credit links Bitstamp $\rightarrow$ Charles and GateHub $\rightarrow$ Edward are increased to $101 and $6, respectively. Several paths between sender and receiver can be used in a single transaction [42].

**Gateways.** A *gateway* is a well-known business wallet established to bootstrap credit links to new wallets in an authenticated manner. Gateways are the Ripple counterparts of user-facing banks and loan agencies in the physical world. Their wallets maintain high connectivity. A newly created Ripple wallet that does not initially trust any existing wallet can create a credit link to a gateway and thereby interact with the rest of the network before forming direct links to other wallets.

**Market Makers.** A market maker is a wallet in the Ripple network that receives a certain currency on one of its credit links and exchanges it for another currency on another link, charging a small fee. Market makers enable transactions with different currencies.

## 3. DATASETS

**Data Sources.** Our experiments are based on publicly accessible data that can be extracted through the API [9] provided by the Ripple company on their servers $\{s1,\ s2,\ data\}.ripple.com$. We crawl the datasets describing the Ripple network, transactions, gateways, and market makers, and summarize them in the rest of this section. We refer the reader to [8] for further statistics.

**Ripple Network Topology.** We collected all wallets and credit links comprising the Ripple network at the end of each year from December 2013 until December 2016. In each snapshot, we only consider the largest connected component, denoted by multigraph-year for each *year*. In certain experiments, we represent all credit links between any pair of wallets (e.g., one link per currency) as a single link. We denote the thereby processed data as graph-year for each *year*. We observe that graph-2016 consists of $96,953$ wallets and $172,877$

credit links, which represent 97.5% of wallets and 98.9% of credit links of the total Ripple network at that point of time. Such percentages are similar for the rest of snapshots. Therefore graph-{2013-2016} is representative of the Ripple network. Finally, in some experiments, we are only interested in the core of the Ripple network, what we call skeleton. For a graph, say graph-2016, we calculate the corresponding skeleton-2016 by iteratively removing wallets with degree one. skeleton-2016 consists of 37,947 wallets and 113,665 credit links.

**Ripple Transactions.** We extracted the transactions in the Ripple network in the period Jan 13 – Dec 16, which is a total of 27,406,877 transactions. We pruned this dataset according to the following criteria. First, we discarded 18,266,173 XRP payments. Second, we discarded 1,530,107 anomalous transactions (e.g., spam) considered outliers by previous studies [42]. Finally, we discarded 1,014,121 transactions that are not performed among wallets included in graph-2016. Our final set of transactions contains a total of 6,596,475 transactions. We refer to this as txs-{2013-2016}.

**Gateways.** We crawled the list of gateways from the Ripple API, and added the gateways identified by the Ripple community [7]. As a result, we obtained a list of 101 gateways and 119 wallets associated with them. We will denote this dataset by gateways-2016.

**Ethical Considerations.** Our Ripple network analysis solely uses publicly available data. We only mention the names of gateways that are well-known in the Ripple community and publicly advertised on websites.

# 4. NETWORK STRUCTURE & EVOLUTION

In this section, we analyze our datasets to understand the structure and evolution of the Ripple network.

## 4.1 Network Topology

**Network Growth and Structure.** graph-2013 contains 14,744 wallets and 23,418 credit links, while graph-2016 contains 96,953 wallets and 172,877 credit links. Therefore, in three years, the Ripple network has 6.6 times more wallets and 7.4 times more links. This trend shows that wallets and credit links grow at a similar rate and new wallets enter the Ripple network by connecting to a few existing wallets.

Table 1 shows the evolution of standard graph metrics for graph-{2013-2016}. We observe that most properties remain stable over the Ripple network lifetime except density, which has continuously decreased. Since the ratio $\mathbb{E}/\mathbb{V}$ has been constant in the Ripple network, the density grows as $\frac{2}{|\mathbb{V}|-1}$, and therefore decreases as the number of wallets increases. Thus, the Ripple network is a sparse graph.

**Network Motifs.** Recent work [21, 50] shows that high-order connectivity patterns or *motifs* (i.e., a subgraph composed of three nodes connected via a certain pattern of two or three edges) are important to understand the structure of a graph. We classify wallets into three groups: gateways, market makers and users, and color them accordingly in graph-2016. Using the *FANMOD* tool [51] on the colored graph-2016, we find that the motif depicted below is the most frequently occurring (88.07%) (out of 152 total possible 3-node motifs). This shows that the Ripple network has gateways as key players, which is consistent with the low clustering coefficient and disassortativity properties in Table 1.

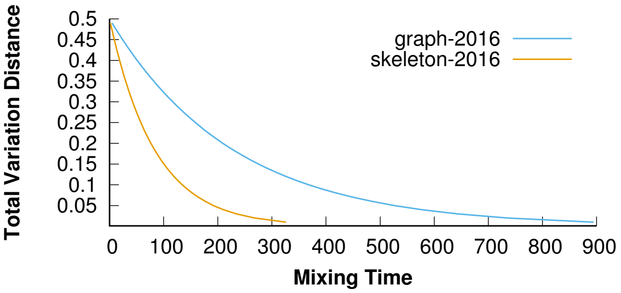

Figure 2: Lower bound of the mixing time for graph-2016 and skeleton-2016.

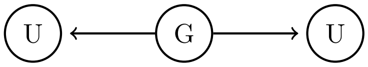

**Mixing Time.** We compute a lower bound on the mixing time of the Ripple network using the second largest Eigenvalue of the transition matrix for the graph as described by Mohaisen et al. [40] (Figure 2). We make two observations. First, the lower bound on the mixing time for an $\epsilon = 0.10$ is 380. This slow-mixing property is similar to the one observed for social networks [40], which is not surprising given the small clustering coefficient of the Ripple network. Second, the mixing time decreases if we consider skeleton-2016, a phenomena also observed for social networks [40].

## 4.2 Transaction Statistics

We count the number of intermediate wallets used in transactions included in txs-{2013-2016} in Figure 3. Paths with 0 to 3 intermediate wallets comprise 95% of all paths. Transactions with zero hops are used by

Table 1: Graph metrics for the Ripple network topology for different snapshots.

| | Dec'13 | Dec'14 | Dec'15 | Dec'16 |
|---|---|---|---|---|
| Avg Degree | 3.12 | 3.52 | 3.43 | 3.53 |
| Clustering | 0.07 | 0.07 | 0.07 | 0.12 |
| Assortativity | −0.49 | −0.35 | −0.33 | −0.43 |
| Density | $2 \cdot 10^{-4}$ | $8 \cdot 10^{-5}$ | $5 \cdot 10^{-5}$ | $3 \cdot 10^{-5}$ |

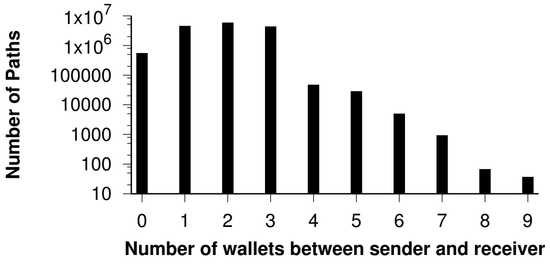

Figure 3: Distribution of number of wallets included in each path for the transactions in txs-{2013-2016}.

a wallet to open and close a credit link, and by a cold wallet to top-off a hot wallet [1, 42]. The remaining transactions require at least one hop as a wallet does not have a direct credit link to every possible receiver. As the credit available on a path is the minimum credit among the links in the path, short paths are preferred for transactions. We also count the number of paths used in transactions and observe that 99% of transactions use between 1 and 5 paths. Interestingly, 19% of transactions are circular (i.e., sender and receiver are the same wallet). Circular transactions can contain paths with market makers whose exchange offers make the overall transaction profitable for the sender [4].

## 4.3 Community Structure

**Community Detection and Localization.** We first extract communities using the Louvain community detection algorithm [22] as implemented in the Gephi software [20]. We use graph-2016 as input, and set the resolution parameter to 0.45. The resolution parameter determines the granularity in the search for communities: lower values of this parameter result in smaller communities that trivially form around a single gateway, and higher values result in larger communities containing several gateways that may be geographically located far apart. Second, we map each gateway included in gateways-2016 to its geographical location based on the information included in their corresponding websites. Finally, we map the communities extracted in the first step to their geographical location: we associate a community with the location of the gateway(s) contained in the community, since several gateways require users to provide identity and address verification documents before they populate links to them.

Figure 4 (column Dec16) depicts the results of this experiment. Mixed-1 refers to Latin America, Singapore, Indonesia and Canada and Mixed-2 refers to Europe, Latin America and China. As of Dec '16, the largest community centers around the gateway RippleIsrael located in Israel, followed by key gateways in Europe and Asia. Note that we have only shown communities that we were able to geo-locate. We observe, nonetheless, that communities that were omitted are much smaller.

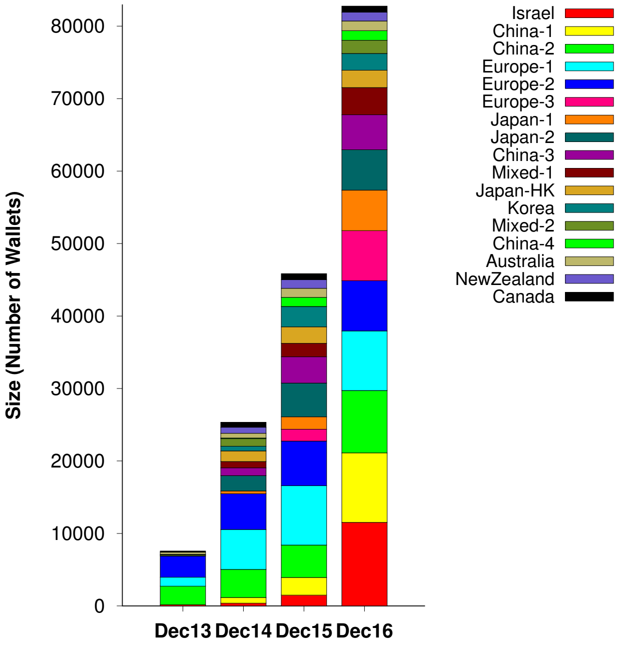

Figure 4: Distribution of communities over time. Each stack shows the size of the community in that snapshot. Various communities located in the same region are additionally labeled with a number.

To validate our results, we repeat the community detection process using the Walktrap Detection algorithm [43] as implemented in python-igraph [24]. We observe that 72% of the wallets were associated with the same community as in Figure 4 (column Dec 16).

**Temporal Evolution of Communities.** We repeat the process described above for each graph in graph-{2013-2016}. The geographical distribution of communities in each year is depicted in Figure 4. We observe two distinct trends. First, several communities (e.g., Israel, China-1 and China-2) have significantly grown over time. The growth of these communities is due to the addition of newly created wallets every year rather than a shift of wallets from other communities. In general, we find that wallets tend to stay in the community they initially joined. Second, several of the smallest communities have not grown over time (e.g., Canada, New Zealand and Australia).

## 4.4 Summary

Ripple user communities form by connecting to gateways in the same geographical region. This is a result of the identity verification process enforced by many gateways. Despite the pseudonymous nature of Ripple wallet identities, this geography of communities can simplify identification tasks for regulation and law-enforcement authorities. However, the identification process before a new credit link is created and funded reduces the number of credit links in the Ripple network. This results in a slow-mixing, unclustered, disassortative network. The slow-mixing property is similar to other networks where link creation requires physical interaction [28].

## 5. SECURITY AND RESILIENCE

In this section, we analyze the network liquidity, the resilience to faulty wallets and the effect of unexpected balance shifts (so-called "rippling").

### 5.1 Ripple Liquidity

We say that a pair of wallets (sender, receiver) has *liquidity* if the amount of credit that can be transferred between them is only bounded by the credit available in either the sender or receiver credit links. To compute this, first, we prune from multigraph-2016 the credit links associated with a currency other than {USD, CNY, BTC, JPY, EUR}, and convert the balance on the remaining credit links to USD using publicly available exchange rates. We select these five currencies since they comprise 62% of the original credit links. We denote the largest connected component of the thereby constructed subgraph as pruned-multigraph-2016.

Second, we transform pruned-multigraph-2016 to denote how much credit can be transferred between wallets instead of how much credit one wallet owes to its counterpart, as described by Dandekar et al. [25]. For example, the credit link Gatehub → Edward with balance \$5 and limit \$100 in Figure 1, results in two credit links: Gatehub → Edward with value \$95 and Gatehub ← Edward with value \$5. Following this approach, Figure 1 can be transformed into Figure 5. We denote this transformed graph as liquidity-multigraph-2016.

To check the liquidity in liquidity-multigraph-2016, we randomly pick $10,000$ pairs of wallets avoiding repetitions, and for each pair ($w_1$, $w_2$), calculate the max-flow from $w_1$ to $w_2$. We observe that 92% of the pairs of wallets have liquidity. In other words, the max-flow value between wallets is determined by the credit value available on either $w_1$'s credit links or $w_2$'s credit links. Therefore, the core of the Ripple network provides high liquidity and the bottleneck for transactions is the credit links from the users. In terms of liquidity, the Ripple network is similar to the current banking system, where the major banks hold more credit than their customers.

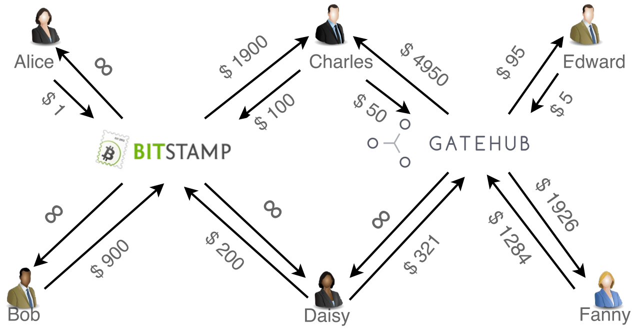

Figure 5: Example graph for liquidity experiment. The edge weight shows the credit that can be transferred from the source of an edge to its destination.

### 5.2 Resilience to Faulty Wallets

If an adversary compromises a highly connected and active wallet (e.g., a gateway) and disables rippling on all its credit links (as we explain in Section 5.3), transactions through the compromised wallet will no longer be possible. This could not only severely affect the liquidity of the network, but also lead to monetary losses to the neighboring wallets. Here, we study the resilience of the network to such adversarial wallet compromises.

We select 100 candidate faulty wallets from graph-2016 according to two different criteria: (i) wallets with highest degree (100-deg) and (ii) wallets involved in most of the transactions (100-ftx). We assess the most disruptive set of wallets by removing them from graph-2016 and observing how the network connectivity is affected. Figure 6 depicts the size of the largest connected component after removing the wallets in 100-deg and 100-ftx. Intuitively, the smaller the component, the fewer transactions are possible, as only wallets in the same component can transact with each other. From this experiment, we conclude that wallets included in 100-deg have a more profound impact on the connectivity of the Ripple network (and therefore on the transactions) than wallets included in 100-ftx. Therefore, we use 100-deg in the rest of this section.

We define the resilience factor, rsl-factor, as the ratio between the component size in the most disruptive splitting of the network after removing a wallet (i.e., splitting the network in two components of equal size) and the size of the actual largest component after removing a wallet. Therefore, the rsl-factor can take values in the range $[0.5, 1]$. Values close to 1 indicate that

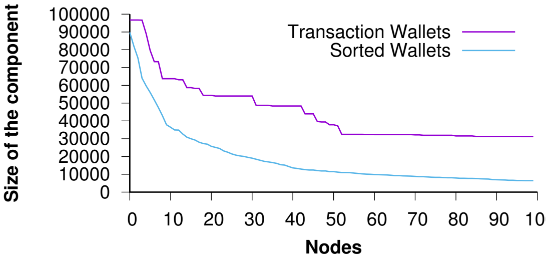

Figure 6: Size of the largest connected component in the Ripple network after removing nodes according to two different criteria.

the network has a low resilience as the removal of a wallet resulted in a component with (close to) half of the wallets of the network. Conversely, values close to 0.5 indicate that the network has a high resilience, as the largest component after removing a wallet is (close to) the entire graph. We observe that rsl-factor in the Ripple network is maintained in the range $(0.5, 0.6)$ after the removal of each wallet in 100-deg, demonstrating that the core of the Ripple network has high resilience.

In conclusion, we can divide the Ripple network into: (1) A small network core of 20% of the wallets that includes the key wallets with high connectivity. This core is highly resilient to the removal of highly connected wallets, and (2) A large set of wallets that can be easily disconnected from the network after removal of key wallets, forming components of very small sizes.

### 5.3 The Effect of Unexpected Balance Shifts

In the Ripple community, *rippling* denotes the redistribution of credit on the credit links for each intermediate wallet as a consequence of a transaction [14]. For instance, in Figure 1 consider a transaction from Bob to Edward through Charles for a value of $40. Among other changes, this transaction increases the balance of the credit link Bitstamp → Charles to $140 and reduces the balance of the credit link Charles ← Gatehub to $10, so that $40 are shifted between the links of Charles.

Although rippling maintains the net balance of intermediate wallets, its use is not innocuous for intermediate wallets [3]. The main issue is that the actual market value and stability of the credit depends on the issuer of such credit. In our running example, Charles may trust the credit from Gatehub more than Bitstamp. Therefore, a transaction involving rippling can induce a redistribution of credit from a more valuable to a less valuable issuer without the specific consent of the involved wallet's owner. We expect gateways to allow rippling; however, less active users may wish to avoid balance shifts not initiated by them.

As a countermeasure, each credit link is associated with a flag no_ripple. When no_ripple is set, the corresponding credit link cannot be part of a rippling operation. This flag was first added in December 2013, and was updated in March 2015 to establish its default state as "set," so users could selectively opt-out. Additionally, a wallet has a new flag defaultRipple that, if set, enables rippling among all the wallet's credit links. Gateway wallets, for instance, follow this pattern [13]. We study the effectiveness of this countermeasure here. First, the credit links not including no_ripple flag are tagged as no_ripple = false. Second, for each wallet that has the defaultRipple flag set, we set no_ripple = false (i.e., rippling is allowed) on all its credit links. Third, we use the no_ripple flag for the remainder of the links as specified in the multigraph-2016 dataset.

We say that a wallet is prone to rippling if it has at least two credit links with no_ripple = false (i.e., they allow rippling) and they use the same currency. We find that $21,604$ wallets (22.3%) are prone to rippling and are not associated with well-known gateways. Moreover, $30,870,298$ USD are prone to rippling, counting only the credit links that wallets prone to rippling have directly with gateways, as they are directly associated with real-world deposits. This demonstrates that unexpected balance shifts in the Ripple network can still affect a significant number of wallets, and more importantly, their credit.

We also observe that many wallets prone to rippling maintain credit links with a low balance (even zero), but with upper limit set to a value larger than zero. Those credit links can be used to shift the balances of wallets, and should have their upper limit as zero to void them.

### 5.4 Summary

The core of the Ripple network, consisting of around $10,000$ wallets, provides high liquidity to the remaining wallets and is extremely resilient. Still, around $50,000$ wallets are highly vulnerable to disruption by as few as 10 wallets, and their credit with the gateways (a total of $14,338,105$ USD) is at risk. This result shows that the Ripple network still has a few wallets that are "too big to fail," and, as a countermeasure, it is necessary for many users to increase their connectivity to avoid losses due to the failure of a handful of wallets.

Additionally, rippling among wallets in the Ripple network is not innocuous for intermediate wallets and can jeopardize a total of $30,870,298$ USD on credit links with the gateways that are additionally prone to rippling. As a countermeasure, wallets prone to rippling should appropriately set the limit on their credit links to the expected balance, or to 0 on credit links no longer being used for financial transactions.

## 6. CONCLUSIONS

Although the core of the Ripple network is resilient and not susceptible to rippling and other attacks, a large number of users is vulnerable. At the same time, thanks to the locality of communities, there is hope to tackle these vulnerabilities through geo-political forces. The users tend to stay bound to the same geographical community, elevating the importance of gateways in shaping the Ripple network. Further, though the core of the network is liquid, users can be affected by the disruption of a handful of nodes, and hence are advised to add credit links. Finally, although this work focuses on the Ripple network, we believe that our findings are relevant to other emerging credit networks (e.g., Stellar [12]) and credit network-based systems [34, 39, 44] that leverage similar design principles and may therefore present similar structural patterns and vulnerabilities.

## 7. REFERENCES

[1] Becoming a Ripple Gateway. Ripple blog. http://ripple.github.io/ripple-dev-portal/tutorial-gateway-guide.html#hot-and-cold-wallets.

[2] CryptoCurrencies Market capitalization. https://coinmarketcap.com/.

[3] Discussion on No Ripple Flag. Ripple blog. https://www.xrpchat.com/topic/3124-noripple-flag/#comment-30837.

[4] Example of code for pseudo circular payments. Forum entry. https://www.xrpchat.com/topic/2910-example-of-code-for-pseudo-circular-payments/.

[5] GateHub Announces Ripple Gateway as a Service. Gatehub blog entry. http://blog.gatehub.net/post/121925502092/gatehub-announces-ripple-gateway-as-a-service.

[6] How EarthPort and Ripple are teaming up to make cross-border payments instant. PYMNTS.com Blog. http://www.pymnts.com/in-depth/2015/how-earthport-and-ripple-are-teaming-up-to-make-cross-border-payments-instant/.

[7] Ripple API to get all gateways. https://github.com/ripple/rippled-historical-database/blob/v2.0.4/api/gateways/gateways.json.

[8] Ripple Charts. https://charts.ripple.com/#/.

[9] Ripple Data API v2. https://ripple.com/build/data-api-v2/.

[10] Ripple Network Banks. Ripple blog. https://ripple.com/network/financial-institutions/.

[11] Ripple Website. https://ripple.com.

[12] Stellar Website. https://www.stellar.org/.

[13] Technical Report on Ripple Flag. Ripple blog. https://ripple.com/files/GB-2015-04.pdf.

[14] Understanding the NoRipple Flag. Ripple blog. https://ripple.com/build/understanding-the-noripple-flag/.

[15] Use Case: Corporate Disbursements. Blog entry. https://ripple.com/solutions/corporate-disbursements/.

[16] Use Case: Retail Remittances. Blog entry. https://ripple.com/solutions/retail-remittances/.

[17] Amit. Bank-Wise Analysis of Blockchain Activity. Let's Talk Payments Blog, Aug 2015. http://letstalkpayments.com/bank-wise-analysis-of-blockchain-activity.

[18] Armknecht, F., Karame, G. O., Mandal, A., Youssef, F., and Zenner, E. *Ripple: Overview and Outlook.* Springer International Publishing, Cham, 2015, pp. 163–180.

[19] Barber, S., Boyen, X., Shi, E., and Uzun, E. *Bitter to Better — How to Make Bitcoin a Better Currency.* Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 399–414.

[20] Bastian, M., Heymann, S., and Jacomy, M. Gephi: An Open Source Software for Exploring and Manipulating Networks.

[21] Benson, A. R., Gleich, D. F., and Leskovec, J. Higher-order Organization of Complex Networks. *Science 353*, 6295 (2016), 163–166.

[22] Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment 2008*, 10 (2008), P10008.

[23] Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., and Felten, E. W. Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015* (2015), pp. 104–121.

[24] Csardi, G., and Nepusz, T. The igraph software package for complex network research. *InterJournal Complex Systems* (2006), 1695.

[25] Dandekar, P., Goel, A., Govindan, R., and Post, I. Liquidity in credit networks: A little trust goes a long way. In *Proceedings of the 12th ACM Conference on Electronic Commerce* (New York, NY, USA, 2011), EC '11, ACM, pp. 147–156.

[26] Dandekar, P., Goel, A., Wellman, M. P., and Wiedenbeck, B. Strategic formation of credit networks. *ACM Trans. Internet Technol. 15*, 1 (Mar. 2015), 3:1–3:41.

[27] DeFigueiredo, D. B., and Barr, E. T. Trustdavis: a non-exploitable online reputation system. In *Seventh IEEE International Conference on E-Commerce Technology (CEC'05)* (July 2005), pp. 274–283.

[28] Dell Amico, M., and Roudier, Y. A measurement of mixing time in social networks. In *STM 2009, 5th International Workshop on Security and Trust Management, September 24-25, 2009, Saint Malo, France* (Saint Malo, FRANCE, 09 2009). http://www.eurecom.fr/publication/2900.

[29] Donet Donet, J. A., Pérez-Solà, C., and Herrera-Joancomartí, J. *The Bitcoin P2P Network.* Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 87–102.

[30] Fugger, R. Money as IOUs in social trust networks & a proposal for a decentralized currency network protocol. *Hypertext document. Available electronically at http://ripple.*

*sourceforge. net*.

[31] Higgins, S. Bitstamp to Launch New Ripple Trading Pairs. CoinDesk blog entry. http://www.coindesk.com/bitstamp-launch-new-ripple-trading-pairs/.

[32] Higgins, S. US Banks Announce Ripple Protocol Integration. CoinDesk - Blog Entry. http://www.coindesk.com/us-banks-announce-ripple-protocol-integration/.

[33] Hunt, C. How Marco Montes is Empowering Migrant Workers. Ripple Blog. https://ripple.com/blog/how-marco-montes-is-empowering-migrant-workers/.

[34] Kakhki, A. M., Kliman-Silver, C., and Mislove, A. Iolaus: securing online content rating systems. In *22nd International World Wide Web Conference, WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013* (2013), pp. 919–930.

[35] Liu, A. Santander: Distributed Ledger Tech Could Save Banks $20 Billion a Year. Blog entry. https://ripple.com/insights/santander-distributed-ledger-tech-could-save-banks-20-billion-a-year/.

[36] Long, M. Santander Becomes the First U.K. Bank to Use Ripple for Cross-Border Payments. Ripple blog, May 2016. https://ripple.com/insights/santander-becomes-first-uk-bank-use-ripple-cross-border-payments/.

[37] Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. A fistful of bitcoins: Characterizing payments among men with no names. In *Proceedings of the 2013 Conference on Internet Measurement Conference* (New York, NY, USA, 2013), IMC '13, ACM, pp. 127–140.

[38] Meyer, D. More Banks Are Trying Out Blockchains For Fund Transfers. Fortune - Blog entry. http://fortune.com/2016/06/23/ripple-blockchain-banks/.

[39] Mislove, A., Post, A., Druschel, P., and Gummadi, P. K. Ostra: Leveraging trust to thwart unwanted communication. In *5th USENIX Symposium on Networked Systems Design & Implementation, NSDI 2008, April 16-18, 2008, San Francisco, CA, USA, Proceedings* (2008), pp. 15–30.

[40] Mohaisen, A., Yun, A., and Kim, Y. Measuring the Mixing Time of Social Graphs. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement* (New York, NY, USA, 2010), IMC '10, ACM, pp. 383–389.

[41] Moreno-Sanchez, P., Ruffing, T., and Kate, A. PathShuffle: Mixing Credit Paths for Anonymous Transactions in Ripple. *Proceedings on Privacy Enhancing Technologies 2017*, 3 (July 2017). To appear.

[42] Moreno-Sanchez, P., Zafar, M. B., and Kate, A. Listening to whispers of ripple: Linking wallets and deanonymizing transactions in the ripple network. *Proceedings on Privacy Enhancing Technologies 2016*, 4 (2016), 436–453.

[43] Pons, P., and Latapy, M. Computing communities in large networks using random walks. In *International Symposium on Computer and Information Sciences* (2005), Springer, pp. 284–293.

[44] Post, A., Shah, V., and Mislove, A. Bazaar: Strengthening user reputations in online marketplaces. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation* (Berkeley, CA, USA, 2011), NSDI'11, USENIX Association, pp. 183–196.

[45] Rizzo, P. Fidor Becomes First Bank to Use Ripple Payment Protocol. CoinDesk - Blog Entry. http://www.coindesk.com/fidor-becomes-first-bank-to-use-ripple-payment-protocol/.

[46] Rizzo, P. Royal Bank of Canada Reveals Blockchain Trial With Ripple. CoinDesk - Blog Entry. http://www.coindesk.com/royal-bank-canada-reveals-blockchain-remittance-trial-ripple/.

[47] Rizzo, P. Japan's SBI Holdings Teams With Ripple to Launch New Company. CoinDesk - Blog entry, Jan 2016. http://www.coindesk.com/sbi-holdings-ripple-new-company/.

[48] Ron, D., and Shamir, A. *Quantitative Analysis of the Full Bitcoin Transaction Graph*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 6–24.

[49] Southurst, J. Australia's Commonwealth Bank Latest to Experiment With Ripple. CoinDesk - Blog entry, May 2015. http://www.coindesk.com/australia-commonwealth-bank-ripple-experiment/.

[50] Wernicke, S. A faster algorithm for detecting network motifs. In *International Workshop on Algorithms in Bioinformatics* (2005), Springer, pp. 165–177.

[51] Wernicke, S., and Rasche, F. Fanmod: a tool for fast network motif detection. *Bioinformatics 22*, 9 (2006), 1152–1153.