Short Paper: Towards Characterizing Sybil Attacks in Cryptocurrency Mixers

Mikerah Quintyne-Collins

HashCloak Inc

Abstract. Sybil attacks [5] are a well-studied problem in peer-to-peer networking systems. However, their relevance to cryptocurrency mixers has received little attention in the literature, with only a few papers [2, 3, 7] in recent times aiming to design mixers that are resistant to Sybil attacks. A lot of the research has been primarily driven by independent cryptocurrency enthusiasts [1, 6, 16]. We attempt to provide a few characterizations of Sybil attacks as they pertain to mixers and provide mitigations based on economics in order to disincentive Sybil attacks against mixers. In doing so, we highlight that the security of mixers need not only be analyzed through the use of cryptographic techniques but also with the use of economic techniques. Moreover, we provide future research directions in determining heuristics for detecting Sybil identities in mixers.

1 Introduction

Mixers [11] are a way to increase privacy in otherwise pseudonymous cryptocurrencies like Bitcoin. This was primarily driven by work presented in [9] that showed how links between Bitcoin transactions and addresses can be used to create profiles of Bitcoin users. Mixers [11] are services that enable cryptocurrency users to mix their coins together in order to increase the fungibility of their coins. They can be provided by a centralized server or done through a peer-to-peer networking protocol. Many proposals have been introduced and some are running in production such as variations of Coinjoin [4, 13, 16] on Bitcoin, for example.

As these mixers are a form of distributed system, they are susceptible to the well-studied Sybil attack [5]. Sybil attacks are attacks where an adversary creates multiple identities in order to exert some form of influence over the network. There have been a few mixer designs that emphasize their protection against Sybil attacks [3, 7, 16], most notably Xim [2]. The methods these proposals use in order to mitigate Sybil identities are expanded upon in this work.

In this paper, we take a closer look at Sybil attacks within the context of mixers in an attempt to characterize such attacks. First, we provide some background on the core concepts behind mixers and Sybil attacks, after which we will expand on the various types of Sybil attacks that affect mixers. These attacks are theoretical in nature but there is the potential that these attacks might

occur in the future. Furthermore, we will expand on some potential mitigations that mixers can apply to make Sybil attacks more costly. Finally, we provide directions for future work for making mixers more Sybil resistant and in the detection of Sybil identities.

2 Background

2.1 Mixers

Cryptocurrency mixers are services that enable users of a cryptocurrency to jointly mix their coins in order to obfuscate the flow of their transactions [11]. Mixers provide a way to add privacy to blockchains that don't have built-in privacy capabilities without having to modify the base protocol.

Roughly, mixers can be categorized on two axes:

- **Centralized vs Decentralized Mixers:** Centralized mixers are services in which users send their cryptocurrency funds to a centralized server that then mixes all the funds for a fee and sends mixed funds to their respective users at new addresses. This type of mixer removes a lot of complexity for the mixer participants. However, there is a risk of theft by the mixing service [22]. Decentralized mixers enable users to come together in a decentralized fashion in order to mix their coins. These mixers don't require the use of a third party and don't require users to trust a centralized server. However, users need to be able to coordinate amongst themselves in order to effectively establish mixes according to a particular protocol.
- **Obfuscation-based vs Zero-Knowledge-based:** Obfuscation-based mixers [15], also known as decoy-based mixers [10] use techniques that hide a user's transaction graph. However, with enough resources, an adversary can reconstruct the transaction graph through various techniques [9, 10]. On the other hand, zero-knowledge-based mixers make heavy use of advanced cryptographic techniques such as zero knowledge proofs in order to completely eliminate the transaction graph. The main downside to this approach is that they require heavy cryptography usage that may hinder scalability.

2.2 Sybil Attacks in Peer-to-Peer Networks

Sybil attacks [5] are attacks where an adversary in a peer-to-peer, permissionless network creates and controls multiple identities in order to exert some influence in the network. This is a problem that affects many peer-to-peer protocols like BitTorrent and Tor and is known to affect cryptocurrencies as well. It is a well-studied problem in distributed systems with varied solutions [8, 12].

3 Types of Sybil Attacks

Now, we describe Sybil attacks that are of relevance to mixers. In the case of mixers, addresses will serve as the identity of a user. As addresses cost nothing to

create, users can own multiple addresses. The first attack presented, the Sybil-based de-anonymization attack, aims to link addresses and transactions that are participating in the mixer. This attack is dependent on whether the mixer is decoy/obfuscation-based or zero-knowledge-based. The second attack presented, Sybil-based DoS attack, aims to disrupt the functioning of the mixer and worsen the other participants user experience through increased fees and longer round times. This attack is dependent on whether the mixer is centralized or decentralized. In both of these attacks, the cost of taking over the anonymity set increases linearly as the size of the anonymity set grows.

3.1 Sybil-based De-anonymization attacks

In Sybil-based de-anonymization attacks also known as Sybil-based linking [2], an adversary creates multiple funded addresses on a blockchain and joins a mixer. In a decentralized mixer, the adversary would need to be paired with the same user(s) multiple times in order to build a profile of that user's transactions in an attempt to de-anonymize them. In the case of a centralized mixer, the mixer server itself can de-anonymize users, even without the need for Sybil identities.

For decoy-based mixers, Sybil-based de-anonymization attacks are a problem. Since the main way these mixers work is by having multiple decoy addresses hide other addresses [10], a Sybil attack simply needs to create as many addresses as possible in order to increase their chances of being paired with target users. For example, in a CryptoNote-like mixer [20], the Sybil attacker will need to make up the majority of the ring in which a targeted user is in.

On the other hand, for zero-knowledge-based mixers, this type of attack doesn't actively de-anonymize users in the mix. This is due to the complete elimination of the transaction graph through advanced cryptographic technique [10]. It does, however, artificially inflate the anonymity set. This attack in such a mixer would give users the false impression that the mixer is much safer than it actually is.

3.2 Sybil-based DoS attacks

Sybil-based Denial-of-Service (DoS) attacks [2] are attacks in which an adversary will create multiple identities in order to disrupt the normal functioning of the mixer. This kind of attack is mainly done at the networking layer.

For decentralized mixers, an adversary simply creates multiple identities and participates in the mixing protocol. During the pairing part of the protocol where each participant exchanges funds in an attempt to mix them, identities controlled by an adversary can easily decide to not pair up with other participants. This can incur fees for the mixer participants and increase the time to get funds mixed in the mixer.

In centralized mixers, Sybil-based DoS attacks aren't a concern because the mixer operator is in control of the mixing. This is because there is no way in which the mixer operator can create identities in order to disrupt the functioning of the mixer. In fact, it is much simply for the mixer operator to perform a

DoS attack without the need for Sybil identities. We do note, however, that centralized mixers are of more susceptible to regular DoS attacks due to the reliance on a central authority.

4 Mitigations

In [5], it was shown that the only defense against Sybil attacks is through the use of a centralized authority in charge of managing identities. However, this approach is in conflict with the goals of mixers in adding extra privacy to cryptocurrency users. Thus, to get around the use of a central authority, researchers have proposed the usage of techniques in which it is difficult for an adversary to scale through the use of external resources. For mixers in particular, these techniques can be classified broadly into 3 categories, based on solutions presented in [8]: those that build on social constraints, those that build on computational constraints and those that build on economic/monetary constraints. First, we will expand on why the first two classes of techniques are not feasible for use in mixers, after which, we will propose a few solutions that build on monetary constraints.

Social constraint-based Sybil resistance techniques [8] rely on some notion of trust between the peers in a network or in our case, users in a mixer. They typically rely on the fact that building trust between an adversary and honest users is hard to do. An adversary would have to create multiple identities, in which, each identity is trusted, according to some metric, by other honest users in the network. Sybil resistance solutions based on social constraints are inadequate for mixers because there are no existing trust relationships between mixer participants. Any attempts at defining a notion of trust between mixer participants would defeat the purpose of using a mixer as the main goal is to increase anonymity.

On the other hand, computational constraint-based Sybil resistance techniques [8] make it such that an adversary creating Sybil identities needs to spend computational resources proportional to the number of identities created. Peer-to-peer networks like Bitcoin [14] used this mechanism in order to limit the number of blocks produced through the use of Proof-of-work (PoW). This technique is usually employed when the cost of joining a network is low. Even though the cost of creating a funded address to participate in a mixer is low, there are very low computational requirements needed in order to manage these addresses. Thus, for our purposes, Sybil resistance solutions employing the use of computational constraints are not feasible.

In short, existing Sybil defenses for peer-to-peer networks are not sufficient for the needs of cryptocurrency mixers. Their unique circumstance require techniques that leverage the economic playground that a blockchain offers. We propose potential mitigations that require Sybil attackers to own multiple coins. Note that each of these mitigations can be employed independently or combined with another technique. We introduce each mitigation in turn.

4.1 Burning

Burning [17, 21] involves making a deposit that is non-refundable in order to make creating a cryptographic identity expensive to make. Practically, it involves sending some amount of cryptocurrency to an unspendable address on the blockchain. If priced accordingly, burning can make it expensive for a Sybil attacker to create various identities as long as one needs to burn tokens for each new identity. Thus, the cost of creating identities would be linear in the number of identities created. The main drawback of such an approach is that, due to pricing, it would necessitate large amounts of a currency to be burned which can be prohibitively expensive for honest users of the mixer. In the case where these honest users have mostly "clean" coins, the funds in the mixer would most likely be derived from illicit sources.

4.2 Time Locking

Time locking [23] is the process of restricting the usage of funds until a specified time in the future. Here, if the time parameter is set to be long enough, this could discourage a potential Sybil attacker. However, this mechanism is dependent on the opportunity of the adversary. If the adversary is economically motivated, there is lost opportunity cost in which they could be using the same funds for something else, instead of disrupting the mixer. On the other hand, if the adversary is simply trying to disrupt the functioning of the mixer, this technique may not be enough to deter them.

4.3 Fidelity bonds

Fidelity bonds [6] are a financial instrument in which one can sacrifice coins in order to receive a bond that is redeemable in the future. It is implemented through the use of burning and time locks. They were first introduced in [6] and their usage applied to mixers was first proposed in [6, 19]. More recently, [1] introduced fidelity bonds to combat Sybil attacks in JoinMarket. This solution is mainly of relevance to adversary that have long-term holdings of the coin. The assumption is that the adversary wouldn't want to Sybil attack a mixer which would hurt the long-term value of their holdings through making the coin less private. However, if an adversary doesn't care about the long-term value of their holdings, this mitigation may not be enough to prevent them from attacking.

4.4 Coin-age

The concept of coin-age [18] is only of relevance to UTXO-based blockchains. It refers to the age of transaction inputs where the age is measured in blocks. The coin-age of transaction inputs can be calculated as follows

$$A \times \texttt{average age of coins}$$

where A is the amount of coins.

Using the coin-age to prevent Sybil attacks in mixer involves restricting the use of coins that are under a certain coin age. The assumption is that it takes a long time for an adversary to acquire enough coins that are old enough to be used in a mixer.

There are however issues with using coin-age as a Sybil resistance method for mixers. First, as coins need to be "old" enough to be used in the mixer, this would discourage so-called "clean" coins from being used in the mixer. Clean coins are coins that are newly-minted or that do not come from illicit sources. Thus, they have a short coin-age. This is problematic because clean coins are the type of coins that one wants to have in a mixer in order to mask the not so clean coins. This would increase the proportion of coins that come from illicit sources which would defeat the purpose of using mixers.

5 Discussion and Conclusion

We re-formulated the problem of Sybil identities within the context of cryptocurrency mixers. We presented two different types of Sybil attacks that are of particular interest to mixers, one that can link users together, the other disrupting the functioning of the mixer. Moreover, we provided an analysis for why Sybil attacks cannot be mitigated with the use of social or computational constraints and provide a new categorization based on monetary constraints of the adversary.

The framing of Sybil attacks in mixers as an economics problem as opposed to a cryptography problem opens the door to how one might analyse the security of mixer proposals from an economic perspective. Our goal with this paper is to provide a stepping stone to further analyze various other economic attacks that affect mixers in order to assess the economic security of a proposed scheme. We believe that this will lead to a more well-rounded approach to analyzing the security of such systems due to the combining both economics and cryptographic techniques to provide arguments of security. Moreover, the detection of Sybil identities in mixers is another area of research worth pursuing. This is due to the fact that mixers are anonymity favoring environments and as such, the detection of Sybil identities isn't completely possible. Coming up with a set of heuristics that can aid in the detection of Sybil identities is another research direction we would like to highlight. Thus, we hope to stir more research in the field of Sybil resistance in anonymity favoring environments.

Even with the emergence of so-called privacy coins, billions of dollars of value are still secured by non-privacy favoring blockchains. The importance of mixers will be more pronounced as these cryptocurrencies gain more popularity.

References

[1] Chris Belcher. *[bitcoin-dev] Improving JoinMarket's resistance to sybil attacks using fidelity bonds*. E-mail. July 25, 2019. URL: https://lists.

linuxfoundation.org/pipermail/bitcoin-dev/2019-July/017169.html (visited on 09/13/2019).

[2] George Bissias et al. "Sybil-Resistant Mixing for Bitcoin". In: *Proceedings of the 13th Workshop on Privacy in the Electronic Society - WPES '14*. the 13th Workshop. Scottsdale, Arizona, USA: ACM Press, 2014, pp. 149–158. ISBN: 978-1-4503-3148-7. DOI: 10.1145/2665943.2665955. URL: http://dl.acm.org/citation.cfm?doid=2665943.2665955 (visited on 09/14/2019).

[3] Joseph Bonneau et al. "Mixcoin: Anonymity for Bitcoin with Accountable Mixes". In: *Financial Cryptography and Data Security*. Ed. by Nicolas Christin and Reihaneh Safavi-Naini. Vol. 8437. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 486–504. ISBN: 978-3-662-45471-8 978-3-662-45472-5. DOI: 10.1007/978-3-662-45472-5_31. URL: http://link.springer.com/10.1007/978-3-662-45472-5_31 (visited on 09/15/2019).

[4] *CoinJoin: Bitcoin privacy for the real world*. URL: https://bitcointalk.org/?topic=279249 (visited on 09/13/2019).

[5] John R. Douceur. "The Sybil Attack". In: *Peer-to-Peer Systems*. Ed. by Peter Druschel, Frans Kaashoek, and Antony Rowstron. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 251–260. ISBN: 978-3-540-45748-0.

[6] *Fidelity-bonds for trust free mixers*. URL: https://bitcointalk.org/index.php?topic=172047.0 (visited on 09/15/2019).

[7] Ethan Heilman et al. "TumbleBit: An Untrusted Bitcoin-Compatible Anonymous Payment Hub". In: *Proceedings 2017 Network and Distributed System Security Symposium*. Network and Distributed System Security Symposium. San Diego, CA: Internet Society, 2017. ISBN: 978-1-891562-46-4. DOI: 10.14722/ndss.2017.23086. URL: https://www.ndss-symposium.org/ndss2017/ndss-2017-programme/tumblebit-untrusted-bitcoin-compatible-anonymous-payment-hub/ (visited on 09/15/2019).

[8] Brian Neil Levine, Clay Shields, and N Boris Margolin. "A Survey of Solutions to the Sybil Attack". In: (), p. 11.

[9] Sarah Meiklejohn et al. "A fistful of bitcoins: characterizing payments among men with no names". In: *Proceedings of the 2013 conference on Internet measurement conference - IMC '13*. the 2013 conference. Barcelona, Spain: ACM Press, 2013, pp. 127–140. ISBN: 978-1-4503-1953-9. DOI: 10.1145/2504730.2504747. URL: http://dl.acm.org/citation.cfm?doid=2504730.2504747 (visited on 09/18/2019).

[10] Ian Miers. *Blockchain Privacy: Equal Parts Theory and Practice*. The Zcash Foundation. URL: https://www.zfnd.org/blog/blockchain-privacy/ (visited on 09/17/2019).

[11] *Mixing service - Bitcoin Wiki*. URL: https://en.bitcoin.it/wiki/Mixing_service (visited on 09/18/2019).

[12] Aziz Mohaisen and Joongheon Kim. "The Sybil Attacks and Defenses: A Survey". In: *arXiv:1312.6349 [cs]* (Dec. 22, 2013). arXiv: 1312.6349. URL: http://arxiv.org/abs/1312.6349 (visited on 09/15/2019).

[13] Malte Moser and Rainer Bohme. "Join Me on a Market for Anonymity". In: (), p. 21.

[14] Satoshi Nakamoto. “Bitcoin: A Peer-to-Peer Electronic Cash System”. In: (), p. 9.

[15] Arvind Narayanan and Malte Mfffdfffdser. “Obfuscation in Bitcoin: Techniques and Politics”. In: (), p. 7.

[16] nopara73. *nopara73/ZeroLink*. original-date: 2017-07-28T06:01:08Z. Sept. 10, 2019. URL: https://github.com/nopara73/ZeroLink (visited on 09/13/2019).

[17] *Proof of burn - Bitcoin Wiki*. URL: https://en.bitcoin.it/wiki/Proof_of_burn (visited on 09/15/2019).

[18] *Proof of Stake - Bitcoin Wiki*. URL: https://en.bitcoin.it/wiki/Proof_of_Stake (visited on 09/15/2019).

[19] *Purchasing fidelity bonds by provably throwing away bitcoins*. URL: https://bitcointalk.org/index.php?topic=134827.0 (visited on 09/15/2019).

[20] Nicolas van Saberhagen. “CryptoNote v 2.0”. In: 2013.

[21] *Spam resistant block creator selection via burn auction*. Ethereum Research. July 21, 2019. URL: https://ethresear.ch/t/spam-resistant-block-creator-selection-via-burn-auction/5851 (visited on 09/15/2019).

[22] *The lagest Bitcoin mixer is about to stop working*. URL: https://bitcointalk.org/index.php?topic=2042470.0 (visited on 09/18/2019).

[23] *Timelock - Bitcoin Wiki*. URL: https://en.bitcoin.it/wiki/Timelock (visited on 09/15/2019).