

An Analysis of Anonymity in the Bitcoin System

Fergal Reid

Clique Research Cluster
University College Dublin, Ireland
fergal.reid@gmail.com

Martin Harrigan

Clique Research Cluster
University College Dublin, Ireland
martin.harrigan@ucd.ie

Abstract—Anonymity in Bitcoin, a peer-to-peer electronic currency system, is a complicated issue. Within the system, users are identified by public-keys only. An attacker wishing to de-anonymize its users will attempt to construct the one-to-many mapping between users and public-keys and associate information external to the system with the users. Bitcoin frustrates this attack by storing the mapping of a user to his or her public-keys on that user's node only and by allowing each user to generate as many public-keys as required. In this paper we consider the topological structure of two networks derived from Bitcoin's public transaction history. We show that the two networks have a non-trivial topological structure, provide complementary views of the Bitcoin system and have implications for anonymity. We combine these structures with external information and techniques such as context discovery and flow analysis to investigate an alleged theft of Bitcoins, which, at the time of the theft, had a market value of approximately half a million U.S. dollars.

I. INTRODUCTION

Bitcoin is a peer-to-peer electronic currency system first described in a paper by Satoshi Nakamoto (probably a pseudonym) in 2008 [1]. It relies on digital signatures to prove ownership and a public history of transactions to prevent double-spending. The history of transactions is shared using a peer-to-peer network and is agreed upon using a proof-of-work system [2], [3].

The first Bitcoins were transacted in January 2009 and by June 2011 there were 6.5 million Bitcoins in circulation among an estimated 10,000 users [4]. In recent months, the currency has seen rapid growth in both media attention and market price relative to existing currencies. It has featured in both *The Economist* and *Forbes* magazine. At its peak, a single Bitcoin traded for more than US\$30 on popular Bitcoin exchanges. At the same time, U.S. Senators and lobby groups in Germany, such as *Der Bundesverband Digitale Wirtschaft* (BVWD) or the *Federal Association of Digital Economy*, have raised concerns regarding the untraceability of Bitcoins and their potential to harm society through tax evasion, money laundering and illegal transactions. The implications of the decentralized nature of Bitcoin for authorities' ability to regulate and monitor the flow of currency is as yet unclear.

Many users adopt Bitcoin for political and philosophical reasons, as much as pragmatic ones. While there is an understanding amongst Bitcoin's technical users that anonymity is not a prominent design goal of the system, we believe that this awareness is not shared throughout the community. For example, WikiLeaks, an international organization for anonymous

whistleblowers, recently advised its Twitter followers that it now accepts *anonymous* donations via Bitcoin (see Fig. 1) and states that¹:

“Bitcoin is a secure and anonymous digital currency. Bitcoins cannot be easily tracked back to you, and are a [sic] safer and faster alternative to other donation methods.”

They proceed to describe a more secure method of donating Bitcoins that involves the generation of a one-time public-key but the implications for those who donate using the tweeted public-key are unclear. Is it possible to associate a donation with other Bitcoin transactions performed by the same user or perhaps identify them using external information? At present, there is little detailed work on Bitcoin anonymity in the public domain – the extent to which this anonymity holds in the face of determined analysis remains to be tested.



Fig. 1. Screen capture of a tweet from WikiLeaks announcing their acceptance of ‘anonymous Bitcoin donations’.

This paper is organized as follows. In Sect. II we consider some existing work relating to electronic currencies and anonymity. The economic aspects of the system, interesting in their own right, are beyond the scope of this work. In Sect. III we present an overview of the Bitcoin system; we focus on three features that are particularly relevant to our analysis. In Sect. IV we construct two network structures, the transaction network and the user network using the publicly available transaction history. We study the static and dynamic properties of these networks. In Sect. V we consider the implications of these network structures for anonymity. We also combine information external to the Bitcoin system with techniques such as flow and temporal analysis to illustrate how various types of information leakage can contribute to the de-anonymization of the system's users. Finally, we conclude in Sect. VI.

¹<http://wikileaks.org/support.html> – Retrieved: 22-07-2011

A. A Note Regarding Motivation and Disclosure

Our motivation for this analysis is not to de-anonymize individual users of the Bitcoin system. Rather, it is to demonstrate, using a passive analysis of a publicly available dataset, the inherent limits of anonymity when using Bitcoin. This will ensure that users do not have expectations that are not being fulfilled by the system.

In security-related research, there is considerable tension over how best to disclose vulnerabilities [5]. Many researchers favor full disclosure where all information regarding a vulnerability is promptly released. This enables informed users to promptly take defensive measures. Other researchers favor limited disclosure; while this provides attackers with a window in which to exploit uninformed users, a mitigation strategy can be prepared and implemented before public announcement, thus limiting damage, *e.g.* through a software update. Our analysis does not show a vulnerability of the Bitcoin system *per se*, but does illustrate some potential risks and pitfalls with regard to anonymity. However, there is no central authority which can fundamentally change the system's behavior. Furthermore, it is not possible to mitigate analysis of the existing transaction history.

There are also two noteworthy features of the dataset when compared with, say, contentious social network datasets, *e.g.* the Facebook profiles of Harvard University students [6]. Firstly, the delineation between what is considered public and private is clear: the entire history of Bitcoin transactions is publicly available. Secondly, the Bitcoin system does not have a usage policy. After joining Bitcoin's peer-to-peer network, a client can freely request the entire history of Bitcoin transactions; there is no crawling or scraping required.

Thus, we believe the best strategy to minimise the threat to user anonymity is to be descriptive about the risks of the Bitcoin system. We do not identify individual users – apart from those in the case study – but we note that it is not difficult for other groups to replicate our work. Indeed, given the passive nature of the analysis, other parties may already be conducting similar analyses.

II. RELATED WORK

The related work for this paper can be categorized into two fields: electronic currencies and anonymity.

A. Electronic Currencies

Electronic currencies can be technically classified according to their mechanisms for establishing ownership, protecting against double-spending, ensuring anonymity and/or privacy, and generating and issuing new currency. Bitcoin is particularly noteworthy for the last of these mechanisms. The proof-of-work system [2], [3] that establishes consensus regarding the history of transactions also doubles as a minting mechanism. The scheme was first outlined in the B-Money Proposal [7]. We briefly consider some alternative mechanisms. Ripple [8] is an electronic currency where every user can issue currency. However, the currency is only accepted by peers who trust the issuer. Transactions between arbitrary pairs of users

require chains of trusted intermediaries between the users. Saito [9] formalized and implemented a similar system, i-WAT, in which the the chain of intermediaries can be established without their immediate presence using digital signatures. KARMA [10] is an electronic currency where the central authority is distributed over a set of users that are involved in all transactions. PPay [11] is a micropayment scheme for peer-to-peer systems where the issuer of the currency is responsible for keeping track of it. However, both KARMA and PPay may incur a large overhead when the rate of transactions is high. Mondex is a smart-card electronic currency [12]. It preserves a central bank's role in the generation and issuance of electronic currency. Mondex was an electronic replacement for cash in the physical world whereas Bitcoin is an electronic analog of cash in the online world.

The authors are not aware of any studies of the network structure of electronic currencies. However, there are such studies of physical currencies. The community currency Tomamae-cho was introduced into the Hokkaido Prefecture in Japan for a three-month period during 2004–05 in a bid to revitalize local economy. The Tomamae-cho system involved gift-certificates that were re-usable and legally redeemable into yen. There was an entry space on the reverse of each certificate for recipients to record transaction dates, their names and addresses, and the purposes of use, up to a maximum of five recipients. Kichiji and Nishibe [13] used the collected certificates to derive a network structure that represented the flow of currency during the period. They showed that the cumulative degree distribution of the network obeyed a power-law distribution, the network had small-world properties (the average clustering coefficient was high whereas the average path length was low), the directionality and the value of transactions were significant features, and the double-triangle system [14] was effective. There also exist studies of the physical movement of currency: 'Where's George?' [15] is a crowd-sourced method for tracking U.S. dollar bills where users record the serial numbers of bills in their possession, along with their current location. If a bill is recorded sufficiently often, its geographical movement can be tracked over time. Brockmann et al. [16] used this dataset as a proxy for studying multi-scale human mobility and as a tool for computing geographic borders inherent to human mobility.

B. Anonymity

Previous work has shown the difficulty of maintaining anonymity in the context of networked data and online services which expose partial user information. Backstrom et al. [17] considered privacy attacks which identify users using the structure of the network around them and discussed the difficulty of guaranteeing user anonymity in the presence of network data. Crandall et al. [18] infer social ties between users where none are explicitly stated by looking at patterns of 'co-incidences' or common off-network co-occurrences. Narayanan and Shmatikov [19] de-anonymized the Netflix Prize dataset

using information from IMDB² which had similar user content, showing that statistical matching between different but related datasets can be used to attack anonymity. Puzis et al. [20] simulated the monitoring of a communications network using strategically-located monitoring nodes and show that, using real-world network topologies, a relatively small number of nodes can collaborate to pose a significant threat to anonymity. All of this work points to the difficulty in maintaining anonymity where network data on user behaviour is available and illustrates how seemingly minor information leakages can be aggregated to pose significant risks.

III. THE BITCOIN SYSTEM

The following is a simplified description of the Bitcoin system; see Nakamoto [1] for a more thorough treatment. Bitcoin is an electronic currency with no central authority or issuer. There is no central bank or fractional reserve system controlling the supply of Bitcoins. Instead, they are generated at a predictable rate such that the eventual total number will be 21 million. There is no requirement for a trusted third-party when making transactions. Suppose Alice wishes to ‘send’ a number of Bitcoins to Bob. Alice uses a Bitcoin client to join the Bitcoin peer-to-peer network and makes a public transaction or declaration stating that one or more identities that she controls (which can be verified using public-key cryptography), and which previously had a number of Bitcoins assigned to them, wish to re-assign those Bitcoins to one or more other identities, at least one of which is controlled by Bob. The participants of the peer-to-peer network form a collective consensus regarding the validity of this transaction by appending it to the public history of previously agreed-upon transactions (the longest *block-chain*). This process, known as *mining*, involves the repeated computation of a cryptographic hash function so that the digest of the transaction, along with other pending transactions, and an arbitrary nonce, has a specific form. This process is designed to require considerable computational effort, from which the security of the Bitcoin mechanism is derived. To encourage users to pay this computational cost, the process is incentivized using newly generated Bitcoins and/or transaction fees.

In this paper, there are three features of the Bitcoin system that are of particular interest. Firstly, the entire history of Bitcoin transactions is publicly available. This is necessary in order to validate transactions and prevent double-spending in the absence of a central authority. The only way to confirm the absence of a previous transaction is to be aware of all previous transactions. The second feature of interest is that a transaction can have multiple inputs and multiple outputs. An input to a transaction is either the output of a previous transaction or a sum of newly generated Bitcoins and transaction fees. A transaction frequently has either a single input from a previous larger transaction or multiple inputs from previous smaller transactions. Also, a transaction frequently has two outputs: one sending payment and one returning change. Thirdly, the

payer and payee(s) of a transaction are identified through public-keys from public-private key-pairs. However, a user can have multiple public-keys. In fact, it is considered good practice for a payee to generate a new public-private key-pair for every transaction. Furthermore, a user can take the following steps to better protect their identity: they can avoid revealing any identifying information in connection with their public-keys; they can repeatedly send varying fractions of their Bitcoins to themselves using multiple (newly generated) public-keys; and/or they can use a trusted third-party mixer or laundry. However, these practices are not universally applied.

The three features above, namely the public availability of Bitcoin transactions, the input-output relationship between transactions and the re-use and co-use of public-keys, provide a basis for two distinct network structures: the *transaction network* and the *user network*. The transaction network represents the flow of Bitcoins between *transactions* over time. Each vertex represents a transaction and each directed edge between a source and a target represents an output of the transaction corresponding to the source that is an input to the transaction corresponding to the target. Each directed edge also includes a value in Bitcoins and a timestamp. The user network represents the flow of Bitcoins between *users* over time. Each vertex represents a user and each directed edge between a source and a target represents an input-output pair of a single transaction where the input’s public-key belongs to the user corresponding to the source and the output’s public-key belongs to the user corresponding to the target. Each directed edge also includes a value in Bitcoins and a timestamp.

We gathered the entire history of Bitcoin transactions from the first transaction on the 3rd January 2009 up to and including the last transaction that occurred on the 12th July 2011. We gathered the dataset using the Bitcoin client³ and a modified version of Gavin Andresen’s *bitcointools*.⁴ The dataset comprises 1 019 486 transactions between 1 253 054 unique public-keys. We describe the construction of the corresponding transaction and user networks and their analyses in the following sections. We will show that the two networks are complex, have a non-trivial topological structure, provide complementary views of the Bitcoin system and have implications for the anonymity of users.

IV. THE TRANSACTION AND USER NETWORKS

A. The Transaction Network

The transaction network \mathcal{T} represents the flow of Bitcoins between *transactions* over time. Each vertex represents a transaction and each directed edge between a source and a target represents an output of the transaction corresponding to the source that is an input to the transaction corresponding to the target. Each directed edge also includes a value in Bitcoins and a timestamp. It is a straight-forward task to construct \mathcal{T} from our dataset.

²<http://www.imdb.com>

³<http://www.bitcoin.org>

⁴<http://github.com/gavinandresen/bitcointools>

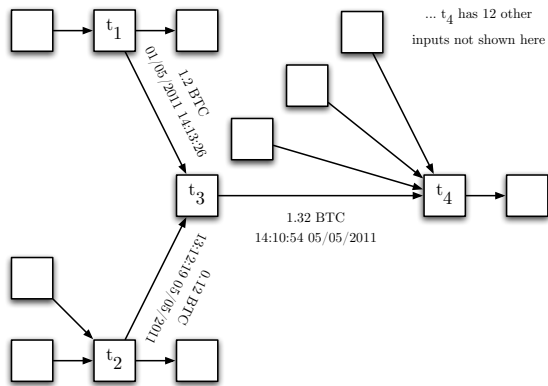


Fig. 2. An example sub-network from the transaction network. Each rectangular vertex represents a transaction and each directed edge represents a flow of Bitcoins from an output of one transaction to an input of another.

Figure 2 shows an example sub-network of \mathcal{T} . t_1 is a transaction with one input and two outputs.⁵ It was added to the block-chain on the 1st May 2011. One of its outputs assigned 1.2 BTC (Bitcoins) to a user identified by the public-key pk_1 .⁶ The public-keys are not shown in Fig. 2. Similarly, t_2 is a transaction with two inputs and two outputs.⁷ It was accepted on the 5th May 2011. One of its outputs sent 0.12 BTC to a user identified by a different public-key, pk_2 .⁸ t_3 is a transaction with two inputs and one output.⁹ It was accepted on the 5th May 2011. Both of its inputs are connected to the two aforementioned outputs of t_1 and t_2 . The only output of t_3 was redeemed by t_4 .¹⁰

\mathcal{T} has 974 520 vertices and 1 558 854 directed edges. The number of vertices is less than the total number of transactions in the dataset because we omit transactions that are not connected to at least one other transaction. These correspond to newly generated Bitcoins and transactions fees that are not yet redeemed. The network has neither multi-edges (multiple edges between the same pair of vertices in the same direction) nor loops. It is a directed acyclic graph (DAG) since the output of a transaction can never be an input (either directly or indirectly) to the same transaction.

Figure 3(a) shows a log-log plot of the cumulative degree distributions: the solid red curve is the cumulative degree distribution (in- and out-degree); the dashed green curve is the cumulative in-degree distribution; and the dotted blue curve is the cumulative out-degree distribution. We fitted power-law distributions, $p(x) \sim x^{-\alpha}$ for $x > x_{min}$, to the three distributions by estimating the parameters α and x_{min} using

⁵The transactions and public-keys used in our examples exist in our dataset. The unique identifier for the transaction t_1 is 09441d3e52fa0018365fcd2949925182f6307322138773d52c201f5cc2bb5976. You can query the details of a transaction or public-key by examining Bitcoin’s longest block-chain using, say, the Bitcoin Block Explorer (<http://www.blockexplorer.com>).

⁶13eBhR3oHFD5wkE4oGtrLdbdi2PvK3ijMC

⁷0c4d41d0f5d2aff14d449daa550c7d9b0eaf35d81ee5e6e77f8948b14d62378

⁸19smBSUoRGmbH13vif1Nu17S63TnmG7h9n

⁹0c034fb964257ecbf4eb953e2362e165dea9c1d008032bc9ece5cebbbc7cd4697

¹⁰f16ece066f6e4cf92d9a72eb1359d8401602a23990990cb84498cddb93026402

a goodness-of-fit method [21]. Table I shows the estimates along with the corresponding Kolmogorov–Smirnov goodness-of-fit (GoF) statistics and p-values. We observe that none of the distributions for which the empirically-best scaling region is non-trivial have a power-law as a plausible hypothesis ($p > 0.1$). This is likely due to the fact that there is no preferential attachment [22], [23]: new vertices are joined to existing vertices whose corresponding transactions are not yet fully redeemed.

There are 1 949 (maximal weakly) connected components in the network. Fig. 3(b) shows a log-log plot of the cumulative component size distribution. There are 948 287 vertices (97.31%) in the giant component. This component also contains a giant biconnected component with 716 354 vertices (75.54% of the vertices in the giant component).

Variable	\tilde{x}	\bar{x}	s	α	x_{min}	GoF	p-val.
Degree	3	3.20	6.20	3.24	50	0.02	0.05
In-Degree	1	1.60	5.31	2.50	4	0.01	0.00
Out-Degree	1	1.60	3.17	3.50	51	0.05	0.00

TABLE I

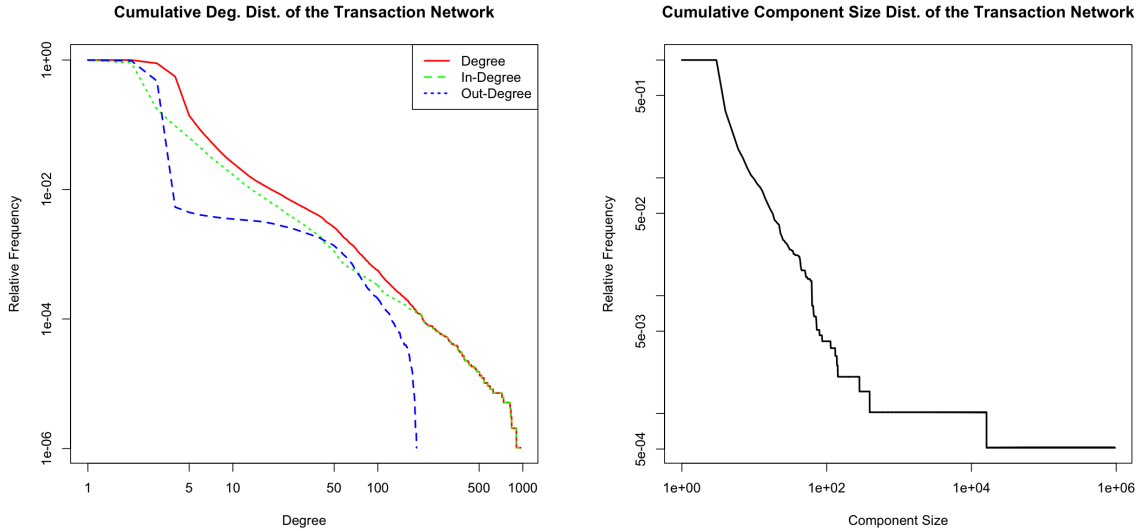
THE DEGREE, IN-DEGREE AND OUT-DEGREE DISTRIBUTIONS OF \mathcal{T} .

We also performed a rudimentary dynamic analysis of the network. Figures 3(c), 3(d) and 3(e) show the edge number, density and average path length of the transaction network on a monthly basis. These measurements are not cumulative. The network’s growth and sparsification are evident. We also observe some anomalies in the average path length during July and November 2010.

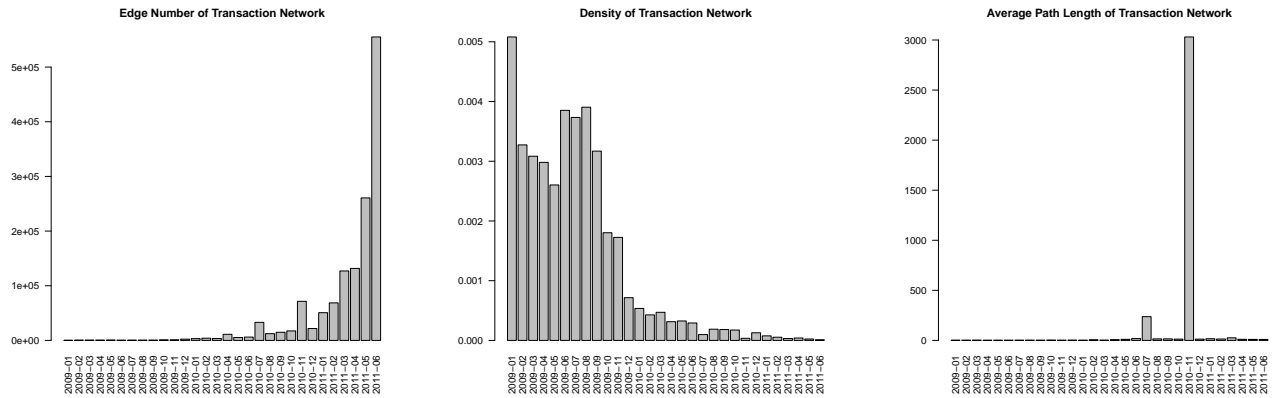
B. The User Network

The user network \mathcal{U} represents the flow of Bitcoins between *users* over time. Each vertex represents a user and each directed edge between a source and a target represents an input-output pair of a single transaction where the input’s public-key belongs to the user corresponding to the source and the output’s public-key belongs to the user corresponding to the target. Each directed edge also includes a value in Bitcoins and a timestamp.

We need to perform a preprocessing step before we can construct \mathcal{U} from our dataset. Suppose \mathcal{U} is, at first, imperfect in the sense that each vertex represents a single public-key rather than a user and that each directed edge between a source and a target represents an input-output pair of a single transaction, where the input’s public-key corresponds to the source and the output’s public-key corresponds to the target. In order to perfect this network, we need to contract each subset of vertices whose corresponding public-keys belong to a single user. The difficulty is that public-keys are Bitcoin’s mechanism for ensuring anonymity: ‘the public can see that someone [identified by a public-key] is sending an amount to someone else [identified by another public-key], but without information linking the transaction to anyone.’ [1]. In fact, it is considered good practice for a payee to generate a new public-private key-pair for every transaction to keep transactions from



(a) A log-log plot of the cumulative degree distributions. (b) A log-log plot of the cumulative component size distribution.



(c) A temporal histogram showing the number of edges per month. (d) A temporal histogram showing the density per month. (e) A temporal histogram showing the average path length per month.

Fig. 3. The degree distributions, component size distribution and monthly edge number, density and average path length of the transaction network.

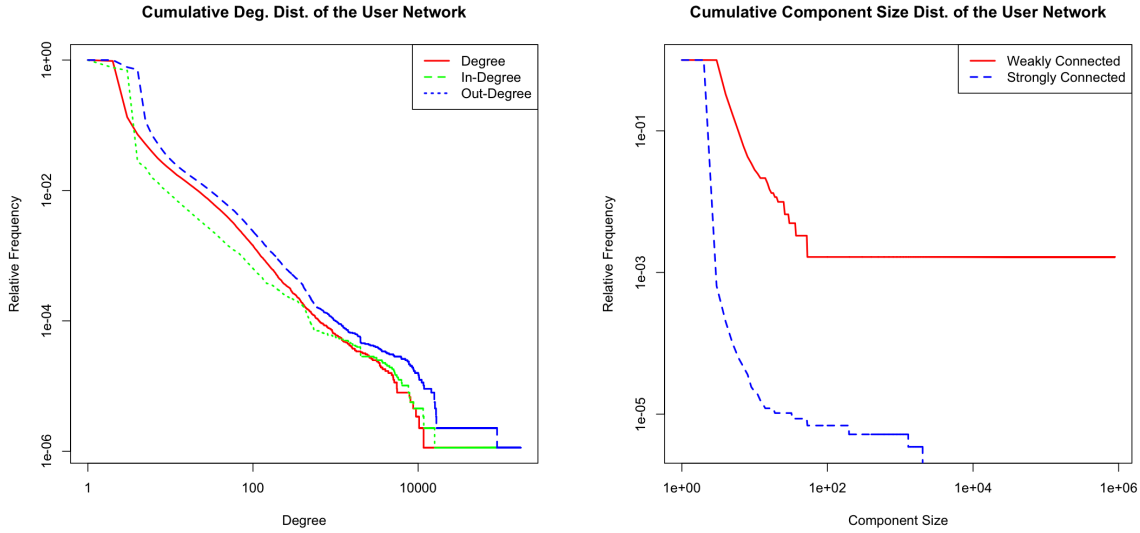
being linked to a common owner. Therefore, it is impossible to completely perfect the network using our dataset alone. However, as noted by Nakamoto [1],

“Some linking is still unavoidable with multi-input transactions, which necessarily reveal that their inputs were owned by the same owner. The risk is that if the owner of a key is revealed, linking could reveal other transactions that belonged to the same owner.”

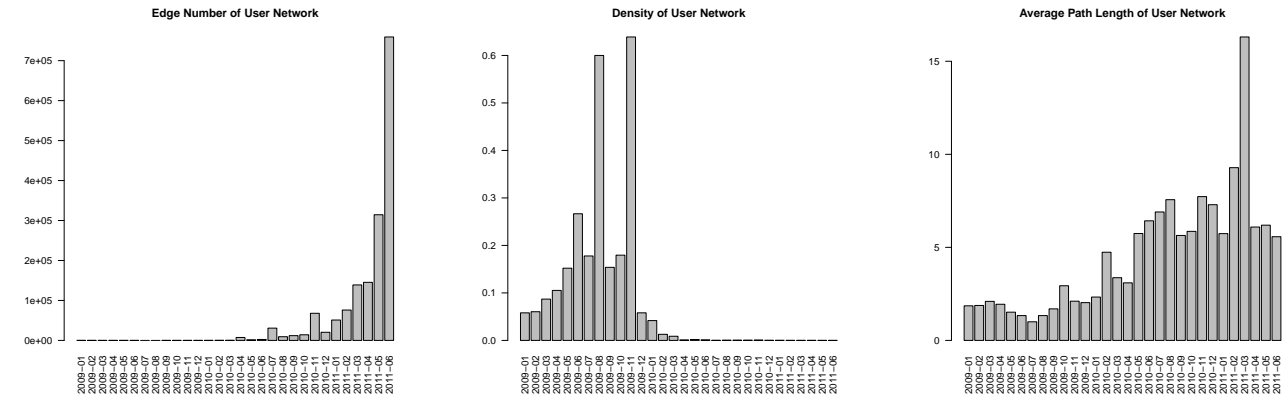
We will use this property of transactions with multiple inputs to contract subsets of vertices in the imperfect network. We construct an ancillary network where each vertex represents a public-key and each undirected edge between two end points represents a pair of inputs of a single transaction whose public-keys correspond to the end points. Using our dataset, this network has 1 253 054 vertices (unique public-keys) and 4 929 950 edges. More importantly, it has 86 641

non-trivial maximal connected components. We deduce that each maximal connected component corresponds to a user and each component’s constituent vertices correspond to that user’s public-keys.

Figure 5 shows an example sub-network of the imperfect network overlaid onto the example sub-network of \mathcal{T} from Fig. 2. The outputs of t_1 and t_2 that were eventually redeemed by t_3 were sent to a user whose public-key was pk_1 and a user whose public-key was pk_2 respectively. Figure 6 shows an example sub-network of the user network overlaid onto the example sub-network of the imperfect network from Fig. 5. pk_1 and pk_2 are contracted into a single vertex u_1 since they correspond to a pair inputs of a single transaction. In other words, they are in the same maximal connected component of the ancillary network (see the vertices representing pk_1 and pk_2 in the dashed grey box in Fig. 6). A single user owns both public-keys. We note that the maximal connected



(a) A log-log plot of the cumulative degree distributions. (b) A log-log plot of the cumulative component size distribution.



(c) A temporal histogram showing the number of edges per month. (d) A temporal histogram showing the density per month. (e) A temporal histogram showing the average path length per month.

Fig. 4. The degree distributions, component size distribution and monthly edge number, density and average path length of the user network.




Fig. 5. An example sub-network from the imperfect network. Each diamond vertex represents a public-key and each directed edge between diamond vertices represents a flow of Bitcoin from one public-key to another.

component in this case is not simply a clique; it has a diameter of four indicating that there are at least two public-keys belonging to that same user that are connected indirectly via three transactions. The sixteen inputs to transaction t_4 result in the contraction of a further sixteen public-keys into a single vertex u_2 . The value and timestamp of the flow of Bitcoins from u_1 to u_2 is derived from the transaction network.

After the preprocessing step, \mathcal{U} has 881 678 vertices (86 641 non-trivial maximal connected components and 795 037 isolated vertices in the ancillary network) and 1 961 636 directed edges. The network is still imperfect. We have not contracted all possible vertices but it will suffice for our present analysis. Unlike \mathcal{T} , \mathcal{U} has multi-edges, loops and directed cycles.

Figure 4(a) shows a log-log plot of the network's cumulative degree distributions. We fitted power-law distributions to the three distributions and calculated their goodness-of-fit and statistical significance as in the previous section. Table II




Fig. 6. An example sub-network from the user network. Each circular vertex represents a user and each directed edge between circular vertices represents a flow of Bitcoins from one user to another. The maximal connected component from the ancillary network that corresponds to the vertex u_1 is shown within the dashed grey box.

shows the results. We observe that none of the distributions have a power-law as a plausible hypothesis.

There are 604 (maximal) weakly connected components and 579 355 (maximal) strongly connected components in the network; Fig. 4(b) shows a log-log plot of the cumulative component size distribution for both variations. There are 879 859 vertices (99.79%) in the giant weakly connected component. This component also contains a giant weakly biconnected component with 652 892 vertices (74.20% of the vertices in the giant component).

Variable	\tilde{x}	\bar{x}	s	α	x_{\min}	GoF	p-val.
Degree	3	4.45	218.10	2.38	66	0.02	0.00
In-Degree	1	2.22	86.40	2.45	57	0.05	0.00
Out-Degree	2	2.22	183.91	2.03	10	0.22	0.00

TABLE II
THE DEGREE, IN-DEGREE AND OUT-DEGREE DISTRIBUTIONS OF \mathcal{U} .

Our dynamic analysis of the user network mirrors that of the transaction network in the previous subsection. Figures 4(c), 4(d) and 4(e) show the edge number, density and average path length of the user network on a monthly basis. These measurements are not cumulative. The network's growth and sparsification are evident. We note that even though our dynamic analysis of the user network is on a monthly basis, the preprocessing step is performed using the ancillary network of the entire imperfect network. This enables us to resolve public-keys to a single user irrespective of the month in which the linking transactions occur.

V. ANONYMITY ANALYSIS

Prior to performing the analyses above, we expected the user network to be largely composed of disjoint trees representing Bitcoin flows between one-time public-keys that were not linked with other public-keys. However, our analyses reveal

that the user network has considerable cyclic structure. We now consider the implications of this structure, coupled with other aspects of the Bitcoin system, for anonymity.

There are several ways in which the user network can be used to deduce information about Bitcoin users. We can use global network properties, such as degree distribution, to identify outliers, *e.g.* users with unusually high activity in a time-window following an event of interest. We can use local network properties to examine the context in which a user operates by observing the egocentric network of a particular user and the users with which he or she interacts with either directly or indirectly. The dynamic nature of the user network also enables us to perform flow and temporal analyses. We can examine the significant Bitcoin flows between groups of users over time. We will now discuss each of these threats in more detail, and provide a case study to demonstrate their potential.

A. Integrating Off-Network Information

There is no user directory for the Bitcoin system. However, we can attempt to build a partial user directory associating Bitcoin users (and their known public-keys) with off-network information. If we can make sufficient associations and combine them with the network structures above, a potentially serious threat to anonymity emerges.

Many organizations and services such as on-line stores that accept Bitcoins, exchanges, laundry services and mixers have access to identifying information regarding their users, *e.g.* e-mail addresses, shipping addresses, credit card and bank account details, IP addresses, etc. If any of this information was publicly available, or accessible by, say, law enforcement agencies, then the identities of users involved in related transactions may also be at risk. To illustrate this point, we consider a number of publicly available data sources and integrate their information with the user network.

1) *The Bitcoin Faucet*: The Bitcoin Faucet¹¹ is a website where users can donate Bitcoins to be redistributed in small amounts to other users. In order to prevent abuse of this service, a history of recent give-aways are published along with the IP addresses of the recipients. When the Bitcoin Faucet does not batch the re-distribution, it is possible to associate the IP addresses with the recipient's public-keys. This page can be scraped over time to produce a time-stamped mapping of IP addresses to users.

We found that the public-keys associated with many of the IP addresses that received Bitcoins were contracted with other public-keys in the ancillary network, thus revealing IP addresses that are somehow related to previous transactions. Fig. 7(a) shows a map of geolocated IP addresses belonging to users who received Bitcoins over a period of one week. Fig. 7(b) overlays the user network onto a sample of those users. An edge between two geolocated IP addresses indicates that the corresponding users are linked by an undirected path of length at most three in the user network; the path must not contain the vertex representing the Bitcoin Faucet itself.

¹¹<http://freebitcoins.appspot.com>

These figures serve as a proof-of-concept from a small publicly available data source. We note that large centralized Bitcoin service providers are capable of producing much more detailed maps.

2) *Voluntary Disclosures*: Another source of identifying information is the voluntary disclosure of public-keys by users, for example, when posting to the Bitcoin forums¹². Bitcoin public-keys are typically represented as strings approximately thirty-three characters in length and starting with the digit one. They are indexed very well by popular search engines. We identified many high-degree vertices with external information using a search engine alone. We proceeded to scrape the Bitcoin Forums where users frequently attach a public-key to their signatures. We also gathered public-keys from Twitter streams and user-generated public directories. It is important to note that in many cases we are able to resolve the ‘public’ public-keys with other public-keys belonging to the same user using the ancillary network. We also note that large centralized Bitcoin service providers can do the same with their user information.

B. Egocentric Analysis and Visualization of the User Network

There are several pieces of information we can directly derive from the user network regarding a particular user. We can compute the balance held by a single public-key. We can also aggregate the balances belonging to public-keys that are controlled by a particular user. For example, Fig. 8(a) and Fig. 8(b) show the receipts and payments to and from WikiLeaks’ public-key in terms of Bitcoin and transaction volume respectively. The donations are generally small donations and are forwarded to other public-keys periodically. There was also a noticeable spike in donations when the facility was first announced. Figure 8(c) shows the receipts and payments to and from the creator of a popular Bitcoin trading website aggregated over a number of public-keys that are linked through the ancillary network.

An important advantage of deriving network structures from the Bitcoin transaction history is our ability to use network visualization and analysis tools to investigate the flow of Bitcoins. For example, Fig. 9 shows the network structure surrounding the WikiLeaks’ public-key in the imperfect user network. Our tools resolve several of the vertices with identifying information gathered in Sect. V-A. These users can be linked either directly or indirectly to their donations. The presence of a Bitcoin mining pool (a large red vertex) and a number of public-keys between it and WikiLeaks’ public-key is interesting.

C. Context Discovery

Given a number of public-keys or users of interest, we can use network structure and context to better understand the flow of Bitcoins between them. For example, we can examine all shortest paths between a set of vertices or consider the maximum number of Bitcoins that can flow from a source




Fig. 9. An egocentric visualization of the vertex representing WikiLeaks’ public-key in the imperfect user network. The size of a vertex corresponds to its degree in the entire imperfect user network. The color denotes the volume of Bitcoins – warmer colors have larger volumes flowing through them. The large red vertices represent a Bitcoin mining pool, a centralized Bitcoin wallet service and an unknown entity.

to a destination given the transactions and their ‘capacities’ in an interesting time-window. For example, Fig. 10 shows all shortest paths between the vertices representing the users we identified using off-network information in Sect. V-A and the vertex that represents the MyBitcoin service¹³ in the user network. We can identify more than 60% of the users in this visualization and deduce many direct and indirect relationships between them.

Case Study – Part I: We analyse an alleged theft of 25 000 BTC reported in the Bitcoin Forums¹⁴ by a user known as *allinvain*. The victim reported that a large portion of his Bitcoins were sent to pk_{red} ¹⁵ on 13/06/2011 at 16:52:23 UTC. The theft occurred shortly after somebody broke into the victim’s Slush pool account¹⁶ and changed the payout address to pk_{blue} ¹⁷. The Bitcoins rightfully belonged to pk_{green} ¹⁸. At the time of the theft, the stolen Bitcoins had a market value of approximately half a million U.S. dollars. We chose this case study to illustrate the potential risks to the anonymity of a user (the thief) who has good reason to remain anonymous.

We consider the imperfect user network before any contractions. We restrict ourselves to the egocentric network surrounding the thief: we include every vertex that is reachable by a path of length at most two ignoring directionality and all

¹³<http://www.mybitcoin.com>

¹⁴<http://forum.bitcoin.org/index.php?topic=16457.0>


¹⁵1KPTdMb6p7H3YCwsyFqrEmKGmsHqe1Q3jg

¹⁶<http://http://mining.bitcoin.cz>


¹⁷15iUDqk6nLmav3B1xUHPQivDpfMruVsu9f

¹⁸1J18yk7D353z3gRVcdbS7PV5Q8h5w6oWWG

¹²<http://forum.bitcoin.org>

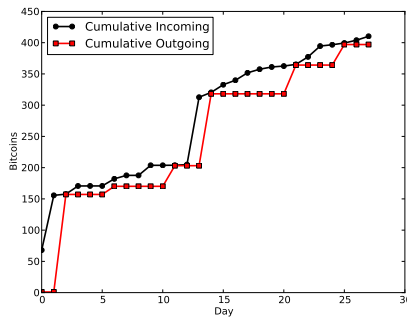


(a) A map of geolocated IP addresses associated with users receiving Bitcoins from the Bitcoin Faucet during a one week period.




(b) A map of a sample of the geolocated IP addresses in Fig. 7(a) connected by edges where the corresponding users are connected by a path of length at most three in the user network that does not include the vertex representing the Bitcoin Faucet.

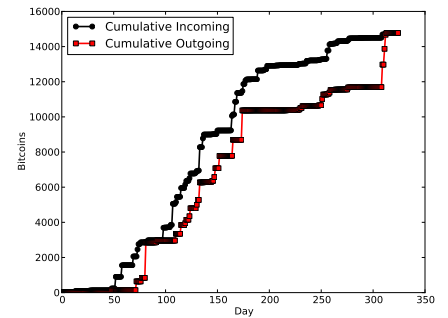
Fig. 7. We can use the Bitcoin Faucet to map users to geolocated IP addresses.



(a) The receipts and payments to and from WikiLeaks' public-key over time.



(b) The number of transactions involving WikiLeaks' public-key over time.



(c) The receipts and payments to and from the creator of a popular Bitcoin trading website aggregated over a number of public-keys.

Fig. 8. Plots of the cumulative receipts and payments to and from Bitcoin public-keys and users.

edges induced by these vertices. We also remove all loops, multiple edges and edges that are not contained in some biconnected component to avoid clutter. In Fig. 11, the red vertex represents the thief who owns the public-key pk_{red} and the green vertex represents the victim who owns the public-key pk_{green} . The theft is the green edge joining the victim to the thief. There are in fact two green edges located nearby in Fig. 11 but only one directly connects the victim to the thief.

Interestingly, the victim and the thief are joined by paths (ignoring directionality) other than the green edge representing the theft. For example, consider the sub-network shown in Fig. 12 induced by the red, green, purple, yellow and orange vertices. This sub-network is a cycle. We contract all vertices whose corresponding public-keys belong to the same user. This allows us to attach values in Bitcoins and timestamps to the directed edges. We can make a number of observations. Firstly, we note that the theft of 25 000 BTC was preceded by a smaller theft of 1 BTC. This was later reported by the victim in the Bitcoin forums. Secondly, using off-network data, we have identified some of the other colored vertices: the purple vertex represents the main Slush pool account and the orange vertex

represents the computer hacker group known as LulzSec.¹⁹ We note that there has been at least one attempt to associate the thief with LulzSec²⁰. This was a fake; it was created after the theft. However, the identification of the orange vertex with LulzSec is genuine and was established before the theft. We observe that the thief sent 0.31337 BTC to LulzSec shortly after the theft but we cannot otherwise associate him with the group. The main Slush pool account sent a total of 441.83 BTC to the victim over a 70-day period. It also sent a total of 0.2 BTC to the yellow vertex over a two day period. One day before the theft, the yellow vertex also sent 0.120607 BTC to LulzSec.

The yellow vertex represents a user who is the owner of at least five public-keys.²¹ Like the victim, he is a member of the Slush pool, and like the thief, he is a one-time donator

¹⁹<http://twitter.com/LulzSec/status/76388576832651265>

²⁰<http://pastebin.com/88nGp508>

²¹1MUpbAY7rjWxvLtUwLkARViqSdzypMgVW4
13ts9ukW294Q7f6zRjr3VmLq6zp1C68EK
1DcQvXMD87MaYcFZqHzDZyH3sAv8R5hMZe
1AEW9ToWWwKoLFYsSsLkPqDyHeS2feDVsvZ
1EWASKF9DLUCgEFqgrNaHzp3q4oEgjTsF




Fig. 10. A visualisation of all users identified in Sect. V-A and all shortest paths between the vertices representing those users and the vertex representing the MyBitcoin service in the user network.




Fig. 11. An egocentric visualization of the thief in the imperfect user network. For this visualization, vertices are identified through their colors in the text, edges are colored according to the color of their sources and the size of each vertex is proportional to its edge-betweenness within the egocentric network.

to LulzSec. This donation, the day before the theft, is his last known activity using these public-keys.

D. Flow and Temporal Analyses

In addition to visualizing egocentric networks with a fixed radius, we can follow significant flows of value through the network over time. If a vertex representing a user receives a large volume of Bitcoins relative to their estimated balance, and, shortly after, transfers a significant proportion of those Bitcoins to another user, we deem this interesting. We built a




Fig. 12. An interesting sub-network induced by the thief, the victim and three other vertices. The notation is the same as in Fig. 11.

special purpose tool that, starting with a chosen vertex or set of vertices, traces significant flows of Bitcoins over time. In practice we have found this tool to be quite revealing when analyzing the user network.

Case Study – Part II: To demonstrate this tool we reconsider the Bitcoin theft described earlier. We note that the victim has developed their own tool to generate an exhaustive list of public-keys that have received some portion of the stolen Bitcoins since the theft²². However, this list grows very quickly and, at the time of writing, contained more than 34 100 public-keys. Figure 13 shows an annotated visualization produced using our tool. We observe several interesting flows in the aftermath of the theft. The initial theft of a small volume of 1 BTC is immediately followed by the theft of 25 000 BTC. This is represented as a dotted black line between the relevant vertices, magnified in the left inset.

In the left inset, we can see that the Bitcoins are shuffled between a small number of accounts and then transferred back to the initial account. After this shuffling step, we have identified four significant outflows of Bitcoins that began at 19:49, 20:01, 20:13 and 20:55. Of particular interest are the outflows that began at 20:55 (labeled as ‘1’ in both insets) and 20:13 (labeled as ‘2’ in both insets). These outflows pass through several subsequent accounts over a period of several hours. Flow 1 splits at the vertex labeled *A* in the right inset

²²<http://folk.uio.no/vegardno/allinvain-addresses.txt>

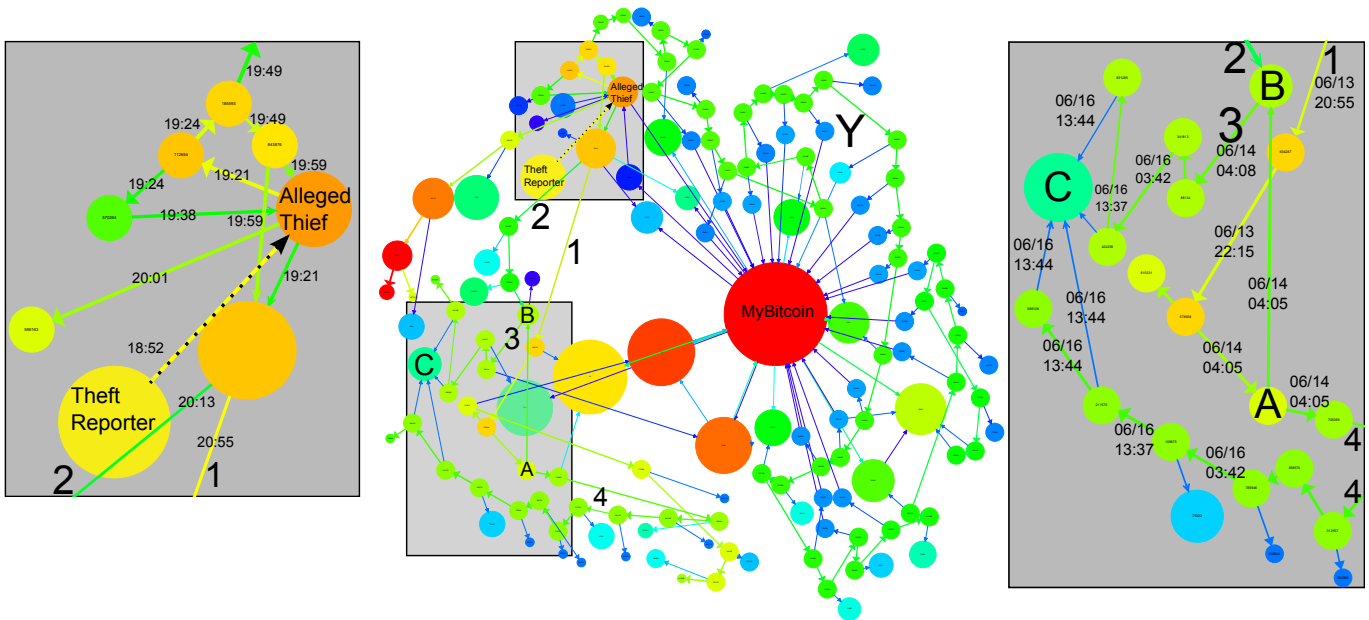


Fig. 13. Visualisation of Bitcoin flow from the alleged theft. The left inset shows the initial shuffling of Bitcoins among accounts close to that of the alleged thief, during which all transfers happen within a few hours of the incident. The right inset shows detail on the events of several subsequent days, where Bitcoin flows split, and then later merge back into each other, validating that the flows found by the tool are probably still controlled by a single party.

at 04:05 the day after the theft. Some of its Bitcoins rejoin Flow 2 at the vertex labeled *B*. This new combined flow is labeled as ‘3’ in the right inset. The remaining Bitcoins from Flow 1 pass through several additional vertices in the next two days. This flow is labeled as ‘4’ in the right inset.

A surprising event occurs on 16/06/2011 at approximately 13:37. A small number of Bitcoins are transferred from Flow 3 to a heretofore unseen public-key pk_1 .²³ Approximately seven minutes later, a small number of Bitcoins are transferred from Flow 3 to another heretofore unseen public-key pk_2 .²⁴ Finally, there are two simultaneous transfers from Flow 4 to two more heretofore unseen public-keys: pk_3 ²⁵ and pk_4 .²⁶ We have determined that these four public-keys, pk_1 , pk_2 , pk_3 and pk_4 – which receive Bitcoins from two separate flows that split from each other two days previously – are all contracted to the same user in our ancillary network. This user is represented as *C* in Fig. 13.

There are several other examples of interesting flow. The flow labeled as *Y* involves the movement of Bitcoins through thirty unique public-keys in a very short period of time. At each step, a small number of Bitcoins (typically 30 BTC which had a market value of approximately US\$500 at the time of the transactions) are siphoned off. The public-keys that receive the small number of Bitcoins are typically represented by small blue vertices due to their low volume and degree. On 20/06/2011 at 12:35, each of these public-keys makes a transfer to a public-key operated by the the MyBitcoin

service.²⁷ Curiously, this public-key was previously involved in another separate Bitcoin theft.²⁸




Fig. 14. The Bitcoins are transferred between public-keys along the highlighted paths very quickly.

We also observe that the Bitcoins in many of the above flows are transferred between public-keys very quickly. Fig. 14 shows two flows in particular where the intermediate parties

²³ 1FKFiCYJSFqxT3zkZntHjfu47SvAzauZXN

²⁴ 1FhYawPhWDvkZCJVBrDfQoo2qC3EuKtb94

²⁵ 1MJZZmmSrQZ9NzeQt3hYP76oFC5dWAF2nD

²⁶ 12dJo17jcr78Uk1Ak5wfgYXtciU62MzcEc

²⁷ 1MAazCWMydsQB5ynYXqSGQDjNQMn3HFmEu

²⁸ <http://forum.bitcoin.org/index.php?topic=20427.0>

waited for very few confirmations before re-sending the Bitcoins to other public-keys.

Naturally, much of this analysis is circumstantial. We cannot say for certain whether or not these flows imply a shared agency in both incidents. There is always the possibility of drawing false inferences. However, it does illustrate the power of our tool when tracing the flow of Bitcoins and generating hypotheses. It also suggests that a centralized service may have further details on the user(s) in control of the implicated public-keys.

The same tool that we use here to examine outflows of Bitcoins from users of interest can also be used to examine inflows of Bitcoins to users and organizations accepting Bitcoins as payment or as donations.

E. Other Forms of Analysis

There are many other forms analysis that can be applied in order to de-anonymize the workings of the Bitcoin system:

- Order books for Bitcoin exchanges are typically made available to support trading tools. As orders are often placed in Bitcoin values converted from other currencies, they often have a precise decimal value with eight significant digits. It may be possible to find transactions with corresponding amounts and thus map public-keys and transactions to the exchanges.
- It is conceivable that over an extended time period, several public-keys, if used at particular times, may belong to the same user. We could construct and cluster a co-occurrence network to help deduce further mappings between public-keys and users.
- Finally, there are far more sophisticated forms of attack where the attacker actively participates in the network, for example, using marked Bitcoins or by operating a laundry service.

VI. CONCLUSIONS

For the past half-century futurists have heralded the advent of a cash-less society [24]. Many of their predictions have been realized, *e.g.* Anderson et al.'s [24]'s 'on-line real-time' payment system and bank-maintained 'central information files'. However, cash is still a competitive and relatively anonymous means of payment. Bitcoin is an electronic analog of cash in the online world. It is decentralized: there is no central authority responsible for the issuance of Bitcoins and there is no need to involve a trusted third-party when making online transfers. However, this flexibility comes at a price: the entire history of Bitcoin transactions is publicly available. In this paper we investigated the structure of two networks derived from this dataset and their implications for user anonymity.

Using an appropriate network representation, it is possible to map many users to public-keys. This is performed using a passive analysis only. Active analyses, where an interested party can potentially deploy marked Bitcoins and collaborating users can discover even more information. We also believe that large centralized services such as the exchanges and wallet

services are capable of identifying considerable portions of user activity.

Technical members of the Bitcoin community have cautioned that strong anonymity is not a prominent design goal of the Bitcoin system. However, casual users need to be aware of this, especially when sending Bitcoins to users and organizations they would prefer not to be publicly associated with.

VII. ACKNOWLEDGEMENTS

This research was supported by Science Foundation Ireland (SFI) Grant No. 08/SRC/I1407: Clique: Graph and Network Analysis Cluster. Both authors contributed equally to this work. It was performed independently of any industrial partnership or collaboration of the Clique Cluster.

REFERENCES

- [1] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008. [Online]. Available: <http://bitcoin.org/bitcoin.pdf>
- [2] C. Dwork and M. Naor, "Pricing via Processing or Combatting Junk Mail," in *Proceedings of the 12th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO'92)*. Springer, 1992, pp. 139–147.
- [3] A. Back, "Hashcash – A Denial of Service Counter-Measure," 2002. [Online]. Available: <http://www.hashcash.org/papers/hashcash.pdf>
- [4] The Economist, "Digital Curriencies – Bits and Bob," June 2011. [Online]. Available: <http://www.economist.com/node/18836780>
- [5] H. Cavusoglu, H. Cavusoglu, and S. Raghunathan, "Emerging Issues in Responsible Vulnerability Disclosure," in *Proceedings of the 4th Annual Workshop on Economics of Information Security (WEIS'05)*, 2005.
- [6] K. Lewis, J. Kaufman, M. Gonzalez, A. Wimmer, and N. Christakis, "Tastes, Ties, and Time: A New Social Network Dataset using Facebook.com," *Social Networks*, vol. 30, pp. 330–342, 2008.
- [7] W. Dai, "B-Money Proposal," 1998. [Online]. Available: <http://www.weidai.com/bmoney.txt>
- [8] R. Fugger, "Money as IOUs in Social Trust Networks A Proposal for a Decentralized Currency Network Protocol," 2004. [Online]. Available: <http://www.ripple-project.org/decentralizedcurrency.pdf>
- [9] K. Saito, "i-WAT: The Internet WAT System – An Architecture for Maintaining Trust and Facilitating Peer-to-Peer Barter Relationships," Ph.D. dissertation, Keio University, 2006.
- [10] V. Vishnumurthy, S. Chandrakumar, and E. Sirer, "KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing," in *Proceedings of the 1st Workshop on Economics of Peer-to-Peer Systems*. [Online]. Available: <http://www2.sims.berkeley.edu/research/conferences/p2pecon/index.html>
- [11] B. Yang and H. Garcia-Molin, "PPay: Micropayments for Peer-to-Peer Systems," in *Proceedings of the 10th ACM Conference on Computer and Communication Security (CCS'03)*, V. Atluri and P. Liu, Eds. ACM Press, 2003, pp. 300–310.
- [12] F. Stalder, "Failures and Successes: Notes on the Development of Electronic Cash," *The Information Society (TIS)*, vol. 18, no. 3, pp. 209–219, 2002.
- [13] N. Kichiji and M. Nishibe, "Network Analyses of the Circulation Flow of Community Currency," *Evolutionary and Institutional Economics Review*, vol. 4, no. 2, pp. 267–300, 2008.
- [14] M. Nishibe, "Chiiki Tuka No Susume (in Japanese)," Hokkaido Shokoukai Rengou, 2004.
- [15] "Where's George?" <http://www.wheresgeorge.com>.
- [16] D. Brockmann, L. Hufnagel, and T. Geisel, "The Scaling Laws of Human Travel," *Nature*, vol. 439, no. 26, pp. 462–465, 2006.
- [17] L. Backstrom, C. Dwork, and J. Kleinberg, "Wherefore Art Thou r3579x?: Anonymized Social Networks, Hidden Patterns, and Structural Steganography," in *Proceedings of the 16th International Conference on World Wide Web*. ACM, 2007, pp. 181–190.
- [18] D. Crandall, L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg, "Inferring Social Ties from Geographic Coincidences," *Proceedings of the National Academy of Sciences*, vol. 107, no. 52, p. 22436, 2010.

- [19] A. Narayanan and V. Shmatikov, "Robust De-anonymization of Large Sparse Datasets," in *Proceedings of the 29th Symposium on Security and Privacy*. IEEE, 2008, pp. 111–125.
- [20] R. Puzis, D. Yagil, Y. Elovici, and D. Braha, "Collaborative Attack on Internet Users' Anonymity," *Internet Research*, vol. 19, no. 1, pp. 60–77, 2009.
- [21] A. Clauset, C. Shalizi, and M. Newman, "Power-Law Distributions in Empirical Data," *SIAM Review*, vol. 51, no. 4, pp. 661–703, 2009.
- [22] H. Simon, "On a Class of Skew Distribution Functions," *Biometrika*, vol. 42, pp. 425–440, 1955.
- [23] A. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [24] A. Anderson, D. Cannell, T. Gibbons, G. Grote, J. Henn, J. Kennedy, M. Muir, N. Potter, and R. Whitby, *An Electronic Cash and Credit System*. American Management Association, 1966.