

An Inquiry into Money Laundering Tools in the Bitcoin Ecosystem

Malte Möser
Department of Information Systems
University of Münster
Münster, Germany
malte.moeser@uni-muenster.de

Rainer Böhme
Department of Information Systems
University of Münster
Münster, Germany
rainer.boehme@uni-muenster.de

Dominic Breuker
Department of Information Systems
University of Münster
Münster, Germany
dominic.breuker@uni-muenster.de

Abstract—We provide a first systematic account of opportunities and limitations of anti-money laundering (AML) in Bitcoin, a decentralized cryptographic currency proliferating on the Internet. Our starting point is the observation that Bitcoin attracts criminal activity as many say it is an anonymous transaction system. While this claim does not stand up to scrutiny, several services offering increased transaction anonymization have emerged in the Bitcoin ecosystem – such as Bitcoin Fog, BitLaundry, and the Send Shared functionality of Blockchain.info. Some of these services routinely handle the equivalent of 6-digit dollar amounts. In a series of experiments, we use reverse-engineering methods to understand the mode of operation and try to trace anonymized transactions back to our probe accounts. While Bitcoin Fog and Blockchain.info successfully anonymize our test transactions, we can link the input and output transactions of BitLaundry. Against the backdrop of these findings, it appears unlikely that a Know-Your-Customer principle can be enforced in the Bitcoin system. Hence, we sketch alternative AML strategies accounting for imperfect knowledge of true identities but exploiting public information in the transaction graph, and discuss the implications for Bitcoin as a decentralized currency.

Index Terms—Bitcoin, Money Laundering, Deanonimization

I. INTRODUCTION

The recent crackdown on Liberty Reserve, supposedly the largest case of cross-border money laundering in history [1], has put a spotlight on other digital currencies as potential vehicles for money laundering. This includes Bitcoin, a decentralized cryptographic currency which emerged over the past three years and has been regarded with suspicion for its allegedly anonymous and irreversible transactions [2], its popularity in underground markets [3], and its association with several cases of investment fraud [4], [5]. As Bitcoin is not controlled by any central entity, the core system defies regulation and enforcement efforts, which adds to the worries of financial regulators and cybercrime fighters. Indeed, Bitcoin is very different from many other digital currencies.

Liberty Reserve was essentially a financial intermediary offering online banking to international customers from its base in Costa Rica (though without having a bank license there), and interfacing with various payment systems. Customer accounts were denominated in dollar, euro, or gold, all relabeled to appear like new currencies by adding the prefix “Liberty”. What made it susceptible to money laundering was the – intentionally, some say – lax enforcement of the Know-

Your-Customer (KYC) principle [6], a requirement mandating financial service providers to validate the identity of account holders. KYC was tightened in the US Patriot Act (with most other jurisdictions following suit) in order to strengthen efforts of anti-money laundering (AML) and combating the financing of terrorism (CFT). However, KYC is only one cornerstone in achieving these ends. It must be complemented by risk assessment, monitoring, reporting and enforcement measures. Once identities are established via KYC, they become the identifiers enabling the downstream activities. Standard procedures include suspicious activity reports filed with financial intelligence units (FIUs) or automatic cross-checks against blacklists maintained by financial crime fighters, such as the US Office of Foreign Assets Control. In simple terms, AML in conventional payment systems relies on known identities and does not require a full picture of all transactions.

Bitcoin, by contrast, is designed with pseudonymous identities. Account numbers are public keys of a specific asymmetric encryption system. Account ownership is established by knowing the corresponding private key. Everyone with a computer can create valid key pairs from large random numbers and thus open one or many Bitcoin accounts. Although the relation between Bitcoin accounts and civil identities of their owners is a priori unknown, Bitcoin transactions are not anonymous. A simple abstraction for Bitcoin is to think of it as a *public* distributed ledger which records all transactions between valid Bitcoin accounts. This information is securely stored in a constantly growing data structure called block chain and remains visible for everyone, forever. Bitcoin attracted media attention for allegedly being an anonymous digital currency (e. g., [7]), especially when organizations like WikiLeaks described it as a “secure and anonymous digital currency”, that “cannot be easily traced back to you” [8]. However, because all transactions are recorded in the public ledger, the anonymity of a user relies on the pseudonym (i. e., the account number or public key) not being associated with his or her true identity. Meanwhile also the Bitcoin community officially acknowledges that “the current implementation is not very anonymous” [9]. The lesson for cybercrime fighters is: AML in Bitcoin has to deal with imperfect knowledge of identities, but may exploit perfect knowledge of all transactions. This calls for new strategies.

In principle, if KYC could be enforced at the edges of the Bitcoin system, that is where bitcoins are exchanged for conventional currencies or products and services, then it becomes possible to identify suspicious activities in the public transaction ledger and hold the perpetrators accountable when and where they interact with the real world. Recent initiatives towards better customer identification by the major Bitcoin exchanges are encouraging signals to this end. Mt.Gox, a popular exchange, for instance, requires scans of national identity documents for transactions involving conventional currencies [10]. However, the endeavor towards stronger identities at the periphery of the Bitcoin system is thwarted by intermediaries who offer services to anonymize the relation between senders and recipients of transactions within the Bitcoin system. These services are also known as Bitcoin mixes. The term alludes to David Chaum's [11] concept of mixes for anonymous communication systems (although they are technically different, as we will explain below). Some of these services trade under indicative names, such as "BitLaundry."

Of course, other use-cases for Bitcoin mixes than money laundering are conceivable. For example, donors may have legitimate interest in financial privacy. They may use transaction anonymizers to evade being observed by attackers monitoring incoming transactions to publicly known Bitcoin addresses of organizations who advocate, say, a political mission. In the context of this work, however, we focus on the threat of money laundering and options to track it down.

The purpose of this paper is to give a systematic account of the available money laundering tools in the Bitcoin ecosystem, to understand their modes of operation, and to draw conclusions on the effectiveness of AML efforts in Bitcoin. More specifically, we evaluate whether currently available Bitcoin mixes can increase the anonymity of its users; and if so, at which cost and risk. We test three services and try to trace the anonymized transactions in the public ledger. We find that while the service BitLaundry does not provide sufficient anonymity, both Bitcoin Fog and Blockchain.info make it impossible for us to find any direct connections in the transaction graph.

The remainder of this paper is organized as follows. Section II recalls essentials of the Bitcoin protocol with special emphasis on anonymization services and their implications for the traceability of money flows. Section III documents our experimental approach and the selection of services under analysis. Section IV presents the results, Section V briefly reviews related work; and the final Section VI concludes with a critical discussion of our results, suggested regulatory options to aid AML, as well as their implications for Bitcoin as a whole.

II. TRACEABILITY OF BITCOIN TRANSACTIONS

A. *The Bitcoin Protocol*

Bitcoin is a digital, distributed, cryptographic currency developed by an open source community. The idea behind it was proposed in 2008 by an unknown author using the pseudonym Satoshi Nakamoto [2]. In short, Bitcoin can be described as a decentralized accounting system in which accounts are associated with public keys of an asymmetric

encryption scheme. Knowledge of the corresponding private key allows account holders to create digital signatures, thereby proving their eligibility to access that account.

A core concept of the Bitcoin system is a transaction, useful to transfer money between accounts. Every transaction consists of a list of outputs, that are tuples of monetary amounts and public keys identifying the destination account; and a list of inputs, which are references to outputs of previous transactions. The semantics of a transaction is that the inputs completely consume the referenced outputs (from previous transactions). Inputs lower the balance of the sender's account(s), whereas outputs (of the current transaction) increase the balance of the receiver's account(s). To ensure that only the owner can withdraw from an account, any transaction is combined with digital signatures corresponding to the public keys referenced in the inputs.

The protocol stipulates that any output referenced by an input is used up and cannot be referenced again. This prevents users from double-spending their money by referencing an output in two different transactions. A difficulty in enforcing this rule is that Bitcoin is a distributed system. Hence, a recipient of a transaction may not be aware that the sender has referenced a particular output in another transaction before.

For this reason, Bitcoin maintains a (probabilistically) consistent record of all transactions, the so-called block chain. Blocks are data structures encapsulating transactions as well as a reference to the previous block, thereby forming a chain. Conflicts, such as forks (i. e., blocks referenced by two or more blocks), are resolved using a proof-of-work scheme. Blocks are considered valid only if all of their transactions are valid and if they come with a solution of a computational intensive problem parameterized by this block. While the solution is hard to find, its validity can easily be verified.

A large number of so-called miners permanently try to find such solutions at the expense of computing power. The difficulty of the problem regularly adapts such that all miners taken together find a solution on average every ten minutes. The finder of a solution receives a monetary reward that is paid out in form of a new transaction that has no inputs: a coinbase transaction. This reward is halved every 210,000 blocks and amounts to 25 BTC at the time of writing. In addition to this reward, the finder can also claim transaction fees for all transactions he included in the block (by increasing the coinbase transaction's value). Whenever the combined monetary value of all a transaction's outputs is lower than that of its inputs, the difference determines the transaction fee.

Any user of the Bitcoin system may maintain a local copy of the block chain and will resolve conflicts by believing in the longest chain. The rationale is that the longest chain must have been created by the majority of all miners. Hence, no single attacker can gain control over the block chain unless he manages to gain control over more computing power than all other miners taken together.

In the experiments conducted in this paper, we use the public information in the Bitcoin block chain to explore anonymization services. In particular, we will explore parts of the Bitcoin

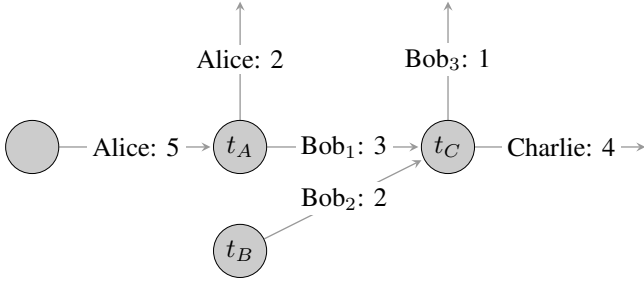


Figure 1. Example of a partial transaction graph

transaction graph. To illustrate these graphs, we use the notation of Figure 1. Nodes correspond to transactions. A directed edge from one node to another denotes that an output of the source transaction is referenced by an input of the target transaction (i.e., if bitcoins from the source transaction are spent in the target transaction). Directed edges not having a target node correspond to outputs not yet referenced by an input. Edges may be annotated with addresses and/or values if needed.

The example in Figure 1 should be read as follows. Alice sends 3 bitcoins to Bob using transaction t_A . To do so, she references an output with a value of 5 from an (unnamed) previous transaction and creates two outputs of t_A , one sending 3 bitcoins to Bob, the other one returning 2 bitcoins to her. Returning bitcoins to oneself is common practice. The amount of bitcoins referenced by all inputs will usually not equal the amount one actually wants to send. As inputs may only be used once, a new output must be created to return the change.

Continuing the description of Figure 1, Bob now wants to send 4 bitcoins to Charlie. Instead of one, he actually owns three different accounts and received 3 bitcoins from Alice to his first: Bob₁. In order to send the 4 bitcoins to Charlie, he needs to create a transaction t_C with two inputs. One references the output of t_A , in which Bob₁ received 3 bitcoins, the other one references an output of another transaction t_B , in which Bob₂ received 2 bitcoins. Together, the referenced outputs sum up to 5, allowing Bob to send 4 of them to Charlie and the change back to him (by creating corresponding outputs in t_C). These outputs are not referenced yet. Hence, it can be seen that Bob owns at least 1 bitcoin, and Charlie owns at least 4.

B. Transaction Anonymization Services

Although Figure 1 may suggest that Bitcoin addresses can be identified with actual individuals, it is in general not so easy. Any user may create as many addresses as he wants. Thus, situations such as Bob using three addresses are pervasive. Without further information, nobody can link the three addresses Bob₁, Bob₂, and Bob₃ to his real identity. At first sight, they may even belong to many instead of one person. Identifying the people behind addresses can nevertheless be possible.

Consider again the example of Figure 1 and assume Alice is a Bitcoin exchange requiring its business partners to identify themselves. Hence, she knows that address Bob₁ belongs to

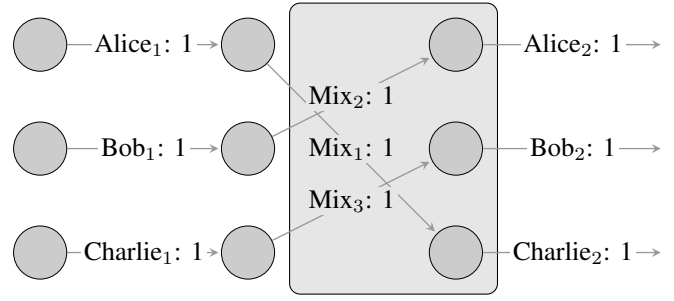


Figure 2. Block diagram of a hypothetical Bitcoin mixing service

Bob, as she transferred 3 bitcoins to it. The fact that outputs belonging to both Bob₁ and Bob₂ are referenced as inputs in the same transaction t_C could be interpreted as evidence that Bob₁ and Bob₂ belong to the same person. Hence, Alice can say with high certainty that also Bob₂ is owned by Bob. Moreover, the fact that a transaction usually has two outputs – the actual output and the change – suggests that one of t_C 's outputs belongs to the same person who owns Bob₁ and Bob₂. As bitcoins from two addresses have been combined to finance a larger output of 4 bitcoins, chances are good that the small output of only 1 bitcoin is the change. Hence, Alice may conclude that also Bob₃ belongs to Bob.

This simple example demonstrates that context information may be useful to reason about identities behind Bitcoin addresses. More comprehensive attempts to identify users have been undertaken in [12], [13]. Hence, it is wrong to state that Bitcoin is an anonymous digital currency. As a consequence, services offering more anonymous transactions are proliferating in the Bitcoin ecosystem. They are often referred to as mixing services.

The idea of such a service can be sketched as follows (cf. Figure 2). Suppose Alice, Bob, and Charlie all have 1 bitcoin at addresses Alice₁, Bob₁, and Charlie₁. They all fear they have been identified and prefer to use a mixing service. Each of them generates a new address (Alice₂, Bob₂, and Charlie₂), sends the bitcoin to an address of the service, and asks the service to send a bitcoin to their respective new address. The service has now 3 bitcoins at three addresses: Mix₁, Mix₂, and Mix₃. Alice, Bob, and Charlie do not care which address the mixing service uses to send a bitcoin back. Hence, the service may choose them at random. As long as it keeps this information private, no external observer can know the persons behind addresses after mixing, even if the observer knew them before mixing. Given a particular person, e.g., Alice, all addresses Alice₂, Bob₂, and Charlie₂ appear equally likely.

Currently, a number of different mixing services for bitcoins exist. As they disguise the origin of bitcoins, they can be perceived as money laundering tools. Also, suggestive names such as “BitLaundry” indicate the application scenarios the service operators have in mind. The primary motivation for providing these services is making profit. Hence, they send back to their users only what has been paid in minus a fee.

This fee varies between the services and may consist of a fixed and a variable part.

In terms of the anonymity literature, and with respect to the terminology put forth by Pfizmann and Köhntopp [14], the anonymity these services try to provide can be characterized as relationship anonymity. The goal is to make sender and recipient of mixed bitcoins unlinkable. Unlinkability means that all pairs of senders and recipients must as likely be related after the attacker makes his observations than they were before the observation. The mixing service must not leak any information about their relatedness.

Note that this terminology is borrowed from the literature on mix networks designed for anonymous communication. Consequently, the items of interest in this literature are messages, not transactions. These services may provide relationship anonymity, but may also offer stronger forms of anonymity. Sender anonymity refers to a service rendering attackers unable to link senders to their messages. Conversely, recipient anonymity prevents the identification of a message's recipient. For instance, the latter can be achieved by broadcasting a message to many parties. The message must be prepared in such a way that only the intended recipient recognizes that the message is for him, e. g., by using cryptography.

Clearly, an anonymization service for Bitcoin transaction cannot provide these stronger notions of anonymity. A transaction is valid only if it is documented correctly in the public block chain, including addresses of both senders and recipients. Transactions are required to send bitcoins in and out of the mixing service. Hence, there is no hope of completely hiding a relation, e. g., between a sender and the corresponding transaction he uses to send bitcoins to a mixing service.

C. Attacking Transaction Anonymizers

In analogy to mixes for communication networks (e. g., [15]), a number of attack scenarios on transaction anonymization services for the Bitcoin system are conceivable. First and foremost, an attacker could take over the service or may even set it up. As the service has to keep logs for a certain time span in order to route the bitcoins through the system, full knowledge about all relationships between senders and recipients could be acquired. A remedy is to use not just one but multiple independent services in a row. In this scenario, a user does not have to trust a single service. A single trustworthy service among all services used would be sufficient to provide anonymity. Unfortunately, this solution comes at higher cost and risk as each service will charge a fee and the availability of the chain depends on its weakest link.

Another requirement for anonymization is a sufficient number of independent users. Anonymity means that a user is anonymous with respect to a set of users, the anonymity set [14]. That is, the user is not identifiable and could be any of the users in the anonymity set. If this set is small, the degree of anonymity will be small too. Consequently, if the number of users is very small, they may either have to wait long until sufficiently many users have been found, or accept low degrees of anonymity. An attacker could also try to make heavy use

of an anonymization service, possibly by using multiple fake identities [16]. While such a sybil attack would be costly, it could give other users a false sense of anonymity.

Mixes for messages must ensure that the timing of incoming and outgoing messages does not reveal their relationships. If the mixing service would send all messages to their recipients in the same order the senders have provided them to the mix, the relationships between senders and recipients would be obvious. One way of avoiding this is to delay message forwarding by random amounts of time. Transaction anonymization services will have to provide similar features to ensure that the timing of incoming and outgoing transactions provides as little information as possible about relations between transactions.

Yet another threat to anonymity is a transaction's value. Similar to the unique size of a message in a communication network, the transaction value in a transaction system like Bitcoin could serve as a fingerprint, revealing the origin of a transaction. For example, if an attacker monitors the addresses of a user and knows how many bitcoins he transferred into the service, he could try to search for an output transaction of equal size (minus the predictable fee) in the subsequent blocks. This is why many services advise their users not to pay out the full amount of bitcoins they previously paid in [17]. Furthermore, users are encouraged to split the outgoing transaction into multiple smaller transactions and to spread them over a period of time, making it harder for an attacker to link them together.

Another weakness of these services could be the communication between users and the service. A user must provide the service with all information regarding the bitcoins he wants to pay in and out of the service. This information includes the addresses. Hence, if the traffic can be intercepted, an attacker would get all information he needs.

In this paper, we focus only on attacks that can be done ex-post, i. e., attacks based only on publicly available information from the Bitcoin block chain. That means, we do not consider any form of attacks involving taking over mixing services, monitoring their communication infrastructure, etc.

D. Measuring Anonymity

While the literature offers a number of approaches to quantify the anonymity a mixing service can provide (e. g., [18], [19]), their adaptation for transaction systems like Bitcoin leave many open research questions. We are not aware of any rigorous model to quantify the degree of anonymity Bitcoin mixes provide.

One tool that can help to evaluate the anonymity is the taint analysis¹ of Blockchain.info. As previously mentioned, a Bitcoin mix can swap bitcoins between different users and thereby remove the link between the identity and the transactions. By following the transaction graph, the taint analysis tool analyzes which bitcoin addresses have been used in previous transactions leading to the current one and thus might be an origin of the bitcoins. The higher the taint value,

¹http://blockchain.info/de/taint/_ADDRESS_

the more likely is the connection between the transaction and an address. We can use this metric to evaluate the anonymity by searching for addresses related to our input transactions. If we find a match, this means that a service did not successfully unlink the bitcoins from our identity and thus does not reliably increase our anonymity.

Similar to taint analysis, we try to identify direct connections in the transaction graph. However, the transaction graph can only reveal connections between transactions, not between addresses. The better a mixing service obfuscates the relation between input and output transaction the more anonymity it provides. An ideal transaction anonymizer would make it impossible for us to find any connection.

While the taint analysis tool aims at measuring the “correlation” between two addresses, there is another notion of *taint* in the Bitcoin community which refers to the percentage of bitcoins, that come from a known theft or scam and have been blacklisted by popular exchange markets. For example, in 2012 the bitcoin exchange Mt.Gox froze accounts of customers, who owned bitcoins that could be directly related to such an incident [20].

III. METHOD

A. Procedure

The mixing services evaluated in this study are either directly accessible via the Internet or require a connection through the Tor² network, a popular anonymous communication system based on a peer-to-peer mix network. Depending on the functionality offered, users either have to create an account or interact with the service via a web interface, where they fill in all necessary information and receive a (sometimes single-use) address to send bitcoins to. The account-based websites typically work like a virtual wallet, allowing to deposit and withdraw bitcoins. In order to analyze the services, we pay in small amounts of bitcoins that can be payed out once the input transaction is confirmed. Parameters that can be specified include the amount of bitcoins to withdraw, one or more destination addresses, the number of output transactions, and a time period over which the transactions should be spread. Each process of mixing bitcoins in one of the services will be referred to as an experiment.

For each experiment, we use one or multiple newly generated destination addresses belonging to our own private Bitcoin wallet. Once we receive the payment, we gather the relevant block chain data using the API of Blockchain.info³. We reconstruct the transaction graph by following the inputs of the outgoing transaction, as described in Section II-A, and visualize it using the open source software Gephi⁴. Inspecting the transaction graph, we try to understand how the service works and to identify patterns or special characteristics. Furthermore, we try to find direct connections between the input and output transactions using both a local search as well as the taint analysis tool presented in Section II-D.

²<https://www.torproject.org/>

³http://blockchain.info/api/blockchain_api

⁴<https://gephi.org/>

B. Services

As of July 2013, we are aware of the following transaction anonymizers offering their services in the Bitcoin ecosystem:

- OnionBC⁵ is an online Bitcoin wallet accessible via Tor only. It offers the functionality to send transactions anonymously, for which it takes a fee of 3%⁶ with a minimum transaction size of 0.5 BTC (48 USD on 2013-07-12)⁷. Furthermore, it offers an escrow service which can be used to delay Bitcoin payments for goods bought online until the goods have been delivered.
- Bitcoin Fog⁸ is another service only accessible via Tor. It allows generating up to 5 addresses for depositing bitcoins and takes a (random) fee between 1–3% of the transaction value. Bitcoins can be withdrawn to a maximum of 20 addresses, spread over a timespan of 6–96 hours with a minimum total of 0.2 BTC.
- BitLaundry⁹ is a simple mixing service, that, in contrast to OnionBC and Bitcoin Fog, does not allow to deposit bitcoins into a virtual wallet. Instead, the destination addresses, the number of outgoing transactions, and a time span have to be specified. A single-use address is generated to which the user must send at least 0.25 BTC. The mixing fee for BitLaundry is split into two parts. The first is 2.49% of the total, the second is 0.00249 BTC per outgoing transaction.
- Blockchain.info offers a service called Send Shared¹⁰ that uses a shared wallet to swap bitcoins between different users. It takes a mixing fee of 0.5%, making it the cheapest service in this comparison. Its minimum transaction size is 0.2 BTC.
- On 13 April 2013, the Bitcoin forum user BlindMixrDR announced a mixing service¹¹ that would combine Bitcoin with a blind signature scheme. Unfortunately, the service and detailed information about the system are not available anymore.

We have selected three services for our analysis: Bitcoin Fog, BitLaundry and the Send Shared functionality of Blockchain.info. Table I displays a comparison of key features.

We exclude OnionBC from the analysis due to concerns regarding the trustworthiness of the service. Its minimum deposit requirement of 0.5 BTC is rather high, and as we were not able to find any positive reviews of it on the Bitcoin boards, we decided to avoid the risk of falling for a scam.

IV. RESULTS

We report the results of our experiments for each tested service and then summarize our findings in Section IV-D.

⁵<http://6fgd4t0gcynxylb.onion>

⁶While the frontpage states a fee of 2%, the transaction view says 3%.

⁷For the exchange rate, we use the weighted market price of Mt.Gox from <http://bitcoincharts.com>.

⁸<http://bitcoinfo.com>

⁹<http://app.bitlaundry.com>

¹⁰<https://blockchain.info/de/wallet/send-shared>

¹¹<https://bitcointalk.org/index.php?topic=175959.0>

Table I
OVERVIEW OF MIXING SERVICES ANALYZED IN THIS STUDY

Service	Input		Output				
	No. of addresses	Online wallet	Fee	Multiple transactions	Time span	Minimum transaction size	
Bitcoin Fog	5 per Account	yes	1–3%	1–20	6–96 h	0.20 BTC	
BitLaundry	1 per Tx	no	2.49% + 0.00249 per Tx	1–10 per day	1–10 days	0.25 BTC	
Blockchain.info	unlimited	yes	0.5%	no	no	0.20 BTC	

A. Blockchain.info Send Shared

Experiment: The Send Shared functionality of Blockchain.info offers – in contrast to Bitcoin Fog and BitLaundry – neither the option to split the transaction into multiple smaller ones nor allows to spread payments over time. Users can work around this limitation by manually splitting a single transaction into multiple.

We send 0.40012345 BTC into our online wallet and, 6 minutes later, use the shared wallet feature to send them to another address. As we cannot detect any special patterns in the transaction graph, we create eleven additional transactions in order to increase the chance of, for example, receiving multiple coins from the same address.

We are not able to find any direct connections between the input and output transactions. However, instead of twelve there are only eight separate graphs (cf. Figure 3), meaning that there are connections between multiple outputs. Furthermore, there exist hubs where a large number of transactions are bundled into one transaction. We only find a few coinbase transactions, which indicates that mainly the coins of other users are used. The transactions that are connected to multiple output transactions suggest that transactions are bundled into larger ones and then split again for payouts. One example is shown in Figure 4, where red nodes represent output transactions and green nodes represent transactions connected to multiple output transactions. Following the left green transaction, we find an address¹² that bundles transactions to a total size of 2,000 BTC (247,640 USD on 2013-05-18), which is then split into eight transactions with a size of around 250 BTC each, and then again into smaller transactions.

Results: Although our input has been used by the service, it is not possible to find any direct connections between the input and output transactions. The service bundles a large number of small transactions into larger ones, which are then split again, making it difficult to infer the bitcoins’ origin.

B. Bitcoin Fog

Experiment: After creating an account for the service Bitcoin Fog, we obtain a newly generated address for deposits. As a first attempt, we send 0.3 BTC (43 USD on 2013-04-29) to this address (cf. Table III). As of 28 June 2013, almost two months later, these bitcoins have not been moved, yet. After the deposit is confirmed by Bitcoin Fog, we withdraw the whole

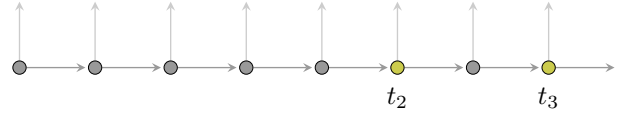


Figure 5. Chain of payout transactions characteristic for Bitcoin Fog

amount using three destination addresses, of which only two receive a transaction later on (i. e., one address we provided the service with has not been used).

We can now analyze the transaction graph of the outgoing transactions. Building the graph reveals an interesting pattern: both transactions t_2 and t_3 have only one large input transaction with a size of about 474 BTC (66,298 USD on 2013-04-30). The time difference between the transactions is only 15 minutes, and as the graph in Figure 5 shows, there is only one transaction between them.

We extend this graph, trying to identify the origin of the large transaction. After 1445 single-input transactions, there is a transaction¹⁴ that took place on 20 April 2013 and combines five big transactions with a total of 6,013 BTC (745,833 USD on 2013-04-20). Following these 5 transactions and using a community detection algorithm [21], we can identify five big communities in which a large number of transactions are bundled into one (cf. Figure 6). On the right side, they are connected by a few single-input chains. Probably, these chains are used to pay out bitcoins to other users. We cut off the graph at the edges of the communities.

In one of the communities we find a transaction¹⁵ with a size of 44,039 BTC. The coins origin from an even larger transaction¹⁶ that bundles a large number of inputs to a total of 50,000 BTC (613,500 USD on 2012-09-21). While we do not know with certainty if these belong to the same service, the transactions show the same pattern of a long single-input chain paying small amounts to many different Bitcoin addresses.

We take a closer look at the first chain of single-input transactions. By comparing the size of a transaction with the size of the previous one, we calculate the amount of bitcoins that has been payed out in each transaction. The minimum payout amounts to 0.04239 BTC (6 USD on 2013-04-25), the

¹⁴Unique transaction ID: e315f8c1cb7a85762d07511d41c7e621bcd83000185ed51443a9a72370346667

¹⁵Unique transaction ID: 5fe155fd1b72eb8acca41cc03bf6abc13083c990613907c8a8bc15bd750d1ba3

¹⁶Unique transaction ID: 443d8f0511ec1f77132b06c739b6b6f29f008dc58a373fa511ab1b182390c4fe

¹²Bitcoin address: 13udyfBcdA2PUDCFM69VYDEHRRFnqjEkx

¹³We withdrew a partial amount, about 0.1 BTC remain in the online wallet.

Table II
TRANSACTIONS RECORDED IN THE BLOCKCHAIN.INFO SEND SHARED EXPERIMENT

	Time	Type	Value	Unique transaction ID (hash value)
t_{17}	2013-05-27 16:09	In	0.40012345	c8536ce1809f296d9ed82c37a406a5cb01b63c780aa5b76324a2f26c1a7063cd
t_{18}	2013-05-27 16:15	Out	0.39713345	7fa8bf0c9c346a3e1b57ce15409473693427411729ac5664487ce6f811016517
t_{19}	2013-05-27 16:18	In	0.21212121	e72bf981bdf893a0acf55f9c54cab361c476a2bdf131d5127cc03ce105e79702
t_{20}	2013-05-28 15:55	In	0.41	10ce8832084bb1625d180d71eafc79cdea46c24dd647e44e2a50c9309182892d
t_{21}	2013-05-28 16:15	Out	0.2	c70237e203a5d3d70d1b92ced9253240810228e7b947ac73afc4e75ab34393e1
t_{22}	2013-05-28 16:17	Out	0.2	6c4c0a974999c0f83fc2f4a581da223d3cc26f7b2eacccc85ebcf5a302e18f90
t_{23}	2013-05-28 16:19	Out	0.2	b45d9a2a45c9985a9e1236aaff70d6865c562c2d7184303ebadb4303c8246d2c
t_{24}	2013-05-28 20:02	In	0.63	c2319a47c5811aaa00575343030e80b31fa482f243b297a650dfc8b12b6b660e
t_{25}	2013-05-28 20:05	Out	0.21	a3b0226c4fb44bbf0829c0be13b4dd4613daa517dd0c3616c651c04a3c06f43b
t_{26}	2013-05-28 20:08	Out	0.21	f5c3c844d9c1b7f48c45826059df7608af532d3528e05b60d9fd28c2aca3b78e
t_{27}	2013-05-28 20:13	Out	0.21992121	aab4d3d66f4a08c713e71becdd3c28cf9bf8fb34a29bf5f8d96dceb26bdecebe5
t_{28}	2013-05-28 20:52	In	0.5	1fca72c0fe447c35a5db1cc6381cc9fde7439847354b01de773053e413ae9404
t_{29}	2013-05-28 20:55	Out	0.204191	d0cf1c9fdcd2e4ac3e0421e8bd5f81ce85aled1e7ebc6cb78980e4c0b52b9e4b
t_{30}	2013-05-28 20:57	Out	0.203799	985bd5a528e2992820f5a5a1b64d537b518e29dabd40651662e5fbc09b8caf49
t_{31}	2013-05-28 21:07	In	0.6	b12e7bb024ab1a98dfe27375eb4b378cbb5e316751cee7faf5cc2c70cd5b738a
t_{32}	2013-05-28 21:13	Out	0.2110955	4de6e9651f3801bfa110dce3e1c3d01c129dcfc87ad098909e508122014fc18f
t_{33}	2013-05-28 21:15	Out	0.21336685	c2bd5ab1a52621684150ad3d4d087c131d9bbbd17d38d0db523da85ab5406bb2
t_{34}	2013-05-28 21:30	Out	0.25707765	e490ad336994f2c570a5d28edc85c80316ed00f4d8cfd0a99a86a5a224ba127a




Figure 3. Partial transaction graphs of the Blockchain.info Send Shared experiment

maximum to 717.94096 BTC (94948 USD on 2013-04-26). The average payout size is 3.8328 BTC (548 USD on 2013-04-25) with a standard deviation of 24.5344 BTC (3510 USD on 2013-04-25). The distribution of the payout sizes is shown in the left part of Figure 7. Most transactions have a size between 0.1 and 5 BTC, with a median of 0.80111 BTC. The large difference between median and mean can be ascribed to a few large transactions. As the anonymity set for large transactions is small, it can be easier to trace those.

A week after the first experiment, we make a second deposit of 0.31 BTC (33 USD on 2013-05-07). This time we withdraw 0.21 BTC (23 USD on 2013-05-15), spread over two transactions and two days. Again, we create the transaction graph of the inbound transactions and see a long chain of single input transactions. It originates from a transaction¹⁷ that, similar to the communities in the first experiment, combines

¹⁷Unique transaction ID: d7cfafaba42d952fee3ec4617f07d40808bc52fd14e507cd7fb2e0e168d40635

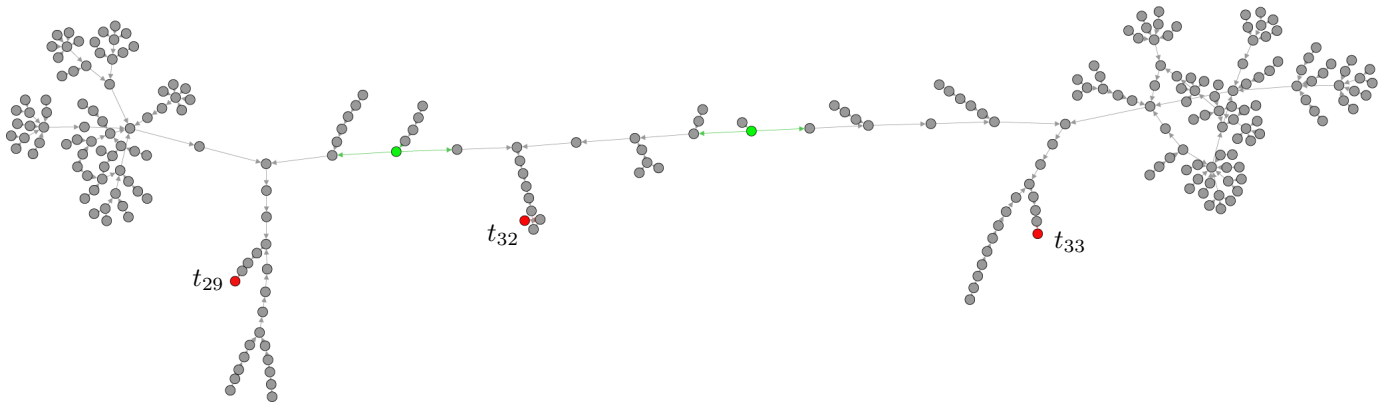


Figure 4. Transactions in the Blockchain.info Send Shared graph

Table III
TRANSACTIONS RECORDED IN THE BITCOIN FOG EXPERIMENT

	Time	Type	Value	Unique transaction ID (hash value)
t_1	2013-04-29 07:23	In	0.3	97e723ded27cd1e4f9954689c503d092fe5a1b79747d6c45b18ad8f90bf61c62
t_2	2013-04-30 08:45	Out	0.2052473	56a4f35b4a2fb5eb15549befdb1285e831a5dd67bc1b559c1b2ef8e145627856
t_3	2013-04-30 09:00	Out	0.08804699	8f4bf3e95c00025d42fc2c6a9f28e66c7ed75eb08560b7675c712accb1d75b2c
t_4	2013-05-07 20:13	In	0.3141593	ac8d82b3c3088a633fc4b48562e8c5794f502acbfbec360b406958e0acc92451
t_5	2013-05-14 08:36	Out	0.1104155	18ee1ea93a9c84dd5f1e7bd758410368e545a45a989aafcd78584f51c3da4566
t_6	2013-05-15 20:22	Out	0.1019295 ¹³	a95e2fea5498dae5ec3419d8d5c62dea23b09d69923eb15e829a562a6975a962

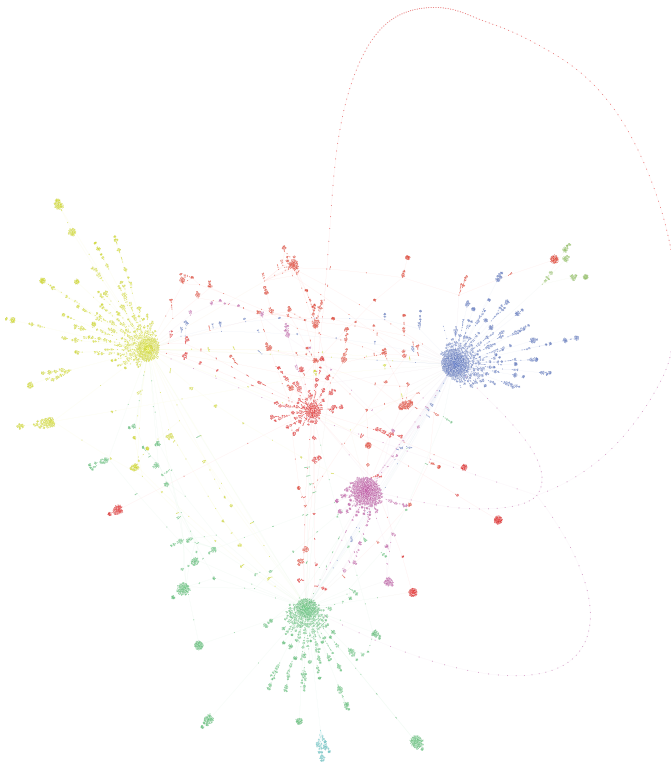


Figure 6. Communities in the first Bitcoin Fog transaction graph

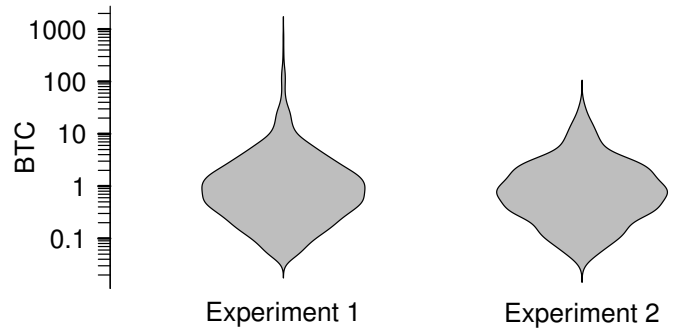


Figure 7. Kernel density estimation of Bitcoin Fog payout sizes

Figure 8, there are 30 coinbase transactions with a total value of 683 BTC. If we increase the depth of the graph, this size increases. However we cannot determine whether they belong to the service or not.

The transaction sizes range between 0.04 and 36.83 BTC (cf. right side of Figure 7), with an average transaction size of 1.89 BTC and a standard deviation of 3.72. Again, the median of 0.745 BTC is lower than the average due to some large output transactions. Similar to the first experiment, our input transaction has not been spent yet, making it impossible to find connections in the transaction graph.

Results: The service Bitcoin Fog bundles a large number of small transactions into a small number of large transactions, which are then used to create all outgoing transactions. The input transactions, however, remain untouched for a long

multiple transactions into one, with a total value of 942.88 BTC (110,336 USD on 2013-05-13). In the graph, shown in

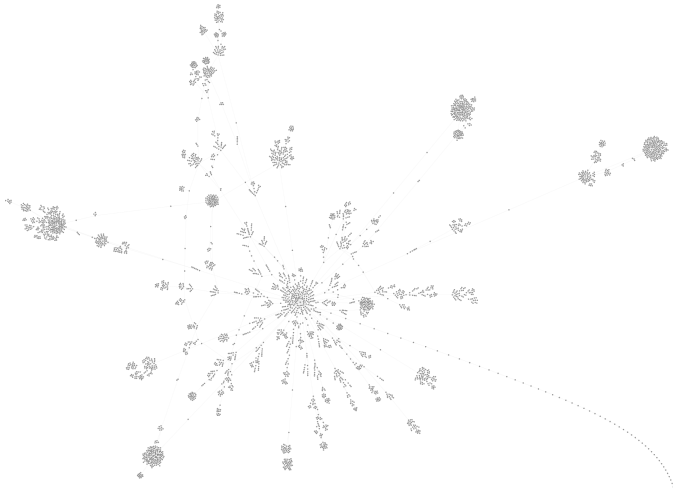


Figure 8. Transaction graph of the second Bitcoin Fog experiment

time. This way, the service prevents us from detecting any direct connections between the input and output transaction in the transaction graph. Even though the clear structure of the service might leave a possibility to decrease the anonymity of transactions using additional context information. This finding enables us to estimate the empirical distribution of the size of outgoing transactions. Cautious users can use it to choose a transaction size that fits the distribution, for example by random sampling, in the hope to defeat attackers who try to link transactions by matching amounts. Note that this defense can be difficult or impossible under budget constraints, i. e., if the user has only a small amount of bitcoins available.

C. BitLaundry

Experiment: The last service analyzed is BitLaundry. On 13 May 2013 we deposit 0.33 BTC (39 USD on 2013-05-13) in order to be transferred to a single address, split up into two transactions over a period of 2 days. Instead of two, we receive four transactions (cf. Table IV). The first observation is that the payouts seem to take place at 10:45 p.m. and 12:15 a.m., which suggests the service does batch processing at specific times.

Figure 9 shows the transaction graph visualizing the flow of the incoming transactions. The transactions are colored as follows: t_8 = red, t_9 = yellow, t_{10} = green, t_{11} = blue. In contrast to our experiment with Bitcoin Fog, we also find our deposit transaction t_7 in the graph (colored black). All five transactions are connected to each other.

A large part of the input transaction is forwarded to an address¹⁸ that received and sent about 18.45 BTC (2415 USD on 2013-05-25) over a timespan of 14 days. A small amount of 0.0244 BTC, 7.79% of the total amount, goes directly into t_{11} , which means that there is a direct connection between our input transaction and one of the output transactions in the transaction graph. The taint analysis of Blockchain.info also reports a strictly positive taint value.

¹⁸Bitcoin address: 1KdPv6GWpg6eoj6cxcV65uc1NwufvhtGGQ

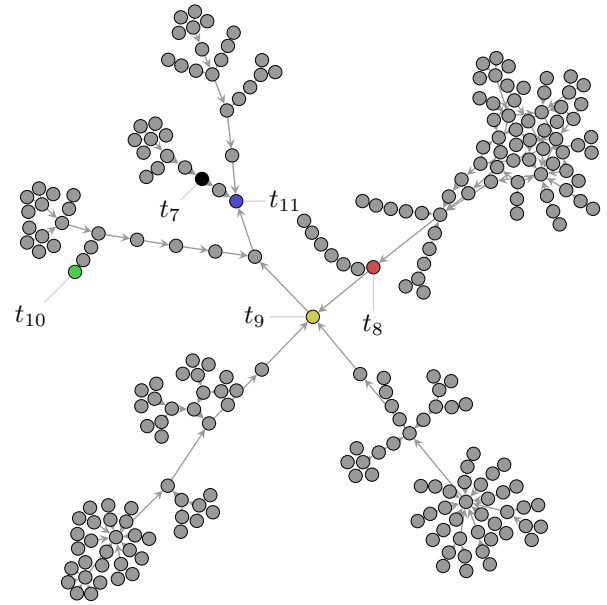


Figure 9. Transaction graph of the first BitLaundry experiment

Several hubs combining multiple transactions can be found in the graph. Initially, we observe only one coinbase transaction. If we increase the the depth of the graph, the number of coinbase transactions increases. However, we do not know whether they belong to the service. And we conjecture that the service cannot draw on a large number of coinbase transactions, which would offer perfect transaction anonymity.

We conduct two more experiments with BitLaundry in order to find out if, for example, low usage of the service might lead to another connection in the transaction graph. Therefore, we first create a transaction with a size of 0.31415 BTC, to be paid back in one transaction within one day. This experiment's transaction graph does not reveal direct connections between input and output.

After that, we pay 0.332211 BTC into the service, in order to be transferred within one day, spread over two transactions. This time, we do find a direct connection between the input t_{14} and the first output t_{15} in the transaction graph (cf. Figure 10), which results in a high taint value. The second output t_{16} is not connected to our input.

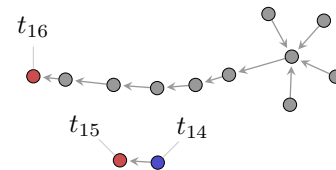


Figure 10. Transaction graph of the third BitLaundry experiment

Results: In the first experiment, we were able to find a connection between one output and our input. Although the direct link makes up only a small part of the transaction size, it is evidence for imperfect anonymization. The second experiment did not reveal any connections. However, in the last

Table IV
TRANSACTIONS RECORDED IN THE BITLAUNDRY EXPERIMENT

	Time	Type	Value	Unique transaction ID (hash value)
t_7	2013-05-13 20:04	In	0.33158651	3f574ac9026d265250fb987468346dc84a339d6ae3741356940aed723579aab5
t_8	2013-05-13 22:45	Out	0.09387001	50b78013f4e5a7acea29e721179e9ead6742bc9c9993b41d26c95fd13591f210
t_9	2013-05-14 12:09	Out	0.0818	bbb6320539a61abfde853e1ee684ec9430d19fa40926b1abe27a22bcfa7daf16
t_{10}	2013-05-14 22:43	Out	0.0782	529f930f65001a6ef519c54c7c5ad463db864cce5656fdd706ab4c5d91792845
t_{11}	2013-05-15 12:22	Out	0.0595	5fcf3ea2565672a65a389de99653a9672fa06a0c7ad90c17231bd354d2422767
t_{12}	2013-06-22 20:45	In	0.31415	9809ab21a659724b1c52cdd22427c83420f486df3935f17b0c1e3c0a1fc7b38a
t_{13}	2013-06-23 01:18	Out	0.30383767	2f917d2a38e68b99d87c47b8a78db1c8f4d7310840c23c9f9e84239dabae8cdd
t_{14}	2013-06-24 15:56	In	0.332211	6078f4779354d3cd8902be6703c0f5bb2b13417f43c60e62ed0f6375acd66a09
t_{15}	2013-06-25 00:56	Out	0.16055895	06e5b3c0d5e3be98abd8f1cd18fc91370f6e9161e4085184f11680d35ffd8af8
t_{16}	2013-06-25 16:26	Out	0.1584	238e5c60fbb09a92f8c4b6e0c94ca03658f6177ca6d50a68468aba5c4453f35d

experiment the service directly used half of the input transaction to create an output transaction. Although our sample is not very large, it suggests that this service does not provide very good anonymity. A reason for this could be a low usage of the service as well as a lack of technical measures to ensure that users do not receive their input coins back.

D. Result Summary

We can conclude that both Bitcoin Fog and Blockchain.info make it hard for an attacker to relate input and output transaction. In our analysis of Bitcoin Fog, we found a clear structure which allowed us to understand how the service works. This might make it easier for an attacker who has additional information available to detect outgoing transactions. We were not able to find any direct connections in the transaction graph of Blockchain.info. As the service provider also offers a popular online Bitcoin wallet, in which the Send Shared functionality is integrated, the large user base probably adds to the anonymity of the mixing service. In the analysis of the service BitLaundry, we found direct connections in the transaction graph in two of our three experiments. In the last one, we directly received half of the coins we paid in. Thus, BitLaundry cannot be considered to reliably increase anonymity.

On a more general note, two things are remarkable. First, some Bitcoin anonymizing services seem to deal with values in the order of 6-digit dollar amounts. This is substantial and indicates a viable – though questionable – business model as well as ample demand for transaction anonymization. Second, a particularity of transaction anonymization over conventional mixes is the possibility to keep incoming transactions untouched for a very long time (provided the service has sufficient free capital), thereby stretching the anonymity set over time. This is a clear advantage of anonymous transaction systems over anonymous communication systems where the tolerable message latency imposes an upper bound on the batch size.

E. Combinations and Costs

All three services pose the risk that the operator itself is an attacker or colludes with one. In order to reduce this risk, it would be possible to combine multiple services in a cascade. However, this comes at the cost of additional delay and higher fees. We calculate the fees for using Bitcoin Fog, BitLaundry

and Send Shared to anonymize a transaction. The output O can be calculated by multiplying the input I with the fees of the single services, minus the number of outgoing transactions m of BitLaundry, minus our initial transaction cost. In this case we end up with a total cost of about 5% for using these three services.

$$O = I * 0.98 * 0.995 * 0.9751 - m * 0.00249 - 0.0005 \quad (1)$$

The risk-adjusted costs are even higher, because the risk that one of the services goes bankrupt (e. g., it gets hacked and all coins are stolen) or offline (stealing all coins that are in the system) as well as the risk that transactions are not included in the block chain have to be priced in as well. There have been multiple incidents where large services were accused of stealing their user’s bitcoins, e. g., the shutdown of MyBitcoin in July 2011 [22] resulted in a loss of at around 78,740 BTC (1,063,777 USD on 2011-07-29) [22] and the scam of Bitcoin Savings & Trust, which turned out to be a pyramid scheme [5], in a loss of at least 200,000 BTC (2,328,000 USD on 2012-08-12). A detailed list of major thefts and scams can be found at [22].

Our expenses for conducting this experiment have been rather small. Altogether we spent around 0.08 BTC (8 USD on 2013-07-12) on service and transaction fees.

V. RELATED WORK

Bitcoin is not the first cryptographic cash system. In 1985, Chaum [23] proposed cryptographic cash that allowed anonymous payments using blind signatures. Another idea for electronic currencies are credit networks, e. g., iOwe [24]. Payments are made with digital bonds that represent a pledge to deliver a certain value or good in the future. Digital credit networks heavily rely on trust models to prevent double spending and sybil attacks. The success of Bitcoin has spurred many derivatives. Zerocoin, proposed by Miers, Garman, Green, *et al.* [25], promises complete anonymity. It uses zero-knowledge proofs to deposit and withdraw special transactions with unlinkable inputs and outputs.

Several authors have studied the anonymity of Bitcoin. Ron and Shamir [26] analyze statistical properties of the Bitcoin transaction graph. Reid and Harrigan [12] analyze the network of Bitcoin users by combining addresses that are inputs of

multi-input transactions and therefore must belong to the same sender. They use publicly available data, like forum posts including Bitcoin addresses, to identify users. Ober, Katzenbeisser, and Hamacher [27] analyze structural aspects of the transaction graph and draw implications on the anonymity of transactions. Androulaki, Karame, Roeschlin, *et al.* [28] study privacy implications of multi-input transactions and shadow addresses generated by the Bitcoin client for receiving change. They could identify 40% of the users in an artificial transaction graph based on (simulated) behavior. Meiklejohn, Pomarole, Jordan, *et al.* [13] identify a large number of intermediaries by interacting with them and using a change address heuristic to identify addresses of the same user. Using this dataset they analyze popular thefts and are able to relate payouts to popular exchanges. To the best of our knowledge, this paper documents the first study of Bitcoin transaction anonymizers.

The idea of establishing anonymity by mixing messages of multiple users is due to Chaum [11]. A popular anonymous communication system is The Onion Router (Tor), which anonymizes applications and users communication on the TCP-layer [29]. Other relevant systems include AN.ON/JonDo [30]. Attacks on mix networks are often performed using context or linkability information [31]–[34].

VI. DISCUSSION AND CONCLUSION

Our starting point for this paper was to recall Bitcoin's principle of pseudonymous accounts. Contrary to common belief, it permits certain AML measures by imposing regulation (such as the KYC principle) on the intermediaries who offer financial or real services in exchange for bitcoins. However, this approach is thwarted by the existence of transaction anonymizers. Intermediaries of this special kind operate within the Bitcoin system and are therefore hard to locate and presumably even harder to regulate. The core contribution of this study is a systematic analysis of three popular transaction anonymizers based on the transaction graph extracted from the block chain. Although our deanonymization attempts were not exactly the strongest conceivable (see Section VI-A below), the results lend support to the notion that budget-constrained cybercrime fighters are effectively set back by two of the three tested services. In the remainder of this discussion, we will point out ways forward along three dimensions: more powerful deanonymization (Section VI-A), new directions for regulation (Section VI-B), and some considerations on the viability of Bitcoin as a decentralized currency (Section VI-C).

A. Limitations of the Study

Our study uses reverse-engineering methods to understand the modes of operation of three popular transaction anonymizers. We have carried out a series of experiments with probe transactions and tried to establish relationships between inputs and outputs using public information in the transaction graph. The insights gained from these experiments are not necessarily generalizable because the analyzed services may change their mode of operation at any point in time. Moreover, due to limited number of experiments, our results are of exploratory nature

and need further substantiation with evidence from quantitative measurement studies.

Our results likely overstate the level of anonymity provided by the analyzed services. For one thing, we do not consider auxiliary information from traffic analyses that would require realtime interception of (parts of) the communication network. It is likely that such measures are available for law enforcement of serious crimes. Possible targets include the communication with the transaction anonymizer, e. g., account setup and control of payout structure (to prevent this, Bitcoin Fog requires Tor as anonymous control channel), as well as the communication in Bitcoin's peer-to-peer network. The latter might reveal IP address ranges and timing information related to the publication of suspect transactions before they are validated and included in the persistent block chain. This level of detail may help to identify devices, locations, and eventually people. In addition, we have not systematically exploited timing information and relations between Bitcoin transaction, addresses, and blocks; although this information is persistent and publicly available without access to realtime surveillance.

All this suggests that our statements on the anonymity offered by the services should be interpreted as upper bounds. This view is supported by the general warning that users must fully trust the anonymizing services regarding the confidentiality (ideally, immediate deletion) of input–output relations and the willingness to return the temporarily entrusted values. In practice, criminals might have a hard time to identify trustworthy transaction anonymizers between other criminals setting up fake services and cybercrime fighters operating their own or taking over existing services as decoys.

It remains an open research question if user-verifiable transaction anonymizers with provable anonymity and security guarantees can be built for the existing Bitcoin protocol.

B. Regulatory Options

Recall from the introduction that KYC is only a first step that enables downstream activities such as blacklisting suspicious account holders. As Bitcoin accounts have weak identities at best, but all transaction records are public, for AML to be effective, it should blacklist transaction histories (i. e., bitcoins) rather than accounts or account holders. A good offline analogy is registering serial numbers of bank notes used to pay ransoms. There is some precedence in Bitcoin. In 2012, the Mt.Gox exchange began to reject bitcoins looted in major thefts or frauds [20], [35]. Unsurprisingly, this behavior is fiercely debated in the Bitcoin community (e. g., [36]) because the blacklists are governed by powerful entities, making a dent in the idea of complete decentralization of control. At the same time, this offers a backdoor for regulation. Entities operating in the Bitcoin ecosystem could be mandated by law to observe official blacklists of transaction prefixes. Combine this with systematic undercover test purchases, which offer a relatively cheap and effective way of enforcement in each jurisdiction where a service is provisioned in exchange for bitcoins.

It is easy to see that this regulation would affect the business model of transaction anonymizers. In the long run,

coinbase transactions are rare, so every output of a transaction anonymizer has a history. Either transaction anonymizers reject blacklisted coins, the intended outcome, or all users sending clean (i. e., not blacklisted) bitcoins take the risk of getting blacklisted bitcoins in return. This leads to adverse selection: only owners of blacklisted bitcoins have nothing to lose and might keep using transaction anonymizers. Yet, if all inputs are blacklisted, so are all outputs. There remains nothing to gain from such services. Indeed, it seems that the only way to escape such regulation is to change the design towards a cryptographic currency that offers transactions without history.

C. Implications for Bitcoin at Large

Considering the long-term perspective of Bitcoin is relevant because it appears that regulators are currently facing the dilemma of finding the right level of intervention. While a tough policy against Bitcoin might indeed hamper cybercrime (or drive it elsewhere), it comes at the risk of rashly eradicating a platform that may unleash future innovation as legitimate businesses and consumers adopt (see [37] for payment innovations in general). Bitcoin's relationship with the underground has striking parallels to another financial innovation: the first automatic teller machine was set out of operation in 1939 after just six months in service because it attracted clients who did not want to look into a clerk's eyes – gamblers and prostitutes. It took until 1969 for the second attempt that thrived [38, p. 333]. However, if the potential of Bitcoin for legitimate purposes is foreseeably limited from the outset, the opportunity costs of a tough policy are low.

Now turning to the implications: Bitcoin set out to be a decentralized currency. Unlike earlier and commercially unsuccessful proposals for electronic cash [23], [39], [40], calling Bitcoin a “currency” can be justified because the system enforces an upper bound on money creation, which makes bitcoins scarce and therefore valuable. More generally, economists define money by three functions: medium of exchange, store of value, and unit of account. Critics of Bitcoin debate the store of value function on grounds of the high volatility observable on Bitcoin exchanges and the built-in deflationary tendency. The latter may impair the medium of exchange function, because investors in bitcoin have incentives to hoard rather than spend [41].

Our observation calls into question the last function standing. Bitcoins are not alike. Every transaction has a different history. The fact that transaction anonymizers exchange the history for a fee implies that bitcoins with different histories have different value. Coinbase transactions, for example, should be valued highest because they are scarcest. They are an important resource required as input to make transactions provably untraceable. This effect is aggravated under the above-described policy of blacklisting transactions for AML, where “virgin” coinbase transactions have the lowest risk of being blacklisted at the time of conversion or spending. By contrast, bitcoins with known blacklisted transaction prefixes should have very little value, as they can only be used in the underground. One can complete this example by speculating

that collectors might ascribe higher value to bitcoins with a famous transaction history. For enthusiasts, a bitcoin involved in the 10,000 BTC pizza purchase in 2010 might be as valuable to own as, say, a dollar note provably spent by James Dean to buy his Porsche. We believe that it is just a matter of time until price spreads between bitcoins of different provenance appear in the marketplace. The bottom line is that the uniqueness of every bitcoin thwarts the very idea of money as a homogeneous commodity to serve as *unit* of account. Inventors of future cryptographic currencies should take note.

ACKNOWLEDGEMENTS

The authors are grateful to Raimo Radczewski for insightful discussions and to the APWG for a generous travel stipend.

REFERENCES

- [1] US Department of Justice, *Manhattan U.S. attorney announces charges against Liberty Reserve*, 2013. [Online]. Available: <http://www.justice.gov/usao/nys/pressreleases/May13/LibertyReservePR.php> (visited on 05/28/2013).
- [2] S. Nakamoto, *Bitcoin: A Peer-to-Peer Electronic Cash System*, 2008.
- [3] N. Christin, “Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace”, in *Proceedings of the 22nd International World Wide Web Conference*, Rio de Janeiro, 2013, pp. 213–224.
- [4] M. Geuss, *Bitcoinica users sue for USD 460k in lost Bitcoins*, 2012. [Online]. Available: <http://arstechnica.com/tech-policy/2012/08/bitcoinica-users-sue-for-460k-in-lost-bitcoins/> (visited on 07/12/2013).
- [5] J. Mick, *Pirateat40 Makes Off 5.6M USD in BitCoins From Pyramid Scheme*, 2012. [Online]. Available: <http://www.dailytech.com/pirateat40/article25538.htm> (visited on 07/05/2013).
- [6] M. Santora, W. K. Rashbaum, and N. Perlroth, *Online currency exchange accused of laundering \$6 billion*, May 2013. [Online]. Available: http://www.nytimes.com/2013/05/29/nyregion/liberty-reserve-operators-accused-of-money-laundering.html?pagewanted=all&_r=0.
- [7] O. Solon, *A simple guide to Bitcoin*, 2013. [Online]. Available: <http://www.wired.co.uk/news/archive/2013-05/7/bitcoin-101> (visited on 05/23/2013).
- [8] *Donate to WikiLeaks*. [Online]. Available: <http://shop.wikileaks.org/donate> (visited on 05/30/2013).
- [9] Bitcoin Wiki, *Anonymity*. [Online]. Available: <https://en.bitcoin.it/wiki/Anonymity> (visited on 05/23/2013).
- [10] Mt.Gox, *Statement Regarding Account Verifications*, 2013. [Online]. Available: https://mtgox.com/press_release_20130530.html (visited on 05/31/2013).
- [11] D. Chaum, “Untraceable electronic mail, return addresses, and digital pseudonyms”, *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, Feb. 1981.

- [12] F. Reid and M. Harrigan, “An Analysis of Anonymity in the Bitcoin System”, in *Security and Privacy in Social Networks*, Y. Altshuler, Y. Elovici, A. Cremers, N. Aharony, and A. Pentland, Eds., New York: Springer, 2013, pp. 197–223.
- [13] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage, “A Fistful of Bitcoins: Characterizing Payments Among Men with No Names”, 2013.
- [14] A. Pfitzmann and M. Köhntopp, “Anonymity, unobservability, and pseudonymity—a proposal for terminology”, in *Designing privacy enhancing technologies*, ser. Lecture Notes in Computer Science, H. Federrath, Ed., vol. 2009, Berlin Heidelberg: Springer, 2001, pp. 1–9.
- [15] J. Raymond, “Traffic analysis: Protocols, attacks, design issues, and open problems”, in *Designing Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, H. Federrath, Ed., vol. 2009, Berlin Heidelberg: Springer, 2001, pp. 10–29.
- [16] J. R. Douceur, “The Sybil Attack”, in *Peer-to-peer Systems*, ser. Lecture Notes in Computer Science, P. Druschel, F. Kaashoek, and A. Rowstron, Eds., vol. 2429, Berlin Heidelberg: Springer, 2002, pp. 251–260.
- [17] *Bitcoin Fog*. [Online]. Available: <http://bitcoinfo.com/> (visited on 06/22/2013).
- [18] C. Diaz, S. Seys, J. Claessens, and B. Preneel, “Towards Measuring Anonymity”, in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, R. Dingleline and P. Syverson, Eds., vol. 2482, Berlin Heidelberg: Springer, 2003, pp. 54–68.
- [19] A. Serjantov and G. Danezis, “Towards an information theoretic metric for anonymity”, in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, R. Dingleline and P. Syverson, Eds., vol. 2482, Berlin Heidelberg: Springer, 2003, pp. 41–53.
- [20] *Mt.Gox thinks it's the Fed. Freezes acc based on "tainted" coins. (unlocked now)*, 2012. [Online]. Available: <https://bitcointalk.org/index.php?topic=73385.0> (visited on 07/12/2013).
- [21] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, “Fast Unfolding of Communities in Large Networks”, *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2008, no. 10, P10008, 2008.
- [22] *List of Major Bitcoin Heists, Thefts, Hacks, Scams, and Losses*. [Online]. Available: <https://bitcointalk.org/index.php?topic=83794> (visited on 07/05/2013).
- [23] D. Chaum, “Security Without Identification: Transaction Systems To Make Big Brother Obsolete”, *Communications of the ACM*, vol. 28, no. 10, pp. 1030–1044, 1985.
- [24] D. Levin, A. Schulman, K. LaCurts, N. Spring, and B. Bhattacharjee, “Making Currency Inexpensive with iOwe”, in *Proceedings of the Workshop on the Economics of Networks, Systems, and Computation (NetEcon)*, San Jose, 2011.
- [25] I. Miers, C. Garman, M. Green, and A. D. Rubin, “Zero-coin: Anonymous Distributed E-Cash from Bitcoin”, in *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, Berkeley, 2013, pp. 397–411.
- [26] D. Ron and A. Shamir, “Quantitative Analysis of the Full Bitcoin Transaction Graph”, in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, A.-R. Sadeghi, Ed., vol. 7859, Berlin Heidelberg: Springer, 2013, pp. 6–24.
- [27] M. Ober, S. Katzenbeisser, and K. Hamacher, “Structure and Anonymity of the Bitcoin Transaction Graph”, *Future Internet*, vol. 5, no. 2, pp. 237–250, 2013.
- [28] E. Androulaki, G. O. Karame, M. Roeschlin, T. Scherer, and S. Capkun, “Evaluating User Privacy in Bitcoin.”, in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, A.-R. Sadeghi, Ed., vol. 7859, Berlin Heidelberg: Springer, 2013, pp. 34–51.
- [29] R. Dingleline, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router”, in *Proceedings of the 13th conference on USENIX Security Symposium*, San Diego, 2004.
- [30] O. Berthold, H. Federrath, and S. Köpsell, “Web MIXes: a system for anonymous and unobservable Internet access”, in *Designing Privacy Enhancing Technologies*, H. Federrath, Ed., ser. Lecture Notes in Computer Science, vol. 2009, Berlin Heidelberg: Springer, 2000, pp. 115–129.
- [31] M. Franz, B. Meyer, and A. Pashalidis, “Attacking Unlinkability: The Importance of Context”, in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, N. Borisov and P. Golle, Eds., vol. 4776, Berlin Heidelberg: Springer, 2007, pp. 1–16.
- [32] S. Berthold, R. Böhme, and S. Köpsell, “Data retention and anonymity services”, in *The Future of Identity in the Information Society*, V. Matyáš, S. Fischer-Hübner, D. Cvrček, and P. Švenda, Eds., Brno, Czech Republic: Springer, 2009, pp. 92–106.
- [33] S. Schiffner and S. Clauß, “Using linkability information to attack mix-based anonymity services”, in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, I. Goldberg and M. J. Atallah, Eds., vol. 5672, Berlin Heidelberg: Springer, 2009, pp. 94–107.
- [34] S. Zhioua, “Anonymity attacks on mix systems: a formal analysis”, in *Information Hiding*, ser. Lecture Notes in Computer Science, T. Filler, T. Pevný, S. Craver, and A. Ker, Eds., vol. 6958, Berlin Heidelberg: Springer, 2011, pp. 133–147.
- [35] V. Buterin, *Mt.Gox: What the largest exchange is doing about the Linode theft and the implications*, 2012. [Online]. Available: <http://bitcoinmagazine.com/mtgox-the-bitcoin-police-what-the-largest-exchange-is-doing-about-the-linode-theft-and-the-implications/> (visited on 07/12/2013).
- [36] *Taint checker list*, 2012. [Online]. Available: <https://bitcointalk.org/index.php?topic=67383.0>.

- [37] R. Anderson, “Risk and Privacy Implications of Consumer Payment Innovation”, 2012.
- [38] —, *Security Engineering*, 2nd ed. Indianapolis: Wiley, 2008.
- [39] D. Chaum, A. Fiat, and M. Naor, “Untraceable electronic cash”, in *Advances in Cryptology — CRYPTO’ 88*, ser. Lecture Notes in Computer Science, S. Goldwasser, Ed., vol. 403, New York: Springer-Verlag, 1990, pp. 319–327.
- [40] *Digicash.com*. [Online]. Available: <http://digicash.com> (visited on 07/12/2013).
- [41] P. Krugman, *Golden Cyberfettters*, 2011. [Online]. Available: <http://krugman.blogs.nytimes.com/2011/09/07/golden-cyberfettters/> (visited on 07/13/2012).