Contents lists available at ScienceDirect

Forensic Science International: Digital Investigation

journal homepage: www.elsevier.com/locate/fsidi

DFRWS EU 2025 - Selected Papers from the 12th Annual Digital Forensics Research Conference Europe

Unmixing the mix: Patterns and challenges in Bitcoin mixer investigations

Pascal Tippe [*], Christoph Deckers

FernUniversität in Hagen, Germany

ARTICLE INFO

Keywords:
Digital currency forensics
Bitcoin mixer forensics
Bitcoin mixer investigation
Custodial mixing services
Cryptocurrency investigation

ABSTRACT

This paper investigates the operational patterns and forensic traceability of Bitcoin mixing services, which pose significant challenges to anti-money laundering efforts. We analyze blockchain data using Neo4j to identify unique mixing patterns and potential deanonymization techniques. Our research includes a comprehensive survey of 20 currently available mixing services, examining their features such as input/output address policies, delay options, and security measures. We also analyze three legal cases from the U.S. involving Bitcoin mixers to understand investigative techniques used by law enforcement. We conduct two test transactions and use graph analysis to identify distinct transaction patterns associated with specific mixers, including peeling chains and multi-input transactions. We simulate scenarios where investigators have partial knowledge about transactions, demonstrating how this information can be leveraged to trace funds through mixers. Our findings reveal that while mixers significantly obfuscate transaction trails, certain patterns and behaviors can still be exploited for forensic analysis. We examine current investigative approaches for identifying users and operators of mixing services, primarily focusing on methods that associate addresses with entities and utilize off-chain attacks. Additionally, we discuss the limitations of our approach and propose potential improvements that can aid investigators in applying effective techniques. This research contributes to the growing field of cryptocurrency forensics by providing a comprehensive analysis of mixer operations and investigative techniques. Our insights can assist law enforcement agencies in developing more effective strategies to tackle the challenges posed by Bitcoin mixers in cybercrime investigations.

1. Introduction

Bitcoin has become an increasingly prominent digital asset and payment system since its inception in 2008. As its popularity grows, so does the demand for enhanced privacy measures, leading to the emergence of Bitcoin mixers. These services offer users increased anonymity by obfuscating the trail of transactions on the blockchain. While Bitcoin mixers provide legitimate privacy benefits for individuals and businesses, they have also become an issue in the realm of anti-money laundering efforts. The rise in cybercrime has made cryptocurrency tracing a crucial tool for investigators seeking to track illicit activities (Meiklejohn et al., 2013). However, Bitcoin mixers pose a significant challenge to these efforts by obscuring money flows and complicating the detection of criminal operations. These developments underscore the importance of understanding the mechanisms and patterns associated with Bitcoin mixing services. While non-custodial mixing protocols use known protocols, custodial mixers usually use their own techniques to conceal funds. We focus on custodial mixing services as recent high-profile cases (U.S. District Court for the District of Columbia, 2021) show their high relevance. Our research draws inspiration from the work of Möser et al. (2013), which employed test transactions to analyze mixing services. Building upon this approach, our study aims to

- Examine the features offered by currently operating custodial Bitcoin mixing services.
- Conduct test transactions on two selected Bitcoin mixing services to gather ground truth data.
- Analyze transaction data to identify patterns that can reveal the activity or connected addresses of these mixing services.
- Explore and evaluate possible investigative methods used to identify users and operators of Bitcoin mixing services.

By combining these objectives, we seek to contribute to the growing body of knowledge on cryptocurrency forensics and provide valuable insights for both researchers and law enforcement agencies tackling the challenges posed by Bitcoin mixers in the evolving landscape of digital

[*] Corresponding author.
E-mail address: pascal.tippe@fernuni-hagen.de (P. Tippe).

https://doi.org/10.1016/j.fsidi.2025.301876

Available online 24 March 2025

financial crime.

2. Related work

2.1. Bitcoin protocol

Bitcoin, introduced in the whitepaper *Bitcoin: A Peer-to-Peer Electronic Cash System* by the pseudonymous Nakamoto (2008), represents a significant development in digital currency. The Bitcoin network, launched in 2009 based on Nakamoto's reference implementation, initiated the development of decentralized digital currencies and blockchain technology. The Bitcoin system operates on a distributed public ledger, known as the blockchain, which records all transactions. This ledger is maintained by a peer-to-peer network that ensures the distribution and consistency across participating nodes and relays transactions and blocks. The blockchain's structure allows for the traceability of all Bitcoin transactions to their origin, contributing to the system's transparency. Bitcoin's security model relies on public-key cryptography. Users can generate a pair of cryptographic keys: a public key serving as the receiving address, and a private key for transaction authorization. To initiate a transfer, the sender digitally signs a hash of the previous transaction and the recipient's public key. Transactions are broadcast to the network for validation and inclusion in the blockchain. Users interact with the network through digital wallets that manage private keys and facilitate transactions.

A key component of the Bitcoin blockchain is its time-stamping mechanism, which provides proof of data existence at specific time points. This is implemented through a proof-of-work system, requiring the computation of a block hash meeting specific criteria. The proof-of-work process, known as mining, is performed by network participants called miners. Successful miners are rewarded with transaction fees and a predetermined number of newly created bitcoins, a mechanism that also controls currency supply. The proof-of-work system introduces computational complexity that increases with the difficulty of the block hash criteria. A variable called a nonce is adjusted incrementally until a matching hash is found. This process contributes to the security and consensus mechanism of the Bitcoin network. Bitcoin addresses have evolved over time to incorporate new features, and there are currently three main types of Bitcoin addresses. Legacy addresses begin with the number *1* and were the original address format. P2SH addresses start with *3* and allow more complex transactions and are often used for multi-signature wallets. Bech32 addresses, also known as native SegWit addresses, start with *bc1* and offer improved efficiency. Each address type has different characteristics in terms of transaction size, fees, and compatibility with older wallets. The choice of address type can impact transaction patterns and may be relevant in forensic analysis of Bitcoin transactions.

It is important to note that while Bitcoin transactions do not directly reveal the identities of the parties involved, the system is not anonymous but rather pseudonymous. Each Bitcoin address acts as a pseudonym for its owner. All transactions associated with an address are publicly visible on the blockchain, creating a permanent and traceable record. This pseudonymity presents both challenges and opportunities for investigators. By analyzing transaction patterns, clustering related addresses, and correlating blockchain data with external information, investigators can potentially link addresses to real-world identities (Androulaki et al., 2013). For instance, if an address interacts with a regulated exchange, the stored identification information can provide a starting point for identification. Additionally, techniques such as transaction graph analysis and heuristic clustering can reveal connections between addresses and identify patterns of behavior associated with specific entities or individuals. However, the effectiveness of these investigative methods can be hindered by the use of privacy-enhancing techniques such as Bitcoin mixing services, which aim to obfuscate the transaction trail (de Balthasar and Hernandez-Castro, 2017).

2.2. Mixing services

In response to growing concerns over transaction traceability and deanonymization in Bitcoin, the cryptocurrency community and researchers have developed various methods to enhance transaction privacy. These efforts have led to the creation of privacy-focused cryptocurrencies like Monero (koe et al., 2020), as well as the development of Bitcoin mixing services. Mixing services, also known as tumblers, aim to obfuscate the trail of transactions by pooling funds from multiple sources and redistributing them, making it difficult to trace the origin of specific coins. Two main types of Bitcoin mixing services are distinguishable: Custodial and non-custodial mixers. Custodial mixers are centralized services where users send their bitcoins to be mixed with others. The mixer temporarily takes custody of the funds, mixes them, and then returns different bitcoins of equivalent value to the users. While effective, these services require users to trust the mixer with their funds, introducing potential risks of theft or compromise. Non-custodial mixers are decentralized protocols that allow users to mix their coins without giving up control of their funds. The most prominent example of this is the CoinJoin protocol, which enables multiple users to combine their transactions into a single transaction, making it difficult to determine which inputs correspond to which outputs (Maxwell, 2013). Wu et al. (2021) introduced an abstraction model by dividing the mixing process into three phases: Taking inputs, performing mixing, and sending outputs. The anonymity is mainly achieved through mixing mechanisms, which can be further categorized into swapping and obfuscating. Swapping mechanisms swap inputs and outputs from different participants, while obfuscating mechanism are designed to protect relationship anonymity by breaking the matching procedure between participant inputs and outputs.

To avoid the detection of mixing transactions, mixers often employ peeling chains. A peeling chain consists of a series of transactions that distribute outputs, resembling normal user transactions with two outputs. This technique makes mixing transactions hard to distinguish from normal user transactions. A mixing transaction in a peeling chain typically consists of one output designated to the payment to the user and another output used for change, which then becomes the input of the next transaction in the chain. Some mixers introduce random time delays between receiving and sending funds to further obfuscate the transaction trail or use variable transaction sizes to make pattern recognition more difficult (Wu et al., 2021).

2.3. Bitcoin analysis techniques

Following the invention of Bitcoin, researchers began developing methods to trace and deanonymize transactions. According to the definitions provided by Pfitzmann and Köhntopp (2001), privacy in the context of an attack is any attempt to obtain additional knowledge about sender(s), recipient(s), or amount(s) of at least one transaction. These attacks often employ heuristics that may not be optimal but produce results in a reasonable timeframe. The effectiveness of these heuristics and attacks depends on underlying assumptions, which significantly influence the validity of their results. Deuber et al. (2022) critically review existing blockchain-related heuristics and introduce a taxonomy classifying them based on their underlying assumptions. They categorize these assumptions into different types, with user behavior and statistical assumptions being particularly relevant for Bitcoin tracing. User behavior assumptions are typically based on common patterns observed in cryptocurrency usage, often stemming from standard wallet software implementations. Statistical assumptions, on the other hand, are based on probabilistic models of transaction behavior. Several key heuristics have emerged in Bitcoin transaction analysis (Deuber et al., 2022; Reid and Harrigan, 2011):

Multi-input heuristic: This heuristic assumes that all inputs in a transaction are controlled by the same user or entity, as standard Bitcoin wallets do not support different users participating in a single transaction. However, CoinJoin, a decentralized mixing protocol, is specifically designed to render this heuristic ineffective. Some researchers argue that the low likelihood of multiple inputs in a single transaction originating from different users is sufficient to justify continued application of this heuristic. It's crucial to note that a single transaction can link two relevant addresses or clusters in an analysis.

Change-address heuristic: Bitcoin requires all inputs to a transaction to be completely spent, with any excess value sent to a change address. In its basic form, this heuristic states that for every transaction with two output addresses, if exactly one address was never used before, then that address is a change address. It assumes that every transaction pays to only one user. Meiklejohn et al. (2013) propose a refined definition to account for cases like gambling sites or mining pools with multiple payouts: An output address of a non-coinbase transaction is the change address if it is the only address in the outputs appearing for the first time, and there is no output address that also appeared in the inputs.

Address reuse heuristic: This heuristic assumes that an address that has been used before is more likely to be a payment address rather than a change address.This is based on the observation that change addresses are typically generated automatically by wallet software, while payment addresses are often reused by human users.

Temporal heuristic: This assumes that transactions occurring close together in time are more likely to be related. However, mixing services often introduce time delays to counteract this assumption.

There are other additional heuristics like the consistent use of the same address type which might be caused by the wallet implementation or assumption about user behavior like them favoring round numbers. While these heuristics rely on reasonable assumptions, mixing services are continually evolving to counteract these heuristics. Similarly, advanced wallet software may implement strategies to confuse change address detection, such as creating multiple change addresses or deliberately reusing addresses.

2.4. Bitcoin mixer tracing

Several studies have conducted test transactions to analyze Bitcoin mixing services, each focusing on different aspects of the mixing process and its implications. Wu et al. (2021) conducted test transactions as part of their comprehensive study on Bitcoin mixing services. Their approach primarily aimed at determining whether mixing services employed swapping or obfuscating mechanisms. While their analysis provided valuable insights into the operational methods of mixing services, it did not delve deeply into specific transaction patterns. Instead, their focus extended to estimating the revenue generated by these services and identifying associated addresses. This broad approach offered a general understanding of the mixing ecosystem. Möser et al. (2013) also performed test transactions in their research, with a particular emphasis on revealing links between input and output transactions. Their work highlighted the importance of tracing funds through the mixing process and identified addresses that held significant amounts of bitcoin. This approach provided valuable insights into the potential for deanonymization of mixed transactions and the concentration of funds within the mixing ecosystem. Without using test transactions, Gong et al. (2023) analyzed peeling chains in the blockchain data and focused on common patterns like the peeling amount and peeling percentage that are likely attributed to mixing services.

3. Analysis of current bitcoin mixing landscape

3.1. Currently available bitcoin mixing services

To gain a comprehensive understanding of the current Bitcoin mixing service landscape, we conducted a search across popular cryptocurrency forums. Our primary sources were the Bitcointalk forum and Reddit, where we used the keywords *B. mixer* and *Bitcoin tumbler* to identify relevant discussions and service mentions. For services mentioned without specific URLs or Tor addresses, we employed search engines to locate their online presence. It is important to note that the information presented here is based on forum posts and the websites of the mixing services themselves, and may not be entirely accurate or up-to-date. We also noted a significant number of fraudulent activities in this space, including website clones with similar domain names and imitation communication channels (e.g., Telegram) designed to deceive users. Distinguishing between legitimate and cloned services is further complicated by some genuine mixers operating multiple domains for redundancy. Our search yielded 20 active mixing services. To avoid inadvertently endorsing any particular service, we have opted not to disclose their names in this paper.

We observed that none of the identified services required user registration or identity verification. Only one mixer allowed multiple input addresses for sending bitcoin to the service. 17 out of 20 services permitted output payments to at least two separate addresses. Ten services offered customizable delay options for output payments, which affects the anonymity set of input transactions. Advertised delay ranges for the output payments varied from immediate to 168 h, with nine services offering delays of 8 h or less, and eight offering delays exceeding 24 h 19 services maintained a clearnet domain directly accessible via the internet, with 13 utilizing Cloudflare's reverse proxy service to obfuscate their direct address. All but one service provided an Onion Service on the Tor network, enhancing user anonymity and concealing server locations. 15 services offered a *letter of guarantee* or *signed warranty*, as described by Bonneau et al. (2014), allowing users to publicly expose non-compliant services.

3.2. Legal cases involving bitcoin mixing services

We reviewed relevant U.S. court cases involving mixing services to gather additional information on their operational methods and the investigative techniques used to identify operators. Our research uncovered three cases in the United States where law enforcement agencies successfully identified operators of Bitcoin mixing services. Notably, in all three cases, while law enforcement conducted test transactions, these did not directly lead to operator identification. The ChipMixer case (U.S. District Court for the Eastern District of Pennsylvania, 2023) provided the most detailed technical information about service operations. Crucially, in all instances, the identification relied on information external to the transactions conducted by the mixing service. ChipMixer, a prominent Bitcoin mixer, derived its name from the *chips* users received post-mixing. According to its announcement on the Bitcointalk forum, ChipMixer created Bitcoin addresses called *chips* and funded them with bitcoins in denominations ranging from 0.001 to 4.096 BTC (ChipMixer, 2017). This approach, utilizing 0.001 BTC multiplied by powers of 2, facilitated the merging and splitting of chips. The service advertised pre-funded chips to ensure no link between incoming and outgoing transactions can be established. Users could further obfuscate the origin of their bitcoins by donating, merging, and splitting chips manually on the platform. The key breakthrough in this case was the FBI's identification of the IP address of one of ChipMixer's Tor Onion Service servers, leading to the tracing of the server and subsequent acquisition of user account details, ultimately revealing the operator's identity. The second Bitcoin mixer, Helix (U.S. District Court for the District of Columbia, 2019), advertised its ability to conceal transactions from law enforcement by providing customers with new bitcoins unlinked to the darknet and employing new addresses for each transaction. Helix partnered with the darknet marketplace *AlphaBay* to offer Bitcoin mixing services to AlphaBay customers. While specific technical details of Helix's operation were not publicly disclosed, evidence suggests that the operator's identification was based on information external to the Bitcoin blockchain. In the third case, Bitcoin Fog, announced in 2011, required users to register accounts and promised payouts from addresses different from those used for deposits. To enhance anonymity, the service claimed to delete logs after one week and charged variable fees between 1 % and 3 %. The administrator of Bitcoin Fog was identified by tracing bitcoins used to pay for the hosting service of *bitcoinfog.com*. Although Bitcoin tracing techniques were employed, the crucial information stemmed from the service's infrastructure rather than its mixing operations (U.S. District Court for the District of Columbia, 2021). In all three cases, law enforcement conducted test transactions, but these did not directly lead to the identification of the operators. The successful identifications were primarily based on information external to the Bitcoin blockchain, highlighting the importance of traditional investigative techniques in combating cryptocurrency-based money laundering operations. These findings underscore the complexity of investigating Bitcoin mixing services and the need for a multifaceted approach that combines blockchain analysis with conventional investigative methods.

4. Bitcoin mixer analysis setup

To analyze Bitcoin mixing services effectively, we required a comprehensive setup to trace and cluster blockchain transactions. Our initial approach considered using BlockSci, an open-source blockchain analysis platform (Kalodner et al., 2020). However, BlockSci had not been updated since 2020, potentially limiting its ability to process more recent blockchain data and transaction types. Given these limitations, we opted for a more flexible solution and utilized a tool called btc-csv, developed by Sommer (2019). This tool allowed us to extract Bitcoin blockchain data and format it for importing it into the Neo4j graph database. Neo4j is a graph based database that allows to generate queries to analyze complex transaction patterns (Neo4j, Inc., 2024). The btc-csv tool processes raw blockchain data and generates CSV files containing information about transactions, addresses, and their relationships. After cleaning up duplicates and incomplete data, the CSV files were then imported into Neo4j using its bulk import feature, which significantly accelerates the data ingestion process compared to individual transaction insertions. However, we encountered a challenge with btc-csv: it uses an older library for decoding Bitcoin addresses that was last updated in 2020. As a result, the parser was unable to process Taproot transactions, which were introduced in the Bitcoin protocol more recently. To address this limitation, we modified the btc-csv program and manually ensured that the peeling chains from the mixers include all addresses and transactions.

Overall, we imported more than 2,000,000,000 nodes representing addresses, blocks and transactions. Fig. 1 shows the data model of the imported data from btc-csv. Blue nodes represent mined blocks, purple nodes Bitcoin addresses and orange nodes transactions. Addresses are linked through transactions that receive and send bitcoins from connected addresses. Blocks contain the transactions and are linked the predecessor, however transactions linked to the same block are not linked semantically, but are only confirmed in the same block. For easier visualization, we only display the ID numbers that Neo4J assigned to the nodes automatically. During our analysis, we used the full addresses. For conducting the actual analysis on the imported blockchain data, we used Neo4j's integrated web interface with the related Cypher query language. Our initial step was to find the nodes for input and output payment addresses and then expand the graph into the comprehensive transaction graph. For visualization purposes, we trimmed parts of the generated Figure to focus on the relevant parts.

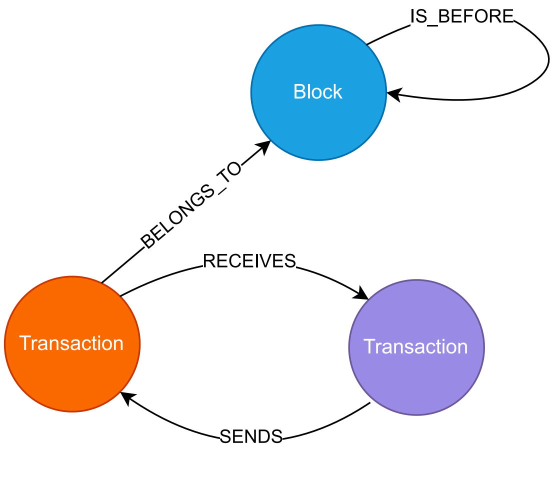

Fig. 1. Data schema showing all node and relationship types.

5. Methodology

To gain insights into the operational patterns of current Bitcoin mixing services and identify potential forensic artifacts, we conducted a series of test transactions and subsequent blockchain analyses. We selected two representative services from the market survey in subsection 3.1 for detailed investigation based on their features and user accessibility. We initiated our test transactions by using a cryptocurrency exchange wallet to send approximately 0.004 BTC to the first mixer, collecting the output from a single address. Subsequently, we reused this output address to send the remaining funds to the second mixer, specifying two output addresses for this transaction. For both mixers, we utilized the default fee and delay settings provided on their respective websites. The delays were 24 h or less. To mitigate potential side effects from wallet software, we used separate wallet software for input and output addresses throughout all transactions. Following the completion of these transactions, we performed a detailed analysis of the blockchain data. We developed a Cypher query, the query language in Neo4j, to determine the shortest path between input and output addresses, allowing us to identify potential links between them. To account for the delay mechanisms employed by mixing services, we conducted our primary analysis on the output address 24 h after the initial transactions, as the selected services indicated delay ranges below this duration. We periodically observed the input address and started analysing it when we noticed significant movement. Moreover, we simulated an attack scenario where an adversary possesses partial information about a user's transaction like timing and amount and constructed additional Cypher queries to calculate the number of candidate addresses by counting transactions occurring between specified block heights within defined value ranges. The Cypher query used for transaction value analysis is detailed in Appendix A. This approach enabled us to explore how varying transaction amounts could influence the likelihood of linking input and output addresses. To enhance our analysis further, we utilized two external anti-money laundering services: CrystalBlockchain (Crystal Blockchain B.V., 2024) and AMLBot (Safelement Limited, 2024). CrystalBlockchain was selected for its current daily allowance of 15 free checks and its presence in literature (Makarov and Schoar, 2021). AMLBot served as a secondary service to cross-verify results from CrystalBlockchain as they do not provide detailed information about their underlying data sources or methodologies. Both services assign risk scores ranging from 0 % to 100 % to addresses, aiding in the detection of suspicious activities. Our overall analysis focused on identifying distinct patterns that could potentially be used to trace mixing transactions and understand the activities of mixing service operators. Apart from the presented results, we also checked for other patterns like CoinJoin transactions that would be easy to spot as the input and output payments are of equal size, but we didn't find any patterns in the immediate vicinity of the input and output addresses beyond the presented results.

6. Mixer 1

Mixer 1 only allowed us to select one output address and the mixer claimed to not use cryptocurrency exchanges for mixing. We used the default fee settings and received a signed letter of guarantee. We couldn't verify the public key linked to a Bitcoin address published on their website because the verification was invalid. We analyzed the graph by loading the input address IN of the mixer and the output address OUT and expanded it as shown in Fig. 2. The input and output addresses are visibly not linked as there is no direct connection in the mixing period defined as the time between input payment and output payment. A transaction with two input addresses B and C sends bitcoins to two output addresses E and F. From E the direct payment to the output address OUT is made. C is part of a preceding peeling chain starting from address A and D is an address with multiple transactions.

6.1. *Peeling chain*

The peeling chain between addresses A and C in Fig. 2 starts with two transactions sending bitcoins to A 1509 and 1534 blocks before the input address was paid. The subsequent peeling chain transactions follow no obvious timing periods as the minimum time between transactions is three blocks and the maximum is 918 with a standard deviation of 339.7 blocks. An additional multi-input address is linked to this peeling chain via the second transaction and the second output address below C. This address is not included in the picture to provide a more focused overview, as this multi-transaction address only receives bitcoins from two addresses on the peeling chain but does not send any bitcoins to the peeling chain. All five addresses forming the peeling chain appear the first time in the blockchain and start with *bc1q* while only three out of the five peeled off addresses start with *bc1q*, one starts with *1* and the other with *3*. All appear to belong to the participants in the mix. This understanding is derived from the fact that the addresses forming the horizontal line feed into a transaction with two input (B and C) and two output addresses (E and F), from which ultimately a payment is made to the output address M.

6.2. *Multi-transactions addresses and multi-address transactions*

Address D is part of 40 transactions over multiple blocks and sends bitcoins to Address B which, with one intermediate hop, feeds the output address OUT. A total of 14 addresses are sending bitcoins to transaction G. This includes the address F which received bitcoins from the mixer's peeling chain. The two addresses H and I receive the bitcoins from transaction G, while the other addresses are sending bitcoins. This pattern suggests that the mixer is pooling bitcoins received from its users in the transaction G and mixes them in the address D.

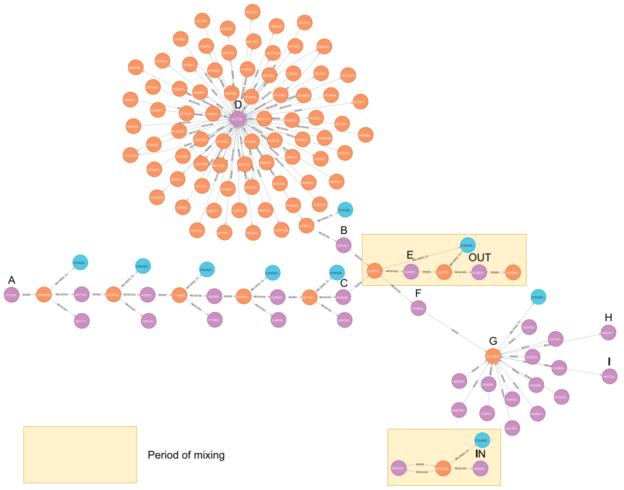

Fig. 2. Transaction graph showing relevant activities from Mixer 1 for the output address OUT.

6.3. *Input and output linkage analysis*

To test whether the anonymization of Mixer 1 was successful, the graph is checked for a path between the input and output address. This path does not exist. Even two months after the input transaction, the input address still held the bitcoins. Therefore, no link can be established between the input and output address. It therefore can be concluded that the mixer for this test transaction successfully prevents a direct linking of the bitcoins sent and received through analysis of the blockchain.

6.4. *Transaction value analysis*

With our Cypher query, we searched for transactions in the following 24 h after the input transactions with a value of the input amount minus the mixing fees. It returned two transactions of which one was the actual transaction to the output address OUT. This shows that the target group of transaction is easily traceable for third parties. Running the same query applying a range of output values based on the range of possible fees applied by the mixer results in a transaction count of 4453. This represents the range of potential output values if the exact fee level applied is not known. Running the query with a value range between 0 and the expected amount based on the minimum fee advertised by the mixer returns 461,140 transactions. These transactions represent potential outputs that combined could form two payouts by the mixer. This population would need to be further analyzed for inputs that result in exact matches for the payout. Nevertheless, this analysis illustrates how much more complicated the linking of transactions or addresses based on input and output values becomes when two output addresses or separate transactions are used instead of one.

6.5. *Taint analysis*

Table 1 shows the results of the taint analysis of key addresses. Both services correctly identify the cryptocurrency exchange used to buy the bitcoins for our test transaction. The taint results are very similar for both services and no owner is determined for any address in the transaction graph. This seems to confirm the statements from Mixer 1 to not use cryptocurrency exchanges during the mixing process. The bitcoins received by the mixer show a higher taint percentage than the bitcoin sent to the mixer but still at a low percentage of 29 and 30 %. Even address D, involved in 40 transactions, and transaction G, which presumably pools bitcoin from users, have low taint ratings.

7. Mixer 2

Mixer 2 offers up to two output addresses. The mixer claims to source bitcoins for payout from cryptocurrency exchanges. The mixing is initiated by submitting an order via its website on the Tor network. We used the default fee settings and two output addresses. To test Mixer 2, we used the output address from Mixer 1 to send approximately 0.00038 bitcoin into Mixer 2. We received a letter of guarantee signed with a PGP key. The referenced PGP key was not provided on the website or key servers and therefore, we could also not validate it. The output was

Table 1
Taint results for selected Mixer 1 addresses.

Address	AMLBot		CrystalBlockchain	
	Taint	Owner	Taint	Owner
B	40 %	not defined	40 %	not defined
C	40 %	not defined	40 %	not defined
IN	10 %	correctly identified	10 %	correctly identified
D	30 %	not defined	29 %	not defined
OUT	30 %	not defined	29 %	not defined

received 15 and 20 blocks later.

Fig. 3 shows the expanded graph in Neo4j. The input address IN is on the right side, below the two output addresses OUT1 and OUT2. Two peeling chains starting from the addresses N and H are visible with 8 and 12 peeling steps. The input and output address are visibly not linked as there is no direct connection in the mixing period, defined as the time between input payment and output payment.

7.1. Peeling chains

The first peeling chain, starting with address A, consists of eight peeling transactions. The Addresses linking the transactions on the peeling chain consistently start with the number *3*. Six of the eight addresses to which bitcoins are peeled off start with *bc1q*. The other 2 start with the number *1*. Aside from a consistent use of the same type of address along the peeling chain, this understanding of control is also based on the fact that the chain ultimately sends bitcoins back to the first output address provided to the mixer. The second peeling chain has twelve peeling transactions beginning from address H. Again, the addresses leading to the OUT2 address all start with *3*. Eight of the twelve addresses to which bitcoins are sent from the peeling chain start with *bc1*, three addresses start with the number *1* and one address starts with the number *3*. The timing between the transactions in the peeling chains seems to follow a more predictable pattern. The minimal delay is 4 blocks and the maximum delay is 118 blocks, while the average delay for the first and second peeling chain is around 58 blocks and the standard deviation is 34.7 and 32.8 blocks. While the number of data points is relatively low with only 20 transactions overall in the peeling chains, the consistent values could hint to default timing values that the mixing service uses.

7.2. Multi-transactions addresses

Address A is linked to multiple transactions and sends bitcoins to the first peeling chain. It sends bitcoins to and receives bitcoins from overall 47 transactions. Six multi-transaction addresses B, C, D, E, F and G send bitcoins to the second peeling chain in one common transaction. Overall, seven addresses are included in this transaction as address B is included twice. Also, address D is also an output address of this transaction. Table 2 shows these addresses and their respective number of receiving and sending transactions. Addresses B and D show a significantly higher number of transactions compared to the rest, this could hint to an online service with a shared wallet for multiple users like a cryptocurrency exchange.

7.3. Input and output linkage analysis

To test whether the anonymization of Mixer 2 was successful, the graph is checked for a path between the input and output address. During the period of mixing as shown in the rectangle in Fig. 3, there is no direct link between IN and OUT1 or OUT2. Therefore, the mixer prevents a direct linking of the sent and received bitcoins. Applying our Cypher query to calculate the shortest path between the input and output addresses reveals a connection between them. While the visualization shows that during the mixing period is no direct connection between them, the mixer reuses addresses that are already connected from previous transactions. This can provide insights into how bitcoins received from users are subsequently used. We ran the Cypher query for the output addresses OUT1 and OUT2 separately and extracted the

Table 2
Number of transactions linked to selected multi-transaction addresses.

Address	Number of transactions	
	Receiving	Sending
A	47	47
B	9712	9649
C	15	14
D	6153	6114
E	25	25
F	26	26
G	27	27

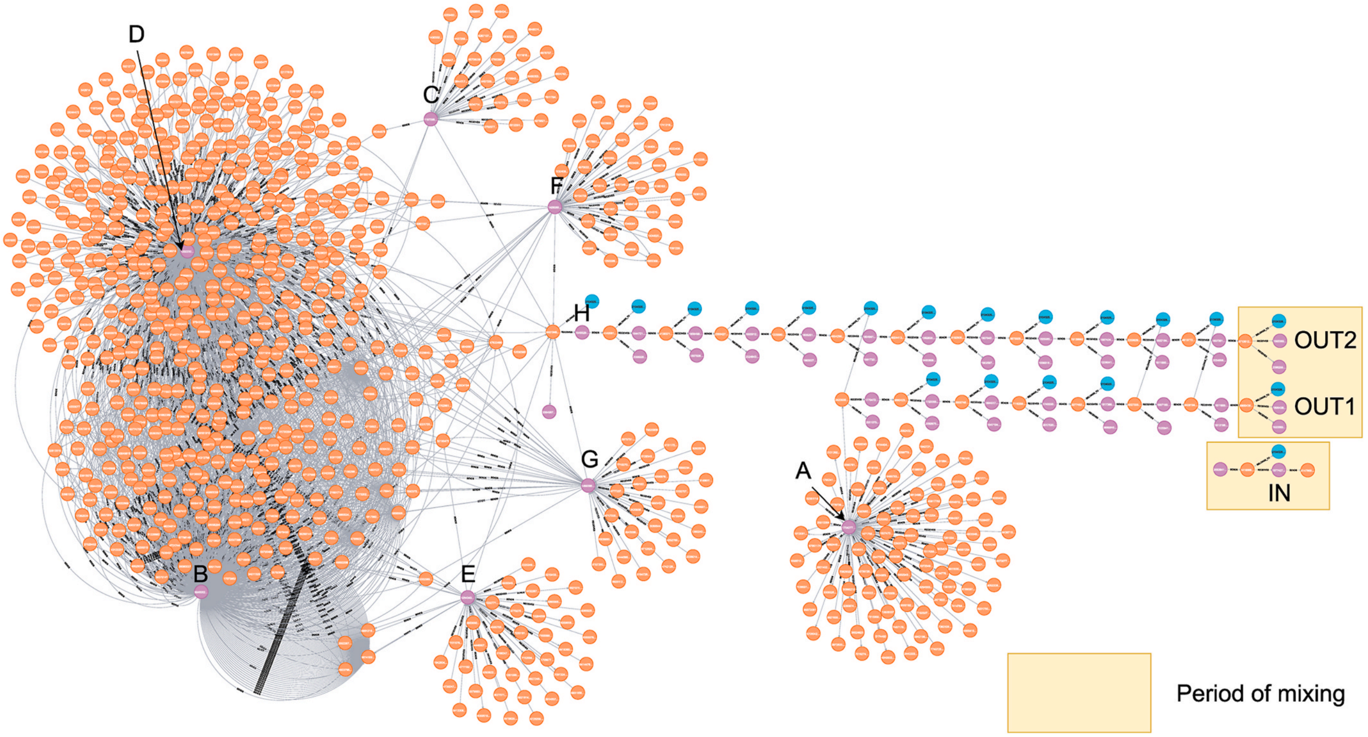

Fig. 3. Transaction graph showing relevant activities from Mixer 2 for the output addresses OUT1 and OUT2.

common subgraph shown in Fig. 4.

The transaction in Fig. 4 form a chain of subsequent transactions, enabling a clear tracing of the bitcoins. From the input address IN and 94 other addresses, 7.7 bitcoins are pooled in the transaction P. P sends 5 bitcoins to address Q and 2.7 bitcoins to address R. R sends, together with two other addresses, 9.7 bitcoins to a transactions that sends a part consisting of 5 bitcoins to the address Q. Two blocks later two addresses send 3.7 and 3.8 to an transaction that sends 4.9 bitcoins to address Q. Then all bitcoins are moved to address S and then subsequently to address D. Afterwards, the link to the output addresses is based on transactions in previous blocks. While it is hard to confirm any hypothesis without multiple test transactions over a prolonged time, it seems likely that transaction P pools bitcoins sent from its users. Also, the mixer seems to use a number of storage addresses like S and D to mix the bitcoin and reuses them.

7.4. Transaction value analysis

Again, we searched in the blockchain data in the following 24 h for transactions within the expected returned value range. There was no transaction with the combined value of the output transactions. There are 150 transactions returning the same value that was received to output address OUT1 and three transactions with the same value received at output address OUT2. Running the same query applying a range of output values assuming a single output payment based on the range of possible fees applied by the mixer results in a transaction count of 1422. Running the query with a value range between 0 and the expected amount based on the minimum fee advertised by the mixer returns 391,998 transactions. This again illustrates the benefit for anonymity of using two output addresses compared to only one.

7.5. Taint analysis

Table 3 shows the result of the taint analysis. Looking at the addresses A to G in Fig. 3, both services provide similar results that are all rated as low risk transactions with a 10 % taint value. The owners of these addresses that mix the bitcoins from different sources are attributed to the HTX cryptocurrency exchange, which indicates that participants use their accounts to pay out proceeds to the customers of the

Table 3
Taint results for selected Mixer 2 addresses.

Label	AMLBot		CrystalBlockchain	
	Taint	Owner	Taint	Owner
IN	blacklisted	not defined	26 %	not defined
OUT1	23 %	correctly identified	25 %	correctly identified
OUT2	16 %	not defined	16 %	not defined
R	73 %	not defined	30 %	not defined
Q	blacklisted	not defined	17 %	not defined
A	10 %	HTX	10 %	HTX
B	10 %	HTX	10 %	HTX
C	10 %	HTX	10 %	HTX
D	10 %	HTX	10 %	HTX
E	10 %	HTX	10 %	HTX
F	10 %	HTX	10 %	HTX
G	10 %	HTX	10 %	HTX
S	10 %	HTX	10 %	HTX

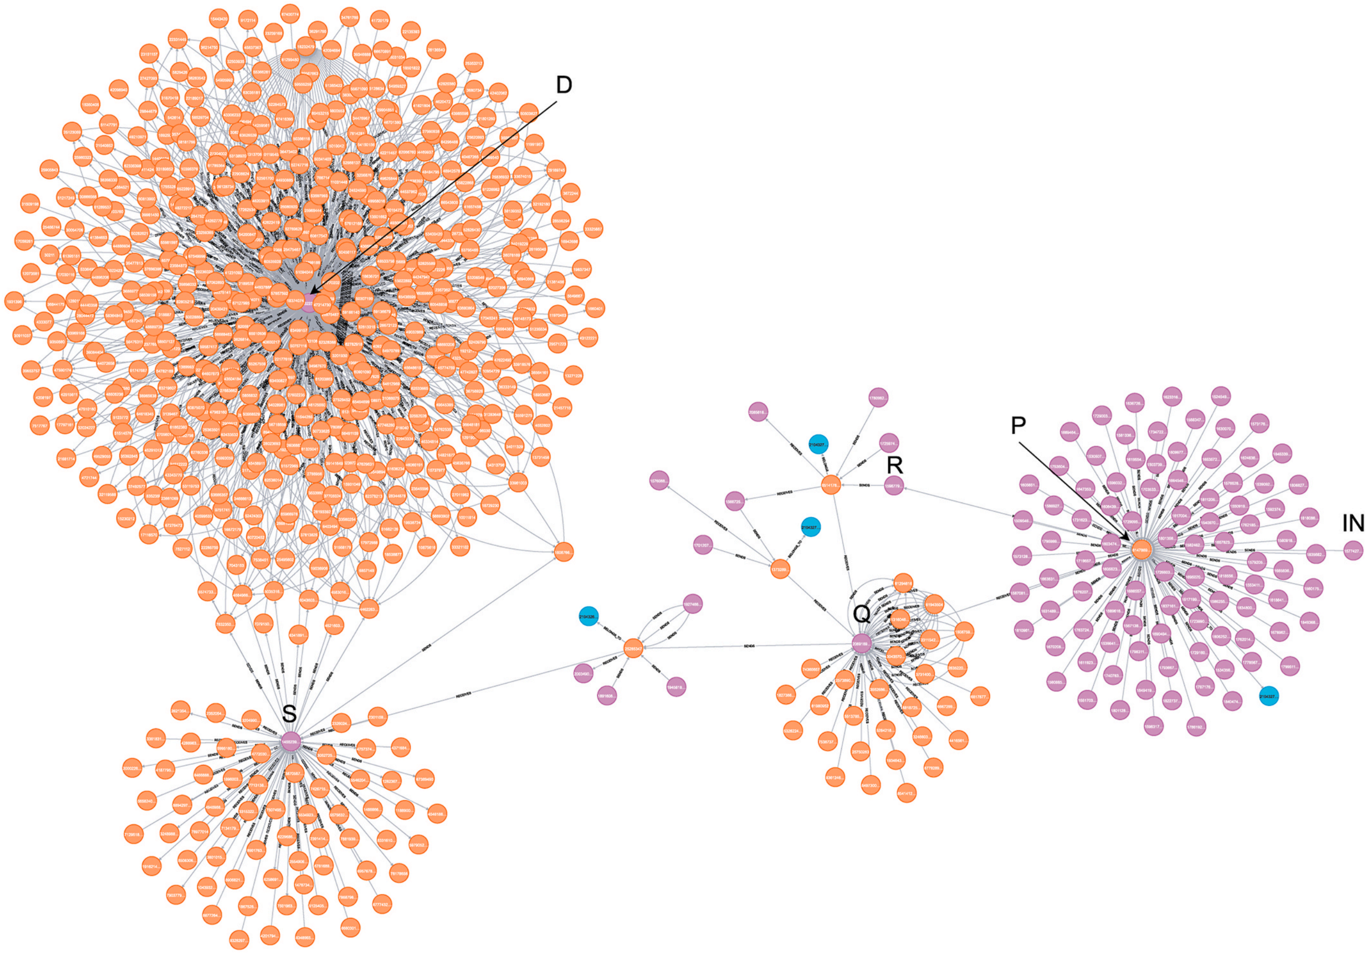

Fig. 4. Transaction graph showing relevant activities from Mixer 2 for the input address IN.

mixing service. The payout addresses OUT1 and OUT2 receive a lower taint rating (between 16 % and 25 %) than the bitcoins originally sent to the mixer that received a rating of 29 % and 30 % (see Table 1). However, the bitcoins received in return are more tainted than the Bitcoin addresses of the cryptocurrency exchange used to source bitcoins for paying back users. This indicates that the services use internal heuristics that might recognize the peeling chain pattern between the addresses A and OUT2, and H and OUT2. The input related addresses in Fig. 4 on the other hand show discrepancy in the risk scoring, the input address IN and address Q are even blacklisted in one service. Also, the address R receives a high rating of 73 % from one service. It is likely that other users of the mix used its services with bitcoins from problematic sources and therefore, tainted the input address IN. While the blacklisting carries over the address R, the other connected address Q is not blacklisted but received a high taint rating. This difference might be due to the different number of transactions with various taint levels that paid into the addresses.

8. Patterns for the analyzed mixing services

Both mixing services use peeling chains to send the bitcoins to the payout address and the transactions leading to the payout addresses all use a consistent type of address. We argue that these analyzed peeling chains are controlled by the mixing service as it's unreasonable for someone else to make the payment to the output address on behalf of the mixer, unless the mixer is able to swap such payments with other participants in an organized and automated manner. The time and cost to find willing participants for the target amount would likely be too high. As normal transactions between users exhibit the same one input to two output addresses pattern, different wallet implementations might use different Bitcoin address types. Therefore, this criteria narrows down the number of candidate peeling chains when searching through the blockchain. The generated graphs for both mixers in Figs. 2 and 3 also vary in complexity, indicating that Mixer 2 uses a more sophisticated way of mixing the proceeds. This aligns with the fact that Mixer 2 involves addresses related to cryptocurrency exchanges and by default uses two output addresses.

8.1. Mixer 1

Fig. 2 shows that the mixing service sourced bitcoins from address J that is involved in 40 transactions. The owner of this address could not be identified by AMLBot or CrystalBlockchain. As this address is used for a prolonged time, it might be connected to an external service that the mixing service uses but the low number of connections rather indicates that this address is directly operated by the mixing service. Performing additional test transactions on Mixer 1 might show whether additional links can be traced to this address from other payouts. The output pattern shows an interesting anomaly as the transaction merging payments from the addresses F and G does not directly send bitcoin to the output address, but uses an intermediary address L that in the same block sends the value to the output address. This suggests that these transactions are performed by the same user or entity. Manually coordinating two transactions to enter one block is rather challenging given the time target of about 10 min to mine a block. The first of the two transactions must be validated before the second transaction can be validated as otherwise the UTXO used in the second transaction would not yet exist. Thus, at first sight this extra step looks like it may be a characteristic of Mixer 1. The property of offering one output address makes this mixing service very susceptible to the transaction value analysis in subsection 6.4. The fact that the input address didn't conduct any transaction that we could follow up on, prevents an analysis of input patterns.

8.2. Mixer 2

Mixer 2 sources its bitcoins for payout to participants from the cryptocurrency exchange HTX. One address attributed to HTX sends bitcoins to peeling chain 1 while six addresses are sending bitcoins to peeling chain 2 in a single transaction. At first sight, this looks like a distinctive characteristic. However, as the sending addresses belong to the cryptocurrency exchange HTX and not Mixer 2 it is rather likely that this is a pattern of HTX making payouts to its customers rather than Mixer 2. 1498 blocks after the payment to Mixer 2 for the test transaction, the mixer pools bitcoins in the transaction P with 95 input addresses and 2 output addresses Q and R. Over multiple addresses, the bitcoins are sent to the cryptocurrency exchange, which should allow investigators to determine the identified owner of the associated accounts and continue to trace the flow of bitcoins by checking further transactions from the accounts. The timing patterns for the peeling chains could help to detect peeling chains that are likely connected to this mixing service. In combination with known input amounts, the transaction value analysis can narrow down the potential number of output address.

8.3. Investigative angles

The identified patterns can help to identify other users, especially for Mixer 1 as it uses mixing addresses that are not related to a cryptocurrency exchange or other known organization. Associated accounts at cryptocurrency exchanges that pay into peeling chains for Mixer 2 can help to identify the operator but also other users as it seems rather unlikely that the mixing service will have a separate account for each new mixing transaction and the address reuse also hints to this. Quickly blacklisting addresses associated with mixing services might not identify the operator or its users, but could potentially incur significant damage to the operation, eventually making it unprofitable. This should be implemented in cryptocurrency exchanges and potentially even in wallet software to protect users from unknowingly receiving blacklisted bitcoins. If a Bitcoin mixer is connected to an illegal marketplace, investigators might be able to determine the time and amount for a limited number of transactions by either conducting test transactions and tracking where the money appears linking to the seller, or by using provided information like reviews or public order status information. In this case, our simulated attack scenario could be practically applied by investigators. Apart from this, the fact that many mixing services surveyed in section 3.1 use clearnet addresses would allow law enforcement agencies to send lawful interception requests to wiretap the traffic and image the server. The first helps to identify potential users of this mixing service and the latter allows to search for further hints to penetrate the infrastructure and identify the operator. With a higher risk of being detected as the certificate changes, it is possible to force the Certification Authority to first release subscriber information and then issue a new TLS certificate that allows investigators to set up a Man-in-the-Middle proxy that can see the entire decrypted traffic which links input and output addresses clearly. This attack is described in a blog post about an alleged attack on a XMPP instance (ValdikSS, 2023).

9. Discussion

Our analysis of Bitcoin mixing services reveals several key findings that contribute to the understanding of their operational patterns and the challenges they pose to forensic investigations. The test transactions we conducted demonstrated distinct characteristics of the mixing services, aligning with previous research by Möser et al. (2013) and Wu et al. (2021). We observed the use of peeling chains and multi-input transactions, confirming the findings of Gong et al. (2023). However, our analysis using Neo4j provided a more granular view of these patterns, allowing for better visualization and potential identification of mixer-specific behaviors. In contrast to previous studies that primarily focused on determining whether mixers employed swapping, obfuscating, or general mechanisms, our research concentrated on specific transaction patterns and their potential for deanonymization. Our contributions demonstrate how concrete patterns for selected mixers can be derived to enhance blockchain analysis by narrowing down the search space for potentially suspicious transactions. The resulting queries can identify and isolate transaction patterns indicative of specific mixer activity, highlighting distinctive patterns such as sending bitcoins to output addresses over two intermediary addresses within one block. By applying these queries, investigators can efficiently flag transactions that warrant further scrutiny, thereby reducing the overall complexity and volume of data that needs to be analyzed. For tracing individuals, transaction value analysis is also a suitable technique to unmix selected mixing transactions in certain cases. By focusing on these narrowed-down transaction sets, investigators can leverage additional information such as off-chain data or traditional investigative techniques to significantly enhance the accuracy and effectiveness of their analyses. This approach bridges the gap between theoretical models of mixer operations and practical forensic techniques. The examination of U.S. legal cases involving Bitcoin mixers together with the large analyzed transaction graphs revealed that while blockchain analysis plays a crucial role, the identification of operators often relies on traditional investigative techniques and off-chain information. This underscores the importance of combining blockchain analysis with conventional law enforcement methods. The cases of ChipMixer, Helix, and Bitcoin Fog demonstrate that mixer operators can be identified through various means, including tracing infrastructure payments and exploiting operational security mistakes. These findings suggest that while mixers can effectively obfuscate individual transactions, they may still leave traces that can be exploited by law enforcement. Also, as one mixer relies on cryptocurrency exchanges, assisting organizations in detecting and denying questionable transactions mitigates the impact of the underground economy on legitimate businesses.

9.1. Limitations and future research

While our test transactions provided valuable insights, they represent only a small sample of mixer operations. Law enforcement agencies would need to conduct more extensive testing to reliably approximate how mixing services operate. However, this approach faces several challenges as mixers can adapt their techniques rapidly, potentially requiring investigators to analyze a significant portion of all incoming transactions to cover all potential methods. Our reliance on open-source tools limited our ability to handle recent blockchain updates, potentially missing nodes in multi-input transactions or multi-transaction addresses. Additionally, the evolving nature of the Bitcoin protocol and user behaviors (e.g., new wallet software) can disrupt previously reliable identification patterns. Therefore, future research should focus on developing more robust and adaptable analysis tools that can keep pace with blockchain updates and mixer innovations. Using external sources to identify address owners can help to uncover how the mixer operates and provide valuable information about the operators.

9.2. Ethical considerations

While the operation of Bitcoin mixing services is questionable as these organizations typically operate without licenses or company registrations to protect themselves from law enforcement agencies or other actors, their usage is not inherently unethical. Legitimate users may seek enhanced privacy for various reasons, including protection from authoritarian governments when making donations to NGOs. Our research, including the test transactions, was conducted with ethical considerations in mind. We used legally acquired bitcoins from a cryptocurrency exchange and limited our transactions to minimize impact on the mixer ecosystem. To protect user anonymity, we anonymized all node identifiers in our graph analysis and refrained from specifying exact transaction amounts or block heights.

10. Conclusion

This study analyzes Bitcoin mixing services, their operational patterns, and the challenges they pose to forensic investigations. Our research includes a survey of 20 currently available mixing services and an analysis of three U.S. legal cases for investigative techniques and operation details. Using Neo4j for blockchain data analysis and conducting test transactions, we identified unique transaction patterns associated with two specific mixers, including peeling chains and multi-input transactions. Our simulations demonstrated how partial transaction knowledge could be leveraged to trace funds through mixers, highlighting that while these services significantly obfuscate transaction trails, certain patterns and behaviors can still be exploited for forensic analysis. The examination of legal cases underscores the importance of combining blockchain analysis with traditional investigative techniques, particularly off-chain attacks and methods for associating addresses with entities. As cryptocurrency adoption and its use in cybercrime continue to grow, the need for advanced forensic tools and methods becomes increasingly crucial. Our research contributes to the field of cryptocurrency forensics by offering insights into mixer operations and potential avenues for improving traceability. We discuss the limitations of current approaches and propose potential improvements that can aid investigators in applying effective techniques. Future work should focus on developing quick techniques to identify addresses associated with mixers while considering the ethical implications of such advancements. Ultimately, this research aims to assist law enforcement agencies in developing more effective strategies to tackle the challenges posed by Bitcoin mixers in cybercrime investigations.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Cypher Query for transaction value analysis

Listing 1: Count the number of transactions between two block heights transferring bitcoins within a defined value range

```
MATCH            (b:Block)<−[:BELONGS_TO]−
                 (t:Transaction)−[r:RECEIVES]
                 −>(a:Address)
WHERE            r.value <= <max_value>
AND              r.value >= <min_value>
AND              b.height >= <start_block>
AND              b.height <= <end_block>
RETURN COUNT(r)
```

This Cypher query is designed to quantify the number of transactions within a specified block height and value range. It focuses on identifying potential output addresses used by a mixer by counting transaction values received. By replacing the last line with ‘RETURN a‘, the query can instead return the addresses receiving the specified transaction value.

An analyst could use this query to identify addresses that receive amounts similar to those sent to a mixer address, minus any fees, within a defined time range. This population of addresses can then be further analyzed to establish potential links between mixing input and output addresses.

References

Androulaki, E., Karame, G.O., Roeschlin, M., Scherer, T., Capkun, S., 2013. Evaluating user privacy in bitcoin. In: Financial Cryptography and Data Security. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 34–51.

de Balthasar, T., Hernandez-Castro, J., 2017. An analysis of bitcoin laundry services. In: Secure IT Systems. Springer International Publishing, Cham, pp. 297–312.

Bonneau, J., Narayanan, A., Miller, A., Clark, J., Kroll, J.A., Felten, E.W., 2014. Mixcoin: anonymity for bitcoin with accountable mixes. In: Financial Cryptography and Data Security. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 486–504.

ChipMixer, 2017. ChipMixer.com - bitcoin mixer/Bitcoin tumbler - mixing reinvented. https://bitcointalk.org/index.php?topic=1935098.0. (Accessed 2 October 2024).

Crystal Blockchain, B.V., 2024. Blockchain intelligence & crypto compliance platform. https://crystalintelligence.com/. (Accessed 2 October 2024).

Deuber, D., Ronge, V., Rückert, C., 2022. Sok: assumptions underlying cryptocurrency deanonymizations. Proc. Priv. Enhancing Technol. 2022, 670–691. https://doi.org/10.56553/POPETS-2022-0091.

Gong, Y., Chow, K.P., Yiu, S.M., Ting, H.F., 2023. Analyzing the peeling chain patterns on the bitcoin blockchain. Forensic Sci. Int.: Digit. Invest. 46, 301614. https://doi.org/10.1016/j.fsidi.2023.301614.

Kalodner, H., Möser, M., Lee, K., Goldfeder, S., Plattner, M., Chator, A., Narayanan, A., 2020. BlockSci: design and applications of a blockchain analysis platform. In: 29th USENIX Security Symposium (USENIX Security 20). USENIX Association, pp. 2721–2738. https://www.usenix.org/conference/usenixsecurity20/presentation/kalodner.

Makarov, I., Schoar, A., 2021. Blockchain Analysis of the Bitcoin Market. National Bureau of Economic Research. https://doi.org/10.3386/w29396. Working Paper 29396.

Maxwell, G., 2013. Coinjoin: bitcoin privacy for the real world. https://bitcointalk.org/index.php?topic=279249. (Accessed 20 September 2024).

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G.M., Savage, S., 2013. A fistful of bitcoins: characterizing payments among men with no names. In: Proceedings of the 2013 Conference on Internet Measurement Conference. Association for Computing Machinery, New York, NY, USA, pp. 127–140. https://doi.org/10.1145/2504730.2504747.

Möser, M., Böhme, R., Breuker, D., 2013. An inquiry into money laundering tools in the bitcoin ecosystem. In: 2013 APWG eCrime Researchers Summit, pp. 1–14. https://doi.org/10.1109/eCRS.2013.6805780.

Nakamoto, S., 2008. Bitcoin: a peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf. (Accessed 20 September 2024).

Neo4j, Inc., 2024. Neo4j graph database & analytics | graph database management system. https://neo4j.com/. (Accessed 2 October 2024).

Pfitzmann, A., Köhntopp, M., 2001. Anonymity, unobservability, and pseudeonymity — a proposal for terminology. In: International Workshop on Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Unobservability. Springer-Verlag, Berlin, Heidelberg, pp. 1–9.

Reid, F., Harrigan, M., 2011. An analysis of anonymity in the bitcoin system. In: 2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing, pp. 1318–1326. https://doi.org/10.1109/PASSAT/SocialCom.2011.79.

Safelement Limited, 2024. Comprehensive AML compliance solutions for crypto. https://amlbot.com/. (Accessed 2 October 2024).

Sommer, D., 2019. Processing bitcoin blockchain data using a big data-specific framework. Bachelor’s thesis. University of Zurich. https://files.ifi.uzh.ch/CSG/staff/scheid/extern/theses/BA-D-Sommer.pdf.

U.S. District Court for the District of Columbia, 2019. United States V. HARMON. 1:19-cr-00395.

U.S. District Court for the District of Columbia, 2021. United States V. Sterlingov. 1:21-cr-00399.

U.S. District Court for the Eastern District of Pennsylvania, 2023. United States v. QUOC NGUYEN. 2, 528, 23-mj.

ValdikSS, 2023. Encrypted traffic interception on Hetzner and Linode targeting the largest Russian XMPP (Jabber) messaging service. URL: https://notes.valdikss.org.ru/jabber.ru-mitm/. (Accessed 2 October 2024).

Wu, L., Hu, Y., Zhou, Y., Wang, H., Luo, X., Wang, Z., Zhang, F., Ren, K., 2021. Towards understanding and demystifying bitcoin mixing services. In: Proceedings of the Web Conference 2021. Association for Computing Machinery, New York, NY, USA, pp. 33–44. https://doi.org/10.1145/3442381.3449880.