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

The Ripple credit network has emerged as a payment backbone with key advantages for financial institutions and the remittance industry. Its path-based IOweYou (IOU) settlements across different (crypto)currencies conceptually distinguishes the Ripple blockchain from cryptocurrencies (such as Bitcoin and altcoins), and makes it highly suitable to an orthogonal yet vast set of applications in the remittance world for cross-border transactions and beyond.

This work studies the structure and evolution of the Ripple network since its inception, and investigates its vulnerability to devilry attacks that affect the IOU credit of linnet users' wallets. We find that about 13M USD are at risk in the current Ripple network due to inappropriate configuration of the rippling flag on credit links, facilitating undesired redistribution of credit across those links. Although the Ripple network has grown around a few highly connected hub (gateway) wallets that constitute the core of the network and provide high liquidity to users, such a credit link distribution results in a user base of around 112,000 wallets that can be financially isolated by as few as 10 highly connected gateway wallets. Indeed, today about 4.9M USD cannot be withdrawn by their owners from the Ripple network due to PayRoutes, a gateway tagged as faulty by the Ripple community. Finally, we observe that stale exchange offers pose a real problem, and exchanges (market makers) have not always been vigilant about periodically updating their exchange offers according to current real-world exchange rates. For example, stale offers were used by 84 Ripple wallets to gain more than 4.5M USD from mid-July to mid-August 2017. Our findings should prompt the Ripple community to improve the health of the network by educating its users on increasing their connectivity, and by appropriately maintaining the credit limits, rippling flags, and exchange offers on their IOU credit links.

KEYWORDS

Ripple credit network; IOweYou (IOU); credit devilry; rippling; faulty gateways; stale exchange offers

ACM Reference Format:
Pedro Moreno-Sanchez, Navin Modi, Raghuvir Songhela, Aniket Kate, Sonia Fahmy. 2018. Mind Your Credit: Assessing the Health of the Ripple Credit Network. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3178876.3186099

WWW 2018, April 23–27, 2018, Lyon, France

ACM ISBN 978-1-4503-5639-8/18/04.
https://doi.org/10.1145/3178876.3186099

1 INTRODUCTION

The Ripple network [1, 12, 20, 34] conceptually differs from the plethora of flourishing cryptocurrencies because it simultaneously allows transactions across traditional fiat currencies, cryptocurrencies as well as user-defined currencies over IOU credit paths. Its inherent capability to perform cross-currency transactions in a matter of seconds for a small fee in a publicly verifiable manner paves the way for reducing costs of financial institutions and the remittance industry by billions of dollars [23]. Given that, early embracers of Ripple [13, 41] have been recently followed by a wave of financial institutions worldwide [24, 35, 37, 42, 43, 45], including 12 of the world's top 50 banks [28], remittance institutions [15, 40] and online exchange services for cryptocurrencies [14, 16].

Among early academic efforts, Armknecht et al. [1] and Di Luzio et al. [25] present basic statistics of the Ripple network usage such as transaction volume, and consider the centralized nature of the Ripple blockchain consensus process respectively. Moreno-Sanchez et al. [33, 34] focus on deanonymization attacks and privacy enhancing solutions for users. Nevertheless, the Ripple network is yet to get its due attention similar to Bitcoin [2, 6, 10, 27, 44] from the academic community. This is critical because the Ripple IOweYou credit network and path-based transactions over credit links clearly set it apart (structurally and functionally) from cryptocurrencies.

Security of the credit in the Ripple network has not been studied thus far. Yet, it is crucial at this juncture to determine how users are handling their credit in the Ripple network and, more importantly, identify potential vulnerabilities, and determine countermeasures and best practices for future usage. By analyzing the collected Ripple network data which includes 181,233 wallets and 352,420 credit links, as well as 29,428,355 transactions during the period Jan '13 – Aug '17, we make the following key contributions.

This work presents the first extensive, longitudinal study of the Ripple network and its transactions throughout its *complete* lifetime up to August 2017, shedding light on its evolution and analyzing its security. We characterize the Ripple network graph (Section 4). We show that the number of wallets and credit links has grown at a steady rate through 2016 with a sudden spike in 2017, in tune with wide adoption over the second quarter of 2017. The ratio between wallets and credit links has however remained constant and hence the network density is decreasing. The network is slow-mixing, unclustered and disassortative. We identify *gateway* nodes as the key players in the network today. Gateways are highly connected bootstrapping wallets trusted to set up links to new users. We show that wallets are dynamically grouped into geographically demarcated communities, where each community is defined by (on average) two gateway wallets. We find that the core of the Ripple network provides enough liquidity for transactions from other wallets.

We assess the security of the credit held by users in the Ripple network in three ways. First, we investigate the effect of undesired redistribution of credit, i.e., rippling (Section 5). We show that more than 11,000 wallets in the Ripple network are prone to rippling among their credit links if they are used in a transaction as intermediate wallets. We observe that credit links at risk are associated with more than 13M USD.

Second, we study the resilience to disruptive wallets in the Ripple network (Section 6) and observe that although the core of the Ripple network, composed of around 65,000 wallets, is resilient to disruptive wallets, there exists a large user base of more than 112,000 wallets that is prone to disconnection by as few as 10 highly connected gateway wallets, and their credit (currently about 42M USD) is at risk of being no longer connected to the main component of the Ripple network and thus of being stagnated. In fact, we delve into the effect caused by a disruptive wallet by analyzing the case of PayRoutes, a gateway tagged by the Ripple community as faulty [38]. We observe that as of Aug '17, more than 600 wallets still have credit issued by PayRoutes for around 4.9M USD that is stagnated (and cannot get transferred) as PayRoutes does not provide the rippling option for those wallets.

Finally, we study the effect that stale exchange offers have on the credit of market makers and users (Section 7). In particular, we observe that during a period of ten days in 2013, market makers put at risk around 250,000 USD due to stale offers, and 24 wallets were able to gain more than 7,500 USD by taking advantage of those offers. We find that this effect, caused by stale offers, not only continues in the current Ripple network but is amplified. In particular, during a period of one month in 2017, market makers put at risk at least 500,000 USD due to stale offers, and cunning users gained more than 4.5M USD.

Our work motivates the Ripple community to 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. Additionally, we encourage the market makers to frequently update their offers in the Ripple network, according to the corresponding exchange rates in the real world.

2 BACKGROUND

The Ripple blockchain has emerged in the landscape of financial networks as an alternative settlement backbone for financial institutions and the remittance industry. Ripple adoption is fueled by potential savings of more than 20 billion dollars per year [23]. At the time of writing, Ripple's market capitalization is third, only behind Bitcoin and Ethereum.

The Ripple network. With its roots in IOweYou credit networks [1, 8, 12, 34], 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. By default, the credit on a link is upper-bound by ∞, but the wallet owner (u_2 in our example) can customize it. Additionally, each wallet is associated with a non-negative amount of XRP. XRP is the native currency in Ripple, initially conceived perhaps for users to pay a small fee per transaction towards curbing denial of service attacks and unbounded wallet creation (or Sybil attacks).

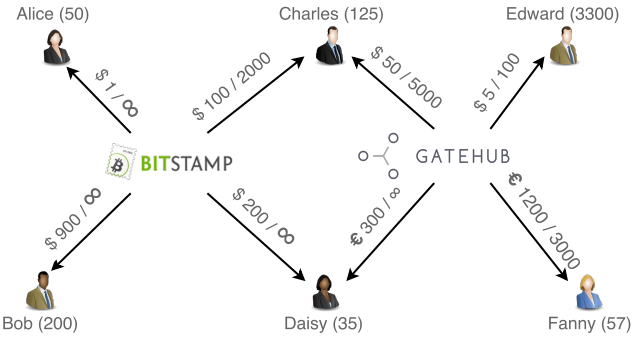

Figure 1: Illustrative example of the Ripple network. Each credit link is tagged with two values a / b, where a denotes the current credit and b denotes the upper bound. The edge lower bound is always zero. Numbers within parentheses denote the amount of XRP owned by each user.

Figure 1 depicts an excerpt of the Ripple network. Here, the credit link *Bitstamp* → *Alice* denotes that Bitstamp owes Alice 1 USD, and there is no upper-bound for such credit link. The credit link *Gatehub* → *Edward* denotes that Gatehub owes Edward 5 USD, and such credit can increase only up to 100 USD.

Ripple wallets. A wallet is governed by a pair of signing and verification keys from ECDSA or Schnorr signature scheme. An encoded version of the hash of the verification key identifies the wallet. Any operation associated with a wallet is only valid if it is signed by the corresponding signing key. Therefore, whoever holds the signing key for a wallet can do transactions with such wallet set as sender, create exchange offers or update its credit links.

Ripple transactions. Ripple allows two types of transactions: direct XRP payments and path-based settlement transactions. A direct XRP payment exchanges XRP between two wallets, even if they are not connected via a network path. The payment amount is subtracted from the sender's XRP balance and added to the receiver's XRP balance. Direct XRP payments thereby resemble debit payments between users rather than path-based credit settlements, which are the focus of this paper. Therefore, we omit direct XRP payments in our analysis and refer to [39] for further details.

A path-based settlement transaction (or simply a transaction hereby) uses a path of credit links between sender and receiver to settle credit between them. In the example of Figure 1, assume that Alice wants to pay Edward 1 USD. 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 ← Bitstamp → Charles ← GateHub → 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 USD, while reverse credit links are decreased by 1 USD. In the running example, Alice ← Bitstamp and Charles ← GateHub are decreased to 0 and 49 USD, respectively, whereas Bitstamp → Charles and GateHub → Edward are increased to 101 and 6 USD, respectively. Several paths between sender and receiver can be used in a single transaction [34].

Key players: Gateways and market makers. 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. Bitstamp and Gatehub are two examples gateways in the current Ripple network.

A *market maker* is a wallet that receives a certain currency on one of its credit links and exchanges it for another currency on another credit link, charging a small fee. Market makers enable transactions where senders and receivers hold different currencies. For instance, in Figure 1 assume that Bob wishes to pay 100 EUR to Fanny by spending 120 USD. Further, assume that Daisy has published an exchange offer USD/EUR at a rate 1.2 USD = 1 EUR. In this manner, Daisy plays the role of market maker and facilitates the transaction from Bob to Fanny: Bob ← Bitstamp is decreased by 120 USD while Bitstamp → Daisy is increased by 120 USD. Now, Daisy's offer is replenished, Daisy ← Gatehub is decreased by 100 EUR and finally, Gatehub → Edward is increased by 100 EUR.

Key operations: Rippling and exchange offers. In the Ripple community, *rippling* denotes the redistribution of credit on the links for each intermediate wallet as a consequence of a transaction [21]. Rippling can only occur between two credit links that belong to the same wallet and have credit in the same denomination. Nevertheless, several rippling operations can be concatenated to carry out a transaction with several intermediate wallets.

For instance, in Figure 1 consider a transaction from Bob to Edward through Charles for a value of 40 USD. Among other changes, this transaction decreases the balance of the link Bitstamp → Bob to 860 USD and increases the balance of Charles ← Bitstamp to 140 USD, so that 40 USD are shifted between the links of Bitstamp due to rippling. We expect that rippling is allowed by gateways; however, less active users may opt for avoiding balance shifts not initiated by them.

An *exchange offer* is created by a wallet to indicate its willingness to exchange one currency for another. Such wallet is then identified as *market maker*. The typical exchange offers are of the type described above, where Daisy offers an exchange USD/EUR. However, the Ripple network also allows offers that involve XRP. In fact, throughout the Ripple network lifetime, several market makers have included XRP in their offers, later fulfilled by wallets as part of a path-based transaction.

The combination of direct XRP payment and path-based transaction is natively supported in the Ripple network and they are atomically executed as a whole. As they involve the reallocation of credit among wallets, we consider this type of transaction in this work and denote it as a *path-based transaction involving XRP*. We stress, however, that this is different from single direct XRP payments, where only XRP is involved. For instance, a path-based transaction involving XRP can be used by a wallet to pay other wallets for performing transactions on its behalf. Assume now that in Figure 1, Fanny wants to pay Edward 1 USD. However, she only has credit in EUR and Gatehub has not indicated any exchange offer of the form EUR/USD. Instead, assume that Alice publishes an exchange offer of the type XRP/USD. In such a situation, Fanny can pay the amount of XRP corresponding to 1 USD to Alice, who in turn transfers 1 USD in the credit path connecting her to Edward.

3 DATASETS

Data sources. Our experiments are based on publicly accessible data extracted through the API [19] provided by the Ripple network on their servers *{s1, s2, data}.ripple.com*. We crawl the datasets describing the Ripple network topology (wallets and credit links among them), transactions, gateways and market makers. We summarize these datasets in the rest of this section and refer the reader to [18] for further statistics. The scripts for the data crawling and experiments are available at [32].

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, as well as at the end of August 2017. We model each snapshot as a directed graph with multiple edges between wallets (i.e., one edge per currency). For each snapshot, we only consider its largest connected component, and denote it by gr-year. We observe that gr-17 consists of 181, 233 wallets and 352, 420 credit links, which represent 98.6% of wallets and 99.28% of credit links of the total Ripple network at that point of time. Such percentages are similar for other snapshots. Therefore, we believe that gr-{13-17} are representative of the Ripple network snapshots.

Ripple transactions. We extract the transactions in the Ripple network in the period Jan '13 – Aug '17, obtaining a total of 29, 428, 355 transactions. We prune this dataset according to the following criteria. First, we discard 1, 530, 107 anomalous transactions (e.g., spam) considered outliers by previous studies [34]. Second, we discard 16, 180, 972 transactions carried out among wallets not included in gr-17. The majority of these transactions are XRP payments that do not require a credit path. Third, we discard 3, 255, 837 direct XRP payments among wallets in gr-17. We only consider path-based transactions, even if they are extended with a XRP payment. Our final set of transactions contains a total of 8, 461, 439 transactions. We refer to this as tx-{13-17}.

Market makers. We compiled the list of market makers present in gr-17 obtaining a set of 8, 105 wallets with at least one currency-exchange offer. We denote this dataset as mm-17.

Gateways. We crawled the list of gateways from the Ripple API, and added the gateways identified by the Ripple community throughout the Ripple network lifetime. As a result, we obtained a list of 101 gateways and 119 wallets associated with them. We denote this dataset by gw-17.

Ethical considerations. Our Ripple network analysis solely uses publicly available data. Moreover, we do not deanonymize any user that owns a wallet in the Ripple network or include sensitive data about them. We only give the names of gateways that are well-known in the Ripple community, and publicly advertised on websites and forums.

4 GRAPH CHARACTERISTICS OF THE RIPPLE NETWORK

In this section, we dissect our datasets to investigate the structure and evolution of the Ripple network throughout its lifetime.

Ripple network topology. Table 1 shows the Ripple network wallets and credit links as well as the evolution of standard graph metrics for gr-{13-17}. We make two observations. First, apart from the natural spike in the size of gr-14 due to the early stage of the system, wallets and credit links have grown in gr-14-16 at a steady rate of 1.55 ± 0.03 and 1.52 ± 0.07 correspondingly, a trend showing that wallets and credit links grow at a similar rate and new wallets enter the Ripple network by connecting to a few existing wallets. However, these ratios have soared in the first eight months of gr-17.

Second, we observe that most graph properties remain stable over the Ripple network lifetime except density, which has continuously decreased after gr-14. 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. This confirms the fact that the Ripple network is a sparse graph. We have validated these observations when considering the snapshot of the Ripple network every four months, considering thereby each economic quarter.

Ripple transactions. We first separate tx-{13-17} into two groups: (i) Transactions involving XRP and (ii) Transactions not involving XRP, obtaining 1,751,394 transactions in the first group and 6,710,045 transactions in the second group. This shows that although transactions involving XRP are supported in Ripple, they are not the norm.

We make the following observations on Ripple transactions. First, we observe that there exist 2,001,650 circular transactions where a wallet transfers credit to itself. Circular transactions can be used by a user, for instance, to transfer credit from one gateway to another. Alternatively, as we discuss in Section 7, circular transactions are used by cunning users to gain credit from stale offers. Second, we observe that there exist 2,136,387 non-circular and cross-currency transactions that exemplify the use of the Ripple network for remittance. Third, we observe that there exist 2,608,891 transactions that use at least one exchange offer available in the Ripple network. This demonstrates the importance of exchange offers.

Finally, we study the use of intermediate wallets in payment paths. We count 1,285,024 transactions that use 0 intermediate wallets and represent mainly deposit or withdrawal of credit with gateways. Moreover, we count 4,464,027 transactions that use a single intermediate wallet and represent, among others, interactions of gateways with their users following the hot-cold wallet mechanism [34]. Finally, we observe 2,712,388 transactions that use two or more intermediate wallets and represent cross-currency transactions. They exemplify the use of rippling and exchange offers in the Ripple network.

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

	gr-13	gr-14	gr-15	gr-16	gr-17
# wallets	14657	40051	61173	96953	181233
# credit links	26969	82305	119790	190675	352420
Avg degree	3.68	4.11	3.91	3.93	3.88
Clustering	0.08	0.08	0.08	0.13	0.07
Assortativity	−0.23	−0.15	−0.13	−0.16	−0.13
Density	$12 \cdot 10^{-5}$	$5.1 \cdot 10^{-5}$	$3.2 \cdot 10^{-5}$	$2.0 \cdot 10^{-5}$	$1.0 \cdot 10^{-5}$

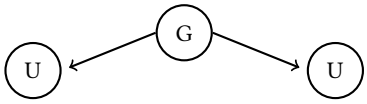

Figure 2: Most frequent motif in the Ripple network. U denotes user and G denotes gateway.

Ripple network structure. Recent work [4, 46] 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 in understanding the structure of a graph. We follow this strategy to study the Ripple network structure. For that, we first classify wallets in gr-17 into gateways, market makers and users, and color them accordingly. Using the *FANMOD* tool [47] on a colored version of gr-17 and parameters set to full enumeration, we find that the motif depicted in Figure 2 is the most frequently occurring with a frequency of 67.8%. 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.

Mixing time. As described by Mohaisen et al. [30], the mixing time in a graph represents how quick a random walk on the graph reaches the stationary distribution. In terms of a payment network as the Ripple network, mixing time intuitively determines how many intermediate wallets are required to reach a receiver wallet from any given sender wallet.

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 in [30] (Figure 3). We make two observations. First, the lower bound on the mixing time for an $\epsilon = 0.10$ is 730. This slow-mixing property is similar to the one observed for social networks [30], which is not surprising given the small clustering coefficient observed. The fact that the Ripple network is slow mixing increases the need for intermediate wallets to perform a transaction between any two wallets. Second, the

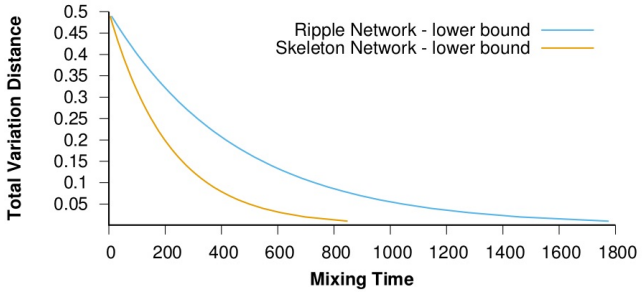

Figure 3: Lower bound on the mixing time for gr-2017 and skeleton-2017.

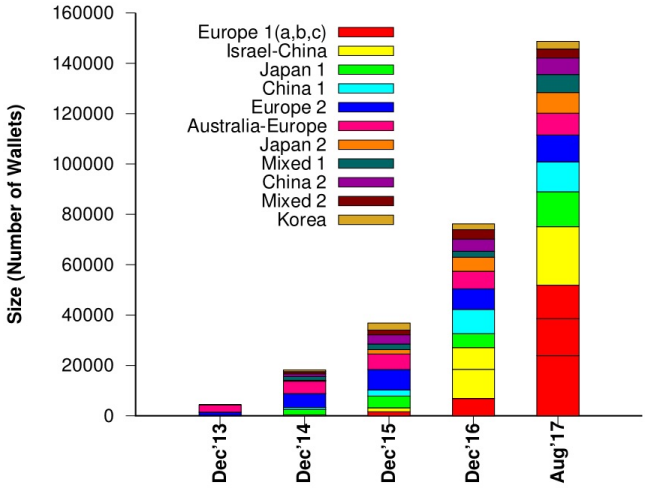

Figure 4: Distribution of communities over time. Each stack shows the size of the community in that snapshot. Communities located in the same region are additionally labeled with a number. Finally, split or merge of communities are represented by split stacks of the same color.

mixing time decreases if we consider the core of the Ripple network (skeleton-17), a phenomenon also observed for social networks [30]. Here, skeleton-17 is obtained by iteratively removing wallets with a single neighbor from gr-17. This shows that the core of the Ripple network has higher connectivity than the periphery.

Communities in the Ripple network. We next consider how wallets group into communities and how those communities have evolved over time. We extract the communities using the Louvain community detection algorithm [5] as implemented in the Gephi software [3] on input gr-17. The Louvain algorithm is parametrized by a *resolution* to determine the granularity in the search for communities. We set this parameter to 0.45 since we observe that lower values of the resolution result in smaller communities that trivially form around a single gateway, whereas higher values result in larger communities containing several gateways that may be geographically located far apart. For the chosen parameter, we have extracted 77 communities of sizes ranging from 3 to 23869 wallets.

We then derive the geographical location for each community to shed light on the community structure of the Ripple network. We opt for geographical location since most gateways require users to provide identity and address verification documents before they populate links to them and may restrict users by geographical location. Towards this goal, we first map each gateway included in gw-17 to its geographical location based on the information included in their corresponding websites; and second, we map a community to the location of the gateway(s) contained in the community, discarding 61 communities that are not associated with a known gateway. The discarded communities are the smallest communities we found, with sizes ranging from 3 to 999 wallets.

Figure 4 depicts the results of this experiment. Mixed-1 refers to Japan, HongKong, Turkey, NewZealand, and Mixed-2 refers to Sanghai, Canada, Indonesia, Singapore, Latin America. As of Aug '17, the largest community centers around Gatehub (Europe 1(a) in Figure 4), a key gateway in Europe, followed by a community represented by PayRoutes and RippleChina.

To understand the evolution of the communities, we repeat the experiment with gr-13-16. We observe that communities are dynamic. In fact, the community tagged as Europe-1 in Dec '16 has been split into three communities Europe1(a,b,c) as of Aug '17, built around three emerging gateways in Europe. Conversely, communities separately built around PayRoutes and RippleChina in Dec '16 have merged together in Aug '17. We believe that this phenomenon corresponds to the growing activity between wallets of these two gateways.

In 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 [9].

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 on either the sender or receiver credit links. We now study whether credit in the Ripple network effectively facilitates transactions among Ripple wallets.

First, we prune from gr-17 the credit links associated with a currency other than {USD, CNY, BTC, JPY, EUR}, extract the largest connected component of the pruned graph, and convert the balance on the remaining credit links to USD using publicly available exchange rates. We select these five currencies since they are the most common, comprising more than 65% of the original credit links. We denote this processed subgraph as pruned-g-17.

Second, we transform pruned-g-17 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. [7]. For example, the credit link Gatehub → Edward with balance

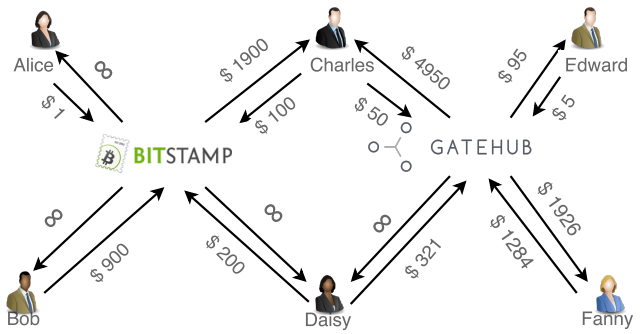

Figure 5: Example graph for liquidity experiment. Each edge weight shows the credit that can be transferred from the source of the edge to its destination.

$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 by liq-g-17.

Finally, we check the liquidity on liq-g-17 by randomly picking a representative sample of the Ripple network consisting of 10,000 pairs of wallets avoiding repetitions, and for each pair (w_1, w_2), we calculate the max-flow from w_1 to w_2. We observe that 92.55% 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.

In conclusion, the core of the Ripple network provides high liquidity and the bottleneck for transactions are 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.

5 RIPPLING AND USERS: THE EFFECT OF UNEXPECTED BALANCE SHIFTS

Although rippling maintains the net balance of intermediate wallets, its use is not innocuous for intermediate wallets. The main issue is that the actual market value and stability of the credit depends on the issuer of such credit. In our illustrative example of Ripple network in Figure 1, Charles may trust the credit from Gatehub more than Bitstamp. Therefore, a transaction involving rippling among the two corresponding credit links 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 have a default state of "set" (i.e., no rippling allowed by default), so users could selectively opt-out and allow rippling. Additionally, a wallet has a new flag called defaultRipple that, if set, enables rippling among all the wallet's credit links. Gateway wallets, for instance, follow this pattern [17].

Goal. In this experiment, we aim to identify wallets other than gateways that allow rippling, and to extract how much credit they put at risk doing so.

Methodology. 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 gr-17 dataset. Now, 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 hold credit in the same currency.

Results. We find that more than 11,000 wallets are prone to rippling and are not associated with well-known gateways. Moreover, more than 13M USD are prone to rippling, counting only the credit links that wallets prone to rippling have directly with gateways, as they are associated with real-world deposits. This gives a lower bound on the amount of credit at risk, and the actual value could be higher, if we count credit at risk with wallets other than the gateways. This result 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. The gap between balance and upper credit limit on these credit links can be used to shift the balances of wallets, thus increasing the risk.

Countermeasures. The users have the possibility of disabling the rippling functionality on their credit links completely. Therefore, less active users may opt for disabling rippling among their credit links to avoid balance shifts not initiated by them. Moreover, more active users can also opt for dynamically adjust the amount of credit prone to rippling and add a rippling fee to it. Finally, users with credit links holding zero balance should reduce their upper limit to effectively void them.

6 RIPPLING AND GATEWAYS: THE EFFECT OF FAULTY GATEWAYS

The gateway wallets are highly connected wallets included in the core of the Ripple network and significantly contribute to the liquidity of the network. A faulty gateway can disable rippling on most credit links of its wallet, ensuring that transactions routed through it are no longer possible and effectively freezing the balance held at credit links of its wallet [39, 48]. This would not only severely affect the liquidity of the network, but also lead to monetary losses to the neighboring wallets, as they no longer can use the credit issued by the compromised wallet.

Goal. We aim to study the effect of faulty gateway wallets (e.g., as a result of adversarial wallet compromise) and the resilience of the Ripple network to them.

Methodology. We select 100 candidate faulty wallets from gr-17 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 gr-17 and observing how the network connectivity is affected. Figure 6 depicts the size of the largest connected component after

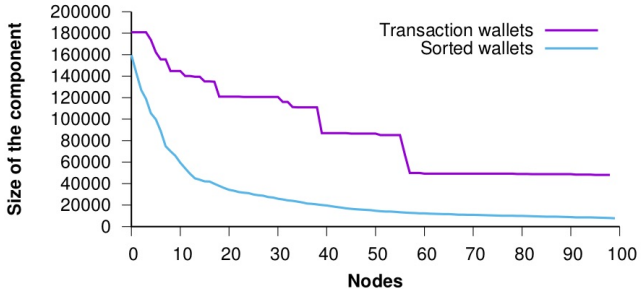

Figure 6: Size of the largest connected component after removing wallets sorted by number of credit links (blue) and number of appearances in transactions (purple).

removing the wallets in 100-deg and 100-ftx. Intuitively, the smaller the component, the fewer the possible transactions, since 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 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.

Results. We observe that the rsl-factor of 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. We conclude that we can divide the Ripple network into: (1) A small network core of around $65,000$ wallets (36% of the total) 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 around $112,000$ wallets that can be easily disconnected from the network after removal of key wallets. Yet, these highly vulnerable wallets have more than 42M USD of credit with the gateways, which is at risk.

Countermeasures. This result shows that the Ripple network still has a few wallets that are "too big to fail." As a countermeasure, it is necessary for many users to increase their connectivity and split their credit among different credit links to avoid losses due to the failure of a handful of wallets.

6.1 A case study: The PayRoutes gateway

While studying the Ripple network communities (see Section 4), we observed that the size of the community created around the PayRoutes gateway suddenly increased in Dec '16. Surprisingly, users in the Ripple community had reported the unresponsiveness of the company running the gateway when contacted regarding the credit issued by it [38]. We also emailed them, but got no answer at the time of this writing. In this state of affairs, we study PayRoutes as an example of a faulty gateway.

Goals. We consider two questions. First, we aim to find the amount of credit in the Ripple network that can only be withdrawn with the cooperation of PayRoutes and, given the unresponsiveness of the gateway, this credit is "stuck" in the Ripple network. Second, we study why wallets with stuck credit obtained it in the first place, even though PayRoutes was already reported as faulty. We describe our methodology and results for each goal separately in the following two sections.

6.1.1 Credit with PayRoutes. We are interested in credit links of the form PayRoutes $\rightarrow w_i$ where PayRoutes has disabled rippling. This implies that the credit on these links can only be used in a withdrawal operation jointly with PayRoutes: w_i sets the credit on the link to 0 to obtain the corresponding amount in the real world from PayRoutes. However, as PayRoutes is a faulty gateway, this operation is no longer available and the credit is stuck. Given that, we first address the question: *how much credit is stuck on credit links with PayRoutes?*

Methodology. From gr-17, we first pick the credit links with PayRoutes as counterparty and positive balance, and derive the status of their rippling flag (as described in Section 5). Then, we classify the neighbor wallets of PayRoutes into two groups as follows. First, we identify those wallets that have a credit link with PayRoutes for which rippling is not allowed, i.e., no_ripple is set to true. We denote this set of wallets by *wallets-no-rippling*. Second, we consider the set of wallets that are not in *wallets-no-rippling* but yet cannot perform a transaction for an amount equal to the balance on their credit link with PayRoutes. We denote this second set as *wallets-rippling-no-tx*. As the wallets in either *wallets-no-rippling* or *wallets-rippling-no-tx* cannot transfer the (entire) credit they have on a credit link with PayRoutes to another wallet in the Ripple network, the only way for them to get their credit back is to contact PayRoutes in the real world and withdraw the corresponding funds. However, as PayRoutes is unresponsive, such credit is "stuck."

Results. We observe that, out of the $2,958$ wallets that have at least one credit link with PayRoutes, there exist 621 wallets in either *wallets-no-rippling* or *wallets-rippling-no-tx*, and therefore with stuck credit. We observe that the stuck credit on these credit links is around 4.9M USD.

Discussion. The PayRoutes case is not typical in the Ripple network. There have been other gateways that have ceased operation during the Ripple network lifetime, but have not caused such an effect. We consider DividendRippler as an example of such a gateway. The difference from PayRoutes is that before shutting down, DividendRippler publicly announced it and mandated its clients to proceed to withdraw the credit available in their credit links with DividendRippler.

We conduct the same above experiment for DividendRippler, and observe that, although 665 wallets have credit stuck with DividendRippler, such credit accounts for around $1,000$ USD only. This is how much DividendRippler currently owes to the rest of wallets. This demonstrates that wallets followed the announcement of the gateway and successfully managed to withdraw most of their credit before the gateway ceased operation.

6.1.2 Obtaining credit from PayRoutes. In this section, we focus on answering the question: *How did wallets with stuck credit obtain such credit in the first place?*

Methodology. We first investigate how new credit links were created with PayRoutes over the lifetime of the Ripple network. We observe a spike of $2,527$ credit links created in Oct '16 from a total of $1,805$ wallets. Out of these, 186 credit links were created by 133 wallets and have balance stuck in PayRoutes. This implies that 21% of the wallets with stuck balance created credit links with PayRoutes during that month. We denote these by stuck-wallets-Oct-16.

Table 2: Summary of the exchange offers between XRP and USD created in the Ripple network during October 2016.

Pay Val	Pay Cur	Get Val	Get Curr	Ratio
1062738.51	XRP	17009.50	USD	62.48 to 1
59678.62	USD	33194.62	XRP	1.78 to 1

Given this unusual behavior, we study how those 133 wallets obtained credit. We identify two possibilities: (i) A path-based transaction from another wallet in the Ripple network; (ii) A circular transaction (i.e., sender and receiver of the transaction are the same wallet), where a wallet pays a certain amount of XRP (or any currency issued by a gateway other than PayRoutes) in exchange for credit issued by PayRoutes on a credit link with it.

Results. We observe that wallets in stuck-wallets-Oct-16 do not receive significant credit from other wallets in the Ripple network during October 2016. In particular, we find only three transactions with credit values of 10 USD, 100 ILS and 5 ILS. Instead, wallets in stuck-wallets-Oct-16 get their credit through circular transactions. We find that 51 wallets perform a total of 286 circular transactions, where these wallets received around 12, 000 USD in exchange for approximately 300 CNY and 12, 000 XRP.

In essence, wallets in stuck-wallets-Oct-16 invested mostly XRP to obtain USD from PayRoutes. We find that the exchange rate XRP/USD in the Ripple network was considerably "better" than in the real world at that time: In the Ripple network at that time, a wallet could get 0.73 USD for 1 XRP on average, with a minimum of 0.14 and a maximum of 2.87 USD using stale offers available in the network. However, in the real world, one could get less than 0.01 USD for 1 XRP at the average exchange rate at that time and up to 0.28 USD for 1 XRP, even considering the best exchange rate over the entire Ripple network lifetime.

The results presented above describe the origin of a small fraction of the credit stuck on credit links with PayRoutes. We repeated the same experiment over the complete Ripple network lifetime and observed similar patterns. First, the credit links with stuck credit are involved in a total of 278 transactions where other wallets in the Ripple network are sending credit to victim wallets at a favorable rate: The receiver gets more credit than actually sent by the sender. Those transactions account for around 158, 000 USD. Second, the highest amount of credit is received as a result of circular transactions that use advantageous offers. In particular, we find that credit links with stuck credit are involved in a total of 16, 469 transactions where they gained more than 63M USD over the complete Ripple network lifetime.

Countermeasures. Although wallets with stuck credit at PayRoutes obtained considerable revenue, a broader perspective reveals that it was a risky operation. For instance, it is possible to check the exchange rates available in the Ripple network at October 2016 to determine how likely it is to get the USD credit back. In particular, we observe that although wallets in stuck-wallets-Oct-16 managed to get "cheap" USDs, the market values were not favorable to get them back: New exchange offers created in the Ripple network in October 2016 (as shown in Table 2) demonstrate this.

7 STALE OFFERS IN THE RIPPLE NETWORK

Exchange offers and rippling are the key operations that enable path-based transactions. The previous two sections investigated the security of rippling, so we now investigate the safety of exchange offers, which are set by the owners of wallets at their own discretion. Naturally, proposed offers should match those of the corresponding currencies in the real world or even be in favor of market makers so that they get credit for their exchange services. Otherwise, cunning users can leverage stale offers to gain credit, while market makers may go bankrupt. This would adversely impact the liquidity and availability of the Ripple network.

Goal. In this experiment, we aim to determine whether there are stale offers in the Ripple network and, if so, study to what extent devilry users have taken advantage of them.

Methodology. We search for sudden changes in a currency's market capitalization. We observed several such changes. We first examine a spike in the price of XRP in late 2013: during a period of ten days (Nov 20th–30th, 2013), the price of 1 XRP with respect to BTC increased by 380%, i.e., 1 XRP was exchanged at 0.00001 BTC at the beginning of the period but within a week, 1 XRP was exchanged at 0.000038 BTC. Given that, we extract from tx-{13-17} the transactions that occurred during this ten-day period, obtaining a total of 1, 932 transactions. We prune this dataset by considering only cross-currency transactions that transfer XRP for BTC or vice versa. We obtain a total of 112 transactions.

We compare the exchange rate between XRP and BTC used in each transaction to the exchange rate in the real world at the same time, as shown in Figure 7. In both (top and bottom) figures, a purple point represents the exchange rate in a Ripple transaction while the corresponding green point denotes the exchange rate in the real world at the same time. For both graphs, if the purple point is higher than the green point (Ripple's offer is more expensive than the real world offer), the market maker made money. In contrast, if the purple point is below the green point, the user who conducted the transaction gained credit.

Results. We analyzed the transactions in which a sender gained credit by exploiting stale offers. We make two observations. First, users could have gained up to around 250, 000 USD by fully exploiting XRP/BTC stale offers during the specified period. In other words, market makers put at risk around 250, 000 USD due to stale offers. Second, 24 different wallets made a monetary benefit of at least 7, 500 USD by exploiting XRP/BTC stale offers (and other offers available in the network at that time). Here, we calculate the USD value by converting the BTC and XRP to their real world exchange rates at the corresponding times. In summary, even in the nascent stages of the Ripple network, when the transaction volume was considerably low, stale offers risked significant loss of credit by market makers.

To confirm these results, we explored another, more recent, substantial change in a currency exchange rate. We found a sudden increase in the price of BTC compared to XRP in 2017, concretely during the period July 16th – August 16th: The value of 1 BTC went from 11, 713 XRP to 25, 735 XRP, that is, an increase of 120%. As before, we extracted the transactions during that period of time and compared the exchange rates of XRP from/to BTC in the Ripple

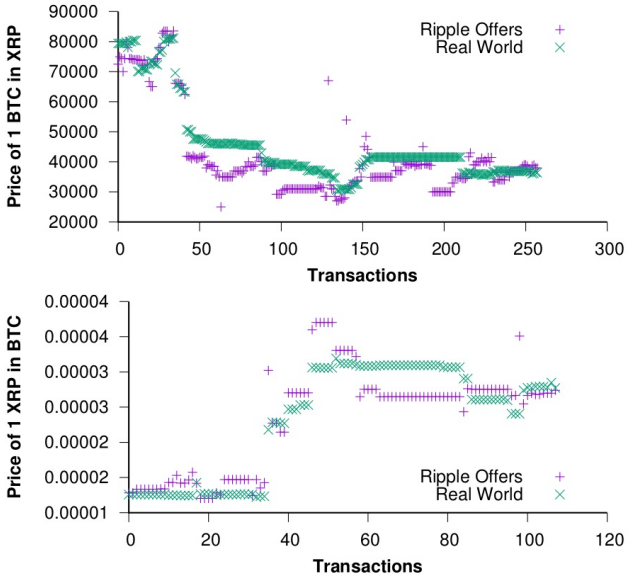

Figure 7: Market maker accepts XRP and pays BTC (top); market maker accepts BTC and pays XRP (bottom). If the purple point (offer in Ripple) is below the green point (offer in real world), the transacting user gained credit. Otherwise, the market maker gained credit. These transactions took place between November 20th and 30th, 2013.

network and in the real world. We observe that market makers put at risk around 500, 000 USD due to stale offers exchanging XRP to BTC and vice versa. Moreover, we observe that 84 wallets exploited these stale offers (and possibly other offers) to gain more than 4.5M USD. These results confirm that stale offers continue to be a risk for market makers. In fact, the effect of stale offers is now amplified given the growth of the Ripple network and transactions.

Countermeasures. A market maker can update a previously offered exchange rate at any time. Therefore, a market maker should continuously monitor the price for the currencies involved in its offers and promptly update its Ripple offers when a sudden change occurs in the real world. The gaps between exchange rates in the Ripple network and real world are thereby reduced, and with them, the windows for cunning users to gain credit.

8 RELATED WORK

Some research work [2, 6, 10, 27, 44] has studied Bitcoin and other cryptocurrencies. Although it is possible to extract lessons from that work, the conceptual differences between cryptocurrencies such as Bitcoin and the Ripple network mandate a dedicated look.

There is limited work studying path-based settlement networks. Moreno-Sanchez et al. [34] present the first detailed study of the Ripple network. In particular, they identify the privacy breaches of the publicly available ledger. Their study links wallets that belong to the same user and deanonymizes the transactions associated with the main gateways. Their follow-up work [33] presents a path-mixing protocol to allow anonymous transactions in the Ripple network, thereby mitigating privacy breaches.

Di Luzio et al. [25] consider two aspects of the Ripple network. They study the evolution of the amount and behavior of participants in the consensus protocol used to add transactions to the ledger during the first three years of the Ripple network. They also propose a novel technique to deanonymize the transactions of a given user, leveraging side-channel information (e.g., the amount of a recent transaction performed by the victim).

Armknecht et al. [1] present an overview of the Ripple network and give statistics about the number of transactions, and types of transactions and exchanges. The work is limited to the first two years of operation of the Ripple network. The work also demonstrates the conditions under which the Ripple consensus protocol fails, leading to a situation where the Ripple ledger might be forked.

In summary, related work studies two dimensions of the Ripple network: privacy and consensus. Some work such as [1, 25] also computes statistics about the network structure during the first few years of the network lifetime. In this work, we consider the consensus protocol and privacy as interesting but orthogonal dimensions to be studied, and instead focus on the evolution of the Ripple network and its vulnerabilities during the *complete* network lifetime through August 2017. We thoroughly study several *security* vulnerabilities and their implications on the Ripple network.

9 CONCLUDING REMARKS

The Ripple network has been gaining momentum, with substantial growth in the number of wallets and credit links in 2017. Yet, new wallets create credit links with only few other key wallets, primarily gateways. This makes the Ripple network slow-mixing, with wallets grouped in demarcated communities. The users tend to stay bound to the same geographical community, elevating the importance of gateways in shaping the Ripple network. The core of the network composed of around 65, 000 wallets provides sufficient liquidity for the remaining wallets.

The key operations in the Ripple network such as rippling and exchange offers pose important security challenges. Although the core of the Ripple network is resilient, a large number of users may be vulnerable to undesirable shift of credit among their credit links. Thanks to the locality of communities, there is hope to tackle these vulnerabilities through geo-political forces. Further, users can be affected by the disruption of a handful of nodes, as demonstrated in the case of PayRoutes, and hence are advised to add credit links. Last but not least, due to the importance of exchange offers in the current Ripple network, market makers are advised to periodically update their offers according to the real-world exchange rates, as they otherwise risk several hundreds of thousands of dollars.

Although this work focuses on the Ripple network, we believe that our findings are relevant to other emerging credit networks (e.g., Stellar [11]) and credit network-based systems [22, 26, 29, 31, 36] that leverage similar design principles and may therefore present similar structural patterns and vulnerabilities.

Acknowledgments. We thank the Ripple forum members and the anonymous reviewers for their feedback. This work is partially supported by an Intel/CERIAS RA and by the National Science Foundation under grants CNS-1319924 and CNS-1717493.

REFERENCES

[1] Frederik Armknecht, Ghassan O. Karame, Avikarsha Mandal, Franck Youssef, and Erik Zenner. 2015. Ripple: Overview and Outlook. In *Trust and Trustworthy Computing*. 163–180.

[2] Simon Barber, Xavier Boyen, Elaine Shi, and Ersin Uzun. 2012. Bitter to Better — How to Make Bitcoin a Better Currency. In *Financial Cryptography and Data Security*. 399–414.

[3] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. 2009. Gephi: An Open Source Software for Exploring and Manipulating Networks. In *Conference on Weblogs and Social Media*.

[4] Austin R. Benson, David F. Gleich, and Jure Leskovec. 2016. Higher-order Organization of Complex Networks. *Science* 353, 6295 (2016), 163–166.

[5] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008), P10008.

[6] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. 2015. SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *Symposium on Security and Privacy*. 104–121.

[7] Pranav Dandekar, Ashish Goel, Ramesh Govindan, and Ian Post. 2011. Liquidity in Credit Networks: A Little Trust Goes a Long Way. In *Conference on Electronic Commerce*. 147–156.

[8] D. B. DeFigueiredo and E. T. Barr. 2005. TrustDavis: a non-exploitable online reputation system. In *Conference on E-Commerce Technology*. 274–283.

[9] Matteo Dell Amico and Yves Roudier. 2009. A measurement of mixing time in social networks. In *Security and Trust Management*.

[10] Joan Antoni Donet Donet, Cristina Pérez-Solà, and Jordi Herrera-Joancomartí. 2014. The Bitcoin P2P Network. In *Financial Cryptography and Data Security*. 87–102.

[11] Stellar Foundation. 2017. Stellar Website. (2017). https://www.stellar.org/.

[12] R. Fugger. 2004. Money as IOUs in social trust networks & a proposal for a decentralized currency network protocol. (2004). http://library.uniteddiversity.coop/Money_and_Economics/decentralizedcurrency.pdf.

[13] Stan Higgins. 2014. US Banks Announce Ripple Protocol Integration. CoinDesk blog entry. (2014). http://www.coindesk.com/us-banks-announce-ripple-protocol-integration/.

[14] Stan Higgins. 2017. Bitstamp to Launch New Ripple Trading Pairs. CoinDesk blog entry. (2017). http://www.coindesk.com/bitstamp-launch-new-ripple-trading-pairs/.

[15] Chloe Hunt. 2015. How Marco Montes is Empowering Migrant Workers. Ripple blog entry. (2015). https://ripple.com/blog/how-marco-montes-is-empowering-migrant-workers/.

[16] Gatehub Inc. 2015. GateHub Announces Ripple Gateway as a Service. Gatehub blog entry. (2015). http://blog.gatehub.net/post/121925502092/gatehub-announces-ripple-gateway-as-a-service.

[17] Ripple Inc. 2015. Technical Report on Ripple Flag. Ripple blog entry. (2015). https://ripple.com/files/GB-2015-04.pdf.

[18] Ripple Inc. 2017. Ripple Charts. (2017). https://xrpcharts.ripple.com/#/.

[19] Ripple Inc. 2017. Ripple Data API v2. (2017). https://ripple.com/build/data-api-v2/.

[20] Ripple Inc. 2017. Ripple Website. (2017). https://ripple.com.

[21] Ripple Inc. 2017. Understanding the NoRipple Flag. Ripple blog entry. (2017). https://ripple.com/build/understanding-the-noripple-flag/.

[22] Arash Molavi Kakhki, Chloe Kliman-Silver, and Alan Mislove. 2013. Iolaus: securing online content rating systems. In *World Wide Web Conference*. 919–930.

[23] Alec Liu. 2015. Santander: Distributed Ledger Tech Could Save Banks $20 Billion a Year. Ripple blog entry. (2015). https://ripple.com/insights/santander-distributed-ledger-tech-could-save-banks-20-billion-a-year/.

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

[25] Adriano Di Luzio, Alessandro Mei, and Julinda Stefa. 2017. Consensus Robustness and Transaction De-Anonymization in the Ripple Currency Exchange System. In *Conference on Distributed Computing Systems*. 140–150.

[26] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, and Matteo Maffei. 2017. SilentWhispers: Enforcing Security and Privacy in Decentralized Credit Networks. In *Network and Distributed System Security Symposium*.

[27] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2013. A Fistful of Bitcoins: Characterizing Payments Among Men with No Names. In *Internet Measurement Conference*. 127–140.

[28] David Meyer. 2016. More Banks Are Trying Out Blockchains For Fund Transfers. Fortune blog entry. (2016). http://fortune.com/2016/06/23/ripple-blockchain-banks/.

[29] Alan Mislove, Ansley Post, Peter Druschel, and P. Krishna Gummadi. 2008. Ostra: Leveraging Trust to Thwart Unwanted Communication. In *Symposium on Networked Systems Design & Implementation*. 15–30.

[30] Abedelaziz Mohaisen, Aaram Yun, and Yongdae Kim. 2010. Measuring the Mixing Time of Social Graphs. In *Conference on Internet Measurement*. 383–389.

[31] Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Kim Pecina. 2015. Privacy Preserving Payments in Credit Networks. In *Network and Distributed System Security Symposium*.

[32] Pedro Moreno-Sanchez, Navin Modi, Raghuvir Songhela, Aniket Kate, and Sonia Fahmy. 2018. Mind Your Credit: Project Website. https://pedrorechez.github.io/Ripple-Credit-Study/. (2018).

[33] Pedro Moreno-Sanchez, Tim Ruffing, and Aniket Kate. 2017. PathShuffle: Credit Mixing and Anonymous Payments for Ripple. *Proceedings on Privacy Enhancing Technologies* 2017, 3 (2017), 110.

[34] Pedro Moreno-Sanchez, Muhammad Bilal Zafar, and Aniket Kate. 2016. Listening to Whispers of Ripple: Linking Wallets and Deanonymizing Transactions in the Ripple Network. *Proceedings on Privacy Enhancing Technologies* 2016, 4 (2016), 436–453.

[35] David Patterson. 2016. Ripple Network Banks. Ripple blog entry. (2016). https://ripple.com/network/financial-institutions/.

[36] Ansley Post, Vijit Shah, and Alan Mislove. 2011. Bazaar: Strengthening User Reputations in Online Marketplaces. In *Conference on Networked Systems Design and Implementation*.

[37] Amit (pseudonym). 2015. Bank-Wise Analysis of Blockchain Activity. Let's Talk Payments blog entry. (2015). http://letstalkpayments.com/bank-wise-analysis-of-blockchain-activity.

[38] Aapeli (pseudonym). 2016. Ripple Explodes Monday Morning Setting All-Time Best Market Cap - Glitch or Gain? XRPChat forum entry. (2016). https://www.xrpchat.com/topic/2187-ripple-explodes-monday-morning-setting-all-time-best-market-cap-glitch-or-gain.

[39] Bitmex Research (pseudonym). 2018. The Ripple story. BitMex blog entry. (2018). https://blog.bitmex.com/the-ripple-story/.

[40] PYMNTS (pseudonym). 2015. How EarthPort and Ripple are teaming up to make cross-border payments instant. PYMNTS.com blog entry. (2015). http://www.pymnts.com/in-depth/2015/how-earthport-and-ripple-are-teaming-up-to-make-cross-border-payments-instant/.

[41] Pete Rizzo. 2014. Fidor Becomes First Bank to Use Ripple Payment Protocol. CoinDesk - Blog Entry. (2014). http://www.coindesk.com/fidor-becomes-first-bank-to-use-ripple-payment-protocol/.

[42] Pete Rizzo. 2014. Royal Bank of Canada Reveals Blockchain Trial With Ripple. CoinDesk blog entry. (2014). http://www.coindesk.com/royal-bank-canada-reveals-blockchain-remittance-trial-ripple/.

[43] Pete Rizzo. 2016. 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/.

[44] Dorit Ron and Adi Shamir. 2013. Quantitative Analysis of the Full Bitcoin Transaction Graph. In *Financial Cryptography and Data Security*. 6–24.

[45] Jon Southurst. 2015. Australia's Commonwealth Bank Latest to Experiment With Ripple. CoinDesk blog entry. (2015). http://www.coindesk.com/australia-commonwealth-bank-ripple-experiment/.

[46] Sebastian Wernicke. 2005. A faster algorithm for detecting network motifs. In *Workshop on Algorithms in Bioinformatics*. 165–177.

[47] Sebastian Wernicke and Florian Rasche. 2006. FANMOD: a tool for fast network motif detection. *Bioinformatics* 22, 9 (2006), 1152–1153.

[48] Joseph Young. 2015. Ripple Directs Bitstamp to Freeze Funds of Former Co-Founder Jed McCaleb. Cointelegraph blog entry. (2015). https://cointelegraph.com/news/ripple-directs-bitstamp-to-freeze-funds-of-former-co-founder-jed-mccaleb.