

Pedro Moreno-Sanchez*, Muhammad Bilal Zafar, and Aniket Kate*

Listening to Whispers of Ripple: Linking Wallets and Deanononymizing Transactions in the Ripple Network

Abstract: The decentralized I owe you (IOU) transaction network Ripple is gaining prominence as a fast, low-cost and efficient method for performing same and cross-currency payments. Ripple keeps track of IOU credit its users have granted to their business partners or friends, and settles transactions between two connected Ripple wallets by appropriately changing credit values on the connecting paths. Similar to cryptocurrencies such as Bitcoin, while the ownership of the wallets is implicitly pseudonymous in Ripple, IOU credit links and transaction flows between wallets are publicly available in an online ledger. In this paper, we present the first thorough study that analyzes this globally visible log and characterizes the privacy issues with the current Ripple network. In particular, we define two novel heuristics and perform heuristic clustering to group wallets based on observations on the Ripple network graph. We then propose reidentification mechanisms to deanonymize the operators of those clusters and show how to reconstruct the financial activities of deanonymized Ripple wallets. Our analysis motivates the need for better privacy-preserving payment mechanisms for Ripple and characterizes the privacy challenges faced by the emerging credit networks.

Keywords: Credit Networks, Ripple, deanonymization, linking wallets, crypto-currencies

DOI 10.1515/popets-2016-0049

Received 2016-02-29; revised 2016-06-02; accepted 2016-06-02.

*Corresponding Author: Pedro Moreno-Sanchez: Purdue University, E-mail: pmorenos@purdue.edu

Muhammad Bilal Zafar: MPI-SWS, E-mail: mzafar@mpi-sws.org

*Corresponding Author: Aniket Kate: Purdue University, E-mail: aniket@purdue.edu

1 Introduction

In recent years, we have observed a rather unexpected growth of IOU transaction networks such as Ripple [36, 40]. Its pseudonymous nature, ability to perform multi-currency transactions across the globe in a matter of seconds, and potential to monetize everything [15] regardless of jurisdiction have been pivotal to their success so far. In a transaction network [54, 55, 59] such as Ripple [10], users express trust in each other in terms of *I Owe You* (IOU) credit they are willing to extend each other. This online approach allows transactions in fiat money, cryptocurrencies (e.g., bitcoin¹) and user-defined currencies, and improves on some of the current banking system drawbacks: transactions are settled in seconds, requiring a small consistent fee world-wide. Banks such as Santander have claimed that adopting Ripple could save them \$20 billion a year [37]. At the time of writing, Ripple has over 170 thousand user wallets, a network value of more than \$790 million and a daily transaction volume over \$1 million [76].

Ripple has the potential to bounce into not only big banks but also into smaller ones as an interesting alternative to avoid large fees charged by intermediate banks while performing world-wide transactions. The Kansas-based CBW Bank and Cross River Bank are the first American banks to adopt Ripple [22]. Recently the Royal Bank of Canada has decided to adopt Ripple after exploring the numerous available blockchain options [43]. In Europe, the German bank Fidor [21] has adopted Ripple as a way of providing its customers with cheaper and faster worldwide transactions. In fact, other banks and financial institutions are adopting Ripple as their payment backbone as well [25, 26, 42, 44]. Companies are also using advantages of Ripple (e.g., fast and low-cost international transactions) to build better cross-border payment services. For example, Earth-

¹ Following the established convention, we use the capitalized term *Bitcoin* when referring to the system, and the lowercase term *bitcoin* (abbreviated BTC), when referring to the unit of currency.

port [28, 41], the largest open network for global banking, has adopted Ripple to perform transactions between banks over more than 60 countries over the world, while Saldo.mx uses the Ripple network to improve cross-border transactions between USA and Mexico [29]. Moreover, Microsoft has partnered with Ripple and is using part of its Azure BaaS to contribute to the execution of the Ripple network [31].

Although Ripple and Bitcoin share the goal of making money easier and faster to move around the world, they are conceptually different from each other. While Bitcoin is a digital cryptocurrency, Ripple is a transaction network that enables transactions in any currency between arbitrary pairs of agents. In fact, Ripple allows Bitcoin transactions: the payer of a transaction can send a certain amount of IOUs in any chosen currency and the payee receives the corresponding amount in bitcoins. Therefore, all Bitcoin merchants now accept transactions from Ripple users using the Ripple network. More than 8,500 merchants, such as Wordpress, are now available on the Ripple network [17].

Ripple, the company responsible for the creation of the Ripple network, has joined standardization organizations to assist with the regulations of transactions on the Web. The W3C Web Payment Interest Group [35], the International Payments Framework Association [33] or the Center for Financial Services Innovation Network [34] are a few examples of such organizations. In general, Ripple is intended to be the *Internet of Value*, facilitating the exchange of value at large scale as information moves today on the Internet [36].

Privacy issues. The Ripple network, at its core, is a replicated, *public* database (called the *Ripple ledger*) that tracks wallets and credit links extended between wallets along with their balances. Anyone can view the Ripple ledger and see a historical record of all activity on the Ripple network [4]. Moreover, Ripple identities are pseudonyms; thus, while not explicitly tied to real-world individuals or organizations, all transactions are completely transparent up to pseudonyms. However, businesses aim at protecting privacy of their financial activities from friends and foes by selectively revealing their pseudonymous identities to their partners [27, 30].

Interesting claims about the privacy for wallets and transactions are made by the Ripple company [8]:

Financial transaction history is recorded on the Ripple ledger; however, transactions are not linked to any identifiable information and cannot be directly associated with any individual account or your financial institution.

Nevertheless, all transactions remain linkable to each other and they are susceptible to deanonymization attacks. For example, users in XRPTalk forum have tracked down a payment carried out by the Fidor bank using Ripple in a few days [20]. The Ripple community has compiled a list with the Ripple wallets and their associated XRP² balances, which has been considered an intrusion of privacy by several Ripple users [9].

In this direction, the Ripple community has started to consider the privacy issues in Ripple [23, 39]. Banks do not wish to have all their transactions published on an open network [27, 30]. There are proposals in the Ripple community to provide privacy [12, 24]. However, they are in an early stage and they have not been implemented yet. Therefore, the current situation of the Ripple network regarding privacy leads to the question: does the public nature of the Ripple ledger lead to any serious privacy issue? Can we measure this?

Privacy studies on Bitcoin [45, 46, 63, 64, 75, 77, 81] have shown privacy issues on the Bitcoin system. However, the Ripple network and its Ripple ledger have not been thoroughly studied yet.

Our contribution. This work aims at improving the understanding of the traceability of Ripple flows and using it to explore the privacy breaches inherent to the public nature of the Ripple ledger. In particular, our goal is to cluster different Ripple wallets belonging to same users. This allows then to recognize previously unlinked transactions performed by the known Ripple users and further deanonymization of businesses performed over Ripple. In this direction, we present two novel heuristics to cluster Ripple wallets attending to transaction patterns and the Ripple network topology.

Our first heuristic is based on a settlement transaction between two wallet owners over their Ripple link to settle their bitcoin exchange. This heuristic enables linking of Bitcoin and Ripple wallets owned by the two involved users. It is not only the first heuristic performing clustering across two different payment networks but also, unlike in the well-known Bitcoin tagging attack [63], this heuristic allows identification of some cryptocurrencies wallet owners from the inherently public information of the linked Ripple wallets. Moreover, it is not restricted to Bitcoin, and enables the clustering of Ripple wallets with other cryptocurrencies wallets (i.e., altcoins), improving the set of clustered wallets.

² The XRP currency is defined in Ripple to protect the network from abuse and DoS attacks. For more details, see Section 2.2.

Our second heuristic leverages the transaction patterns performed by a user when deploying the hot-cold wallet security mechanism [5], which limits her risk profile on the Ripple network by enforcing a separation of roles that promotes stronger security (somewhat similar to master/session key in TLS). This heuristic allows to link several hot and cold Ripple wallets belonging to the same user. The heuristic employs temporal correlations between transactions and the Ripple network topology.

To analyze the efficacy of our heuristics, we crawl the Ripple network (as of December 2015) obtaining a total of *174,738 wallets* and *115,996 credit links*, and extract the complete dataset of transactions from the Ripple network, obtaining *17,645,343 transactions*. We deploy our heuristics over this dataset, resulting in the clustering of *959 Ripple wallets*, *3,113 Bitcoin wallets* and *1,130 Altcoin wallets*, which are involved in *934,484 transactions* in total. Among these clustered transactions, our de-anonymization process identifies the sender or the receiver for more than 78% transactions. Interestingly, we have reconstructed the complete set of transactions of the most widely deployed gateways³, and showed that is indeed bigger than the set of transactions associated to their publicly announced Ripple wallets.

Validation of the heuristics has been a difficult task given the absence of extensive ground truth data. Nevertheless, we have contacted several online services with the list of Ripple wallets linked to them by our heuristics. We have so far received replies from two of them (Bitstamp and RippleFox) confirming our findings.

Finally, we also study the effect of setting a Ripple validator server on the privacy of Ripple wallets. These servers collect transactions from the Ripple users and can tremendously increase the de-anonymization rate from the observed network identifiers (e.g., IP address) of the contacting users. Given the recent selection of commercial players such as Microsoft as validator server [31], we discuss the gravity of these large scale privacy leaks to the validator servers.

Organization. Section 2 presents an overview of the Ripple network. Section 3 describes the collected data. Section 4 defines the heuristic to link Ripple wallets with Bitcoin (and altcoin) wallets and Section 5 describes the heuristic to link Ripple hot and cold wallets. Section 6 shows the de-anonymization process using our clustering results. We discuss related work in Section 7. We conclude this paper in Section 8.

³ A gateway is a highly connected Ripple wallet that exchanges IOU in Ripple for the equivalent value in the outside world.

2 Ripple Overview

In this section we present a brief overview of the Ripple network. We characterize the structure of its IOU graph, describe the two available types of transactions, and compare the Ripple network and cryptocurrencies.

2.1 The Ripple IOU Network

The Ripple network is a weighted, directed graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$. The set \mathbb{V} of vertices represents the wallets in the network. The set \mathbb{E} of weighted edges represents the IOU credit links between wallets.

Similar to cryptocurrencies, a Ripple wallet is initialized with a pair of private (signing) and public (verification) keys. The wallet is then labeled with an encoding of the hashed public key. The wallet owner signs every transaction initiated by him using the private key, and includes the corresponding public key in the transaction. If the private key is stolen, the thief has direct access to the complete funds and extended credits to the wallet. The thief can also issue credit on behalf of the owner of the stolen wallet.

A directed edge $(u, v) \in \mathbb{E}$ is labeled with a dynamic scalar (weight) value α_{uv} indicating the amount of *unconsumed* credit that wallet u has extended to wallet v (i.e., u owes α_{uv} to v). The credit available on an edge is lower-bounded by 0 and is upper-bounded by ∞ by default, while a more strict upper bound can optionally be adopted by the wallet owner (i.e., v in the previous example). Additionally, every wallet has associated with it zero or more XRP (i.e., the Ripple currency). An illustrative example of the Ripple network is shown in Figure 1.

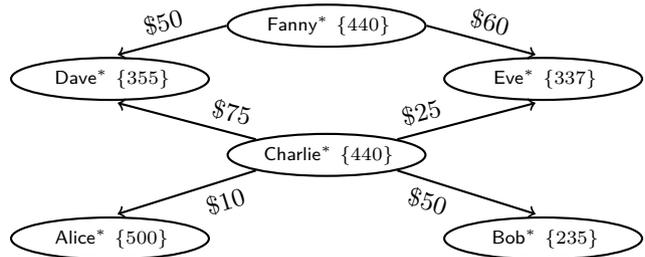


Fig. 1. An illustrative example of the Ripple network. For readability, every Ripple wallet is represented by a name with superscript * instead of a hashed public key used in practice. Values in {} represent the XRP currency balance, and edges represent credit links between pairs of connected nodes. The edge weights show IOU credit values on the edges. The edge weights are lower-bounded by zero and upper-bounded by ∞ . For readability, we show only one currency on the edges.

For ease of explanation, we assume that there is only one IOU currency (i.e., USD) over the Ripple links in the figure. Any other currency works in the same manner and we omit them in the rest of the paper.

A new wallet willing to interact with others in the Ripple network, and not yet having any trusted wallet to interact with, needs to receive some IOUs on a credit link. The Ripple network solves this *bootstrapping* problem by introducing gateways. A *gateway* is a well-known reputed wallet that several wallets in the system can trust to create and maintain a credit line in a correct manner. As wallets for gateways are highly connected nodes in the Ripple network, the thereby created credit line will allow the new wallet to interact with the rest of the Ripple network.

2.2 Transactions in the Ripple Network

Ripple allows two types of transactions: direct XRP payments and path-based settlement transactions. Intuitively, a direct payment involves a transfer of XRP between two wallets which may not have a credit path between them. Path-based settlement transactions transfer any type of credit (fiat currencies, cryptocurrencies and user-defined currencies) between two wallets having a suitable set of credit paths between them.

Direct XRP payments. The XRP currency is defined in Ripple to protect the network from abuse and DoS attacks. A Ripple wallet needs to hold XRP for two reasons: the wallet is considered active only if it has a certain amount of XRP; moreover, the issuer of any transaction must pay a transaction fee in XRP.

The direct XRP payments allow the exchange of XRP between two wallets. Assume that u wants to pay β XRP to v and that u has at least β XRP in her XRP balance. Then β XRP are removed from u 's XRP balance and added to v 's XRP balance. For example, in the illustrative transaction showed in Figure 2, 200 XRP are about to be transferred

Field	XRP payment	Path-based settlement transactions
Sender	Alice*	Dave*
Receiver	Bob*	Eve*
Amount	200 XRP	\$50
Path	–	Dave* ← Fanny* → Eve*
SigningPubKey	Alice*'s public key	Dave*'s public key
Tx Signature	752EF7...3402D1	42EF56...34DDFF

Fig. 2. Ripple transaction examples for both direct XRP payments and path-based settlement transactions. In the direct XRP payment, 200 XRP are sent from Alice* to Bob*. In the settlement transaction, \$50 are transferred from Dave* to Eve* via Fanny*. Irrelevant transaction fields have been omitted.

from Alice* to Bob*. Given that Alice*'s XRP balance is high enough, 200 XRP are taken from Alice*'s XRP balance and added to Bob*'s XRP balance. Notice that this type of transaction does not require the existence of any (direct or indirect) credit line between the sender and the receiver of the transaction. Therefore, the *Path* field of the transaction is not used.

Path-based settlement transactions. The edge weights in the Ripple network represent IOUs for three different types of currencies, namely, *fiat currencies* (e.g., USD), *cryptocurrencies* (e.g., bitcoins) and *user-defined currencies*. In Ripple, these three types of currencies are treated in the same manner. Moreover, there are exchange wallets (market makers in Ripple terms) that receive a certain currency in one of their links and exchange it for another currency in another link, charging a small fee for the exchange. They thereby allow transactions involving different currencies.

Path-based settlement transactions make use of the credit lines available in the Ripple network. Assume that u wants to pay β IOUs to v and that u and v are connected by a path of the form $u - u_1, \dots, u_n - v$. Edges are considered undirected to find a path from the sender u to the receiver v of the transaction. In order to perform the transaction, the credit value on every edge in the path from u to v is updated depending on the direction of the edge as follows: edges in the direction from u to v are increased by β , while reverse edges are decreased by β . For the settlement transaction to be successful, edges weights must always remain non-negative and must not exceed the pre-defined upper bound of the edge (if the upper bound is other than ∞).

In the illustrative settlement transaction shown in Figure 2, assume Dave* wants to pay \$50 to Eve*. The settlement transaction can be routed using the path Dave* ← Fanny* → Eve* (see Figure 1). Since edge Dave* ← Fanny* holds at least \$50 and edge Fanny* → Eve* has no upper bound, the settlement transaction can be performed and credit links are updated as follows: link Dave* ← Fanny* is deleted while link Fanny* → Eve* is increased to \$110.

It is not necessary to find a single path with available credit along each credit link; instead, the settlement transaction can be split across multiple paths such that the sum of credit available on all paths is larger than or equal to β . For example, in the network from Figure 1, assume now that Dave* wants to pay \$70 to Eve*. The settlement transaction now can be split into two settlement transactions with amounts of \$50 and \$20. The \$50 settlement transaction can be performed as ex-

plained earlier, while the \$20 settlement transaction is carried out over the path $\text{Dave}^* \leftarrow \text{Charlie}^* \rightarrow \text{Eve}^*$. In Ripple, it is possible to include the information about the several paths used in a single settlement transaction: the list of paths are included in the *Path* field annotated with the amount of credit used per path. The *Amount* field still indicates the total amount of transacted IOUs.

Settlement transactions are possible between arbitrary pairs of wallets, even if they have not extended a direct credit link between each other. As explained before, the existence of a credit path between a pair of wallets suffices to perform a settlement transaction between them. Therefore, settlement transactions take advantage of the transitive trust among wallets in the credit network to notably improve credit network's usability in several application scenarios [66, 68, 69, 74, 82].

2.3 Comparison with cryptocurrencies

Bitcoin [72] is a fully decentralized digital cryptocurrency, which initiated the wave of online decentralized payments in 2008; over the last few years, we have been observing an unprecedented growth of Bitcoin and its competitors [6, 7, 11]. Moreover, Bitcoin and the competitor cryptocurrencies are being traded on the Ripple network. We use the integration between Bitcoin and Ripple in one of our heuristics.

Although as in Bitcoin, Ripple opted for a ledger-based consensus to demonstrate consistency of transactions through transparency, they are conceptually different from each other. First, Bitcoin and other cryptocurrencies allow any two users to exchange bitcoins by means of a direct payment between them. Instead, Ripple allows to settle a transaction between two users only for which there exists a credit path in the network with enough IOU credit. Second, the Ripple network is composed of IOU links in any user-specified currency (including Bitcoin), and supports same- and cross-currency settlements over those links, a feature conceptually impossible in the currency networks such as Bitcoin.

Given their differences, although ideas showed in Bitcoin heuristics available in the literature [45, 46, 63, 64, 75, 77, 81] can be minimally reused when it comes to studying Bitcoin payments within Ripple, they cannot be directly applied to the fullest extent to the Ripple network. In this work, we notice this ineffectiveness of the Bitcoin heuristics to Ripple, and work towards designing novel clustering heuristics that *particularly* consider the IOU credit network graph.

3 Data Collection

In this section we describe the crawling of transactions from the Ripple network. Then we describe our crawled dataset in terms of number of transactions, number of wallets and observations over the Ripple network.

The complete dataset used in this work has been extracted from public sources available to the interested reader for validation of results. The code to parse the dataset and perform the analysis carried out in this work is available at project website [71].

3.1 Data Sources

Ripple public servers. The Ripple company maintains a set of public servers at *api.ripple.com* and *history.ripple.com*. We connected to them and crawled the Ripple transactions, following the protocol defined in their corresponding API.

Ripple server code. The Ripple company has published the source code of their Ripple server (*rippled*). We have installed our own instance of *rippled* and synchronized it with the rest of the Ripple network, thereby being able to extract information about additional Ripple wallets and their transactions.

Ripplebot. This online tool summarizes the content of the Ripple ledger and opens it up for inspection. It thereby becomes a valuable source for extracting information regarding Ripple transactions. Although this tool is no longer available, we used it while performing the work described in this paper.

Ripple graph. The Ripple company has developed an online tool that allows any user to check the set of credit links associated to any Ripple wallet. We have used it to check the network connectivity and to verify our findings within the Ripple transaction network.

Gateways. There are gateways (e.g., DividendRippler) allowing Ripple users to deposit and to withdraw credit from the network. The information available at these gateways websites differs from gateway to gateway. DividendRippler shows detailed information about every single transaction performed at the gateway. We have used it to carry out one of our heuristics.

Bitcoin blockchain. There are online tools (e.g., BlockExplorer and BlockChain.info) which allow to easily find the details of a given transaction in the Bitcoin blockchain. We employed BlockExplorer to obtain necessary details about the related Bitcoin transactions.

3.2 Collected Data

We extracted the list of Ripple wallets found in a Ripple ledger on the day of December 1, 2015 (ledger number 17,410,130). We found a total of 174,738 wallets in the system and 115,996 credit links between them. Then, we used the public API provided by RippleBot and Ripple public servers to get all the transactions involving these wallets. As of December 1, 2015, we collected a total of 17,645,343 transactions involving 168,422 wallets. We found that 6,316 wallets were dormant, and were not involved in any transaction.

The Ripple network. There were a total of 482 connected components of nodes with at least one link: a prominent component composed of 97.4% of the nodes, while the rest of the components were composed of less than 1% of the nodes in the Ripple network. There were 109,488 disconnected nodes with no links in the Ripple network. We visualize the Ripple network we crawled in Figure 3. In particular, we have decomposed the Ripple network into sets of highly interconnected nodes (i.e., communities) [53]. In order to do that, we have used the Louvain method for community detection in large networks [50] implemented in the software Gephi.

In total there are 470 communities composed only of two or more nodes. The communities composed of bigger number of nodes include the gateways wallets. Finally, communities with 2 to 7 nodes are the most frequent in the Ripple network. This community structure indicates the Ripple network is developing.

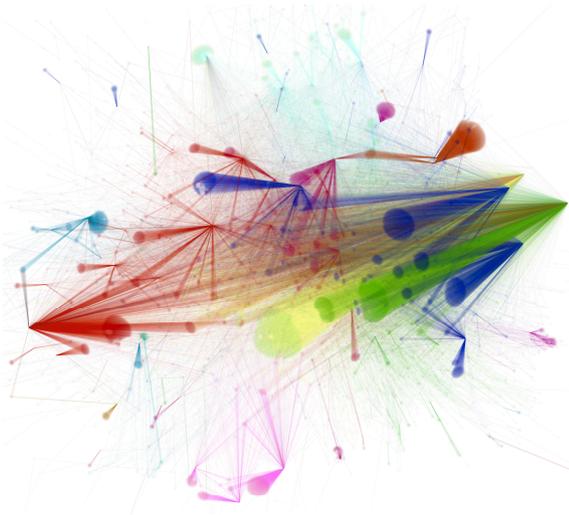


Fig. 3. A visualization of the Ripple network as of December 2015. We show only nodes with at least one link in the network. Different colors represent the communities (as computed by Gephi) in the network.

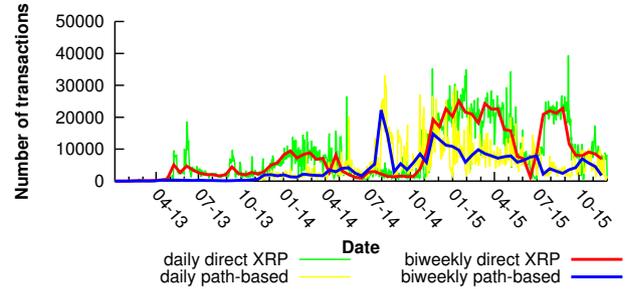


Fig. 4. Evolution of the Ripple transactions in January 2013 – December 2015. For both types of transactions we show the daily number and the median of payments per two weeks. For readability, we show only up to 50,000 transactions per day, thereby omitting the documented outliers.

We observe that the average node's degree is 3.51, the average path length is 6.83 and most nodes have links to a few highly connected nodes, which intuitively correspond to the gateways. The diameter of the network is 13, while the number of hops in transactions ranges from zero to nine.

Ripple payments. We show in Figure 4 the evolution of the number of transactions performed in the Ripple network from January 2013 to December 2015.

We make the following observations. First, direct XRP payments (Figure 4, green and red lines) started earlier than path-based settlement transactions (Figure 4, yellow and blue lines). As explained in Section 2.2, Ripple wallets need to have an XRP balance in order to be active and perform transactions. Therefore, most of the transactions at the early stage of the Ripple network are sending XRP to fund wallets and allow them to perform transactions to other wallets in the Ripple network.

We also analyze anomalous transactions in the Ripple network activity. We observe two spikes in the number of direct payments happening in Jul-13 which correspond to documented spamming attacks in the Ripple network [18]. We also observe two spikes in the number of path-based settlement transactions in Jan-14 and Oct-14. We identify the spikes on Jan-14 as settlement transactions sent by a Ripple wallet transacting CCK currency and the spikes on Oct-14 with a wallet which is sending payments in MTL currency to every other wallet in the Ripple network. Although they have not been documented as spam in the Ripple community, we believe they are outliers and do not reflect the normal activity of the Ripple network. Thus, we discard those to obtain the set of 13,181,194 Ripple transactions employed in our study.

In this characterization of the Ripple network we have only studied some basic graph statistics and the evolution of transactions during three years. Some other aspects of the Ripple network such as the evolution of the trade volume and the number of wallets can be found over *Ripple Charts* [76].

4 Linking Ripple wallets to Bitcoin wallets

In this section we show our heuristic to link Ripple and Bitcoin wallets that belong to the same user. For this purpose, we focus in the interaction of the users with the gateways. In the process we have extracted also other Blockchain-based cryptocurrencies (i.e., altcoins) wallets that can further be linked.

4.1 Heuristic 1: General interaction with the online currency exchanges

Motivation. Assume Alice has certain bitcoins in her Bitcoin wallet. Using the Bitcoin system, she can only pay services which accept payments in BTC. Alice can, however, transfer these bitcoins into the Ripple network, getting thereby the corresponding amount of BTC IOU. In this case, she is able to pay for the service independently of the currency accepted by the service provider. The Ripple network will allow the exchange from BTC IOU into the currency accepted by the service provider using the currency exchanges offered by market makers.

There are gateways (e.g., RippleWise, Bitstamp and DividendRippler) which allow users to transfer bitcoins (or any of the altcoins) into the Ripple network and vice versa. For example, Alice can pay the gateway a certain amount of bitcoins. The gateway, upon reception of the bitcoins, issues the corresponding BTC IOUs to the credit link Alice has previously formed with the gateway. We call this transaction *deposit transaction*. On the other hand, Alice could send (part of) her BTC IOUs to the gateway which in turn, transfers back the corresponding amount of bitcoins to the Alice's Bitcoin wallet. We call this transaction *withdrawal transaction*.

Heuristic algorithm. We use the publicly available information regarding deposit and withdrawal transactions at the gateways to link together Ripple and Bitcoin wallets that belong to the same user. Moreover, if the gateway supports other cryptocurrencies (i.e., altcoins

such as Litecoin [6]), the same heuristic can be used to link a Ripple wallet and a corresponding altcoin wallet.

Heuristic 1. [Deposit and withdrawal at the gateway] *The heuristic for deposit operations to link Bitcoin and Ripple wallets belonging to the same user involves the following steps:*

1. Assume w_g is a Ripple wallet owned by the gateway. Extract the set of all transactions in the Ripple network where w_g is the sender. We denote this set by $T_s(w_g)$. Moreover, for every transaction $t \in T_s(w_g)$, obtain the corresponding Bitcoin transaction. We denote the corresponding Bitcoin transaction by t_b .
2. For every transaction $t \in T_s(w_g)$ create a pair $(w_g, rcv(t_b))$, where $rcv(t_b)$ is the receiver of the Bitcoin transaction t_b corresponding to t . All these pairs thereby created correspond to Ripple, Bitcoin wallets belonging to the gateway. On the other hand, for every transaction $t \in T_s(w_g)$, create a pair $(rcv(t), sdr(t_b))$, where $rcv(t)$ denotes the receiver wallet of the Ripple transaction t and $sdr(t_b)$ denotes the sender wallet of the corresponding Bitcoin transaction. The two wallets of such a pair are owned by the same user.

The heuristic for withdrawals to link together Bitcoin and Ripple wallets belonging to the same user involves the following steps:

1. Assume that w'_g is a Ripple wallet owned by the gateway. Extract the set of all transactions in the Ripple network where w'_g is the receiver. We denote this set by $T_r(w'_g)$. Moreover, for every transaction $t' \in T_r(w'_g)$, obtain the corresponding Bitcoin transaction, which we denote by t'_b .
2. For every transaction $t' \in T_r(w'_g)$ create a pair $(w'_g, sdr(t'_b))$, where $sdr(t'_b)$ is the sender of the Bitcoin transaction t'_b corresponding to t' . All these pairs thereby created correspond to Ripple, Bitcoin wallets belonging to the gateway. On the other hand, for every transaction $t' \in T_r(w'_g)$, create a pair $(sdr(t'), rcv(t'_b))$, where $sdr(t')$ denotes the sender wallet of the Ripple transaction t' and $rcv(t'_b)$ denotes the receiver wallet of the corresponding Bitcoin transaction. The two wallets contained in such a pair are owned by the same user.

In Heuristic 1 we have omitted the necessary steps to link Bitcoin change wallets and thereby for clarity we assume a Bitcoin transaction has only one Bitcoin input and one Bitcoin output wallet. However, when deploying the heuristic in practice, we do take into account the change wallet and Bitcoin mixing transactions. More-

over, Heuristic 1 can be easily extended to take into account Bitcoin heuristics [46, 63, 75, 77, 81].

Figure 5 (top) shows a deposit transaction. Assume Alice wants to get 20 BTC IOU into her Ripple wallet $Alice_1^*$. To achieve that, she first creates a Bitcoin transaction where she transfers 20 BTC from her Bitcoin wallet $Alice_1^B$ to the gateway's Bitcoin wallet Gw_1^B . Once the gateway has checked the validity of the Bitcoin transaction, it creates a Ripple settlement transaction where it issues 20 BTC IOU from its Ripple wallet Gw_2^* to Alice's Ripple wallet $Alice_1^*$. This implies that $Alice_1^B$ and $Alice_1^*$ are owned by Alice while Gw_1^B and Gw_2^* are owned by the gateway. Moreover, following the heuristics regarding Bitcoin change addresses proposed by Meiklejohn et al [63], we can infer that $Alice_2^B$ also belongs to Alice.

Figure 5 (bottom) shows a withdrawal transaction. Assume Alice wants to withdraw 10 BTC IOU from the Ripple network into her Bitcoin wallet. For that, she first sends 10 BTC IOU from her Ripple wallet $Alice_2^*$ to the gateway's Ripple wallet Gw_1^* . Once the gateway has received the BTC IOU, it transfers 10 BTC from its Bitcoin wallet Gw_2^B to Alice's Bitcoin wallet $Alice_3^B$. The withdrawal implies that $Alice_3^B$ and $Alice_2^*$ are owned by Alice while Gw_2^B and Gw_1^* are owned by the gateway. Moreover, as mentioned before, we can infer that Gw_3^B belongs to the gateway.

Heuristic in practice. We have tested the Heuristic 1 in the gateway DividendRippler. In the following we explain how we have extracted the necessary information

Deposit	Bitcoin Transaction		RippleTransaction	
	Input	Output	Field	Value
	$Alice_1^B:30$	$Gw_1^B:20$ $Alice_2^B:10$	Sender	Gw_2^*
		Receiver	$Alice_1^*$	
		Path	$Gw_2^* \leftarrow Gw_1^* \rightarrow Alice_1^*$	
		Amount	20 BTC IOU	

Withdrawal	RippleTransaction		Bitcoin Transaction	
	Field	Value	Input	Output
	Sender	$Alice_2^*$	$Gw_2^B:15$	$Alice_3^B:10$
Receiver	Gw_1^*		$Gw_3^B:5$	
Path	$Alice_2^* \leftarrow Gw_1^*$			
Amount	10 BTC IOU			

Note: Irrelevant transaction fields have been omitted.

Fig. 5. An illustrative example of deposit and withdrawal processes in a gateway. For a deposit, first Alice sends 20 BTC to the gateway and then, the gateway sends 20 BTC IOU in the Ripple network to Alice. For a withdrawal, first Alice sends 10 BTC IOU to the gateway within the Ripple network and then the gateway sends 10 BTC back to Alice in the Bitcoin system.

for the steps defined in Heuristic 1 for the deposit process (i.e., steps 1-2). The heuristic for the withdrawal process has been implemented in a similar manner.

1. The DividendRippler wallet (i.e., w_g) is publicly available at its website. The set $T_s(w_g)$ has been obtained from our crawled Ripple transactions.
2. Every deposit has its own page in the DividendRippler's website. This page details both the Bitcoin (correspondingly the Altcoin) and the Ripple transaction involved. Therefore, the t_b corresponding to every transaction $t \in T_s(w_g)$ can be obtained from it. In Section 4.2, we discuss how to generalize this step to get the Bitcoin transaction corresponding to a Ripple settlement transaction even if the gateway does not publicly show it.
3. For every transaction $t \in T_s(w_g)$, $sdr(t)$ and $rcv(t)$ have been obtained from our Ripple database. The transaction t 's webpage also contains a link to the Bitcoin (correspondingly the altcoin) block where the corresponding t_b is stored. From this block, we have obtained the fields $sdr(t_b)$ and $rcv(t_b)$.

We run the algorithm defined in Heuristic 1. Our heuristic finds out a total of 435 Ripple wallets involved in trading with the gateway DividendRippler. Moreover, we have been able to extract 3,145 Bitcoin wallets and 1,173 altcoin wallets divided into 841 Litecoin wallets, 178 Terracoin wallets and 154 Namecoin wallets.

False positives. While employing the Bitcoin heuristic based on the change wallet, it is necessary to avoid Bitcoin transactions resulted from mixing services [13, 78] to avoid false positives. Mixing services in Bitcoin imply the creation of Bitcoin transactions with several input and output wallets, where the link between each input and the corresponding output is hidden. We avoid them by only considering Bitcoin transactions with at most 2 output wallets (the receiving wallet and the change wallet). In our results we observe 2 Bitcoin wallets that are linked to 15 and 49 Ripple wallets each, while the rest are linked to at most 5 Ripple wallets. Therefore, we consider them as outliers and discard them. In Ripple there are not mixing techniques currently deployed, therefore no extra actions are necessary.

4.2 Discussion

Heuristic verification. As a minimal ground truth, we have extracted Ripple and cryptocurrency wallets published by DividendRippler as its own wallets. We have checked that all of them are linked by our heuristic.

Novelty. This is the first heuristic that link wallets across cryptocurrencies. Additionally, gateways have started to accept transactions from other emerging transaction networks such as Stellar. Thus, this heuristic can also be used to link wallets from these networks. Finally, since gateways must publish one Ripple wallet to its customers, this heuristic allows deanonymization of some cryptocurrencies wallets owners from the inherently public information of the linked Ripple wallets.

Privacy impact. This heuristic enlarges the set of wallets among different cryptocurrencies that can be linked to a given user. This fact has several privacy implications. First, it paves the way to reconstruct the business of a company in a more accurate manner. It is interesting to note that since a business must publicly announce at least one wallet to its customers, the complete (and possible large) set of wallets linked to it are deanonymized. Second, larger sets of linked wallets among different systems affect also to privacy of users. For instance, even if a given user has private wallets in Bitcoin (e.g., she always uses mixing techniques for her transactions), deanonymizing one of her Ripple wallets directly deanonymizes her Bitcoin wallets as well.

Generalization. Although we use a gateway that publishes the Ripple and Bitcoin transactions involved in deposits and withdrawals, our heuristics are also applicable to gateways not publishing this information. In such case, it is possible to collect the Ripple transactions performed by the gateway and link them with high probability to Bitcoin transactions issued in a similar time and transacting the corresponding amount of bitcoins. This approach leads, however, to a probabilistic guarantee on accuracy and might include false positives in the results. Moreover, as mentioned earlier, our heuristic enables to link not only Ripple and Bitcoin wallets, but also wallets corresponding to other transaction networks (e.g., Stellar) and other cryptocurrencies (e.g., Litecoin, Namecoin or Terracoin)

5 Clustering Ripple wallets

Interactions between wallets in Ripple are inherently associated to the (social) topology of the Ripple network. Using the connectivity in the Ripple IOU graph, we have defined a novel heuristic to link Ripple wallets controlled by the same user, which we present in this section.

5.1 Heuristic 2: link Ripple wallets with their cold wallets

Motivation. Users willing to use the Ripple network to attract new business must publicly announce (at least) one of their wallets (i.e., issuing wallet) so that future clients can create credit link with those. From example, gateways publicly advertise their issuing Ripple wallet in their websites. Then, the issuing wallet's owner can issue credit to the clients through the newly created links. However, this practice has two main drawbacks.

First, the issuing wallet becomes an attractive target for an attacker: if the secret key of such wallet gets compromised, the attacker can freely issue an amount of unauthorized IOUs bounded only by the upper bound on these wallet's links. This problem is even more prominent given that upper bounds in the links are set to ∞ by default unless the user changes them. Such an attack has already been observed in the Ripple network and the stolen wallet's owner have gone bankrupt [19]. Second, as the Ripple ledger is publicly available, announcing ownership of a wallet and using it to carry out all the settlement transactions clearly leads to privacy leaks: everybody can track the settlement transactions of the issuing wallet and reconstruct the complete activity of the given user. Nevertheless, current businesses (such as banks and gateways) seek to maintain privacy of their activities while using the Ripple network [27, 30].

In order to overcome these issues, Ripple defines the hot-cold wallet security mechanism to issue IOUs of any currency [5]. The cold wallet is publicly linked to a certain user. However, actual issuing of the IOUs in a credit link extended to the cold wallet is performed by the hot wallet as follows. First, the hot wallet creates a credit link with the cold wallet. Then, when the owner of the cold wallet must extend credit to a user, she uses the hot wallet to extend that credit, using for this settlement transaction the existing path (hot wallet) \leftarrow (cold wallet) \rightarrow (user wallet).

The hot wallet is therefore considered to be online as it is used for daily settlement transactions. For example, the secret key of the hot wallet might be used by a web application to automatically perform settlement transactions to other users when requested. When the credit link between the hot and cold wallet runs out of IOUs, the cold wallet extends extra IOUs. This operation happens, however, less often and can be performed offline (e.g., signing locally the necessary transaction). Thus, the cold wallet is considered offline.

Following this mechanism, if the thief steals the private key of the hot wallet, he can issue a number of unauthorized IOUs bounded by the IOUs extended from the cold wallet to the hot wallet. Two observations are important here. First, this bound is normally notably smaller than the bound on the number of IOUs a cold wallet can issue. Second, the maximum number of IOUs in the link between hot and cold wallet is totally controlled by the owner of the cold wallet. She, however, does not have any control over the upper bound with the credit links created with the rest of the users.

With respect to privacy, we note that a settlement transaction from the hot wallet to any other user’s wallet has the same path structure as a settlement transaction between any two users (i.e., (sender wallet) ← (cold wallet) → . . . → (receiver wallet)). Thus, settlement transactions from the hot wallet to any user cannot be directly linked to cold wallet’s owner.

Heuristic algorithm. Implementing the hot-cold wallet mechanism forces the user to use her Ripple wallets following a pattern that makes it possible to link her wallets together. Intuitively, first our heuristic detects the possible cold wallets. Then, it checks settlement transactions where the cold wallet is the sender. The receivers of these transactions are the possible hot wallets. Finally, our heuristic links together hot and cold wallets that belong to the same user.

Heuristic 2. [Hot and cold wallets]

1. Extract the wallets that only have outgoing credit links in the Ripple network. They form the initial set of potential cold wallets and we denote it by CW . Among the wallets connected to a cold wallet in CW , those that have being paid at least once by such cold wallet are potential hot wallets, which we denote by HW . The rest of the connected wallets (say, a set \overline{HW}) are discarded as they are wallets from users other than cold wallet’s owner.
2. Reduce the set of potential hot wallets HW to those that are paying to other wallets connected to the cold wallet (i.e., the set $HW \cup \overline{HW}$). Let HW' be the thereby reduced set of potential hot wallets. Discarded wallets in this step (i.e., $HW - HW'$) are added to \overline{HW} , obtaining the set \overline{HW}' . This step intuitively ensures that potential hot wallets are being used to issue IOU to other wallets.
3. Reduce the set of potential cold wallets CW to those that have less potential hot wallets than discarded hot wallets. In other words, for each cold wallet $cw_i \in CW$, accept cw_i only if $|\overline{HW}'(cw_i)| <$

$|\overline{HW}'(cw_i)|$. Let CW' be the thereby reduced set of cold wallets. This step ensures there are indeed many wallets demanding IOUs, which are then supplied using a few hot wallets.

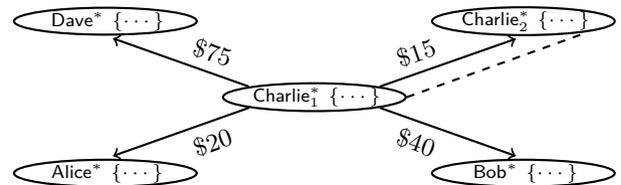
4. For each cold wallet $cw_i \in CW'$, create pairs (cw_i, hw_j) for each hot wallet $hw_j \in HW'(cw_i)$. Here, each pair of wallets thereby created belongs to the same user.

Figure 6 depicts an example of Heuristic 2. The wallet $Charlie_1^*$ is the cold wallet of Charlie as it does not have any incoming link in the Ripple network. In other words, the cold wallet can issue IOUs to other wallets in the network, but no other wallet can issue IOUs to it.

Charlie uses his cold wallet ($Charlie_1^*$) to fund his hot wallet ($Charlie_2^*$) with 80 and 70 credits in two settlement transactions, while no other wallet is paid by the cold wallet. Then, $Charlie_2^*$ is used to issue credit to wallets that have extended a credit line with the cold wallet $Charlie_1^*$, in this example $Alice^*$, Bob^* and $Dave^*$. Interestingly, although Bob^* transfers credit to $Alice^*$, it is not linked to Charlie given that Bob^* does not receive any settlement transaction from Charlie’s cold wallet $Charlie_1^*$.

Our heuristic can thereby derive the fact that $Charlie_1^*$ and $Charlie_2^*$ belong to the same user (i.e., Charlie), even though settlement transactions from $Charlie_2^*$ to other users follow the same path structure as transactions among other users (e.g., settlement transaction from Bob^* to $Alice^*$).

Heuristic in practice. We run the Heuristic 2 algorithm over our Ripple transactions dataset. After discarding false positives, our algorithm results in 261 cold



Ripple Ledger					
Sender	Receiver	Amount	Sender	Receiver	Amount
$Charlie_1^*$	$Charlie_2^*$	\$80	$Charlie_2^*$	Bob^*	\$50
$Charlie_2^*$	$Alice^*$	\$10	$Charlie_2^*$	$Dave^*$	\$75
$Charlie_1^*$	$Charlie_2^*$	\$70	Bob^*	$Alice^*$	\$10

Fig. 6. An illustrative example of Heuristic 2. The arrows show the credit links. The dashed line represents the wallets linked by the heuristic. Cold wallets ($Charlie_1^*$) do not have incoming credit links. Hot wallets ($Charlie_2^*$) receive credit from the cold wallets. XRP balances are omitted as they are not used in this heuristic.

wallets, 268 hot wallets, having a total of 529 Ripple wallets. Although the results of this heuristic in practice has resulted in a low number of clustered wallets, they cover a large number of settlement transactions as we show in Section 6.1.

False positives. The hot-cold wallet mechanism is a rather recent addition to the Ripple network, and it is not yet extensively applied by the Ripple users. Therefore, it is important to avoid false positives while applying this heuristic. In the following, we describe our mechanism to handle false positives.

During our process to handle false positives we apply the principle of being as strict as possible in order to reduce the number of them. Moreover, from our results we observe that false positives fall into two categories: wallets that do not follow the hot-cold wallet mechanism yet and wallets that follow such mechanism but have used the cold wallet to make sporadic payments to wallets other than the hot wallets. We perform the following steps to detect false positives.

First, we calculate the distribution of settlement transactions from cold wallets to potential hot wallets. In the absence of significant ground truth data, we use three gateways (Bitstamp, RippleFox and SnapSwap) well known in the Ripple community for using the hot-cold wallet mechanism, to bootstrap a minimal ground truth for the settlement transaction distributions. Their settlement transaction distributions resemble the Poisson distribution with parameter $\lambda = 1$. We then compute the divergence of each distribution and the Poisson distribution to detect falsely tagged cold wallets.

In detail, we calculate the statistical distance between two distributions using the Kullback–Leibler (KL) divergence [62] as a measure. Then, we flag a cold wallet as false positive if its settlement transaction distribution diverges from Poisson more than a threshold T . We set up T as the maximum divergence value between our ground truth distributions and Poisson with $\lambda = 1$.

This mechanism has flagged as false positives those cold wallets that do not follow the hot-cold wallet mechanism. In such case, the cold wallet is used to transfer IOUs to many other wallets with a somewhat equal probability, thus having a diversion from Poisson greater than T . We believe that these gateways' behavior is transient and that eventually they will correctly apply the hot-cold security mechanism. Otherwise, as it has happened already [19], they risk huge losses and the possibility of even going out of business in case their wallet's key is stolen.

In addition, we observe some wallets following the hot-cold mechanism sporadically paying other wallets other than the hot wallets. We conjecture that these cases represent anomalous settlement transactions. A reason for having anomalous transactions is that, in early stages, users employ the hot-cold wallet mechanism in a non-consistent manner. However, we expect that over the period they will start using this hot-cold wallet mechanism correctly and in a consistent manner; otherwise, they may risk huge credit losses and even bankruptcy as it has been already observed in the Ripple network [19]. Moreover, for known gateways using the hot-cold wallet mechanism, we have observed that percentage of anomalous transactions is fairly small. In order to flag these anomalous cases as false positives, we rely on the fact that cold wallet must refund the hot wallet repeatedly over time.

In detail, we consider 3 months (i.e., an economic quarter) as a time frame. Then, only potential hot wallets that are refunded by the cold wallet at least once per quarter for a period of at least two quarters are flagged as real hot wallets. The rest are flagged as false positives. There is a tradeoff choosing these thresholds. First, enforcing a less frequent refund or a shorter time frame would tag less wallets as false positives, decreasing thus the accuracy of the approach. Enforcing that hot wallets are refunded periodically from when they are created until today would tag real hot wallets as false positives, reducing also the accuracy: Ripple suggests to have several hot wallets [5], so that some cold wallets use one hot wallet for a period of time and then change to another hot wallet. Moreover, thresholds for this criteria have been selected following our design principle of being as strict as possible considering the fact that there are path-based settlement transactions in Ripple only for less than 2 years (see Figure 4).

5.2 Discussion

Heuristic validation. We have contacted several gateways with the list of Ripple wallets linked to them by our heuristic. We have received responses from two of them (i.e., Bitstamp and RippleFox) and both have confirmed the ownership of such wallets. Moreover, these response do not include any wallet missed by our heuristics.

The Ripple gateways publish their cold wallets on their webpages so that they can be used by the gateways' clients to create credit links to them. In order to bootstrap a minimal ground truth data, we have extracted and compare them with our heuristic's results.

The heuristic detects 14% of the cold-hot wallet pairs announced by the 109 gateways included in our study. The rest are not clustered due to our pessimistic parameter selection for the heuristics and the fact that there exist several gateways which do not employ the hot-cold wallet mechanism yet. Interestingly, a gateway publishes the same wallet as hot and cold wallet⁴, which does not lead to any security or privacy improvement.

Novelty. The heuristic we describe in this section is based on relations among different settlement transactions and the IOU network topology for wallets involved in such settlement transactions. In fact, the IOU network topology is a crucial ingredient of our heuristic. This heuristic thus greatly differs from heuristics proposed for Bitcoin. In Bitcoin only direct payments among wallets are possible given that there is not a Bitcoin IOU graph. Thus Bitcoin heuristics cannot take into account the IOU graph topology so that they cannot be successfully be applied to the Ripple network.

Security and Privacy Impact. The hot-cold wallet mechanism has been proposed by Ripple aiming at dissociating settlement transactions from hot wallet and cold wallet so that privacy for cold wallet’s owner is increased. However, our heuristic shows a novel technique to link back hot and cold wallets belonging to the same user, thus allowing to reconstruct the complete business (see Section 6.2). Thus, our heuristic shows that hot-cold wallet mechanism does not increase privacy in practice.

Moreover, linking hot and cold wallets using our heuristic leads to hinder the security supposedly provided by the hot-cold wallet mechanism. Using our heuristic, an attacker can lucratively target the hot wallets belonging to the target business and issue unauthorized IOUs. This forces the attacked business to create new hot wallets. This simple countermeasure however does not help as the attacker can repeat the process and target the newly created wallet.

Generalization. Additionally, our heuristic works for any IOU network following the hot-cold wallet mechanism as described earlier. We focus on the Ripple network as it is currently deployed in practice and several banks intend to adopt it as a transaction backbone in the near future. However, we are observing that the hot-cold wallet mechanism is already being discussed in the recently created transaction network Stellar [2, 38] so that our heuristic will directly apply to it when they grow to the level of Ripple today.

6 De-anonymizing Ripple Users

In this section, we first group our heuristics results, thereby increasing the linking among Ripple and cryptocurrencies wallets: if a wallet in a cluster gets deanonymized, the complete cluster can be deanonymized, independently of the system where the deanonymized wallet belongs. Second, we make concrete instances of deanonymization of such clusters identifying transactions associated to Ripple gateways. Third, we generalize our approach and apply this deanonymization process to our complete clustering. Finally, we discuss how a more capable adversary running a Ripple validator can perform further deanonymization.

6.1 Grouping Heuristics

We have presented two heuristics that enable the finding of a set of Ripple wallets as well as cryptocurrencies wallets which are owned by a certain user. Table 1 shows a summary of our findings.

Clustering in Ripple. To study the implications of our findings, we have grouped the results of the two heuristics. This process has allowed us to reconstruct 561 clusters which in total contain 959 Ripple wallets, 3,113 Bitcoin wallets and 1,130 altcoin wallets. Moreover, Ripple wallets clustered by our heuristics are involved in 161,624 XRP payments and 772,860 settlement transactions. Our clustered wallets are jointly involved in the 7.09% of the transactions in the Ripple network.

In this study, we focus on settlement transactions, which constitute the 35.5% of total transactions in Ripple. The rest are XRP payments. Thus, the inclusion of XRP in our study would lead to more clustered wallets and transactions. Thus, we consider it as an interesting future work. Moreover, the Ripple network is in an early stage. There are many wallets that do not have any credit link or even if they have some, they are not

Heuristic	Ripple		Bitcoin	Altcoins
	Wallets	Transactions	Wallets	Wallets
1	435	96,009	3,145	1,173
2	529	863,614	–	–
Grouped	959	934,484	3,113	1,130

Table 1. Number of wallets clustered in the different heuristics. In Altcoins we consider Litecoin, Namecoin and Terracoin. Finally, for each heuristic and for their grouping, we show the number of Ripple transactions where either the sender or the receiver is a clustered wallet.

⁴ <https://ripple.coinpip.com/ripple.txt>

connected to the main component in the Ripple network (as we show in Section 3). Therefore, these wallets cannot be clustered by our heuristics. Nevertheless, our clustering gathers wallets for most of the main gateways as we show later in this section. The gateways and their associated transactions represent the main activity for the core of the current Ripple network. They are used to transfer value from the real world into Ripple and vice versa, a crucial task to create liquidity in any starting transaction network such as Ripple [1] or Stellar [2].

Clustering among payment systems. We have further studied the results of applying our clustering to link wallets from different payment systems, and we have made the following observations. First, there are a total of 249 clusters composed only of Ripple wallets. These are Ripple wallets clustered in the Heuristic 2, but not appearing in the Heuristic 1. The rest of the clusters include at least one of the studied cryptocurrencies, and therefore they contain information extracted from Heuristic 1.

Second, there are 2 clusters that group wallets from Ripple, Bitcoin and altcoin simultaneously. Obviously, one of these clusters belongs to the gateway DividendRippler, which operate with all the crypto-currencies considered in this study. Third, there are 241 clusters grouping Ripple and Bitcoin wallets, while 40 clusters group together Ripple and altcoin wallets. Therefore, the most common cryptocurrency is Bitcoin, which clearly fits the fact that Bitcoin is nowadays more widespread than the rest of the cryptocurrencies.

6.2 Reconstructing gateways businesses

Services using Ripple as the underlying transaction network strive to ensure the privacy of their transactions from the prying eyes of competitors, authorities, and even customers. For that, gateways and other prominent business do not reveal all their wallets hoping to stop third parties from fully reconstructing their economic activities, and to protect their wallets from the potential break-ins. In contrast, our work shows different mechanisms to question the privacy of a given gateway and reconstruct its complete activity within the Ripple network. This implies that anybody accessing the publicly available Ripple data can reconstruct the total number of transactions carried out by a gateway, and not only transactions associated to the gateway’s public wallets, thereby having a significant privacy breach.

	Total Sent	Total received	Total Balance
Public wallet	1062.29	1064.08	1.79
Clustered wallets	5724.38	5724.41	0.03

Table 2. Deanonymization of Dividendripler Bitcoin business.

Single gateway business. We first consider the deanonymization of business of a single gateway at a time for both DividendRippler and Bitstamp. DividendRippler publicly announces only one Bitcoin wallet. Extracting the transaction history of such wallet from the Bitcoin blockchain, we observe that more than 1000 bitcoins have been transacted. However, this is only a partial view of the gateway’s business. As shown in Table 2, transaction history of Bitcoin wallets linked to the gateway by our heuristics shows that more than 5000 bitcoins have been transacted. These results have been possible given the wallets linked by Heuristic 1.

At the time of writing Bitstamp has only published its cold wallet and one of its hot wallets, for which we observe that there have been 72,042 transactions. However, our Heuristic 2 has flagged another Ripple wallet as belonging to Bitstamp. Using this extra information, it is possible to derive that Bitstamp has instead been involved in 132,543 transactions. Therefore, our heuristics enable the finding of 60,501 extra transactions involving Bitstamp. During our deanonymization process, we consider transactions where either the sender or the receiver is the linked wallet by our heuristic.

It is possible to monitor the gateway’s business even further. Once the clustering is performed, it is possible to monitor the network to notice every time a transaction is received by a given Ripple wallet. Using this approach it is possible to monitor the complete set of wallets in a user cluster, and thereby her full activity in real time.

Several gateways business. We have carried out the reconstruction of the business associated to the most widely deployed gateways [32] in the same manner we did with Bitstamp’s business. We show the most interesting results in Figure 7.

We make the following observations. First, there are gateways for which the numbers of publicly available transactions are different. However, adding up the transactions performed with the wallets resulting from our heuristics (Figure 7, red bar), they have performed the same total amount of transactions. DividendRippler, DYM and Chriswen constitute an example of this observation. We have verified that indeed DividendRippler and DYM are operated by the same owner [14, 16].

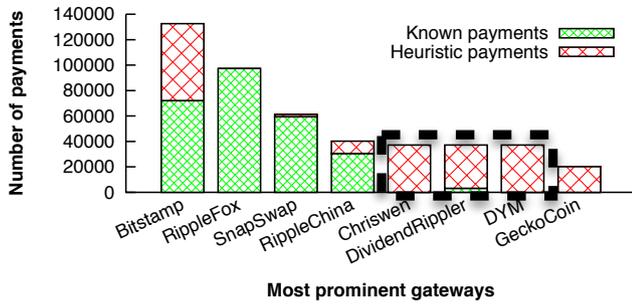


Fig. 7. Comparison of the number of transactions associated to publicly known gateways' wallets (i.e., Known payments) and transactions performed with wallets clustered by our heuristics to those gateways (i.e., Heuristic payments). Dashed line groups gateways sharing an owner.

Chriswen has been linked due to the combined results of both heuristics presented in this work: the hot wallet for Chriswen extracted from Heuristic 2 has been used in DividendRippler and it appears in the cluster for DividendRippler and DYM resulting from Heuristic 1.

Second, there are gateways with a few transactions made by their public wallets. However, when adding the payments associated to wallets clustered to them by our heuristics, the number of transactions increases. This is the case, for example, for GeckoCoin and RippleChina. Finally, we observe that no gateway (except for DividendRippler) publishes its Bitcoin wallets. As our heuristics link Bitcoin and other cryptocurrencies wallets to them, we can further deanonymize their financial activities.

6.3 De-anonymization at large

We have shown in Section 6.2 how to use the results of our heuristics to reconstruct the business carried out by a subset of gateways. In this section, we show our results when applying the de-anonymization process over our complete clustering.

In particular, we have further studied the privacy implications of our heuristics by applying the de-anonymization process over the transactions for which at least a wallet has been clustered by our heuristics (see Table 1). We have deanonymized 85,962 XRP payments and 649,640 settlement transactions, which jointly represent the 78.7% of the total transactions we have considered in our de-anonymization process. These results follow the fact that the probability that a Ripple wallet gets deanonymized is bigger when the wallet is clustered with our heuristics. This is an important privacy breach: we have shown how to use it to reproduce the business of gateways.

Finally, we have studied the interactions between the clusters we have obtained from our heuristics. The results are shown in Figure 8. As expected from our results while reconstructing the gateways' businesses (see Section 6.2), we observe that Bitstamp is the gateway with the largest amount of transactions within our cluster. Moreover, we have deanonymized 98 Ripple wallets belonging to the gateway DividendRippler (Figure 8, blue nodes). We have observed that most of these wallets were clustered to DividendRippler by the Heuristic 1. Interestingly, Bitcoin wallets deanonymized from our heuristic can be further linked to more Bitcoin wallets using the available Bitcoin deanonymization techniques [46, 63, 75, 77, 81]. We focus on privacy attacks on the Ripple network in this work, and leave those improved deanonymization attacks on the Bitcoin wallets as a future exercise.

False positives. As any other study based on heuristics, our analysis might have false positives. A false positive would imply an inaccurate association of financial activities to a gateway. Nevertheless, we have deployed several measures to reduce as much as possible the false positives rate: we have adopted strict configuration values in the deployment and validation of our heuristics and explained the corresponding tradeoffs; we have compared our results with available ground truth data; we have contacted several gateways with our results and received confirmation from two of them. Finally, we have not added our own "ground truth" transactions since they might not represent realistic financial activities and bias our results.

Ethical Principles. We have conducted our privacy analysis of the Ripple network over publicly available data as we discussed in Section 3. Moreover, in our de-anonymization process we mention the names of gateways services that are well known in the Ripple community and publicly advertised in their websites.

6.4 De-anonymization Using a Ripple Server

In the literature, there are several attacks based on maliciously including certain nodes in a network to deanonymize other nodes in the same network. For example, in the case of the Bitcoin network, a series of works [47, 58, 61] have shown that by including a few machines in the Bitcoin network it is possible to link Bitcoin transactions to their corresponding source IP addresses. Our results increase the privacy breach re-

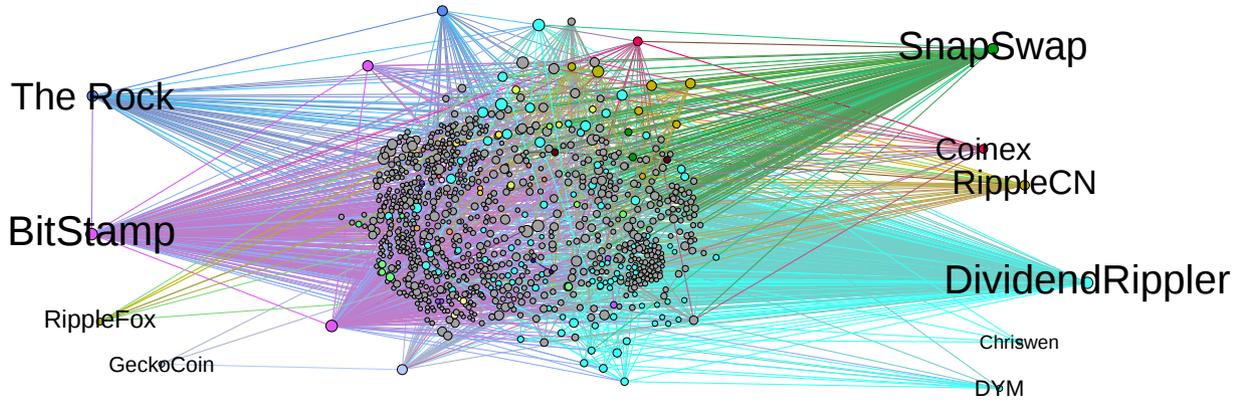


Fig. 8. A visualization of the deanonimization process over our clustered graph. The sizes of the nodes correspond with the number of transactions involving the nodes. Nodes with the same color belong to the same cluster. Gray nodes depict wallets not deanonimized by our heuristics. Links are colored with the color of the sending wallet.

sulting from these techniques since if a Bitcoin wallet is deanonimized, the complete cluster (including Ripple and other cryptocurrency wallets) is deanonimized.

Ripple transactions are collected by Ripple validator servers. Similar to Bitcoin, it is possible to further deanonimize Ripple transactions and wallets by deploying a Ripple validator server. As of today, validator servers are run by the core Ripple team (e.g., `s-west.ripple.com`) and by a few big gateway owners (e.g., SnapSwap). These parties can leverage our heuristics to further deanonimize Ripple wallets, and users are particularly vulnerable to deanonimization by them.

Assume we deploy one Ripple server. Then, a Ripple client can create an IP connection to our deployed server to send us the Ripple transactions. As Ripple transactions are sent in the clear, we can inspect them, and by looking at the *Sender* field (see Figure 2) it is possible to associate the IP address of the incoming connection to the Ripple wallet specified in the *Sender* field.

This privacy breach can be further exploited to link more than one Ripple wallet to a certain IP address. In detail, assume that different connections from the same IP address submits n transactions $\{t_1, \dots, t_n\}$, where t_i has a Ripple wallet w_i specified in the *Sender* field. This assumption is realistic: the currently Ripple web clients (e.g., *RippleTrade*) issue all the transactions by default to the same Ripple server. Given this scenario, it is likely that all the w_i are owned by the same person and we can further associate this cluster of wallets to the IP address used to establish the connection with our Ripple server.

Although the possibility to employ an anonymous communication network (e.g., Tor [52]) to forward the transactions to the transaction collecting server has been explored, such techniques are found to be vulnerable to denial of service and blacklisting attacks [48].

7 Related Work

Since its inception, questions regarding the security and privacy of the Bitcoin system have attracted interest from the research community. Barber et al. [46] observed that Bitcoin exposes its users to the possible linking of their Bitcoin wallets. Thus, recent works [63, 75, 77, 81] have proposed simple heuristics to thwart anonymity in Bitcoin. In a somewhat different direction, other recent works [47, 61] show the possibility of identifying ownership relationships between Bitcoin wallets and IP addresses. As Bitcoin heuristics do not fit transaction networks, our novel heuristics are focused and have special interest for transaction networks such as Ripple, including the integration of several available cryptocurrencies.

The most prevalent approach to improve anonymity for Bitcoin users is the idea of hiding in a group by Bitcoin mixing: the users in the group exchange their coins with each other to hide the relations between users and coins from an external observer. Several Bitcoin mixing approaches have been proposed [3, 13, 46, 49, 51, 57, 78, 83, 86]. Zerocoin [65] and its successor Zerocash [80] propose an alternative approach to provide privacy-preserving payments in cryptocurrencies based on zero-knowledge proofs. While all these approaches tackle privacy in the Bitcoin network, they do not consider the interactions between Bitcoin and other payment networks (e.g., Ripple), a scenario that we use in one of our heuristics.

Social networks (e.g., Facebook or Twitter) are currently used by millions of users, a fact that has attracted the research community. Several research works [60, 73, 84] propose mechanisms to cluster accounts from different social networks that are owned by the same person. These approaches extract static information associated

to anonymized profiles from different social networks and cluster together profiles that match according to a reidentification algorithm. These approaches cannot be applied to credit networks as they do not consider the inherent dynamic nature of the credit networks.

There are several approaches to enhance social networks with privacy [56, 67, 79, 85]. All of these approaches modify the network connectivity so that the privacy of the link is preserved and the loss of system reliability is bounded. However, this loss might not be acceptable in transactions systems like Ripple due to the fact that it implies the loss of money by the users. Moreno-Sanchez et al. [70] formally define privacy for credit networks in the form of transaction receiver privacy and transaction value privacy, and present a provably secure privacy-preserving payment protocol for credit networks in general which does incur any loss in the system reliability. Nevertheless, none of these works show how privacy of users can be thwarted.

8 Conclusions and Future Work

This work characterizes the current state of the Ripple network along with its complete set of transactions. Additionally, we shed light on the gap—due to certain patterns of use and interaction between parties in the network—between the (supposedly) provided privacy available in the Ripple network and the actual privacy achieved by the current Ripple users and, most importantly, their transactions. These heuristics allow us to cluster wallets belonging to the same user, not only from the Ripple network but also from several (publicly verifiable) blockchain-based cryptocurrency systems such as Bitcoin. More interestingly, this clustering has enabled the deanonymization of more than 78% of the clustered transactions, which in turn has allowed us to reconstruct the complete (and not only the publicly published) amount of trade of the most widely deployed gateways in the Ripple network.

Although the Ripple network is still in its early stages, the heuristics described in this work augur promising results when the Ripple network activity takes off; their effectiveness will improve when the Ripple network expands and become more structured. Therefore, our heuristics show a privacy problem directly in the design and the use pattern of credit networks such as Ripple. Our analysis thereby opens the way for several interesting future works. Finally, our analysis characterizes the privacy challenges faced by

the emerging transaction networks, paves the way towards further deanonymization by forensic techniques and motivates the imperative need for better privacy-preserving transactions mechanisms for Ripple and any other emerging transaction network based on the same design principles.

References

- [1] Becoming a Ripple Gateway. Ripple online documentation. <https://ripple.com/build/gateway-guide/#becoming-a-ripple-gateway>.
- [2] Becoming a Stellar Gateway. Stellar online documentation. <https://www.stellar.org/developers/learn/integration-guides/gateway.html>.
- [3] Bitcoin Wiki: Mixing Services. https://en.bitcoin.it/wiki/Category:Mixing_Services.
- [4] Executive Summary for Financial Institutions. Ripple online documentation. <https://ripple.com/integrate/executive-summary-for-financial-institutions/>.
- [5] Hot and Cold Wallets. Ripple online documentation. <https://ripple.com/build/gateway-guide/#hot-and-cold-wallets>.
- [6] Litecoin. <https://litecoin.org/>.
- [7] Namecoin. <http://namecoin.org/>.
- [8] Ripple brochure. <https://ripple.com/files/ripple-brochure.pdf>.
- [9] Ripple names with their balance of XRP. Reddit. https://www.reddit.com/r/XRPTalk/comments/2c66wl/ripple_names_with_their_balance_of_xrp/.
- [10] Ripple Website. <https://ripple.com/>.
- [11] Terracoin. <http://terracoin.sourceforge.net/>.
- [12] Ripple privacy. Ripple Forum, Nov 2012. <https://forum.ripple.com/viewtopic.php?f=1&t=4>.
- [13] CoinJoin: Bitcoin Privacy for the Real World. Post on Bitcoin Forum, Aug. 2013. <https://bitcointalk.org/index.php?topic=279249>.
- [14] Dividend Rippler. Ripple Forum, Jun 2013. <https://forum.ripple.com/viewtopic.php?t=3084>.
- [15] Introducing Goodwill. Ripple Forum, May 2013. <https://forum.ripple.com/viewtopic.php?t=2895&t=2895>.
- [16] Introducing Ripple Currency: DYM. Bitcointalk Forum, Mar 2013. <https://bitcointalk.org/index.php?topic=149533.0>.
- [17] Ripple allows payments to any Bitcoin address straight from its client. Gigaom Blog, Jul 2013. <https://gigaom.com/2013/07/02/ripple-allows-payments-to-any-bitcoin-address-straight-from-its-client/>.
- [18] Are XRP or STR effective at preventing ledger spam?. Ripple Forum, Aug 2014. <https://forum.ripple.com/viewtopic.php?f=1&t=7614>.
- [19] [Closing] Peercover. Ripple Forum, Mar 2014. <https://forum.ripple.com/viewtopic.php?t=5875&p=44441&t=5875&p=44441>.
- [20] Fidor Bank: Testing the 'Auslandsüberweisung über Ripple'. Archive from XRPTalk Forum, Aug 2014. <http://archive.is/BTATe>.
- [21] Fidor Becomes First Bank to Use Ripple Payment Protocol. CoinDesk Blog, May 2014. <http://www.coindesk.com/fidor->

- becomes-first-bank-to-use-ripple-payment-protocol/.
- [22] Ripple Labs Signs First Two US Banks. Ripple Blog, Sep 2014. <https://ripple.com/blog/ripple-labs-signs-first-two-us-banks/>.
- [23] Ripple Privacy - details on proxy payments or alternative? Ripple Forum, Nov 2014. <https://forum.ripple.com/viewtopic.php?f=1&t=8304&p=57936>.
- [24] Using multi-signature transactions to provide privacy. Ripple Forum, Oct 2014. <https://forum.ripple.com/viewtopic.php?f=2&t=8215>.
- [25] Australia's Commonwealth Bank Latest to Experiment With Ripple. CoinDesk Blog, May 2015. <http://www.coindesk.com/australia-commonwealth-bank-ripple-experiment/>.
- [26] Bank-Wise Analysis of Blockchain Activity. Let's Talk Payments Blog, Aug 2015. <http://letstalkpayments.com/bank-wise-analysis-of-blockchain-activity>.
- [27] Breaking the taboo on private Ripple ledgers. Ripple Forum, Jul 2015. <https://forum.ripple.com/viewtopic.php?f=1&t=10597&p=65410>.
- [28] How EarthPort and Ripple are teaming up to make cross-border payments instant. PYMNTS.com Blog, Aug 2015. <http://www.pymnts.com/in-depth/2015/how-earthport-and-ripple-are-teaming-up-to-make-cross-border-payments-instant/>.
- [29] How Marco Montes is Empowering Migrant Workers. Ripple Blog, Feb 2015. <https://ripple.com/blog/how-marco-montes-is-empowering-migrant-workers/>.
- [30] Implementing the Interledger Protocol in Ripple. Ripple blog, Oct 2015. <https://ripple.com/insights/implementing-the-interledger-protocol/>.
- [31] Microsoft Explores Adding Ripple Tech to Blockchain Toolkit. CoinDesk Blog, Dec 2015. <http://www.coindesk.com/microsoft-hints-future-ripple-blockchain-toolkit/>.
- [32] Ripple Gateway List. Ripple online documentation, Nov 2015. https://ripple.com/knowledge_center/gateway-information/.
- [33] Ripple Labs Joins International Payments Framework Association. Ripple Blog, Mar 2015. <https://ripple.com/blog/ripple-labs-joins-international-payments-framework-association/>.
- [34] Ripple Labs joins the Center for Financial Services Innovation. Ripple Blog, Feb 2015. <https://ripple.com/blog/ripple-labs-joins-the-center-for-financial-services-innovation-network-cfsi-network/>.
- [35] Ripple Labs Joins W3C Web Payment Interest Group to Help Set Standards for the Value Web. Ripple online documentation, Feb 2015. https://ripple.com/ripple_press/ripple-labs-joins-w3c-web-payment-interest-group-to-help-set-standards-for-the-value-web/.
- [36] Ripple Labs Named a Technology that by World Economic Forum. Ripple Blog, Aug 2015. <https://ripple.com/blog/ripple-labs-awarded-as-technology-pioneer-by-world-economic-forum-2/>.
- [37] Santander: Distributed Ledger Tech Could Save Banks \$20 Billion a Year. Ripple Blog, Jun 2015. <https://ripple.com/blog/santander-distributed-ledger-tech-could-save-banks-20-billion-a-year/>.
- [38] StellarVerse's Cold Wallet Generator. Stellar Forum, Jan 2015. https://stellarverse.org/cold_wallet_generator/.
- [39] What would you like to see in Ripple? Ripple Forum, Jun 2015. <https://forum.ripple.com/viewtopic.php?f=1&t=8930&p=60341>.
- [40] World Economic Forum Report: The Rise of Non-Traditional Payment Systems. Ripple Blog, Jul 2015. <https://ripple.com/blog/world-economic-forum-report-the-rise-of-non-traditional-payment-systems/>.
- [41] Earthport launches distributed ledger hub via Ripple, 2016. <http://www.bankingtech.com/420912/earthport-launches-distributed-ledger-hub-via-ripple/>.
- [42] Japan's SBI Holdings Teams With Ripple to Launch New Company. CoinDesk Blog, Jan 2016. <http://www.coindesk.com/sbi-holdings-ripple-new-company/>.
- [43] Royal bank of canada teams up with ripple for blockchain remittance system. Coinspeaker Blog, Feb 2016. <http://www.coinspeaker.com/2016/02/25/royal-bank-of-canada-teams-up-with-ripple-for-blockchain-remittance-system/>.
- [44] 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/>.
- [45] ANDROULAKI, E., KARAME, G. O., ROESCHLIN, M., SCHERER, T., AND CAPKUN, S. Evaluating User Privacy in Bitcoin. *Financial Cryptography and Data Security: 17th International Conference, FC 2013*, pp. 34–51.
- [46] BARBER, S., BOYEN, X., SHI, E., AND UZUN, E. *16th International Conference Financial Cryptography and Data Security*. 2012, ch. Bitter to Better — How to Make Bitcoin a Better Currency, pp. 399–414.
- [47] BIRYUKOV, A., KHOVRATOVICH, D., AND PUSTOGAROV, I. Deanonymisation of clients in bitcoin p2p network. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (2014)*, CCS '14, pp. 15–29.
- [48] BIRYUKOV, A., AND PUSTOGAROV, I. Bitcoin over tor isn't a good idea. In *Security and Privacy (SP), IEEE Symposium on (2015)*, IEEE, pp. 122–134.
- [49] BISSIAS, G., OZISIK, A. P., LEVINE, B. N., AND LIBERATORE, M. Sybil-resistant mixing for bitcoin. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society, WPES '14*, pp. 149–158.
- [50] BLONDEL, V., GUILLAUME, J., LAMBIOTTE, R., AND MECH, E. Fast unfolding of communities in large networks. *J. Stat. Mech* (2008), P10008.
- [51] BONNEAU, J., NARAYANAN, A., MILLER, A., CLARK, J., KROLL, J. A., AND FELTEN, E. W. Mixcoin: Anonymity for Bitcoin with accountable mixes. In *Proc. of the 17th International Conference on Financial Cryptography and Data Security, FC'14*, Springer.
- [52] DINGLEDINE, R., MATHEWSON, N., AND SYVERSON, P. Tor: The second-generation onion router. In *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13, SSYM'04*, pp. 21–21.
- [53] FORTUNATO, S., AND CASTELLANO, C. Community Structure in Graphs. In *Encyclopedia of Complexity and Systems Science*. 2009, pp. 1141–1163.
- [54] FUGGER, R. Money as IOUs in Social Trust Networks & A Proposal for a Decentralized Currency Network Protocol, 2004. <http://archive.ripple-project.org/decentralizedcurrency.pdf>.
- [55] GHOSH, A., MAHDIAN, M., REEVES, D. M., PENNOCK, D. M., AND FUGGER, R. Mechanism design on trust networks. In *Proceedings of the 3rd International Conference on*

- Internet and Network Economics*, WINE'07, pp. 257–268.
- [56] HAY, M., MIKLAU, G., JENSEN, D., TOWSLEY, D., AND WEIS, P. Resisting Structural Re-identification in Anonymized Social Networks. *Proc. VLDB Endow.* 1, 1 (2008), 102–114.
- [57] HEILMAN, E., BALDIMTSI, F., AND GOLDBERG, S. Blindly signed contracts: Anonymous on-blockchain and off-blockchain bitcoin transactions. *IACR Cryptology ePrint Archive 2016* (2016), 56.
- [58] KAMINSKY, D. Black Ops of TCP/IP 2011. Black Hat USA 2011. <http://www.slideshare.net/dakami/black-ops-of-tcpip-2011-black-hat-usa-2011>.
- [59] KARLAN, D., MOBIUS, M., ROSENBLAT, T., AND SZEIDL, A. Trust and Social Collateral. *The Quarterly Journal of Economics* 124, 3 (2009), 1307–1361.
- [60] KORAYEM, M., AND CRANDALL, D. J. De-anonymizing users across heterogeneous social computing platforms. In *ICWSM* (2013), E. Kiciman, N. B. Ellison, B. Hogan, P. Resnick, and I. Soboroff, Eds., The AAAI Press.
- [61] KOSHY, P., KOSHY, D., AND MCDANIEL, P. *An Analysis of Anonymity in Bitcoin Using P2P Network Traffic*. Financial Cryptography and Data Security: 18th International Conference, FC 2014. pp. 469–485.
- [62] KULLBACK, S., AND LEIBLER, R. A. On Information and Sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86.
- [63] MEIKLEJOHN, S., AND ORLANDI, C. *Privacy-Enhancing Overlays in Bitcoin*. Financial Cryptography and Data Security: FC 2015 International Workshops, BITCOIN. pp. 127–141.
- [64] MEIKLEJOHN, S., AND ORLANDI, C. *Privacy-Enhancing Overlays in Bitcoin*. Financial Cryptography and Data Security: FC 2015 International Workshops, BITCOIN. pp. 127–141.
- [65] MIERS, I., GARMAN, C., GREEN, M., AND RUBIN, A. D. Zerocoin: Anonymous distributed e-cash from bitcoin. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy* (2013), pp. 397–411.
- [66] 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, Proceedings*, pp. 15–30.
- [67] MITTAL, P., PAPAMANTHOU, C., AND SONG, D. X. Preserving Link Privacy in Social Network Based Systems. In *Network and Distributed System Security 2013*.
- [68] MOHAISEN, A., HOPPER, N., AND KIM, Y. Keep your friends close: Incorporating trust into social network-based Sybil defenses. In *INFOCOM, 2011 Proceedings IEEE* (2011), pp. 1943–1951.
- [69] MOHAISEN, A., TRAN, H., CHANDRA, A., AND KIM, Y. Trustworthy distributed computing on social networks. In *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security* (2013), pp. 155–160.
- [70] MORENO-SANCHEZ, P., KATE, A., MAFFEI, M., AND PECINA, K. Privacy Preserving Payments in Credit Networks: Enabling trust with privacy in online marketplaces. In *Network and Distributed System Security 2015*.
- [71] MORENO-SANCHEZ, P., ZAFAR, M. B., AND KATE, A. Project website. <http://crypsys.mmci.uni-saarland.de/projects/LinkingWallets>.
- [72] NAKAMOTO, S. Bitcoin: A peer-to-peer electronic cash system. Technical report, 2008. <https://bitcoin.org/bitcoin.pdf>.
- [73] NARAYANAN, A., AND SHMATIKOV, V. De-anonymizing social networks. In *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy*, pp. 173–187.
- [74] 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* (2011), NSDI, pp. 183–196.
- [75] REID, F., AND HARRIGAN, M. An analysis of anonymity in the bitcoin system. *Security and Privacy in Social Networks* 2013.
- [76] Ripple Charts. <http://www.ripplecharts.com/>.
- [77] RON, D., AND SHAMIR, A. *Quantitative Analysis of the Full Bitcoin Transaction Graph*. Financial Cryptography and Data Security: 17th International Conference, FC 2013. pp. 6–24.
- [78] RUFFING, T., MORENO-SANCHEZ, P., AND KATE, A. *Coin-Shuffle: Practical Decentralized Coin Mixing for Bitcoin*. Computer Security - ESORICS 2014: 19th European Symposium on Research in Computer Security. pp. 345–364.
- [79] SALA, A., ZHAO, X., WILSON, C., ZHENG, H., AND ZHAO, B. Y. Sharing graphs using differentially private graph models. Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, pp. 81–98.
- [80] SASSON, E. B., CHIESA, A., GARMAN, C., GREEN, M., MIERS, I., TROMER, E., AND VIRZA, M. Zerocash: Decentralized anonymous payments from bitcoin. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy* (2014), pp. 459–474.
- [81] SPAGNUOLO, M., MAGGI, F., AND ZANERO, S. *Bitlodine: Extracting Intelligence from the Bitcoin Network*. Financial Cryptography and Data Security: 18th International Conference, FC 2014. pp. 457–468.
- [82] TRAN, N., MIN, B., LI, J., AND SUBRAMANIAN, L. Sybil-resilient online content voting. Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation, pp. 15–28.
- [83] VALENTA, L., AND ROWAN, B. *Blindcoin: Blinded, Accountable Mixes for Bitcoin*. FC 2015 International BITCOIN Workshops Financial Cryptography and Data Security. pp. 112–126.
- [84] WONDRAČEK, G., HOLZ, T., KIRDA, E., AND KRUEGEL, C. A practical attack to de-anonymize social network users. Proceedings of the 2010 IEEE Symposium on Security and Privacy, pp. 223–238.
- [85] ZHELEVA, E., AND GETOOR, L. Preserving the privacy of sensitive relationships in graph data. Proceedings of the 1st ACM SIGKDD International Conference on Privacy, Security, and Trust in KDD, pp. 153–171.
- [86] ZIEGELDORF, J. H., GROSSMANN, F., HENZE, M., INDEN, N., AND WEHRLE, K. Coinparty: Secure multi-party mixing of bitcoins. In *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy* (2015), pp. 75–86.