Improving the Anonymity of the IOTA Cryptocurrency

Laurence Tennant

2017–10–09

Abstract

IOTA differentiates itself from other cryptocurrencies by being based on a non-blockchain data structure with a highly scalable approach to transaction confirmation. In addition, it exclusively uses post-quantum cryptography. However, as with most cryptocurrencies, IOTA's ledger is currently completely transparent. Constructing an acceptable privacy solution within these parameters is a considerable challenge.

The report begins with a brief introduction to IOTA, followed by a general overview of privacy and anonymity in cryptocurrency. This leads to a review of methods currently used to enhance anonymity in other cryptocurrencies, and an assessment of their effectiveness and applicability to IOTA.

Ultimately, off-ledger mixing using payment channels is found to be the most promising long-term privacy solution. In the meantime, centralised mixing forms a practical way to perform anonymity-enhanced transactions over the IOTA network, and can build a foundation for trustless solutions in future.

1 Introduction

IOTA is a distributed ledger which aims to offer a solution to the issues of scalability and high fees which have afflicted blockchain technology. Created in 2015, it gets its name from its long-term objective to power microtransactions between IoT devices [1], and currently ranks as one of the top ten largest cryptocurrencies by market cap [2]. IOTA's main difference from existing distributed ledger technology is that it is based on a directed acyclic graph consensus structure called the 'Tangle' rather than a blockchain. Instead of requiring special participants—'miners'—to perform computational proof-of-work and validate blocks of transactions in exchange for newly-minted tokens, network participants themselves perform consensus by validating two previous transactions each time they wish to make a transaction [1].

IOTA envisages an open market of devices where resource usage can be billed by the second, opening up new economic possibilities, such as smart

thermostats that share temperature data with weather stations in exchange for micropayments. Canonical are researching the applications of micropayments to telecommunications, and are using IOTA as the accounting system for a profitable smart fountain [3]. Although IOTA is primarily designed with these kinds of use-cases in mind, there is nothing stopping people from using it to exchange value with each other.

Despite IOTA's innovative features, like most cryptocurrencies it is fundamentally a transparent and publicly available ledger. Anyone that a user transacts with can see that user's total balance and parts of their transaction history. This state of affairs is undesirable for much business use, as well as for many ordinary users. Without privacy:

- Firms cannot avoid leaking confidential information to competitors.
- Individual users are at risk of hacking and blackmail.
- IoT transaction activity may be monitored by criminals to plan thefts.
- IOTA's value proposition as digital currency weakens as it loses the property of fungibility.

The IOTA Development Roadmap published in March 2017 listed 'Private Transactions' as a planned feature [4]. This report initiates research and development on this feature, hoping to help prevent the above scenarios from occurring.

2 Privacy & Cryptocurrency

Through the years, the concept of privacy has been framed in relation to distributed ledgers in a variety of ways, and these definitions have implications for the systems that programmers choose to design. Knowledge of practical attacks on cryptocurrency anonymity is key to establishing an approach to attaining privacy.

2.1 Bitcoin & Address Reuse

In the Bitcoin whitepaper, Satoshi Nakomoto noted that despite the inherent transparency of the blockchain, it was possible for users to maintain their privacy by keeping their off-chain identities separate from their on-chain activities. Nakomoto realised that users' privacy—their control over their personal information, such as their purchasing histories—relied heavily upon maintaining pseudonymity. He therefore recommended that fresh addresses be used for every transaction, and that addresses never be linked with personally identifiable information [5, p. 6]. Unfortunately, due to the inconvenience of doing so, many users would follow neither practice.

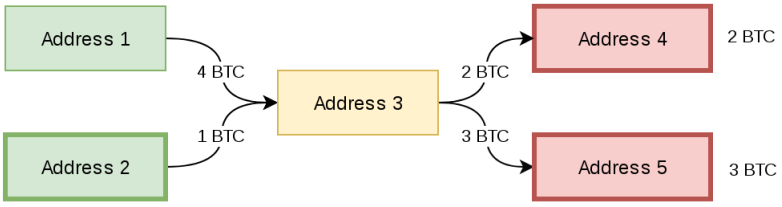

Figure 1: Transaction subgraph

Researchers soon concluded that "strong anonymity is not a prominent design goal of the Bitcoin system" [6], and techniques such as taint and metadata analysis were developed to exploit this.

2.2 Taint Analysis

Taint analysis aims to quantify associations between pairs of addresses. Resistance to taint analysis is central to definitions of cryptocurrency anonymity [7].

Taint analysis works by starting with the graph of all transactions, where each address is a node and each transaction a weighted edge, and calculates the percentage of tokens at a particular address which might have originated from another address [8, p. 3] [9].

For instance, in the transaction subgraph displayed in Figure 1, even though Address 2 never directly pays Address 4, the taint between the two addresses is 20%. Any addresses that Address 4 pays will continue to be tainted by Address 2. Taint analysis is powerful because if any of these addresses can be linked to a known identity (such as through coins purchased on an exchange), that identity's associations with other known entities in the cryptocurrency economy can be constructed.

2.3 Transaction Analysis

Taint analysis considers the broader relationships between addresses on the ledger. It is also worth focussing in on the structure of a transaction itself, and considering what the input and output addresses reveal. IOTA, like Bitcoin, uses change addresses. The funds at input addresses are always used up entirely, and any unspent remainder is sent back to a new address provided by the sender [10]. Change addresses in transactions can usually be distinguished, since the input amounts will typically sum up to a highly precise figure, whereas the payment itself will be rounded to some degree of accuracy (as demonstrated in Figure 2). This leads to three possible scenarios whenever a payment is made, each having different implications for privacy:

1. The sender owns an address with exactly the amount of coins they

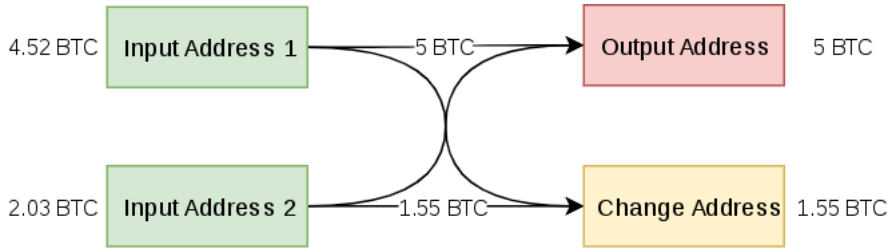

Figure 2: Typical transaction involving a change address

wish to spend. There will be no need for a change address. This situation is unlikely to occur.

2. The sender owns an address containing more coins than they wish to spend. A change address must be created, which will carry the history of the payment.

3. The sender does not own a single address containing enough coins, so must spend from multiple addresses. A permanent link is then made between all these input addresses on the ledger, and a change address will almost certainly have to be created too.

The third situation will eventually arise for users as their total holdings end up divided up amongst numerous change addresses. This means that, over time, unless active measures are taken, that user's activities become easier to tie together via taint analysis. Strategies to counter the growing linkability, such as agreements between senders and recipients to break down payments into multiple smaller transactions matching the amounts in the buyer's input addresses, are difficult to co-ordinate, and just pass the burden on [11].

2.4 Metadata Analysis

Metadata analysis has been much less discussed in the literature, but also deserves consideration [12]. In 2011 Dan Kaminsky pointed out that because Bitcoin nodes must connect to peers in order to broadcast transactions, peers are capable of associating transactions with IP addresses. The proposed solution was to configure the Bitcoin client to run over Tor [13].

However, running a Bitcoin full node eventually became prohibitive in terms of computational resources for many ordinary users, leading to a migration to lightweight clients and web wallets which interact with the blockchain on the user's behalf. It is straightforward for these services to link Bitcoin addresses to IP addresses as well as other metadata, unless users consistently take steps to mask identifying information.

2.5 Measuring Anonymity

Privacy in cryptocurrency is often framed in terms of an 'anonymity set'. This refers to the number of other entities one is indistinguishable from in a system [15]. However, given the extent to which the aforementioned analysis techniques are probabilistic in nature, the concept's utility is limited in this context.[1]

A more suitable formal definition is the 'degree of anonymity' in a system, d, provided by Díaz et al. [17]:

$$d = \frac{H(X)}{H_M}$$

Where $H(X)$ is the entropy of the system taking into account observations made by an attacker:

$$H(X) = -\sum_{i=1}^{N} p_i * \log_2(p_i)$$

N represents the total number of entities in the system; p_i the probability of a message originating from a particular entity.

H_M is the maximum state of entropy in a system, under which every entity is equally likely to be the sender:

$$H_M = \log_2(N)$$

This definition reflects how deanonymisation can become increasingly likely over time as attackers gather information about the relationships in a system more quickly than the maximum entropy of the system is growing, and update the probability distribution accordingly. It can likewise be seen that the behavior of each individual user of a cryptocurrency impacts the anonymity of others, a fact taken advantage of by ledger analysis firms like Elliptic and Chainalysis [18]. Degree of anonymity therefore forms an effective metric by which anonymity-enhancing protocols may be assessed.

2.6 Adversaries

The level of privacy that one wishes to achieve is closely tied to the notion of an adversarial model. Hiding a spending habit from one's family or friends is a rather different proposition from hiding it from a 'global active adversary' which may employ large-scale network and side-channel attacks in addition to ledger analysis. As far as ledger analysis goes, it has been said that a transparent ledger enables anyone to become a global passive adversary [8, p. 12]. At the same time, privacy often incurs significant costs, so the kind

[1] For an extended discussion on the problems of the anonymity set metric and a justification of entropy-based metrics see [16].

of adversaries that users wish to protect their personal information from needs to be considered.

Attaining a suitable level of plausible deniability about one's transaction history may represent sufficient anonymity for most users [14]. Underlying this is the assumption that most users' main adversaries are software startups performing passive ledger analysis, rather than active or coercive adversaries engaged in targeted scrutiny. 'Obfuscation' then is perhaps the best term to describe this level of plausible deniability.

On the other end of the scale, the Monero team proposes an absolutist definition of privacy: it involves not only maximising the degree of anonymity, but making all transactional metadata confidential by default. The aim is that soon, "people snooping the network cannot tell you are even using Monero at all" [19]. Privacy here entails taking as many steps possible to prevent the most capable adversary from uncovering any deanonymising data, to the extent of sacrificing the property of transparency which once defined the blockchain.

Most cryptocurrency users would probably desire a level of privacy somewhere between these two points—strong enough to effectively anonymise their transactions when they need to—but not strong enough to impair their day-to-day activities or attract suspicion.

3 IOTA and Privacy

The previous section outlined a broad framework for the notion of privacy in cryptocurrency. This section will consider the particular design of IOTA and how it fits into this framework.

3.1 IOTA Addresses

IOTA employs a hierarchical deterministic scheme for address generation: the user stores or memorises a single seed (a long random string), and any number of fresh addresses belonging to that user may be generated from it. Due to IOTA's use of one-time signatures, spending from the same address multiple times drastically reduces the security of the funds at that address and is therefore strongly discouraged [20]. Wallets must support automatic generation and handling of new addresses, which makes it very difficult for ordinary users to reuse addresses. This renders IOTA's default taint resistance stronger than that of 'account-based' cryptocurrencies such as Ethereum, which condone address reuse and thus have a transaction graph which is much easier to analyse.

3.2 Post-quantum Cryptography

The IOTA development team has committed to making all cryptography on the ledger quantum-resistant. This is in anticipation of the day when quantum computers are capable of brute forcing discrete logarithms and factoring large primes, the foundations of Elliptic Curve cryptography and the RSA cryptosystem respectively, in superpolynomially quicker time than current machines. This is a situation which some cryptographers do not believe the world is adequately prepared for [21] [22].

The post-quantum commitment, although relating more directly to security than privacy, obviously sets limitations on what kinds of protocols might be implemented in IOTA.

3.3 Masked Authentication Messaging

Since IOTA is feeless, it is possible to freely send messages back and forth over the ledger, with the message occupying the field in the transaction that would otherwise be occupied by the sender's signature. By default, these messages are visible to any observer. The Masked Authentication Messaging module can be used to encrypt messages, providing authentication and integrity with hash-based signatures [23].

This provides a tentative solution to one privacy concern. However, key exchange must occur over a separate channel at present, and tokens cannot be transacted this way.

3.4 Token Origin

All IOTA tokens were created in the genesis transaction. This is something of a blow to privacy, since mining has traditionally been a way to accrue tokens without taint. All tokens in circulation can be traced back to the most recent snapshot, which are special events that prune old transaction data to save space. Although this means that the full history of the ledger is not easily available to users, databases of historic transactions have been recovered.[2]

Ultimately, there are two categories of tokens: those which have never been through an exchange or sent to any identifiable address since IOTA's genesis, and those whose transaction history may be traced back to one of the few exchanges that recently started trading them. Since exchanges tend to hold identifying information on users, only the former category of tokens may be considered potentially untainted by identifiable addresses.

[2]See: `https://github.com/alon-e/iotaWayBack`

Table 1: Connecting to nodes over Tor

Node	Tor Allowed
http://iota.bitfinex.com:80	False
http://service.iotasupport.com:14265	True
http://eugene.iota.community:14265	True
http://eugene.iotasupport.com:14999	True
http://eugeneoldisoft.iotasupport.com:14265	True
http://node01.iotatoken.nl:14265	True
http://node02.iotatoken.nl:14265	True
http://node03.iotatoken.nl:15265	True
http://mainnet.necropaz.com:14500	True
http://wallets.iotamexico.com:80	False
http://5.9.137.199:14265	True
http://5.9.118.112:14265	True
http://5.9.149.169:14265	True
http://88.198.230.98:14265	True
http://176.9.3.149:14265	True
http://node.lukaseder.de:14265	True
https://node.tangle.works:443	True

3.5 Metadata

As far as anonymously connecting to the network goes, IOTA's situation is slightly different to Bitcoin's. It is possible to run a full node, but peer discovery is done manually, and requires a static IP, so is more difficult to route through anonymity technology. On the other hand, there are public full nodes which users can connect to using clients of their choosing.[3] As is shown in Table 1, almost all of the current list of 'light wallet nodes' allow connections from known Tor IPs, making anonymously publishing transactions relatively straightforward in terms of masking an IP address from other nodes.

4 Privacy-Enhancing Protocols

Bearing in mind the potent analysis techniques described so far, and the particular design features of IOTA, we now perform a critical overview of the most important methods that have been proposed over the years to improve the anonymity of cryptocurrency users. These may be divided into three main categories:

1. Protocols that would require *no changes* to the current IOTA core codebase.

[3]See: `http://iotasupport.com/lightwallet.shtml`

2. Protocols that would require *minor or planned changes* to the IOTA core codebase.

3. Protocols that would require *major changes* to the IOTA core codebase.

4.1 No changes

Protocols requiring no changes to the underlying ledger technology, sometimes termed 'overlays' [7], are the most appealing in the context of a short-term project. However, they also tend to possess the weakest privacy properties, being capable only of limited obfuscation of transaction histories rather than provable privacy guarantees.

4.1.1 CoinJoin

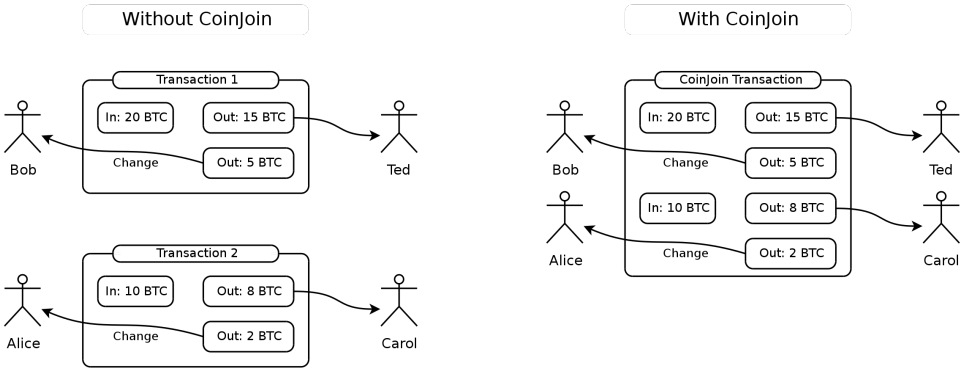

Figure 3: Anatomy of a CoinJoin[4]

CoinJoin, proposed by Gregory Maxwell in 2013, was the first decentralised protocol for mixing coins. It is a method to combine multiple transactions from different users into one large transaction, such that it is not clear to the outside observer which coins have ended up at which address [24]. This ostensibly straightforward protocol is riddled with practical issues:

- All participants must trade the same value, otherwise the mapping is usually trivial to determine, as it would be in Figure 3.[5]

- Historically, it has been a challenge merely to arrange a CoinJoin with enough honest users who want to mix the same value. JoinMarket

[4]Image Credit: MARIODOESBREAKFAST - Image drawn for CoinJoin article, CC BY-SA 3.0, `https://en.wikipedia.org/w/index.php?curid=45648024`

[5]See 'CoinJoin Sudoku' for an analysis of this weakness in Blockchain.info's SharedCoin service: `http://www.coinjoinsudoku.com/advisory/`

tackled this problem by incentivising mix makers to collect fees from the other participants [25].

- There is third party risk in most implementations as somebody must organise the transaction and broadcast it: in a naive version of the protocol, all participants learn the mapping. In JoinMarket, the mix initiator learns the mapping. This leaves the market open to Sybil attacks, where the market is flooded by identities controlled by a single entity initiating a large amounts of mixes, learning all the information the market is aiming to hide, and recovering their funds afterwards [26].

  Worse still, with IOTA's zero fees, in an IOTA JoinMarket analogue a Sybil attack would be more profitable to perform, as there would be no cost to consolidating and reusing funds after performing multiple CoinJoins, or of moving funds in and out of the market.

- CoinJoins are highly distinguishable on the ledger, due to the unusually large number of input and output addresses.

- CoinJoins do not remove taint from coins, since the link from input to output address is not broken. A recent paper takes advantage of this as well as the previous point to describe a 'cluster intersection attack' capable of deanonymising a large number of real world CoinJoins [27].

4.1.2 CoinShuffle

CoinShuffle [28] improves on CoinJoin by removing the need for a particular party to arrange and publish the transaction. Borrowing the idea of layered encryption from decentralised mix networks, CoinShuffle allows participants to randomise the set of output addresses without ever learning the mapping. This is achieved through an initial exchange of public keys and encryption of each address with every key, followed by successive rounds of shuffling encrypted addresses and stripping away a single layer of encryption. By the time the last layer is decrypted, every participant has the means to publish a complete mixing transaction.

The downsides of CoinShuffle include most of those suffered by CoinJoin, along with significant additional complexity. The co-ordination problems of decentralised protocols must not be overlooked. Even after such a protocol gets underway, the lack of co-operation of a single participant (deliberate or otherwise) can disrupt proceedings entirely, and disruption is free to perform in a feeless model. This has to be carefully engineered around, otherwise attacks that reveal an address's signature (see 3.1) become possible. User experience is also an issue; as Mike Hearn puts it, you have to "worry about a random mobile phone the other side of the world driving into a tunnel at the wrong moment" [29].

4.1.3 Centralised Mixers

Unlike the previous two protocols, centralised mixers can completely remove taint by performing mixing over multiple transactions. The user specifies a receiving address and pays in some amount of coins, and the mixer later pays out coins originating from another source. The major disadvantage is the counterparty risk of the mixer absconding with the money or recording the mappings [30, p. 213].

Improvements on the basic centralised mixing protocol have been proposed to somewhat mitigate this risk. Mixcoin introduces the property of accountability: users provide contractual terms when they initiate a mix, and the mixer issues a signed receipt; if the mixer then fails to return the money, the user may produce cryptographic proof of the theft in hopes of damaging the reputation of the service [8].

Blindcoin's additional contribution is a blind signature scheme which hides the mappings from the mixing server itself [31]. However, since output addresses must be linked to the signing key of the mixing server on a public log, Blindcoin sacrifices the indistinguishability of mixing operations from regular transactions on the ledger [31], a property perceived as potentially the biggest advantage of mixing services as it significantly increases the anonymity set of users mixing coins from the point of view of a global passive adversary [32].

4.2 Minor or planned changes

The following protocols rely on some features which IOTA does not possess yet, but may do in the future.

4.2.1 CoinSwap

CoinSwap [33] is similar to centralised mixing services in that it routes payments via a third party. However, it prevents even malicious third parties from absconding with the funds. The protocol involves both the sender and receiver making multisignature transactions with the third party, so the receiver can be paid on the sender's behalf. This alone is not enough to ensure the coins will be received, since the third party may simply never sign the second transaction after signing the first, so an operation of the Bitcoin scripting language informally called a 'hashed timelock' is used. This guarantees transaction atomicity—the second transaction must be redeemable if the first is [34].

One downside of this protocol is that all participants learn the mapping, and if any party attempts to cheat, a sequence of transactions linked by a common hash must be published to the blockchain to recover the funds.

Nonetheless, CoinSwap forms the basis of perhaps the most promising privacy proposal for Bitcoin in recent years, TumbleBit [35]. TumbleBit

combines off-chain RSA computations with CoinSwap-like payment channels to ensure that even the tumbler cannot learn the mappings [36]. It is still in development, but offers some ideas about how to eventually obtain scalable trustless private transactions in IOTA, essentially by performing them off the ledger. Before TumbleBit's concepts can be useful for IOTA, at minimum IOTA would need to establish payment channels and basic scripting capabilities.

4.2.2 Blackbytes

Byteball, a cryptocurrency that, like IOTA, is based on a directed acyclic graph, has a layer of privacy in the form of a secondary currency. 'Blackbytes' are special files doubling as coins which are transmitted directly from one user to another. Spending them requires publication of just two hashes on the ledger: the transaction hash and a 'spending proof' [37, p. 40]. Each file contains its own complete payment history, which is verified by the receiver by checking each payment against published hashes all the way back to the coin's genesis. The 'spending proof' prevents double spends as it depends wholly on data produced by the last spend of the coin, and will be rejected by the receiver if found not to be unique. Blackbytes are therefore like miniature self-contained private ledgers, with observers unable to see the senders, receivers, or values of transactions contained inside.

However, Blackbytes suffer from a number of practical difficulties. They are cumbersome to spend, consisting of fixed denominations, and allow just one input per transaction to prevent payment histories from growing exponentially [37, p. 44]. They are easy to lose as the files must be managed locally by users. If a large amount of them were to eventually pass through a single merchant, that merchant would gain full details of virtually all past 'private' payments. Furthermore, Blackbyte payment histories are ultimately traceable back to particular Bitcoin transactions, as this is how they are issued, which may be deanonymising in itself. The IOTA team have made clear that they will not mint additional coins, so a system like Blackbytes, which introduces a rather shaky notion of privacy, and has seen little adoption or testing so far, would be controversial.

4.2.3 Merge Avoidance

Merge Avoidance [29] is an intriguing idea which, although now inapplicable to Bitcoin due to its high fees, is worth re-evaluating here. It revolves around the notion of a payment protocol [38], and is described here in a manner adapted to a typical use-case of IOTA.

Imagine that a smart thermostat wants to collect a payment in a private manner from a weather station that has made use of its data. The thermostat sends an encrypted payment request to the station, providing multiple

addresses by which to receive the funds, the total requested amount, and perhaps additional information.

Assuming the station respects the payment request, not only will the received funds be sent to addresses initially difficult to link to the thermostat, but, due to being broken down into smaller denominations, will also be easier to spend in a way that does not leak information about the original payment. More importantly, a payment request could be forwarded on to the station's debtors in turn, skipping a link in the chain of payments and making it even harder to determine the true history of a coin [39].

4.3 Major changes

These protocols would require substantial changes to the way IOTA works. However, they offer much stronger privacy guarantees, and as such are each worth a cursory examination in order to establish what it might take for them to be implemented one day.

4.3.1 Zero-Knowledge Proofs

Zero-Knowledge Proofs (ZKPs) allow a party to prove to a verifier that a statement is true without revealing any additional information about that statement [40]. The cryptocurrency Zcash implements a type of ZKP called zk-SNARKs which enable users to hide all data in a transaction: the transaction merely needs to include a ZKP certifying that the internal transaction data is valid [41].

There are a number of reasons why it would be difficult to adapt Zcash-style ZKPs to the IOTA network. First, these proofs require a large amount of computation to produce, and output nonces which must be stored indefinitely by all verifying nodes [42]. Further, verifying Zcash shielded transactions takes orders of magnitude longer than verifying typical transactions [42], in a way inconsistent with the lightweight and decentralised verification underpinning the Tangle. Current implementations exploit numerous non post-quantum cryptographic primitives; it is not clear how they could be made quantum-secure [43]. Further research into ZKPs would have to start by solving the problem of how smaller devices can securely offload the computational burden onto more capable machines.

4.3.2 Ring Signatures

Ring signatures, integrated most notably in the CryptoNote protocol [44], allow a user preparing a transaction to group their own public key with several others, and send a valid transaction containing proof only that they possesses at least one of the keys. In the best case, this can make all senders referenced in the transaction equally likely to be the real sender from the perspective of a blockchain observer, thereby achieving a degree

of unlinkability without needing to involve centralised coin mixers or active participation from other users.

Although they require further research before becoming practical, ring signature implementations based on post-quantum lattice-based cryptography exist [21]. A proposal for 'Anonymous Post-Quantum Cryptocash' along these lines was recently published [45]. The issue with ring signature schemes here is that they involve an established public-key cryptosystem [46], which IOTA does not have.

4.3.3 Others

Several other advanced privacy-enhancing technologies, such as Stealth Addresses [47], Confidential Transactions [48], and Mimblewimble [49], are outside the scope of this report, as they are fundamentally incompatible with IOTA's cryptography.

5 Conclusion

Bitcoin privacy advocate Daniel Krawisz notes that "anonymity in Bitcoin is a hard problem and no single protocol or service is sufficient to provide it" [50]. The research here would suggest that this applies even more so to IOTA, owing to three main design factors:

- IOTA's use of hash-based signatures rules out protocol-enhancing ideas based on elliptic curve and public-key cryptography.
- The requirement for a lightweight and scalable solution further restricts the solution space.
- The fact that IOTA has zero fees makes decentralised protocols hard to build—there is no inherent barrier to a Sybil attack.

Nonetheless, with plenty of upgrades and development on the horizon, IOTA's anonymity has the potential to improve substantially. While an on-ledger solution to transaction privacy in the near future appears difficult to achieve, an off-ledger solution taking advantage of ongoing work on payment channels could make for a compelling alternative.

As mentioned, much work has already been done on an off-chain mixer in Bitcoin by the TumbleBit researchers, who have constructed a 'Classic Tumbler' model with strong privacy properties guaranteed through cryptography. Additional features are also capable, such as trustlessly exchanging coins for other cryptocurrencies. The only disadvantageous aspect is the centralisation of availability. If IOTA gained support for cryptographic opcodes supporting hashed timelocks or an equivalent way to defeat counterparty risk, it would enable this highly scalable approach to privacy.

While this represents an ideal solution, it will take some time before the necessary prerequisites are in place, and development work to adapt a Bitcoin codebase to the IOTA network. In the meantime, centralised mixing forms the most practical means of improving the anonymity of IOTA. Centralised mixing can make tracing payment histories in what was previously a fully transparent ledger far more difficult and doubtful, thereby adding to the degree of anonymity of all users of the cryptocurrency. Although token mixing has clear downsides, it is a first step, and as has recently been pointed out, more simple anonymity-enhancing techniques based on obfuscation have been widely adopted in Bitcoin, whereas those requiring ambitious protocol changes have tended to get stuck in development for years [51].

Merge avoidance and Cut-thru payments (see 4.2.3), following the development of a payment protocol, also represent neat concepts for marginally improving the degree of anonymity on the ledger. By imitating the financial practice of paying on someone's behalf and splitting payments up amongst fresh addresses, taint becomes harder to trace.

This report established a foundation for understanding cryptocurrency anonymity, investigating techniques such as taint analysis and metrics such as degree of anonymity. It provided a critical assessment of how existing research into improving cryptocurrency anonymity and privacy fits in with IOTA's novel design. Finally, it highlighted the most promising avenues for further research.

References

[1] Serguei Popov. IOTA: The Tangle, 2016. Available at: `http://iotatoken.com/IOTA_Whitepaper.pdf`.

[2] CryptoCurrency Market Capitalizations, 2017. `https://www.coinmarketcap.com`, accessed 15/9/2017.

[3] Marrten Ectors. The 1M/month revenue generating fountain each smart city should have, 2017. Available at: `https://insights.ubuntu.com/2017/03/09/the-1mmonth-revenue-generating-fountain-each-smart-city-should-have/`.

[4] David Sønstebø. IOTA Development Roadmap, 2017. Available at: `https://blog.iota.org/iota-development-roadmap-74741f37ed01`.

[5] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. Available at: `http://bitcoin.org/bitcoin.pdf`.

[6] F. Reid and M. Harrigan. An Analysis of Anonymity in the Bitcoin System. *ArXiv e-prints*, July 2011. Available at: `http://adsabs.harvard.edu/abs/2011arXiv1107.4524R`.

[7] Sarah Meiklejohn and Claudio Orlandi. Privacy-Enhancing Overlays in Bitcoin. In *Financial Cryptography and Data Security: FC 2015 International Workshops, BITCOIN, WAHC, and Wearable, San Juan, Puerto Rico, January 30, 2015, Revised Selected Papers*, pages 127–141. Springer Berlin Heidelberg, 2015.

[8] Joseph Bonneau, Arvind Narayanan, Andrew Miller, Jeremy Clark, Joshua A. Kroll, and Edward W. Felten. Mixcoin: Anonymity for Bitcoin with accountable mixes. *IACR Cryptology ePrint Archive*, 2014:77, 2014.

[9] M. Möser, R. Böhme, and D. Breuker. An inquiry into money laundering tools in the Bitcoin ecosystem. In *2013 APWG eCrime Researchers Summit*, pages 1–14, Sept 2013.

[10] Bitcoin Wiki: Change. Available at: `https://en.bitcoin.it/wiki/Change`.

[11] Eduardo Ferreira. Anonymity in Bitcoin: Splitting Transactions, 2016. Available at: `https://bitcointalk.org/index.php?topic=1394184.0`.

[12] Matt Corallo and Adam Back. Fungibility overview. Scaling Bitcoin Workshop, Milan. Available at: `http://diyhpl.us/wiki/transcripts/scalingbitcoin/milan/tumblebit/`, 2016.

[13] Dan Kaminsky. Black Ops of TCP/IP, 8 2011. Available at: `https://www.slideshare.net/dakami/black-ops-of-tcpip-2011-black-hat-usa-2011`.

[14] Adlai Chandrasekhar. Joinmarket. Scaling Bitcoin Workshop, Milan. Available at: `https://diyhpl.us/wiki/transcripts/scalingbitcoin/milan/joinmarket/`, 2016.

[15] UC Berkeley School of Information. Privacy Patterns: Anonymity Set. Available at: `https://privacypatterns.org/patterns/Anonymity-set`.

[16] Andrei Serjantov and George Danezis. *Towards an Information Theoretic Metric for Anonymity*, pages 41–53. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

[17] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In *Proceedings of the 2nd International Conference on Privacy Enhancing Technologies*, PET'02, pages 54–68, Berlin, Heidelberg, 2003. Springer-Verlag.

[18] Natasha Lomas. Elliptic takes in 5M for its blockchain forensics tool, 2016. Available at: https://techcrunch.com/2016/03/21/elliptic-takes-in-5m-for-its-blockchain-forensics-tool/.

[19] Monero Project. How does Monero's privacy work? Available at: https://www.monero.how/how-does-monero-privacy-work.

[20] Dominik Schiener. Generating addresses: Learn the basics, 2017. Available at: https://learn.iota.org/tutorial/generating-addresses-learn-the-basics.

[21] Johannes A. Buchmann, Denis Butin, Florian Göpfert, and Albrecht Petzoldt. *Post-Quantum Cryptography: State of the Art*, pages 88–108. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.

[22] Anastasia Marchenkova. How secure will our data be in the post-quantum era?, 2015. Available at: https://medium.com/quantum-bits/how-secure-will-our-data-be-in-the-post-quantum-era-6a7f444ce7d5.

[23] Paul Handy. Masked Authentication Messaging: Overview, 2015. Available at: https://github.com/iotaledger/mam.client.js/blob/master/Overview.md.

[24] Gregory Maxwell. CoinJoin: Bitcoin privacy for the real world, 2013. Available at: https://bitcointalk.org/index.php?topic=279249.0.

[25] Chris Belcher. Joinmarket - Coinjoin that people will actually use, 2015. Available at: https://bitcointalk.org/index.php?topic=919116.0.

[26] Malte Möser and Rainer Böhme. The price of anonymity: empirical evidence from a market for Bitcoin anonymization. *Journal of Cybersecurity*, August 2017. Available at: https://doi.org/10.1093/cybsec/tyx007.

[27] S. Goldfeder, H. Kalodner, D. Reisman, and A. Narayanan. When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies, August 2017. Available at: https://arxiv.org/pdf/1708.04748.pdf.

[28] Tim Ruffing, Pedro Moreno-Sanchez, and Aniket Kate. *CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin*, pages 345–364. Springer International Publishing, Cham, 2014.

[29] Mike Hearn. Merge Avoidance, 2013. Available at: https://medium.com/@octskyward/merge-avoidance-7f95a386692f.

[30] Pedro Franco. *Understanding Bitcoin: Cryptography, Engineering and Economics*. Wiley, 2014.

[31] Luke Valenta and Brendan Rowan. Blindcoin: Blinded, Accountable Mixes for Bitcoin. In *Financial Cryptography and Data Security - FC 2015 International Workshops, BITCOIN, WAHC, and Wearable, San Juan, Puerto Rico, January 30, 2015, Revised Selected Papers*, pages 112–126, 2015.

[32] Sundance. Byzantine Cycle Mode: Scalable Bitcoin Mixing on Unequal Inputs. 2015. Available at: `https://github.com/sundance30203/coin-mix/blob/master/doc/bcm.pdf?raw=true`.

[33] Gregory Maxwell. CoinSwap: Transaction graph disjoint trustless trading, 2013. Available at: `https://bitcointalk.org/index.php?topic=321228`.

[34] Bitcoin Wiki: Hashlock. Available at: `https://en.bitcoin.it/wiki/Hashlock`.

[35] Ethan Heilman & Leen AlShenibr. Tumblebit. Scaling Bitcoin Workshop, Milan. Available at: `http://diyhpl.us/wiki/transcripts/scalingbitcoin/milan/tumblebit/`, 2016.

[36] Ethan Heilman, Foteini Baldimtsi, Leen Alshenibr, Alessandra Scafuro, and Sharon Goldberg. TumbleBit: An Untrusted Tumbler for Bitcoin-Compatible Anonymous Payments. *IACR Cryptology ePrint Archive*, 2016:575, 2016.

[37] Anton Churyumov. Byteball: A Decentralized System for Storage and Transfer of Value, 2016. Available at: `https://byteball.org/Byteball.pdf`.

[38] Gavin Andresen & Mike Hearn. BIP 70: Payment Protocol, 2013. Available at: `https://github.com/bitcoin/bips/blob/master/bip-0070.mediawiki`.

[39] Peter Todd. FAQ on the Payment Protocol: Cut-thru payments, 2013. Available at: `https://bitcointalk.org/index.php?topic=300809.msg3227426#msg3227426`.

[40] S Goldwasser, S Micali, and C Rackoff. The Knowledge Complexity of Interactive Proof-systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, pages 291–304, New York, NY, USA, 1985. ACM.

[41] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. Cryptology ePrint

Archive, Report 2014/595, 2014. Available at: http://eprint.iacr.org/2014/595.

[42] Peter Todd. Cypherpunk Desert Bus: My Role In The 2016 Zcash Trusted Setup Ceremony, 2016. Available at: https://petertodd.org/2016/cypherpunk-desert-bus-zcash-trusted-setup-ceremony.

[43] Daira Hopwood. Post-quantum Zcash, 2016. Available at: https://github.com/zcash/zcash/issues/805.

[44] Nicolas van Saberhagen. Cryptonote v2.0, 2014. Available at: https://cryptonote.org/whitepaper.pdf.

[45] Huang Zhang, Fangguo Zhang, Haibo Tian, and Man Ho Au. Anonymous Post-Quantum Cryptocash. Cryptology ePrint Archive, Report 2017/716, 2017. Available at: http://eprint.iacr.org/2017/716.

[46] Man Ho Au, Sherman S. M. Chow, Willy Susilo, and Patrick P. Tsang. Short Linkable Ring Signatures Revisited. In *Proceedings of the Third European Conference on Public Key Infrastructure: Theory and Practice*, EuroPKI 2006, pages 101–115, Berlin, Heidelberg, 2006. Springer-Verlag.

[47] Nicolas T. Courtois and Rebekah Mercer. Stealth Address and Key Management Techniques in Blockchain Systems. In *Proceedings of the 3rd International Conference on Information Systems Security and Privacy - Volume 1: ICISSP*, pages 559–566. INSTICC, SciTePress, 2017.

[48] Greg Maxwell. Confidential Transactions, 2015. Available at: https://people.xiph.org/~greg/confidential_values.txt.

[49] Andrew Poelstra. Mimblewimble: Private, Massively-Prunable Blockchains. Available at: https://cyber.stanford.edu/sites/default/files/andrewpoelstra.pdf, 2016.

[50] Kyle Torpey. TumbleBit Part 2: How Does This Bitcoin Privacy Improvement Compare with CoinJoin and CoinShuffle?, 2016. Available at: https://coinjournal.net/bitcoin-privacy-improvement-compare-coinjoin-coinshuffle/.

[51] A. Narayanan and M. Möser. Obfuscation in Bitcoin: Techniques and Politics. *ArXiv e-prints*, June 2017. Available at: https://arxiv.org/pdf/1706.05432.pdf.

[52] R. Mercer. Privacy on the Blockchain: Unique Ring Signatures. *ArXiv e-prints*, December 2016. Available at: https://arxiv.org/pdf/1612.01188.pdf.

[53] Georg Becker. Merkle Signature Schemes, Merkle Trees and Their Cryptanalysis. Available at: https://www.emsec.rub.de/media/crypto/attachments/files/2011/04/becker_1.pdf.

[54] Adam Back. Fungibility and Privacy. Israeli Bitcoin Conference. Available at: http://diyhpl.us/wiki/transcripts/bitcoin-adam3us-fungibility-privacy/, 2014.