

Improving Bitcoin Transaction Propagation by Leveraging Unreachable Nodes

Federico Franzoni^{*1} and Vanesa Daza^{†1}

¹Universitat Pompeu Fabra, Barcelona, Spain,
{federico.franzoni,vanesa.daza}@upf.edu

Abstract

The Bitcoin P2P network is at the core of all communications between clients. The reachable part of this network has been explored and analyzed by numerous studies. Unreachable nodes, however, are, in most part, overlooked. Nonetheless, they are a relevant part of the network and play an essential role in the propagation of messages. In this paper, we focus on transaction propagation and show that increasing the participation of unreachable nodes can potentially improve the robustness and efficiency of the network. In order to do that, we propose a few changes to the network protocol. Additionally, we design a novel transaction propagation protocol that explicitly involves unreachable nodes to provide better protection against deanonymization attacks. Our solutions are simple to implement and can effectively bring immediate benefits to the Bitcoin network.

1 Introduction

The Bitcoin P2P network is the channel through which clients exchange transactions and blocks. The characteristics and behavior of the nodes in this network have been extensively studied in the literature, with particular attention to efficiency and security. However, most papers only focus on the reachable portion of the network, leaving unreachable nodes out of their scope [1]–[3]. Nevertheless, these nodes represent more than 90% of the whole network [4], [5]. The reason for this imbalance resides in the inability of measuring tools to create connections towards unreachable nodes, which, by definition, only establish outgoing connections. As a result, the unreachable part of the network is often overlooked, and its relevance is underestimated.

In this paper, we analyze the importance of unreachable nodes in the propagation of messages and show how their participation can be beneficial to the

^{*}This author is partly supported by the Spanish Ministry of Economy and Competitiveness under the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502).

[†]This author was supported by Project RTI2018-102112-B-I00 (AEI/FEDER,UE).

network. In particular, we study the characteristics of these nodes, as emerged from state-of-the-art research, and identify some of their strengths and weaknesses compared to reachable nodes. We then propose changes to the protocol to improve the connectivity of the network as well as the efficiency of message propagation.

Additionally, we show that unreachable nodes are protected from adversaries that connect to the victims. Based on this characteristic, we design a novel transaction propagation protocol that potentially improves security against deanonymization attacks. Our solution explicitly involves unreachable nodes in the propagation pattern and exploits their position in the network to conceal the source of the message. We thoroughly justify our design choices and study the security of our protocol against an eavesdropper adversary [6].

Our contribution includes:

- we study the characteristics of unreachable nodes based on state-of-the-art research, and show have natural protection against a number of attacks;
- we show how unreachable nodes can play an important role in the network and propose changes to the protocol in order to achieve more robustness and improve the efficiency of message propagation;
- we design a new transaction propagation protocol aimed at improving anonymity and study its resilience against powerful adversaries.

2 Background

Unreachable nodes are typically associated with NATs. In fact, Carrier-Grade NATs (CG-NATs) are the primary cause for the unreachability of a node, as most ISPs use this technology to grant access to a larger number of devices. In this section, we provide more details about NAT and describe the current Bitcoin transactions propagation protocol, which we will modify to provide better anonymity.

2.1 NAT and P2P networks

Network Address Translation (NAT) [7] is a method to map IP addresses between incompatible networks. The most common type, known as *Network Address and Port Translation* (NAPT), is often used to connect private networks to the Internet without the need to assign a unique address to each device. NAPT is often regarded as a solution to the IPv4 address exhaustion problem [8], since it allows a large number of devices to connect through a shared IP address. As a side effect, such devices cannot be reached from the Internet, unless they first open a connection.

While this is not a problem in a client-server setting, it is a serious limitation for P2P networks. Notably, it prevents NATted nodes from connecting to each other. To overcome this limitation, *NAT traversal* techniques have been devised [9].

The Bitcoin reference client implements *Universal Plug-and-Play* (UPnP), which, however, is incompatible with CG-NATs, as it needs direct access from the host. Furthermore, the UPnP option is disabled by default due to a known vulnerability in the protocol. As a consequence, NATted nodes only establish outbound connections, which in the reference client are limited to just 8.

2.2 Transaction Propagation and Anonymity

Forwarding a transaction in Bitcoin is a three-step process. First, the sender transmits an inventory (INV) message to advertise the hash of the transaction. The receiver then checks the hash and, if it is unknown, requests the full transaction data with a GETDATA message. Finally, the sender transmits the full transaction in a TX message. The INV-based transmission allows sending the transaction only to those nodes which still have not received it. When a node receives a new transaction, it relays it to its peers following the same process.

The relay step is fundamental in determining the propagation pattern. Since 2015, Bitcoin adopts the *diffusion spreading* protocol, where nodes relay transactions to each neighbor with an independent, exponential delay. Newly-generated transactions are transmitted in the same way by their source.

As shown in [6], the pattern generated by this gossip-like protocol leads to possible deanonymization attacks based on the so-called *rumor centrality* [10]. In simple words, since a transaction spreads symmetrically from each node to its peers, it is possible to determine the origin of the spreading (i.e., the first node that transmitted the transaction) by observing the state of its propagation through the network. Given that the source broadcasts the transaction in the same way as the relays, detecting the origin of the propagation often means identifying the creator of the transaction (in terms of the device address).

Based on this fact, several attacks [6], [11], [12] have been shown where an *eavesdropper adversary* connects to all reachable nodes and applies the so-called *first-spy* estimator, which simply associates a transaction to the first node that relays it. Fanti et al. [13] showed that this type of strategy often has very high levels of accuracy.

3 Bitcoin Unreachable Nodes

Nodes in P2P networks are traditionally divided in research into reachable and unreachable. A node is called *reachable* if it can accept incoming connections from other peers. Otherwise, it is called *unreachable*.

Nodes can be unreachable because they are protected by a firewall, connecting through a proxy or, most typically, being hosted in a private network, behind a NAT device. Less commonly, nodes purposely choose not to accept incoming connections.

Reachable and unreachable nodes are often named servers and clients, respectively, to recall the ability of the firsts to accept connections and the fact that the seconds connect to them. However, there is no client-server relationship

among them, as they follow the same P2P protocol. A more precise classification commonly used in other P2P-related papers, distinguish between *routable* and *non-routable* peers, and calls *unreachable* those peers that are known to other peers but cannot be contacted (e.g., because they are offline or only accept connections from known peers) [14]. In the following, the terms *reachable* and *unreachable* will be used to indicate *routable* and *non-routable*. Furthermore, for the sake of simplicity, *reachable* and *unreachable* nodes will be denoted by *R nodes* and *U nodes*, respectively.

Despite their relevance, U nodes have been marginally covered by state-of-the-art research. Most Bitcoin network-related analyses focus on R nodes [1], [2], [15], [16], leaving U nodes out of scope. The statistics given by these works hardly give a precise account of U nodes, as they do not distinguish between offline nodes and nodes that are actually out of reach. However, virtually all of them show that the number of U nodes is much greater than R nodes (estimates go from 10 [5] to 30 times more [4]). Furthermore, studies showed that regular Bitcoin users tend to use U nodes, while R nodes are mostly run in data centers [4], [17].

As noted in [18], U nodes can contribute to the robustness of the network, as they increase connectivity and they are harder to attack for adversaries without access to core infrastructure. One of the goals of this paper will be to improving the efficiency and robustness of data propagation by increasing the participation of U nodes in the network.

U nodes are often overlooked due to the difficulty of connecting measurement tools to all of them at the same time, which is commonly done with R nodes. The only way to study U nodes is to deploy R nodes and wait for U nodes to connect to them. By adopting this approach, Wang et al. [4] were able to study U nodes in detail. As of today, their work is the only known global analysis of U nodes. Besides it, the only papers focusing on U nodes are those targeting them for deanonymization [19]–[21]. The interest in this kind of attacks stems from the relative difficulty of targeting U nodes at a global level. In fact, like U nodes are hard to study, they are also hard to target by an attacker, which has to deal with the inability to open connections to them. Specifically, U nodes are hard to include in observation-based attacks [6] and unsolicited-message-based attacks [22]. Similarly, U nodes are also immune to many network-level attacks, such as eclipse attacks [23], topology-inferring attacks [2], and partitioning attacks [24].

NATted nodes are also hard to distinguish, since they can share the same IP address. This is why network-wide deanonymization attacks against U nodes usually require fingerprinting techniques for identifying nodes individually [19], [21]. Nonetheless, U nodes are very susceptible to deanonymization attacks in case they directly connect to the adversary. In this case, even a simple first-spy estimator can obtain a very high level of accuracy [4]. In this paper, we propose a new propagation protocol that allows protecting both R and U nodes from deanonymization.

4 Our Solution

We propose some changes to the network protocol, which leverage the specificity of unreachable nodes to improve the efficiency and security of the network. In particular, we propose the following changes:

- Explicitly distinguish reachable and unreachable nodes;
- Increase connections from unreachable nodes;
- Disable advertisement of unreachable addresses;
- Adopt the propagation protocol described in §4.2.

4.1 Network Changes

We first describe the changes to basic network protocol behavior that allow us to improve security and efficiency.

4.1.1 Explicitly distinguish between R and U nodes

Although there is some difference in the behavior of R nodes and U nodes in the reference client, the Bitcoin network protocol does not make any explicit distinction between them. However, our solution is based on the different characteristics shown by the two types of node, as shown in §3. As such, explicitly distinguish between R nodes and U nodes is a necessary step.

Different strategies can be followed by a node to determine its reachability. A naive approach would be to verify if the client accepts incoming connections. However, it might be the case that a node is accepting connections but its address is unreachable from the outside. A better approach is to have the node connect to its own address, as seen by its peers, and set itself reachable, if the attempt succeeds, and unreachable, otherwise.

4.1.2 Increase U nodes connections

The second modification we propose is to increase the number of outbound connections of U nodes. This change has several effects. Firstly, it helps leveling the imbalance of connectivity between R and U nodes. In fact, while R nodes can reach 125 connections, U nodes only maintain up to 8, corresponding to their outbound peers. On the other hand, inbound slots are often underutilized by R nodes [15], which means they can handle a higher number of connections.

Secondly, increasing the number of peers means receiving, and relaying, more transactions per amount of time. Given the great number of U nodes, even a small increase in their connections might produce a significant improvement in the propagation speed of transactions and blocks.

Furthermore, from a security perspective, it has been shown how increasing the number of outbound connections can improve resistance against DoS attacks [18], eclipse attacks [23], and isolation attacks [25].

Finally, a higher number of connections for U nodes can be beneficial for the anonymity of our propagation protocol, as we will show in §4.2.

4.1.3 Do not advertise U nodes addresses

In the current protocol, U nodes, like R nodes, advertise their public address to their peers. These addresses represent 90% of those being spread through the network [2], [4]. However, being unreachable, these addresses are of no use to any other node. At the same time, they increase network traffic [19] and likely produce a high number of failed connection attempts. Additionally, they potentially reduce the availability of reachable addresses, since new (often unreachable) addresses replace old ones in the node database when the address pool is full.

From a security perspective, these addresses enable fingerprinting techniques, which allow for deanonymization attacks [19], [20]. Disabling their advertisement to outbound peers would effectively invalidate the few known deanonymization attacks targeting U nodes.

4.2 Transaction Propagation Protocol

In the following, we discuss and describe our new propagation protocol design. The protocol explicitly leverages U nodes to improve resiliency against deanonymization. Similarly to Dandelion [26], [27], we include two essential concepts in our design: proxied broadcast and transaction mixing. Proxied broadcast consists in delegating the diffusion of a new transaction to another node (called *proxy*), allowing to hide the real origin of the transaction. Mixing consists in sending to the proxy other new (proxied) transactions, received from other peers. This makes it hard for the proxy to distinguish between transactions generated by the sender and others relayed by the sender, but generated by other nodes.

Given what said about U nodes in §3, we want to leverage their protected position in the network to conceal the origin of a transaction. The core idea is to make R nodes, which are more susceptible to deanonymization, use U nodes as proxies for their transactions. This way, such transactions will look as generated by their proxies instead of the actual source. Additionally, we hinder proxies from distinguishing such transactions by mixing them with transactions from other nodes.

Before detailing the propagation protocol, we define our adversary model and motivate our design choices.

Network and Adversary Model To describe our protocol, we model the Bitcoin network as in Figure 1. We call O the R node running the protocol and generating new transactions. Other nodes are denoted by R_i , if reachable and U_i , if unreachable, for $i = 1, 2, \dots$. The adversary A aims at deanonymizing transactions generated by O and can control various nodes, both reachable and unreachable, which we denote by R_i^A and U_i^A , respectively. A can connect

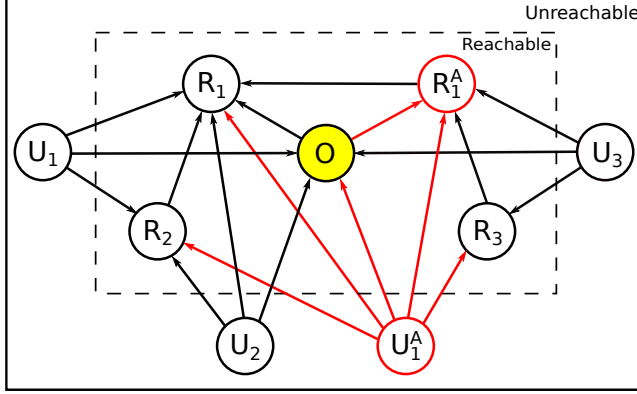


Figure 1: Our view of the Bitcoin network: the origin O of a transaction is connected to R and U nodes. The adversary (colored in red) deploys both R and U nodes and connects to all reachable nodes.

to all reachable nodes and also create multiple connections to the same node (including O). Nonetheless, A cannot directly connect to other U nodes. To that respect, A can only deploy multiple R nodes to increase the chance of having honest U nodes connecting to it. Additionally, the adversary can create and transmit transactions, as well as relay or retain others received from its peers.

Design In our protocol, R nodes leverage U nodes as proxies and use transactions coming from other U peers for mixing. Instead, U nodes use R nodes as proxies and mix new transactions with those coming from other R peers.

This scheme allows protecting both R and U nodes. In fact, U nodes cannot distinguish between transactions generated by their R peers and those proxied by such peers but generated by other U nodes. Similarly, U -generated transactions are indistinguishable to R nodes from those generated by other R nodes and proxied by their U peers.

However, a naive design could lead to easy deanonymization attacks, and also to an ineffective propagation of new transactions through the network. Therefore, we need to define (1) which peers are used for proxying and (2) which transactions are used for mixing.

As for point (1), an R node can select one, all, or a subset of its U peers. Note that an adversary can control a large subset of U peers of R nodes. This increases her probability of being selected as the first proxy for many R -generated transactions, allowing an effective use of the first-spy estimator. At the same time, if we send all transactions to a single proxy, it will be easy for this one to narrow down the set of transactions possibly generated by the sender. As such, we first select a subset of peers to be used as proxy and pick a random one within this subset for every proxy operation. We call this subset the *proxy set*. In order to distribute transactions among all nodes and minimize the risk

of a proxy collecting all new transactions from a node, we change the proxy set at a certain rate. We call *epoch* the time frame in which a proxy set is used.

As for point (2), we first need to identify which transactions are suitable for mixing. Note that transactions received by an R node from other R peers following our protocol have already been diffused, making them unsuitable for mixing (since the adversary might already know them). Similarly, transactions diffused by U peers might have already been received by the adversary. On the other hand, it is easy to see that proxied transactions are the least likely to be known to the adversary, and thus best suited for mixing.

Therefore, we need to identify which transactions are being proxied and which are being diffused. To do so, we mark proxied transactions and distinguish between two propagation phases: the *proxying phase* and the *diffusion phase*. We call transactions in the proxying phase *proxy transactions*. When a new transaction is created is marked as proxying and sent to a node of the proxy set. As for mixing, nodes use proxy transactions coming from their peers. We call the set of proxy transactions used for mixing, the *mixing set* of a node. Transactions in the mixing set are relayed through the same path as newly-generated ones so as to make them indistinguishable from each other.

Ideally, we would like the mixing set to be as large as possible. However, if we used all incoming proxy transactions, they would never be diffused. Instead, we include only a fraction of such transactions in the mixing set, and diffuse the rest. To do so, we need to decide which transactions to diffuse and which to relay. A possible strategy is to select some peers in each epoch, and only use transactions coming from them. However, if an adversary controls many of these peers and also the selected proxy, she could track most transactions in the mixing set of the target, leading to an easy deanonymization. To avoid such a risk, we select proxy transactions from all of our peers and probabilistically include them in our mixing set. In particular, for each proxy transaction, we keep proxying it with a certain probability p and diffuse it otherwise. This way, despite being able to track or inject proxy transactions for a specific node, an adversary cannot affect the number of honest transactions included in its mixing set. A correct choice of p will be fundamental for the effectiveness and efficiency of our protocol.

To further protect R nodes from adversaries controlling many inbound connections, we adopt the *bucketing* strategy used in Bitcoin Core for managing addresses. This mechanism is used to prevent an adversary from filling up the address database with malicious IPs, and it is based on the assumption that the attacker only controls nodes from a limited address space [28]. In particular, each bucket contains addresses from a different subnet. Similarly, we make R nodes select proxies and transactions for the mixing set uniformly at random among peers from different buckets.

Finally, to cope with the risk of a transaction not being diffused, due to a DoS attack by a proxy or to an excessively long proxying phase, each node sets a timeout t for every proxied transaction. When t expires, the node verifies if the transaction has been diffused by checking if the majority of outbound peers have advertised it back to us. We choose to monitor outbound peers

to minimize the risk of an adversary deceiving an R node by relaying proxied transactions from other adversary-controlled U peers. The same rule is applied to both new and relayed transactions, so as to avoid deanonymization due to rebroadcast. In the current protocol, in fact, a rebroadcast is only done by the source of the transaction, and can thus reveal its origin [11]. In our protocol, instead, rebroadcast applies to all proxied transactions, thus leaking no new information.

Protocol Rules To detail the propagation rules of our protocol, we first define the *proxy* operation on a transaction tx as follows:

Algorithm 1: Proxy(tx)

```

Pick a random peer  $P$  from the proxy set;
Send  $tx$  to  $P$  and set a timeout  $t$ ;
When  $t$  expires:
if The majority of outbound peers advertised  $tx$  then
    | Return
else
    | Repeat
end

```

Next, we define the propagation rules for R nodes:

Algorithm 2: R Propagation Rules

```

Divide time into epochs;
if New epoch begins then
    | Select subset  $S$  from U peers uniformly at random from different
    |   buckets;
    | Set  $S$  as the proxy set
end
if Create new transaction  $tx$  then
    | Mark  $tx$  as proxying;
    | Run proxy( $tx$ )
end
if Receive a proxying transaction  $tx_m$  from a U peer then
    | with probability  $p$ , execute proxy( $tx_m$ );
    | otherwise, diffuse( $tx$ )
end

```

U nodes follow the same rules, except they use R peers instead of U peers and do not use buckets.

5 Discussion

5.1 Limitations

Our protocol requires R nodes to have U peers connected to them. However, newly-joined R nodes usually have to wait some time to have other peers connect to them. We address this limitation by having new R nodes use the diffusion protocol until they have a sufficient number of U peers. Additionally, to prevent an adversary from taking advantage from this situation (by filling up all inbound slots), we also adopt the bucketing strategy. Specifically, we make R nodes use our protocol only when enough U peers from different buckets are connected.

5.2 Propagation and anonymity

To better understand our protocol it is useful to depict the propagation pattern of a transaction.

Let us consider an R node O generating a transaction tx . The following sequence of events happens:

1. R selects a proxy P among its proxy set, mark tx as *proxying* and sends it to P ;
2. P receives tx and proxy it with probability p , or diffuses it otherwise;
3. If proxying tx , P selects a node R from its proxy set S and sends it tx ;

Proxying transactions are relayed through a sequence of R and U nodes until it gets diffused. Diffusion can happen at any step, except for the first one. Propagation from an U node follows a similar pattern.

A major risk of proxied broadcast is that a transaction might take too long to diffuse, or not be diffused at all. As for diffusion time, we can statistically guarantee to diffuse every transaction within a reasonable time. Since at every hop, the transaction tx is diffused with probability p , it is possible to tune this value to obtain a target number of hops through which tx is proxied on average. The use of timeouts allows dealing with a transaction not being diffused.

With respect to anonymity, our protocol is designed to be resistant against a first-spy estimator. This type of adversary connects to all R nodes and links each transaction to the first node from which it has been received. As demonstrated in [13], this strategy is very effective with the current propagation model. However, the changes introduced by our protocol make it very unlikely for a node to first receive a transaction from its source. On the contrary, most of the times, transactions will be received by a node different from the origin, thanks to proxying. Furthermore, each transaction is mixed with many others generated by nodes in the proxying path, which are indistinguishable from each other to the receiving node. This means that any claim about the origin of a transaction can be easily denied.

Note that our protocol is designed to resist against very powerful adversaries controlling several nodes and maintaining multiple connections to all reachable

nodes. The adversary can combine information from all of its nodes and coordinate them to influence or track the mixing set of a target node. However, we showed in the previous section how such an adversary has limited capabilities to affect the security of the protocol.

5.3 Ephimerality of U nodes

A possible issue in our design is the short time of connection of many U nodes. In fact, while R nodes are relatively stable [29], U nodes often experience very short-lived connections [4]. This behavior might affect the efficiency of the protocol. However, the timeout mechanism is also meant to deal with this kind of problems and can be fine-tuned independently by each node, depending on the experienced churn.

Moreover, the presence of short-lived proxy nodes, if properly exploited, might serve as an added value to the anonymity level of our protocol, as it makes it harder to track back a transaction to its origin.

5.4 NAT adoption

Another potential limitation of our solution is that it is based on the unreachability of NATted nodes. However, if IPv6 gets adopted by the majority of nodes, it is possible that NATs will cease to be used. The introduction of IPv6, in fact, was mainly intended to deal with the IPv4 address exhaustion problem and remove the need for NATting [30].

Although growing, the adoption rate of IPv6 seems to be variable [31] and not uniform worldwide, with statistics strongly dependent on the adopted metrics [32]. As for the Bitcoin network, Neudecker et al. [17] showed that, unlike IPv4, IPv6 connections have not grown over the past two years.

Either way, most optimistic estimates predict a complete adoption within 7-8 years [33]. In this perspective, our protocol should be considered as a medium-term solution, likely able to work for the next decade.

6 Related Work

6.1 Unreachable nodes

Unreachable nodes have been extensively studied by Wang et al. [4]. In order to perform their analyses, they deployed around 100 nodes, through which they collected information on more than 100 K unreachable peers, which generated more than 2 M transactions. Their findings show that most connections last for less than 60 seconds, while, at the same time, most transaction propagations are sent over long-lived connections (more than 100 seconds), showing a high degree of centrality. Finally, they show a method to deanonymize transactions coming from unreachable peers, with the help of an external listener node. Their results show that unreachable nodes are also susceptible to the first-spy estimator attack. Note that our protocol makes this attack much less effective,

since new transactions are proxied to a single reachable node, reducing the probability that the attacker receives the transaction. At the same time, new transactions are mixed with transactions proxied from its peers, thus reducing the accuracy of the attack.

Other deanonymization attacks also target U nodes globally by means of fingerprinting techniques. Biryukov et al. [19] make use of ADDR messages to uniquely identify U nodes by the set of their peers. Their technique allows linking multiple transactions created by the same node over a single session. In this paper, we proposed a change in the address advertisement by U nodes that would make this attack ineffective.

In [20], Biryukov et al. devise a technique to deanonymize U nodes connecting via Tor, even through multiple sessions. However, their attack is specific to Tor users.

Finally, Mastan et al. [21] exploit block requests patterns to identify U nodes over consecutive sessions. However, their technique only allows linking sessions and thus needs to be used in conjunction with other deanonymization techniques.

6.2 Transaction Propagation Anonymity

Anonymity properties of transaction propagation have been extensively studied in research.

Koshy et al. [11] are among the first ones to show a practical deanonymization technique for Bitcoin, based on transaction propagation analysis. They show that anomalies in the propagation can be exploited to identify the source of transactions.

Neudecker et al. [34] combine observations of the message propagation with Bitcoin address clustering techniques. However, their results show that for the vast majority of users this information does not facilitate deanonymization.

In [13], Fanti et al. thoroughly analyze the anonymity properties of both the former and the actual Bitcoin propagation protocols (Trickle and Diffusion). They theoretically prove that both protocols offer poor anonymity on networks with a regular-tree topology. They also identify the symmetry of current spreading protocols as the main characteristic that allows deanonymization attacks. An alternative protocol, called Dandelion, is then proposed in [26], [27] that specifically addresses this issue. Dandelion breaks the symmetry by proxying the broadcast of new transactions through a network-wide circuit of nodes. Furthermore, they increase anonymity by mixing new transactions over the same path. However, their protocol is somewhat hard to implement and only applies to R nodes. Our protocol, instead, involves both R and U nodes and does not seem to show any difficulties of implementation.

7 Conclusion and Future Work

In this paper, we showed how unreachable nodes are often overlooked despite being a relevant part of the Bitcoin network. We also showed how increasing their connectivity and participation can be beneficial for the robustness and efficiency of the whole network, without introducing overhead. Moreover, we leveraged their peculiar position in the network to design a new transaction propagation protocol to protect all nodes from deanonymization attacks. We analyzed the security of our proposal against powerful adversaries and discussed possible limitations of our approach.

Future work includes implementing our changes and evaluating their effectiveness through experiments. A more formal analysis is also needed to provide better anonymity guarantees. Finally, an analysis of the current participation degree of unreachable nodes in the transaction propagation might shed light on their relevance in the Bitcoin network.

References

- [1] J. A. Donet Donet, C. Pérez-Solà, and J. Herrera-Joancomartí, “The bitcoin P2P network,” in *Financial Cryptography and Data Security*, R. Böhme, M. Brenner, T. Moore, and M. Smith, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 87–102, ISBN: 978-3-662-44774-1.
- [2] A. Miller, J. Litton, A. Pachulski, N. Gupta, D. Levin, N. Spring, and B. Bhattacharjee, “Discovering bitcoin’s public topology and influential nodes,” 2015.
- [3] S. B. Mariem, P. Casas, M. Romiti, B. Donnet, R. Stütz, and B. Haslhofer, *All that glitters is not bitcoin – unveiling the centralized nature of the BTC (IP) network*, 2020. arXiv: 2001.09105 [cs.NI].
- [4] L. Wang and I. Pustogarov, “Towards better understanding of bitcoin unreachable peers,” *CoRR*, vol. abs/1709.06837, 2017. arXiv: 1709.06837. [Online]. Available: <http://arxiv.org/abs/1709.06837>.
- [5] S. Delgado-Segura, S. Bakshi, C. Pérez-Solà, J. Litton, A. Pachulski, A. Miller, and B. Bhattacharjee, “Txprobe: Discovering bitcoin’s network topology using orphan transactions,” *CoRR*, vol. abs/1812.00942, 2018. arXiv: 1812.00942. [Online]. Available: <http://arxiv.org/abs/1812.00942>.
- [6] G. Fanti and P. Viswanath, “Deanonymization in the bitcoin P2P network,” in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., Curran Associates, Inc., 2017, pp. 1364–1373. [Online]. Available: <http://papers.nips.cc/paper/6735-deanonymization-in-the-bitcoin-p2p-network.pdf>.
- [7] P. Srisuresh and M. Holdrege, “IP network address translator (NAT) terminology and considerations,” Tech. Rep., 1999.

- [8] P. Richter, M. Allman, R. Bush, and V. Paxson, "A primer on IPv4 scarcity," *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 2, 21–31, Apr. 2015, ISSN: 0146-4833. DOI: 10.1145/2766330.2766335. [Online]. Available: <https://doi.org/10.1145/2766330.2766335>.
- [9] Z. Hu, "NAT traversal techniques and peer-to-peer applications," in *HUT T-110.551 Seminar on Internetworking*, Citeseer, 2005, pp. 04–26.
- [10] D. Shah and T. Zaman, "Rumor centrality: A universal source detector," *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 1, 199–210, Jun. 2012, ISSN: 0163-5999. DOI: 10.1145/2318857.2254782. [Online]. Available: <https://doi.org/10.1145/2318857.2254782>.
- [11] P. Koshy, D. Koshy, and P. McDaniel, "An analysis of anonymity in bitcoin using p2p network traffic," in *Financial Cryptography and Data Security*, N. Christin and R. Safavi-Naini, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 469–485, ISBN: 978-3-662-45472-5.
- [12] A. Biryukov and S. Tikhomirov, "Deanonymization and linkability of cryptocurrency transactions based on network analysis," in *2019 IEEE European Symposium on Security and Privacy (EuroS P)*, Jun. 2019, pp. 172–184. DOI: 10.1109/EuroSP.2019.00022.
- [13] G. C. Fanti and P. Viswanath, "Anonymity properties of the bitcoin P2P network," *CoRR*, vol. abs/1703.08761, 2017. arXiv: 1703.08761. [Online]. Available: <http://arxiv.org/abs/1703.08761>.
- [14] M. Essaid, S. Park, and H.-T. Ju, "Bitcoin's dynamic peer-to-peer topology," *International Journal of Network Management*, vol. n/a, no. n/a, e2106, 2020, e2106 nem.2106. DOI: 10.1002/nem.2106. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nem.2106>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nem.2106>.
- [15] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *IEEE P2P 2013 Proceedings*, Sep. 2013, pp. 1–10. DOI: 10.1109/P2P.2013.6688704.
- [16] S. Feld, M. Schönfeld, and M. Werner, "Analyzing the deployment of bitcoin's P2P network under an AS-level perspective," *Procedia Computer Science*, vol. 32, pp. 1121–1126, 2014, The 5th International Conference on Ambient Systems, Networks and Technologies (ANT-2014), the 4th International Conference on Sustainable Energy Information Technology (SEIT-2014), ISSN: 1877-0509. DOI: <https://doi.org/10.1016/j.procs.2014.05.542>. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S187705091400742X>.
- [17] T. Neudecker, "Characterization of the bitcoin peer-to-peer network (2015-2018)," Karlsruher Institut für Technologie (KIT), Tech. Rep. 1, 2019, 29 pp. DOI: 10.5445/IR/1000091933.

- [18] T. Neudecker and H. Hartenstein, “Network layer aspects of permissionless blockchains,” *IEEE Communications Surveys Tutorials*, vol. 21, no. 1, pp. 838–857, 2019.
- [19] A. Biryukov, D. Khovratovich, and I. Pustogarov, “Deanonymisation of clients in bitcoin P2P network,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’14, New York, NY, USA: ACM, 2014, pp. 15–29, ISBN: 978-1-4503-2957-6. DOI: 10.1145/2660267.2660379. [Online]. Available: <http://doi.acm.org/10.1145/2660267.2660379>.
- [20] A. Biryukov and I. Pustogarov, “Bitcoin over Tor isn’t a good idea,” in *2015 IEEE Symposium on Security and Privacy*, May 2015, pp. 122–134. DOI: 10.1109/SP.2015.15.
- [21] I. D. Mastan and S. Paul, “A new approach to deanonymization of unreachable bitcoin nodes,” in *Cryptology and Network Security*, S. Capkun and S. S. M. Chow, Eds., Cham: Springer International Publishing, 2018, pp. 277–298, ISBN: 978-3-030-02641-7.
- [22] C. Rossow, D. Andriesse, T. Werner, B. Stone-Gross, D. Plohmann, C. J. Dietrich, and H. Bos, “Sok: P2PWED - modeling and evaluating the resilience of peer-to-peer botnets,” in *2013 IEEE Symposium on Security and Privacy*, 2013, pp. 97–111.
- [23] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, “Eclipse attacks on bitcoin’s peer-to-peer network,” in *24th USENIX Security Symposium (USENIX Security 15)*, Washington, D.C.: USENIX Association, Aug. 2015, pp. 129–144, ISBN: 978-1-931971-232. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/heilman>.
- [24] M. Tran, I. Choi, G. Moon, A. V. Vu, and M. Kang, “A stealthier partitioning attack against bitcoin peer-to-peer network,” in *2020 IEEE Symposium on Security and Privacy (SP)*, Los Alamitos, CA, USA: IEEE Computer Society, May 2020, pp. 515–530. DOI: 10.1109/SP40000.2020.00027. [Online]. Available: <https://doi.ieeeecomputersociety.org/10.1109/SP40000.2020.00027>.
- [25] M. Apostolaki, A. Zohar, and L. Vanbever, “Hijacking bitcoin: Routing attacks on cryptocurrencies,” in *2017 IEEE Symposium on Security and Privacy (SP)*, May 2017, pp. 375–392. DOI: 10.1109/SP.2017.29.
- [26] S. Bojja Venkatakrishnan, G. Fanti, and P. Viswanath, “Dandelion: Redesigning the bitcoin network for anonymity,” *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 1, 22:1–22:34, Jun. 2017, ISSN: 2476-1249. DOI: 10.1145/3084459. [Online]. Available: <http://doi.acm.org/10.1145/3084459>.

- [27] G. Fanti, S. B. Venkatakrishnan, S. Bakshi, B. Denby, S. Bhargava, A. Miller, and P. Viswanath, “Dandelion++: Lightweight cryptocurrency networking with formal anonymity guarantees,” *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 2, no. 2, Jun. 2018. DOI: 10.1145/3224424. [Online]. Available: <https://doi.org/10.1145/3224424>.
- [28] T. Neudecker, P. Andelfinger, and H. Hartenstein, “A simulation model for analysis of attacks on the bitcoin peer-to-peer network,” in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, May 2015, pp. 1327–1332. DOI: 10.1109/INM.2015.7140490.
- [29] Statoshi.info. (2020). “Peers.” (Last accessed: 2020-01-16), [Online]. Available: <https://statoshi.info/dashboard/db/peers>.
- [30] G. Van de Velde, T. Hain, R. Droms, and B. Carpenter, “Local network protection for IPv6,” RFC 4864, May, Tech. Rep., 2007.
- [31] K. McCarthy. (2018). “IPv6 growth is slowing and no one knows why. let’s see if el reg can address what’s going on.” (Last accessed: 2020-07-14), [Online]. Available: https://www.theregister.com/2018/05/21/ipv6_growth_is_slowing_and_no_one_knows_why/.
- [32] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey, “Measuring IPv6 adoption,” *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, 87–98, Aug. 2014, ISSN: 0146-4833. DOI: 10.1145/2740070.2626295. [Online]. Available: <https://doi.org/10.1145/2740070.2626295>.
- [33] L. Howard. (2019). “IPv6 growth.” (Last accessed: 2020-07-14), [Online]. Available: <https://www.retevia.net/ipv6-growth/>.
- [34] T. Neudecker and H. Hartenstein, “Could network information facilitate address clustering in bitcoin?” In *Financial Cryptography and Data Security*, M. Brenner, K. Rohloff, J. Bonneau, A. Miller, P. Y. Ryan, V. Teague, A. Bracciali, M. Sala, F. Pintore, and M. Jakobsson, Eds., Cham: Springer International Publishing, 2017, pp. 155–169, ISBN: 978-3-319-70278-0.