CMixing: An Efficient Coin Mixing Platform to Enhance Anonymity in Cryptocurrency Transactions

Wangze Ni
HKUST
Hong Kong SAR, China
wniab@cse.ust.hk

Yiwei Zhao
PolyU
Hong Kong SAR, China
yiweizhao@polyu.edu.hk

Pengze Chen
HKUST
Hong Kong SAR, China
pchenax@cse.ust.hk

Lei Chen
HKUST (GZ) & HKUST
Guangzhou & Hong Kong SAR, China
leichen@cse.ust.hk

Peng Cheng
ECNU
Shanghai, China
pcheng@sei.ecnu.edu.cn

Chen Jason Zhang
PolyU
Hong Kong SAR, China
jason-c.zhang@polyu.edu.hk

ABSTRACT

Coin mixing methods are widely used to enhance anonymity in cryptocurrency transactions by obfuscating the linkages between recipients and senders. Specifically, coin mixing methods combine several users' transactions into a CoinJoin transaction and decompose the original transactions' outputs into a set of decomposed outputs with similar amounts. However, existing methods have two shortcomings. Firstly, CoinJoin transactions lack anonymity guarantees. Secondly, the number of decomposed outputs is not minimized. To tackle these two shortcomings, we develop a platform named CMixing for mixing transactions with anonymity guarantees and minimal fees. For a CoinJoin transaction obtained by CMixing, the probability of adversaries correctly guessing the original output of a decomposed output does not exceed c, where c is a privacy requirement. Thus, the first shortcoming is solved. Additionally, CMixing uses an approximation algorithm to decompose original outputs, which approximately minimizes the number of decomposed outputs. Thus, the second shortcoming is solved. Our demonstration will showcase how users can use CMixing to make CoinJoin transactions. We will also show the fees saved and the level of anonymity achieved using our algorithm.

PVLDB Reference Format:
Wangze Ni, Yiwei Zhao, Pengze Chen, Lei Chen, Peng Cheng, and Chen Jason Zhang. CMixing: An Efficient Coin Mixing Platform to Enhance Anonymity in Cryptocurrency Transactions. PVLDB, 17(12): 4405 - 4408, 2024.
doi:10.14778/3685800.3685886

PVLDB Artifact Availability:
The source code, data, and/or other artifacts have been made available at https://github.com/HelloGreatWorld/CMixing.

1 INTRODUCTION

Blockchains have gained significant attention from both academic [4, 6] and industry [7] due to their remarkable properties, such as transparency. However, transparency is a double-edged sword, as it also brings privacy concerns [4]. Coin mixing, as a promising privacy protection technique [6], is widely adopted in blockchain systems [2] to enhance anonymity in transactions. The monthly usage of coin mixing methods is worth at least \$214.65 M USD [5].

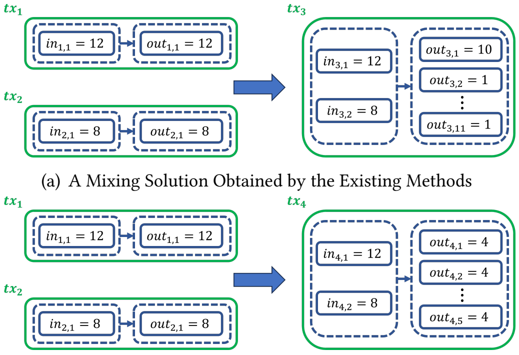

(a) A Mixing Solution Obtained by the Existing Methods

(b) A Better Mixing Solution

Figure 1: Examples of Using a Coin Mixing Method.

Specifically, coin mixing methods combine multiple users' transactions into a CoinJoin transaction, obscuring linkages between inputs and outputs to provide anonymity. Since the outputs of different transactions have varying amounts, for anonymity, these transaction outputs are decomposed into multiple outputs within CoinJoin transactions. We refer to the outputs in the CoinJoin transactions as *decomposed outputs* and the outputs in the original transactions as *original outputs*. The existing coin mixing methods decompose original outputs into decomposed outputs with decimal denominations (e.g., 1, 10) [1]. The receiver of each decomposed output is a different pseudonym. Thus, adversaries cannot infer the linkages between original outputs and decomposed outputs. In Figure 1(a), tx_1 and tx_2 are combined into a CoinJoin transaction tx_3. The output of tx_1 is decomposed into one 10-amount output and two 1-amount outputs. Adversaries cannot accurately distinguish whether $out_{3,2}$ is decomposed from $out_{1,1}$ or $out_{2,1}$.

However, there are two shortcomings associated with existing coin mixing methods. **First**, existing methods fail to ensure the anonymity of the mixing results. For example, in Figure 1(a), adversaries can determine that $out_{3,1}$ is decomposed from $out_{1,1}$, since the amount of $out_{3,1}$ is larger than the amount of $out_{2,1}$. **Second**, existing methods do not effectively minimize the number of outputs

Proceedings of the VLDB Endowment, Vol. 17, No. 12 ISSN 2150-8097.
doi:10.14778/3685800.3685886

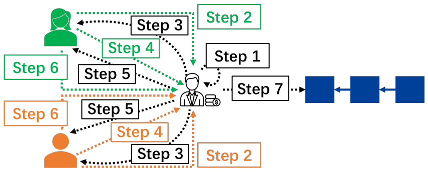

Figure 2: Workflow.

in CoinJoin transactions, resulting in higher transaction fees. For example, there are 11 decomposed outputs in Figure 1(a). Figure 1(b) presents a better mixing solution with only 5 decomposed outputs. Since the transaction fees are proportional to the number of outputs, tx_3 has higher transaction fees than tx_4.

It is challenging to solve these two shortcomings. The transaction data on the blockchain is transparent. Attackers can utilize the transaction data on the blockchain for analysis. Furthermore, attackers may have access to some additional background information [3] (e.g., the original outputs of some decomposed outputs). The background information of attackers are unknown and dynamic. Therefore, it is difficult to design a privacy protection mechanism with an anonymity theoretical guarantee to address the first shortcoming. Additionally, generating a CoinJoin transaction with minimal cost has been proven to be NP-hard [3]. Thus, it is challenging to design an approximate algorithm to quickly obtain a close-optimal CoinJoin transaction to address the second shortcoming.

To address these two shortcomings, we propose a novel platform named CMixing for mixing transactions with anonymity guarantees and minimal fees. CMixing decomposes original outputs by a c-decomposition, which is defined in our prior work [3]. In a c-decomposition, the percentage of x-amount decomposed outputs, originating from the same original output, does not exceed c. In our prior work [3], we theoretically prove the lower bound of the anonymity of a c-decomposition. Thus, **the first shortcoming is solved**. Additionally, CMixing uses an advanced approximation algorithm named Boggart, which was developed in our prior work [3], to obtain a c-decomposition. The number of decomposed outputs in the c-decomposition derived from Boggart is no more than $(\frac{2}{c}+3)$ times the minimum number of decomposed outputs in the optimal solution. Thus, **the second shortcoming is solved**. In summary, our work contributes in two aspects:

- The system provides people with a platform to mix transactions on the blockchain with anonymity guarantees and minimal fees;
- The system is equipped with an advanced approximation algorithm named Boggart, proposed in our prior work [3]. Boggart efficiently makes a mixing solution with a close-optimal cost.

The rest of this paper is organized as follows. § 2 introduce the background. § 3 presents an overview of CMixing. § 4 demonstrates three scenarios for using CMixing. § 5 concludes our paper.

2 BACKGROUND

Coin Mixing. Suppose $OO = \{oo_1, oo_2, \cdots, oo_n\}$ is a set of original outputs, and $DO = \{do_1, do_2, \cdots, do_m\}$ is a set of decomposed outputs from OO. Suppose or_i is the pseudonym that will receive the money in oo_i, dv_i is the amount of do_i, sd_i is the original output that do_i comes from, and dr_i is the pseudonym that will receive the money in do_i. Since the pseudonyms of decomposed outputs are different from the pseudonyms of original outputs, even if attackers know the set of original outputs and their receivers' pseudonyms, attackers cannot know the receivers of decomposed outputs.

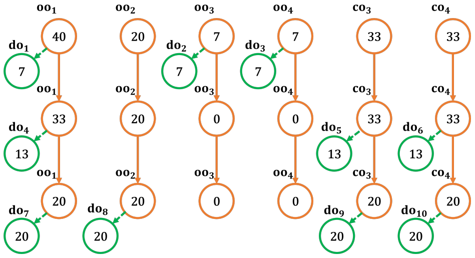

Figure 3: A Running Example of Boggart.

De-anonymization attacks. Attackers can analyze the amount of decomposed outputs to determine which original output the decomposed outputs come from, thereby identifying the receiver of these decomposed outputs and compromising the anonymity of a CoinJoin transaction. As shown in Figure 1(a), attackers know $out_{3,1}$ is decomposed from $out_{1,1}$. Then, the receiver of $out_{3,1}$ must be the receiver of $out_{1,1}$. Thus, the anonymity of tx_3 is compromised.

c-decomposition. In a c-decomposition [3], the percentage of decomposed outputs with an amount of x, originating from the same original output, does not exceed c out of all decomposed outputs with an amount of x, i.e., $\forall do_i \in DO$, $\frac{|\{do_j | dv_j = dv_i, sd_j = sd_i\}|}{|\{do_j | dv_j = dv_i\}|} \leq c$. For example, Figure 1(a) shows a 1-decomposition, since there is only one decomposed output with an amount of 10. In Figure 1(b), the decomposed outputs with an amount of 4 from $out_{1,1}$ constitute 60% of all decomposed outputs with an amount of 4. Thus, Figure 1(b) shows a $\frac{3}{5}$-decomposition. In our prior work [3], we theoretically prove that the probability of attackers correctly guessing the original output of a decomposed output does not exceed c.

3 SYSTEM OVERVIEW

3.1 Workflow

In our system, there are two roles: (1) **Traders** request CoinJoin transactions for anonymity and low costs; (2) **Coordinators** group traders' requests and create CoinJoin transaction plans for each group. Coordinators charge coordination fees based on a commission rate multiplied by the transaction amount [2]. Coordinators may use their own tokens to help traders to make CoinJoin transactions [3]. For instance, when traders' transactions are insufficient to create a CoinJoin transaction that meets the anonymity requirements, coordinators will supplement *compensatory outputs*, transferring tokens from one of their accounts to another.

Figure 2 shows the workflow of CMixing: (1)**Activate coordination services**. The coin mixing service is initiated by a coordinator, who sets the commission rate and the coordination algorithm to assist traders in generating CoinJoin transaction plans. Additionally, the coordinator sets a budget for their own funds allocated to the CoinJoin transaction plans. (2) **Send requests.** A trader fills out a request detailing the amount of an original output and specifying the privacy requirement c, then selects an active coordinator to send the request to. (3) **Make CoinJoin transaction plans.** The coordinator organizes the requests into groups and uses the selected coordination algorithm to create a CoinJoin transaction plan for each group. If the coordinator's funds needed for a CoinJoin transaction plan exceed their budget, the coordinator informs the traders of a failed coordination; otherwise, the coordinator notifies the traders about the generated CoinJoin transaction plan. (4) **Generate Outputs.** Each trader generates the decomposed outputs

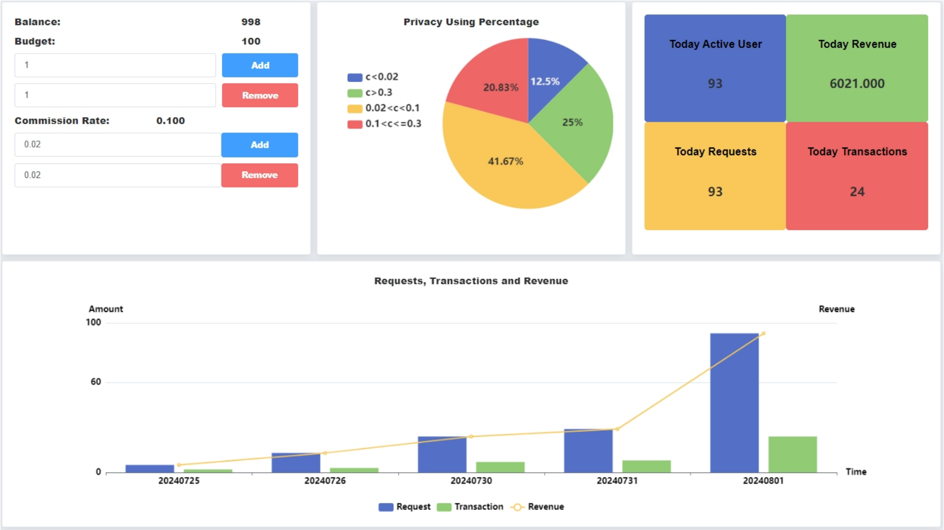

Balance: 998

Budget: 100

1

1

Commission Rate: 0.100

0.02

0.02

Today Active User 93

Today Revenue 6021.000

Today Requests 93

Today Transactions 24

Figure 4: The Coordinator Interface.

according to the CoinJoin transaction plan. Then, each trader submits the outputs and the UTXO will be used in the transaction to the coordinator. (5) **Assemble CoinJoin transactions.** The coordinator puts together the received UTXOs and decomposed outputs from the same CoinJoin transaction plan to create a CoinJoin transaction. Then, the coordinator notifies the traders about the created CoinJoin transaction. (6) **Sign signatures.** After receiving a CoinJoin transaction, if a trader agrees the transaction, she signs the transaction and submit the signed transaction to the coordinator. (7) **Submit CoinJoin transactions.** If a coordinator collects all signatures from the traders in a CoinJoin transaction, they submits this CoinJoin transaction to the system and notifies traders. If the coordinator fails to collect all signatures within the specified time, they inform traders of the coordination failure.

3.2 Implementation

We implemented three algorithms in [3]:

- **Boggart**. Boggart first checks if it needs to add some compensatory outputs. Then, according to the difference between original outputs' amounts, Boggart decomposes original/compensatory outputs round by round until all original outputs are fully decomposed [3]. In the i^{th} round, Boggart obtains decomposed outputs amount of $ov_1 - ov_{i+1}$ from each original output except ov_{i+1}. When an original output's amount is smaller than $ov_1 - ov_{i+1}$, Boggart obtains decomposed output from its compensatory output. Figure 3 shows a running example of Boggart, where we want to get a $\frac{1}{3}$-decomposition for $\{oo_1, oo_2, oo_3, oo_4\}$. Boggart first adds a compensatory output co_3 for oo_3 and co_4 for oo_4. Since $ov_1 - ov_2 = 20$, Boggart tries to decompose each original output except oo_2 by 20. However, since $ov_3 = ov_4 = 7 < 20$, Boggart first obtains decomposed outputs amount of 7 from oo_1, oo_3, and oo_4. Next, Boggart obtains decomposed outputs amount of 13 from oo_1, co_3, and co_4. Finally, Boggart obtains decomposed outputs amount of 20 from each original/compensatory output.
- **DG**. DG adapts existing works (e.g., Dash [1]) to decompose outputs into standard denominations (e.g., 1, 10) from the highest to the lowest. For each denomination x, DG obtain DO', the maximum amount of x from each compensatory and original output. If DO' and the remaining amounts both satisfy a c-decomposition, DG uses DO' to decompose the remaining original/compensatory outputs and turns to the next denomination; otherwise, it tries the next denomination.
- **DR**. DR works like DG except that DR randomly sets DO'.

3.3 Core Technique

Secure. Because coordinators only have the traders' signatures for a particular CoinJoin transaction plan, they cannot create a CoinJoin transaction using a different plan (e.g., moving the traders' funds to the coordinator). Thus, the traders' money are secure.

Anonymity. We have theoretically proved that the anonymity of a c-decomposition is guaranteed in our prior work [3]. Since the Coinjoin transaction by CMixing is a c-decomposition, Coinjoin transactions by CMixing have anonymity guarantees.

Low transaction fees. CMixing employs an approximation algorithm named Boggart, which is proposed in our prior work [3]. Boggart approximately minimizes the number of decomposed outputs in CoinJoin transactions, thereby reducing transaction fees.

4 DEMONSTRATION OVERVIEW

In our demonstration, attendees, acting as traders and coordinators, interactively experience using CMixing to make CoinJoin transactions. This interactive experience allows them to fully immerse themselves in the benefits and advantages of CMixing. Suppose an attendee Alice is a coordinator, and an attendee Bob is a trader.

Scenario 1. Manage the coordination service. Alice effectively manages her coordination service using the coordinator interface, as depicted in Figure 4. The interface provides easy access to crucial information, like the allocated budget for the coordination service. By utilizing the add and remove buttons, Alice can adjust the budget according to her needs. Similarly, Alice also can adjust her commission rate. The graph at the center top of the interface shows the distribution of requests' privacy requirements, providing insights into privacy requirement patterns received by Alice. In the top right corner of the interface, daily statistics summarize the service's performance, helping Alice understand the market and

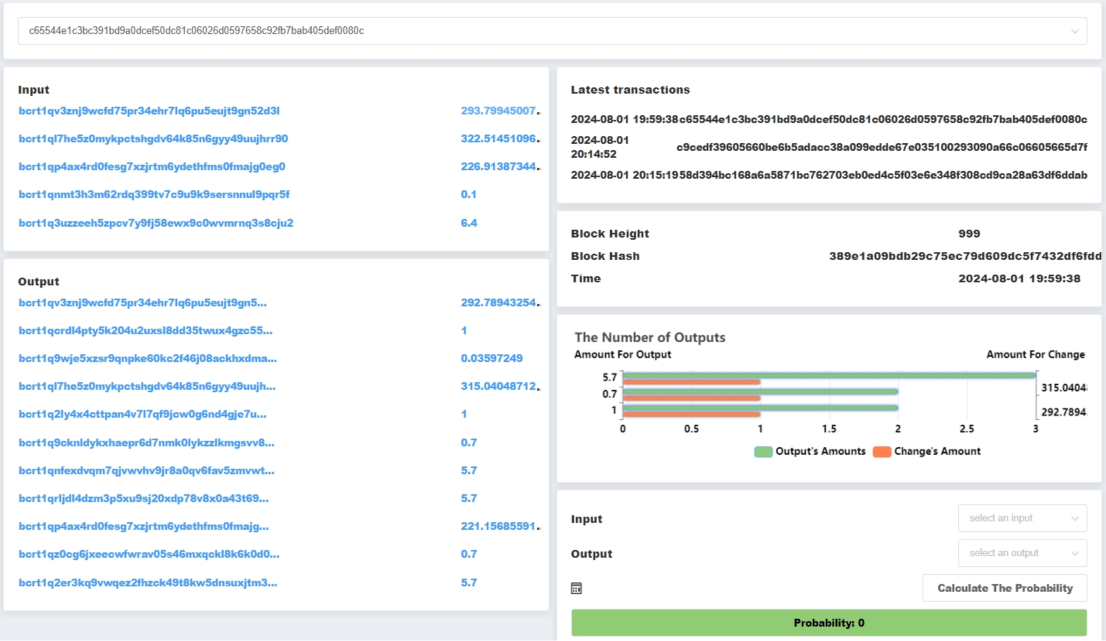

c65544e1c3bc391bd9a0dcef50dc81c06026d0597658c92fb7bab405def0080c

Input

bcrt1qv3znj9wcfd75pr34ehr7lq6pu5eujt9gn52d3l	293.79945007.
bcrt1ql7he5z0mykpctshgdv64k85n6gyy49uujhrr90	322.51451096.
bcrt1qp4ax4rd0fesg7xzjrtm6ydethfms0fmajg0eg0	226.91387344.
bcrt1qnmt3h3m62rdq399tv7c9u9k9sersnnul9pqr5f	0.1
bcrt1q3uzzeeh5zpcv7y9fj58ewx9c0wvmrnq3s8cju2	6.4

Output

bcrt1qv3znj9wcfd75pr34ehr7lq6pu5eujt9gn5...	292.78943254.
bcrt1qcrdl4pty5k204u2uxsl8dd35twux4gzc55...	1
bcrt1q9wje5xzsr9qnpke60kc2f46j08ackhxdma...	0.03597249
bcrt1ql7he5z0mykpctshgdv64k85n6gyy49uujh...	315.04048712.
bcrt1q2ly4x4cttpan4v7l7qf9jcw0g6nd4gje7u...	1
bcrt1q9cknldykxhaepr6d7nmk0lykzzlkmgsvv8...	0.7
bcrt1qnfexdvqm7qjvwvhv9jr8a0qv6fav5zmvwt...	5.7
bcrt1qrljdl4dzm3p5xu9sj20xdp78v8x0a43t69...	5.7
bcrt1qp4ax4rd0fesg7xzjrtm6ydethfms0fmajg...	221.15685591.
bcrt1qz0cg6jxeecwfwrav05s46mxqckl8k6k0d0...	0.7
bcrt1q2er3kq9vwqez2fhzck49t8kw5dnsuxjtm3...	5.7

Latest transactions

2024-08-01 19:59:38	c65544e1c3bc391bd9a0dcef50dc81c06026d0597658c92fb7bab405def0080c
2024-08-01 20:14:52	c9cedf39605660be6b5adacc38a099edde67e035100293090a66c06605665d7f
2024-08-01 20:15:19	58d394bc168a6a5871bc762703eb0ed4c5f03e6e348f308cd9ca28a63df6ddab

Block Height	999
Block Hash	389e1a09bdb29c75ec79d609dc5f7432df6fdd
Time	2024-08-01 19:59:38

Probability: 0

Figure 5: The Query Interface.

adjust her coordination service accordingly. At the bottom, a graph shows important metrics like request frequency, generated plans, and Alice's earnings. *These metrics give a full view of Alice's service performance, helping her track progress and find areas to improve.*

Scenario 2. Send a request. Bob submits his ConJoin transaction request using the request interface. Bob first enters the transaction amount and the privacy requirement. He then selects Alice as the coordinator and sends her the request. Alice, upon receiving the request, provides Bob with a CoinJoin transaction plan. After confirming the plan is correct, Bob clicks "send transaction" to sign and send it back to Alice.

Scenario 3. Query data. After submitting a transaction, Bob uses the query interface to fetch the transaction details from the blockchain, as shown in Figure 5. By entering the hash into the input field at the top, Bob can initiate a search for the corresponding transaction data. If the transaction exists, the interface will present its specifics. The right-side middle section will provide information regarding the number of outputs with varying amounts in the transaction. Finally, in the bottom right corner, Bob can compute the probability that a chosen output comes from a selected input, *allowing him to analyze the linkages between inputs and outputs for a deeper understanding of the transaction's anonymity.*

5 CONCLUSION

Our demonstration introduces CMixing, a new platform for creating anonymous Coinjoin transactions on the blockchain. CMixing makes CoinJoin transactions with guaranteed anonymity and minimized transaction fees. Our demonstration allows the audience to interactively explore how CMixing facilitates coin mixing and compare the transaction fees using Boggart versus baseline algorithms.

ACKNOWLEDGMENTS

Lei Chen's work is partially supported by National Key Research and Development Program of China Grant No. 2023YFF0725100, National Science Foundation of China (NSFC) under Grant No. U22B2060, the Hong Kong RGC GRF Project 16213620, RIF Project R6020-19, AOE Project AoE/E-603/18, Theme-based project TRS T41-603/20R, CRF Project C2004-21G, Guangdong Province Science and Technology Plan Project 2023A0505030011, Hong Kong ITC ITF grants MHX/078/21 and PRP/004/22FX, Zhujiang scholar program 2021JC02X170, Microsoft Research Asia Collaborative Research Grant, HKUST-Webank joint research lab and HKUST(GZ)-Chuanglin Graph Data Joint Lab. Peng Cheng's work is supported by the National Natural Science Foundation of China under Grant No. 62102149. Chen Jason Zhang's work is supported by the Innovation and Technology Fund - Partnership Research Programme (ITF-PRP) (P0040041), Minshang Creative (Donation) (P0046453), Start-up Fund for New Recruits (P0046703), Financial Support for Non-PAIR Research Centres (P0046701), Innovation and Technology Fund - Innovation and Technology Support Programme (ITF-ITSP) (P0048887), Accel Group Holding Limited (Donation) (P0045948), Research Matching Grant Scheme (RMGS) (P0048183, P0048191), and Postdoc Matching Fund Scheme (P0048984). Corresponding author: Peng Cheng.

REFERENCES

[1] Visited 2024. [Online] Dash. https://www.dash.org/.

[2] Visited 2024. [Online] Wasabi Wallet. https://wasabiwallet.io/.

[3] Wangze Ni, Peng Cheng, and Lei Chen. 2022. Mixing transactions with arbitrary values on blockchains. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2602–2614.

[4] Wangze Ni, Peng Cheng, Lei Chen, and Xuemin Lin. 2021. When the recursive diversity anonymity meets the ring signature. In *Proceedings of the 2021 International Conference on Management of Data*. 1359–1371.

[5] Rainer Stütz, Johann Stockinger, Pedro Moreno-Sanchez, Bernhard Haslhofer, and Matteo Maffei. 2022. Adoption and actual privacy of decentralized CoinJoin implementations in bitcoin. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*. 254–267.

[6] Lei Wu, Yufeng Hu, Yajin Zhou, Haoyu Wang, Xiapu Luo, Zhi Wang, Fan Zhang, and Kui Ren. 2021. Towards Understanding and Demystifying Bitcoin Mixing Services. In *Proceedings of the Web Conference 2021*. 33–44.

[7] Xinying Yang, Yuan Zhang, Sheng Wang, Benquan Yu, Feifei Li, Yize Li, and Wenyuan Yan. 2020. LedgerDB: A centralized ledger database for universal audit and verification. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3138–3151.