

Decoupling Bitcoins from Their Transaction History Using the Coinbase Transaction

Bjørn Bjercke <bjorn@bjercke.com>
Keir Finlow-Bates <keir@thinklair.com>

26 March 2020

Abstract

The Bitcoin blockchain is a public open ledger. Therefore all movements of bitcoins are public knowledge, and anyone can submit transactions. As a result, some bitcoins can become “tainted” by passing through an address identified as belonging to a known criminal entity or by association with a bitcoin theft or other malfeasance, and bitcoins are not fully fungible.

Mitigation strategies for overcoming taint exist in the form of mixing services, whereby the taint associated with particular coins can be diluted through the use of mixing transactions.

In the current paper we summarize how Bitcoin transactions and traditional mixing services work, and then present a new method whereby the miner transaction fee property is used to decouple satoshis¹ from their inputs, effectively regenerating them in the Bitcoin coinbase transaction. An enhancement using a collaborating miner is also presented, in which the number of satoshis processed is increased by only forwarding the laundering transaction to the collaborating miner and subsequently revealing it at the same time as a suitable block is found.

Keywords: Bitcoin, cryptocurrency, AML, financial transactions.

¹ A satoshi is the smallest unit of accounting on the Bitcoin blockchain, and 100,000,000 satoshis equal one bitcoin.

Introduction

The Bitcoin blockchain contains a complete record of all coin transactions that have ever been accepted, showing the flow of coins from one address to the next over time. This results in bitcoins being non-fungible, as some coins may have a “taint” attached to them, for example through being used for illicit transactions on Silk Road [7] or as proceeds of a known theft such as the Mt. Gox heist [1]. Tainted coins may therefore have less value to some parties, particularly if cryptoexchanges start to refuse accepting or confiscate tainted coins. As Anti-Money Laundering (AML), Counter Terrorist Financing (CTF) and Know Your Customer (KYC) regulations are tightened to include cryptocurrency transactions [6], the seizure of tainted coins is becoming ever more likely.

The existence of a complete record may seem to suggest that every satoshi can be fully tracked, but this is not actually the case. For example, mixing services allow clean coins and tainted coins to be processed in a manner which reduces the purity of the clean coins and increases the purity of the tainted coins, possibly below a threshold at which subsequent transactions using the mixed coins will be flagged as unacceptable.

However, even though the provenance of an individual satoshi cannot be completely tracked, there is still a transactional link between mixing transactions that can be traced and used to assign a percentage of taint [2] to the outputs of any mixing transaction.

In the current paper we present a new mixing method whereby any number of satothis can be decoupled from their transactional history – effectively destroyed and recreated – and compare and contrast the method with other popular coin laundering systems, examine known limitations, and discuss the ramifications of our research.

Background

The Bitcoin blockchain [10] may be viewed as a ledger documenting the creation of the base units of the Bitcoin cryptocurrency – satothis – and their subsequent transferal from one party (represented on the blockchain by a Bitcoin address) to another over time. The ledger is distributed, so multiple parties maintain records of the ledger, and it is cryptographically secured through the use of hash-linked lists in order to prevent tampering. The elliptic curve digital signing algorithm (ECDSA) is used to authorize transfers, and a proof-of-work consensus system ensures that all parties have one single view of the true state of the ledger, given enough time.

Although the identity of the holder of each balance is pseudonymously² protected through the use of Bitcoin addresses, which are cryptographic hashes of the respective ECDSA public keys, the contents of the ledger itself are public. Anyone with access to a computer and the Internet can download and examine all coin generation events and transactions that have ever taken place since the launch of the system on 3 January 2009, when the “genesis block”, or first Bitcoin data block, was transmitted.

It may therefore seem reasonable to infer that the history of every single satoshi can be tracked to an absolute degree of certainty. However, “mixers”, namely transaction services that take as their input both tainted and non-tainted coins, and output a mix of the inputs, allow for the taint attached to some coins to be diluted.

In order to understand mixing and the new method proposed in this paper, a review of how Bitcoin transactions may be beneficial.

A quick summary of Bitcoin transactions

I. The structure of a bitcoin transaction

Among other things, a Bitcoin transaction consists of “outputs” of satoshis and “inputs” of satoshis.

An “output” is a string of binary digits that specify a sum of satoshis to transfer, the Bitcoin address that currently owns them, and in most cases a Bitcoin address to which they should be transferred. The output is also usually locked with a cryptographic challenge that can only be met by the holder of the private key from which the Bitcoin address receiving the output is derived.

An “input” is a string of binary digits that claim a prior output, and therefore the satoshis transferred by that output. The prior output is most commonly protected with a script, which requires that the holder of a private key corresponding to a public key that hashes to the Bitcoin address specified in the script digitally signs for the release of the satoshis to be transferred.

Each input to a transaction becomes an output to the next transaction, as as a result it is possible to examine the flow of satoshis from one Bitcoin address to the next. However, if a transaction has several inputs from different addresses and several outputs to different addresses, then it is no longer possible to determine exactly which proportion of each input goes to each output.

The sum of the outputs cannot be greater than the sum of the inputs, or the transaction is deemed invalid (as that would involve spending more than was received), except in one special case, namely the “coinbase transaction”, discussed next.

² A pseudonymous identity identifies a “holder” of cryptocurrency, without explicitly linking the identity to a real identity. Nevertheless, due to data leakage there are many forensic techniques that can be used to subsequently link a real identity to the pseudonymous identity.




Figure 1: A sample Bitcoin transaction

II. The coinbase transaction

Bitcoins are created through a special transaction known as the coinbase transaction. It is the first transaction in the transaction list of a mined block, and there are a number of mandatory pieces of data that it must contain for the block to be valid, such as the block height. The miner submitting the block constructs the coinbase transaction to generate the current block reward and to claim transaction fees.

Initially block rewards were set at 50 bitcoins per block, and have since halved roughly every four years. At the time of writing, the block reward is 12.5 bitcoins, but will decrease to 6.25 bitcoins on or around 18 May 2020.

In the same coinbase transaction the miner is also permitted to claim any inputs that are not fully depleted in transactions included in the block – these form the payment of the transaction fees to the miner.

For example, if a transaction contains an input of 1 bitcoin, and two outputs, each of 0.4 bitcoins, there remains an unspent amount of 0.2 bitcoins, which the miner can claim. This constitutes the “transaction fee” or “miner’s fee”. In the example, at this time the coinbase transaction would therefore create and transfer 12.7 bitcoins to the miner’s Bitcoin address.

The transaction fees are not explicitly stated in any of the transactions included in the block. In that sense they are more like a “tip” than a fee – cash left on the table after the diners leave. The miner can choose to pay out the generated reward and transaction fees to multiple addresses, and some mining pools use this to pay out mining shares.

Note that the output(s) of a coinbase transaction cannot be spent until 100 further blocks have been added to the chain, which typically takes about $16\frac{2}{3}$ hours.

Current mixing strategies

The main method for “mixing” coins in order to reduce traceability of prior coin histories involves multiple parties providing inputs to a Bitcoin transaction, and similarly multiple outputs being directed back to those parties. Further obfuscation can be provided by a linked series of multiple transactions. There are several known methods for initiating and coordinating such transactions.

In a centralized laundering service, a laundering operation must necessarily involve a minimum of three transactions: one to supply the laundering service with an input of tainted coins, a second to provide it with clean coins for mixing, and finally a third to mix the coins and return the appropriate proportions of mixed coins to the relevant parties.

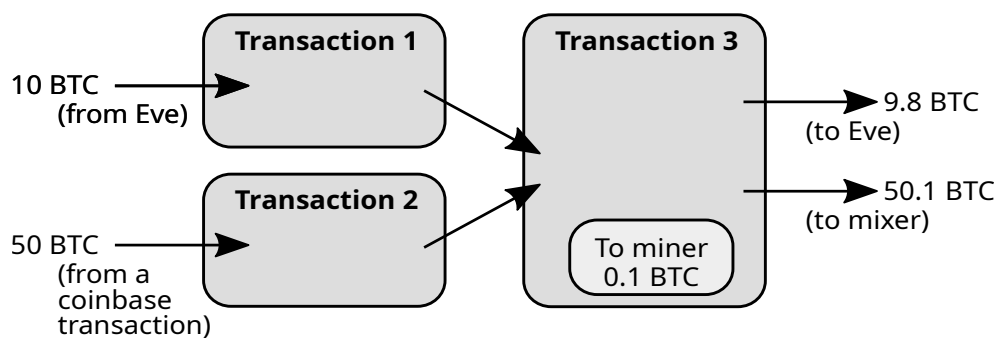


Figure 2: Basic mixing transactions for a mixing service

The users of the mixing service have to take on trust that the service is honest (it can simply run off with their coins), and there is usually a fee charged by the mixing service for the privilege of using it.

It is also possible to enable the decentralized mixing of tainted coins using the CoinJoin system, first proposed by Gregory Maxwell [9]. In CoinJoin, a group of users agree to provide a number of inputs to a transactions, and have those inputs paid back out in equal proportion to each other, but to new, different addresses. The transaction goes round all parties twice:

1. once for its construction, in which every party adds the list of inputs they are using, and the list of outputs that they want the value of their inputs to be sent to, and
2. once again, during which the constructed transaction is signed using each of the private keys that unlock the inputs, in order to make the transaction valid.

If any party is not happy with the transaction that results from step 1, they don’t sign, and the transaction never becomes valid.




Figure 3: A simple CoinJoin transaction

In a simple CoinJoin transaction, Bob and Alice will know of each other's participation, but no other parties will. There are further enhancements to CoinJoin, for example using Chaumian blind signatures [5], which can even hide who provides the inputs and who owns the outputs, but they require more effort and more steps in order to construct the final transaction.

Another mixing method simply involves the use of gambling sites, in which coins can be laundered in a probabilistic manner by repeatedly placing low-odds bets (or if enough bets are placed, even high-odds), and relying on the weak law of large numbers for a convergence to an average expected return [11], although this method runs the risk that the returned coins are just as tainted or even more so than the original ones.

A new method for mixing bitcoins

All of the above methods for mixing coins are based on diluting the taint of one set of satoshis using a clean second set of satoshis by relying on the Bitcoin system not distinctly identifying which individual satoshi in a transaction input is transferred to a specific transaction output. However, the fact remains that in the simplest case only one transaction is used for mixing, and therefore tracing the occurrence of a mixing event is relatively trivial. Even when a web of transactions is used, it is still possible to link the addresses to the mixing activity, as there is always an explicit connection between the transactions.

In the method described below a further level of decoupling is provided by the use of the coinbase transaction, and practical methods are provided for securely maximising the number of satoshis that can safely be decoupled.

I. Decoupling coin transaction histories


The core of our method relies on the observation that unclaimed satoshis from the inputs of a transaction, although described as “transaction fees” or “being claimed by miners”, are actually destroyed, and (optionally) re-created as new satoshis in the coinbase transaction along with the block reward. As a result there is no direct link that can be pointed to between the laundering transaction providing the input and the coinbase transaction cleaning and re-claiming the coins. It can only be inferred from the mined block.

Tainted coins can therefore slowly be laundered by only using them as transaction fees. For example, let us say there is a 1 bitcoin input of tainted coins. The coins are spent to 100 new addresses, each of which receives approximately 0.01 bitcoins.

Subsequently each input from these addresses is used over time to pay transaction fees by including it in normal transactions of clean coins. As there is no explicit link between the laundering transaction and the coinbase transaction the process obfuscates the transaction flow history.

Furthermore, all other transactions included in a block that comprises a transaction using transaction fees to launder tainted coins are unwittingly participate in the laundering activity, as opposed to traditional mixing services that must supply the clean coins for mixing themselves.

Transaction fees are generally not remarkably high. The following chart shows the average daily sum of transaction rewards over a two month period. There was a peak near the end of the period, during which the price of bitcoins dropped rapidly, presumably resulting in more people wishing to quickly transfer their coins and therefore being willing to spend more on their transactions.



*Figure 4: Daily transaction fee returns for the last two months
(figures retrieved from <https://www.blockchain.com/> on 15 March, 2020)*

Over the prior two months the daily sum of average transaction rewards averaged 25.56 bitcoins, and the median was 23.80 bitcoins. For comparison,

the daily block reward total is around 1800 bitcoins at the moment, but will halve to 900 in the near future.

Processing tainted coins by paying transaction fees with them is therefore going to be a slow process, and only of use to those who submit many transactions over time.

However, there is a second step that can be applied to the coinbase mixing process that allows more coins to be safely cleaned and relatively reliably claimed.

II. Increasing the amount processed by not broadcasting the transactions

As mentioned, the problem with submitting a transaction with a large transaction fee to the network is that any miner can pick it up, include it in a block, and claim the unexpectedly large miner's fee.

The solution to this is to refrain from broadcasting the transaction that is using the transaction fee method to launder tainted coins, and only submitting it to a colluding miner.

After the miner has successfully mined a block including this transaction the block (and optionally the laundering transaction) can be broadcast to the system. This prevents other miners from claiming the transaction fee. However, if the laundering fee is set too high, miners may refrain from mining the next block on top of the reward block, and instead try to replace it. This is a variant of the "freeze on transaction" attack, or FRONT problem [8].

The Bitcoin Core code [4] attempts to overcome the problem (and prevent accidentally submitting high fees in error) by imposing a maximum transaction fee, above which transactions that are created will not be transmitted by the wallet included in the code. Constructing a raw transaction using custom code, setting the higher transaction fee and passing it directly to the colluding miner overcomes the Bitcoin Core restriction, as does setting the `-maxtxfee` configuration option higher [3], but does not mitigate the FRONT problem. Earlier versions of the Bitcoin Core code also contained fee checking in the node portion of the code that prevented re-transmission of high fee transactions, but in the most recent release this check has been moved to the wallet portion of the software.

Furthermore, even if the transaction fee is set low enough to reduce the probability of a FRONT mining situation, there is still the possibility that another valid block is uncovered at around the same time as the colluding miner finds one, and for the network to mine on top of that, in which case it is highly likely that the subsequent block will include the laundering transaction and be found by another miner, as shown below:

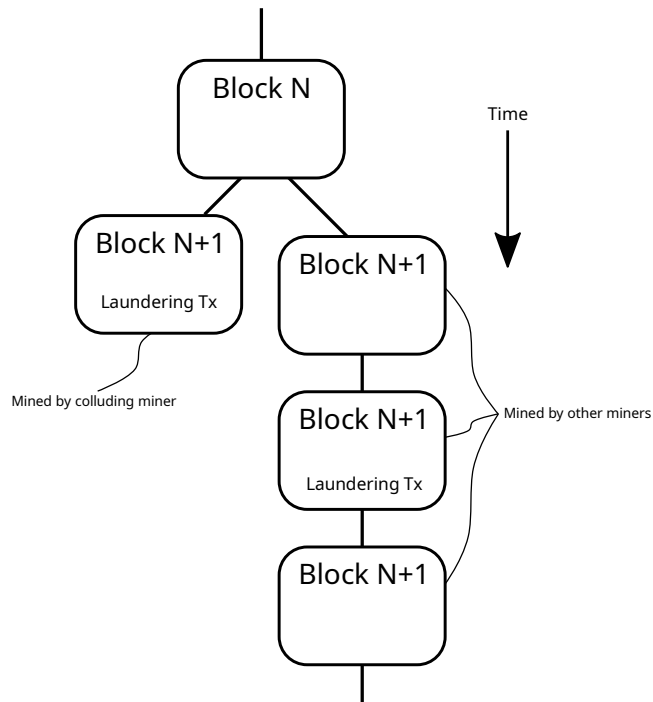


Figure 5: Lost laundering fee due to near-synchronous mining

This results in the coin launderer losing their tainted coins completely.

III. Summary of strengths and known weaknesses

The full process still suffers from a number of weaknesses, which are listed here:

- a significant amount of mining power is required to implement it with relative safety,
- the transactions can still be “linked”, as the laundering transaction and the cleaned coin claiming transaction (in the form of the coinbase transaction) are within the same block,
- furthermore, the presence of a large transaction fee can trigger suspicion that laundering is taking place,
- a risk that the laundering block may be orphaned, and that the laundering transaction fee may subsequently be claimed by another miner exists, and
- the rate at which coins can be laundered may be lower than some parties require, and the laundered coins cannot be spent until about two-thirds of a day has passed.

On the other hand, the strengths of the process compared to other coin mixing strategies are:

- parties submitting legitimate transactions are co-opted into the mixing process involuntarily,
- the process only requires the co-operation between a miner and the party requiring coins to be laundered, and
- there is no explicit link between the laundering transaction and the coinbase transaction, which may offer legal protection.

Our recommendation at this time is that Bitcoin blockchain analytics companies start recording the incidence of “hidden transactions”, that is – transactions that are never received over the peer-to-peer network and entered into the mempool, but are only revealed during the publishing of a block candidate, should be logged as suspicious.

Conclusions

We have presented a method for decoupling the history of Bitcoin transactions through the submission of high transaction fees that are claimed by the coinbase transaction, by concealing the laundering transaction until it is included in a block generated by a colluding miner.

The strengths and weaknesses of the method have been analyzed, and compared with traditional coin laundering approaches.

A recommendation for improving the detection of such coin laundering activity has been proposed – the monitoring of the inclusion of non-broadcast transactions in blocks, especially if such transactions comprise unusually high transaction fees.

Bibliography

- [1] Adelstein, J. (2019); “Solving the world's largest bitcoin heist”, retrieved from <https://www.japantimes.co.jp/news/2019/04/06/national/media-national/solving-worlds-largest-bitcoin-heist/#.XmtHTHUza0k> on 13 March 2020, the Japan Times.
- [2] Apodaca, R. (2015); answer to the question “Procedure for calculating taint?”, retrieved from <https://bitcoin.stackexchange.com/questions/37645/procedure-for-calculating-taint> on 21 March 2020, StackExchange.com.
- [3] Bitcoin.org (2019); “Bitcoin Core version 0.19.0.1 released”, retrieved from <https://bitcoin.org/en/release/v0.19.0.1#updated-rpcs> on 26 March 2020, Bitcoin.org.
- [4] Bitcoin.org (ongoing); “Bitcoin Core integration/staging tree”, retrieved from <https://github.com/bitcoin/bitcoin> on 26 March 2020, Bitcoin.org.
- [5] Chaum, D. (1983); “Blind Signatures for Untraceable Payments”, in Chaum D., Rivest R.L., Sherman A.T. (eds), *Advances in Cryptology*. Springer, Boston, MA.

- [6] FATF (2019); “Virtual Assets and Virtual Asset Service Providers”, retrieved from <https://www.fatf-gafi.org/media/fatf/documents/recommendations/RBA-VA-VASPs.pdf> on 13 March 2020, FATF.
- [7] Greenberg, A. (2015); “Prosecutors Trace \$13.4M in Bitcoins From the Silk Road to Ulbricht's Laptop”, retrieved from <https://www.wired.com/2015/01/prosecutors-trace-13-4-million-bitcoins-silk-road-ulbrichts-laptop/> on 9 March 2020, Wired.
- [8] Lerner, S. D. (2014); “The Bitcoin Freeze on Transaction Attack (FRONT)”, retrieved from <https://bitslog.com/2014/10/05/the-bitcoin-freeze-on-transaction-attack-front/> on 26 March, 2020, Bitslog.
- [9] Maxwell, G. (2013); “CoinJoin: Bitcoin privacy for the real world”, retrieved from <https://bitcointalk.org/index.php?topic=279249.0> on 8 March 2020, Reddit.
- [10] Nakamoto, S. (2008); “Bitcoin: A Peer-to-Peer Electronic Cash System”, retrieved from <https://bitcoin.org/bitcoin.pdf> on 13 March 2020, Bitcoin.org.
- [11] van Wegberg, R., Oerlemans, J., and van Deventer, O. (2018), “Bitcoin money laundering: mixed results? An explorative study on money laundering of cybercrime proceeds using bitcoin”, Journal of Financial Crime, Vol. 25 No. 2, pp. 419-435, retrieved from <https://doi.org/10.1108/JFC-11-2016-0067> on 21 March 2020, Emerald Insight.