An Analysis of Routing Attacks Against IOTA Cryptocurrency

Pericle Perazzo
Information Engineering dept.
University of Pisa
Pisa, Italy
pericle.perazzo@iet.unipi.it

Antonio Arena
Information Engineering dept.
University of Pisa
Pisa, Italy
antonio.arena@ing.unipi.it

Gianluca Dini
Information Engineering dept.
University of Pisa
Pisa, Italy
gianluca.dini@iet.unipi.it

***Abstract*—IOTA is a new type of distributed ledger designed for allowing fee-less and rate-scalable micropayments in Internet of Things applications. Security research on IOTA has focused mainly on attacks involving its cryptographic operations or its consensus algorithm. In this paper, we present a preliminary analysis of the IOTA security with respect to malicious Autonomous Systems (ASes), which can intercept IOTA connections by manipulating routing advertisements (BGP hijacking) or by naturally intercepting traffic. We make the simplifying assumption that the malicious AS can intercept routes between hosts without causing side effects, or without these side effects being noticed by the intercepted hosts. We identify three notable attacks that can lead to permanent money freeze, and to local or global interruptions of the consensus mechanisms. We then analyze the vulnerability of IOTA against malicious ASes on the real Internet topology, and we show that IOTA cryptocurrency is, at the time of writing, pretty susceptible of these attacks because quite centralized from the point of view of BGP routing. We then study the routing-level security of the next version of IOTA (post-coordicide), which has been proposed by the IOTA Foundations to make the cryptocurrency fully distributed.**

***Keywords*-Blockchain, Cryptocurrency, BGP Hijack, BGP Interception, IOTA**

I. Introduction

IOTA [1] is a distributed ledger operating from 2016, whose mechanisms are specifically designed for Internet of Things applications involving micropayments. Its main difference with respect to other distributed ledgers like Bitcoin [2] and Ethereum [3] is that IOTA replaces the classic blockchain data structure with a directed acyclic graph named *Tangle*. The Tangle allows IOTA to avoid two major drawbacks of blockchain-based ledgers: the transaction fees and the poor scalability in terms of global transaction rate. The IOTA distributed ledger has already been proposed for application in many Internet-of-Things scenarios [4], [5], [6]. On the downside, IOTA currently requires a trusted centralized entity, the *coordinator*, in order to finalize the consensus about the status of the distributed ledger. However, IOTA maintainers are designing a major change on the IOTA consensus algorithm called the *coordicide* [7], which will remove any trusted centralized entity from the network, making IOTA fully distributed.

So far, IOTA's security have been investigated mainly from the point of view of the underlying cryptography mechanisms [8], [9], [10], [11] and the consensus mechanisms [12], [13]. However, the possibility of attacking IOTA directly via the Internet routing infrastructure has been neglected so far. By advertising false routing information (BGP hijacking [14], [15]) or by naturally intercepting traffic, malicious Autonomous Systems (ASes) can intercept, drop, or manipulate IOTA traffic, which is currently transmitted in the clear without any integrity protection mechanism.

In this paper, we present a first security analysis of IOTA with respect to Internet routing attacks. In particular we identify three notable attacks that a malicious AS can mount against IOTA, namely *address freeze*, *targeted denial of consensus* and *general denial of consensus*. Such attacks can lead, respectively, to a money loss, an interruption of the consensus mechanisms for a given victim node, and an interruption of consensus for the whole IOTA cryptocurrency. We also reconstruct the current IOTA network topology and we study how much vulnerable such a network is with respect to possible malicious ASes. The results show that IOTA is, at the time of writing, pretty centralized from the BGP routing point of view, and thus susceptible to attacks by ASes. Finally, basing on the last published IOTA coordicide white paper [7], we also analyze the security of *post-coordicide* IOTA with respect to malicious ASes, and we show some possible vulnerabilities that may eventually lead to double spending. Based on the above analyses, we conclude that IOTA could solve many of its routing-level vulnerabilities by supporting secure channels between its nodes. Secure channels do not avoid traffic interception, but they impede malicious ASes to selectively drop messages, which is the key to mount the attacks presented herein.

The rest of the paper is organized as follows. Section II introduces some background about BGP hijacking and IOTA technology, and it analyzes relevant related work. Section III introduces three notable attacks against IOTA carried out by malicious ASes. Section IV reconstructs the current IOTA network topology and studies the vulnerability of such a network with respect to malicious ASes. Section V analyzes the possible routing-level vulnerabilities of post-coordicide IOTA. Finally, Section VI concludes the paper.

II. Background and Related Work

A. BGP Hijacking

BGP is the de-facto standard Internet routing protocol [16]. With BGP, ASes can exchange and propagate information on subnet reachability by means of *advertisements*. Notably, BGP does not check the validity of advertised routes, hence it is reasonably easy for an AS to inject false advertisements of routes that it does not have actually. *BGP hijacking* [14], [15] is the act of advertising a route that the originator actually does not have, with the aim of attracting Internet traffic that was directed to that destination.

The basic effect of BGP hijacking is a denial of service against the hijacked destination subnet, because all the Internet traffic that was destined to that subnet is instead forwarded to the attacker AS, which drops it. However, if the attacker is able to preserve at least one path to the destination, then it can forward all the intercepted traffic to it (*interception attack*) [14]. With an interception attack, the malicious AS can act as an Internet-scale man in the middle for all the traffic destined to the victim subnet. Note that some interception attacks may be ineffective or imprecise in the real life, and they may have detectable side effects, for example interrupting communications of hosts other than the intercepted two. In this paper, we make the simplifying assumption that the malicious AS can intercept routes between hosts *without causing* side effects, or without these side effects being noticed by the intercepted hosts. Every route can in principle be intercepted with BGP hijacking, except those routes that are internal to a single AS. Indeed, intra-AS traffic does not get routed by BGP, but by internal routing protocols (e.g., RIP, OSPF), so it is not affected by malicious BGP advertisements. Of course, all routes, even intra-AS ones, can be always intercepted by a malicious AS that *hosts* one of the two endpoints, or that is included in the route. In these cases, the malicious AS does not have to perform any BGP hijacking to intercept the traffic. In this paper we consider both types of adversaries: an AS that intercepts the traffic after a BGP hijacking, and an AS that *naturally* intercepts the traffic.

B. IOTA Cryptocurrency

The main goal of IOTA is to address most of the perceived inefficiencies of blockchain technologies, such as the poor scalability of the transaction rate, the high transaction fees, and the high energy costs of the consensus algorithm. The first noticeable difference of IOTA is the involved data structure, the Tangle [1]. The Tangle retains similarities to the blockchain while removing the block structure in favor of a Directed Acyclic Graph (DAG) structure. In particular, every block in the Tangle represents a single transaction, called *bundle* in the IOTA terminology. Every transaction must confirm two transactions that are already included in the Tangle. The IOTA transaction format is composed of six mandatory fields: (i) the IOTA source address, i.e., the address representing the node which is spending money; (ii) the IOTA source destination, i.e., the address representing the node which is receiving money; (iii) the quantity of money to be transferred, that can be zero as well; (iv-v) the other two transaction to be confirmed; (vi) the signature of the node that is spending money. As a signature scheme, IOTA employs a quantum-resistant Winternitz one-time signature [11], using the source address itself as the public key.

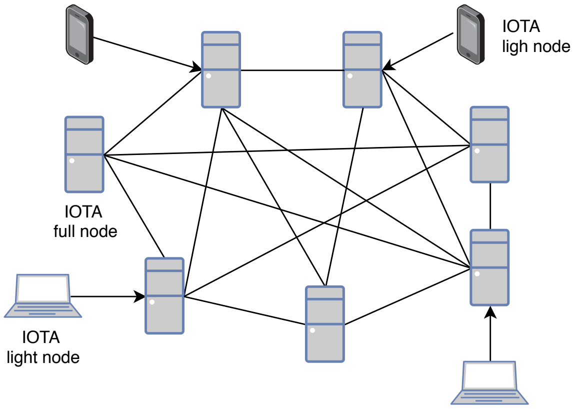

Figure 1. IOTA architecture.

The IOTA architecture is shown in Figure 1. IOTA nodes fall into two classes: *full nodes* and *light nodes*. Full nodes are in charge of maintaining a local copy of the Tangle and performing the consensus algorithm. They connect with a random number of other full nodes for broadcasting new transaction to each other and uniforming each local vision of the Tangle to a global one. On the other hand, light nodes generate transaction, but they do not maintain local copies of the Tangle. For this reason, when generating transactions, a light node queries a full node for obtaining two other transactions to confirm, and then it forwards the generated transaction to one or more full nodes. Full nodes then broadcast new transactions to their neighbors through a *gossip protocol*. With this protocol, a full node wishing to advertise the knowledge of some transactions broadcasts the identifiers of such transactions. Also, a full node wishing to update its local copy of the Tangle (because it learnt of new transactions from its neighbors) must query its neighbors for new transactions by specifying their identifiers. It is important to highlight that, even if the transactions themselves are signed, the current version of the gossip protocol does not offer any kind of transport-layer security. This means that a possible man in the middle can eavesdrop on the entire communication, and also drop, delay, or reorder messages containing transactions.

The coordinator is a special full node that periodically publishes transactions named *milestones*. Only the coordinator can produce milestones, since they are signed with a specific public key, well-known by the other nodes. The consensus on a particular transaction is finalized by means of such milestones. In particular, a transaction can be considered *finalized* only if it has been confirmed by a milestone, or by another finalized transaction.

C. Related Work

Current security research on IOTA distributed ledger has focused mainly on attacks involving IOTA cryptographic operations [8], [9], [10], [11] or IOTA consensus [12], [13], and it proposed solutions to several security problems. However, BGP hijacking and in general routing attacks against IOTA have never been analyzed to the best of the authors' knowledge.

1) Crypto Attacks: Colavita et al. [9] and, independently, Heilman et al. [10] partially cryptanalyzed the Curl-P hash function formerly used by IOTA transaction signatures, showing how it was possible to forge signatures and thus steal money from honest addresses. The disclosure of these attacks caused IOTA to abandon Curl-P for signatures and adopt instead a SHA3-based hash function in 2017. De Roode et al. [8] analyzed the replay attack, which permits an adversary to replicate multiple times a honest transaction in the Tangle, resulting in a money theft. Shafeeq et al. [11] undertaken the security issues affecting the Winternitz one-time signature scheme employed in IOTA. Both the replay attacks and the attacks against the Winternitz signature scheme are possible only in case that a IOTA address is misconfigured to be used multiple times. On the other hand, the attacks we analyze in the present paper are possible even if the full nodes and light nodes are correctly configured.

2) Consensus Attacks: Cullen et al. [12] introduced the parasite chain attack, which permits an adversary to double spend in case the IOTA distributed ledger is used with a proof-of-work consensus. The authors also propose some modifications to the IOTA consensus algorithm to improve the resistance against this attack. We note that currently, IOTA does *not* use a proof-of-work consensus, since the consensus is finalized by a centralized coordinator. Moreover, in the post-coordicide IOTA, consensus will be reached with a PBFT-like mechanism, which is immune to parasite chain attacks. On the other hand, BGP hijacking attacks can have important consequences in the availability of the IOTA mechanisms, and, in the post-coordicide IOTA, also in the consensus mechanisms (see Section V).

3) BGP Hijack Against Bitcoin: For what regards routing-level attacks [17], [18], Apostolaki et al. [15] first analyzed the impact of malicious ASes against a cryptocurrency, namely Bitcoin. The authors found that a BGP hijacker or an AS that naturally intercepts Bitcoin traffic can significantly alter the Bitcoin consensus, which is a proof-of-work one, thus leading to mining power wasting, revenue losses, and double spending. The analysis and the results of [15] *are not applicable* in the IOTA distributed ledger, whose consensus is radically different from the Bitcoin one. Due to the centralized nature of the current IOTA consensus, malicious ASes can cause various forms of denial of service rather than consensus problems.

III. Routing-Level Attacks Against IOTA

As explained in Section II-B, the gossip protocol between two full nodes is neither encrypted nor integrity-protected. This leaves space for man-in-the-middle attacks that break the integrity of the communication. A man-in-the-middle adversary cannot forge new transactions, as these are signed by the original issuer. Nevertheless, the adversary can drop one or more messages carrying transactions, resulting in a *transaction censorship*. In order to censor a transaction to a victim node, the adversary must intercept all the communications with the victim's neighbors, *in at least one of the two directions*. Indeed, if the attacker intercepts the traffic *towards* the victim, she can drop the message carrying the transaction to be censored. Otherwise, if the attacker intercepts the traffic *from* the victim, she can drop the message that requests the transaction to be censored, thus obtaining the same result in practice.

In the following we will introduce and analyze three notable attacks that can be carried out against a single node or the entire IOTA cryptocurrency by means of transaction censorship. All these attacks can be avoided by supporting secure channels between IOTA full nodes. Needless to say, secure channels does not avoid interception attacks themselves, since encrypted traffic can be intercepted as well. However, with a proper integrity-protection between full nodes, it becomes hard for an interceptor to selectively drop single messages to perform transaction censorship. The interceptor can still drop the *entire* communication, but this is more detectable by the victim nodes.

A. Address Freeze

In the *address freeze* attack, the adversary intercepts a victim full node and censors a single transaction issued by such a node. When the victim node noticed that its transaction has not been confirmed by the coordinator, it may try to re-broadcast it. If this happens, the adversary stubbornly censors the transaction again, and so on. The first impact of this attack is of course that the victim node is impeded from publishing its transaction, but the attack may have deeper consequences. Due to the signature scheme employed in IOTA, this may lead to an irreversible freeze of the IOTA source address, resulting in practice to a money destruction. Indeed, the Winternitz one-time signature scheme adopted by IOTA does not permit to use the same public key (i.e., the same IOTA source address) to sign two different transactions. An attempt to use twice the same public key leads to possible signature forging or to a complete private key compromise. Due to this, IOTA addresses must be *single-use*, in the sense that they can spend all the money they contain with a single transaction, and then they can never be used again. However, once a transaction has been signed and transmitted but censored by the adversary, the user controlling the IOTA address cannot change idea and, for example, send the same money to a different IOTA address. In other words, the victim can never *rollback* a failed transaction. The problem gets worse, for example, if the transaction was issued to buy a service with a limited time validity. In this case, the user fails to buy the service, but he also loses his money forever because he cannot use the same money again.

B. Targeted Denial of Consensus

In the *targeted denial of consensus* attack, the adversary intercepts a victim full node and censors all the milestones that its neighbors send to it. In IOTA, receiving the milestones is necessary to finalize the consensus on the state of the ledger. As a consequence, the victim node cannot be sure anymore whether the distributed ledger has reached consensus on the other received transactions. If the victim node is run by some service provider that accepts IOTA payments to provide its service, then its business is completely locked down. Note that the service provider is indeed receiving the payments from its customers, but such payments seem not to be confirmed by the coordinator, so the victim cannot finalize the consensus on them. Note that if the adversary isolates completely the full node by censoring *all* the transactions (instead of censoring only the milestones), she may not obtain the same effect, as it could be easier for the victim to notice the attack.

C. General Denial of Consensus

The *general denial of consensus* is similar to the targeted one, except that the milestones are censored directly at the coordinator level. The adversary intercepts the coordinator itself, and she censors all the milestones it produces, impeding all the full nodes from receiving them. As a consequence, the consensus in the whole IOTA cryptocurrency cannot be finalized. This attacks leverages the centralized nature of the IOTA consensus, which makes the coordinator a single point of failure. Note that, in order to perform this attack, the adversary has to locate the coordinator's subnet. The IP address of the coordinator is unknown and kept secret by the IOTA Foundation for security reasons. However, it could be discovered by a group of colluding full nodes by observing who receives the milestones first, and where they come from. Propagation source identification techniques [19] can help the adversary in this task. We leave as future work the investigation on the effectiveness of such techniques to discover the coordinator's IP address.

IV. Vulnerability of Current IOTA Network Topology

In this section, we reconstruct the current topology of the IOTA full nodes, and we evaluate the impact of malicious ASes that *naturally* intercept IOTA traffic. We do *not* experimentally evaluate the impact of ASes that intercept traffic by means of BGP hijacking in the present paper, as it would require to perform the attack on the real Internet and thus to potentially harm the honest routing functionalities. Therefore, the following results refer to a static evaluation of the reconstructed IOTA full nodes topology in the ASes network.

As the IOTA Foundation and the IOTA community do not provide a tool for reconstructing the topology of the full nodes, we used a recursive algorithm that leverages the possibility to query a full node for its neighbor list. We define S as the set of IOTA full nodes yet to be queried for neighbors. First, we fill S with nodes taken from two websites that regularly publish a list of synced IOTA full nodes, i.e., a list of full nodes with their local Tangle updated to the latest published milestone[12]. Then, at each step we select a full node in S, we query it for its neighbor nodes, and we remove it from S. Every time a full node replies with its neighbor list, we add the newly discovered full nodes in S and save the neighboring relationships between such nodes. The algorithm ends when the set S is empty, that is, when there are no full nodes to be queried any more. Note that there is the possibility that a full node is configured not to answer with its neighbor list. Thus, the reconstructed topology will be necessarily partial. However, for the following analyses it is not important to get a complete topology of the IOTA full nodes, but rather to obtain a representative data set that allows us to estimate the distribution of IOTA full nodes over the different ASes.

Figure 2 shows the topology of the IOTA full nodes, obtained with the above algorithm, represented on a world map.

Each full node has been geo-localized from its IP address, and it is represented as a point in the map. A line between two points represents a neighboring relationship between two full nodes. Nodes represented with the same color belong to the same AS. A single point in the map can represent multiple IOTA full nodes with the same IP prefix, since these are geo-localized at the same coordinates. Note that there are IOTA full nodes which apparently do not have any neighbor. These nodes are not isolated, since they reply to queries on the Tangle status, and they appear to be synchronized with the last milestones. However, they did not reply with their neighbor lists. By analyzing the IOTA full nodes connections and the membership of IOTA full nodes to their ASes, we can observe that 82.5% of IOTA connections is *inter-AS*, and only 17.5% is *intra-AS*. This means that the majority of IOTA connections can in principle be intercepted by BGP hijacking. By inspecting the map, we can also observe that most of the IOTA full nodes are gathered in Europe. Figure 2

[1] IOTA Nodes website: https://iota-nodes.net.

[2] IOTA Dance website: https://iota.dance.

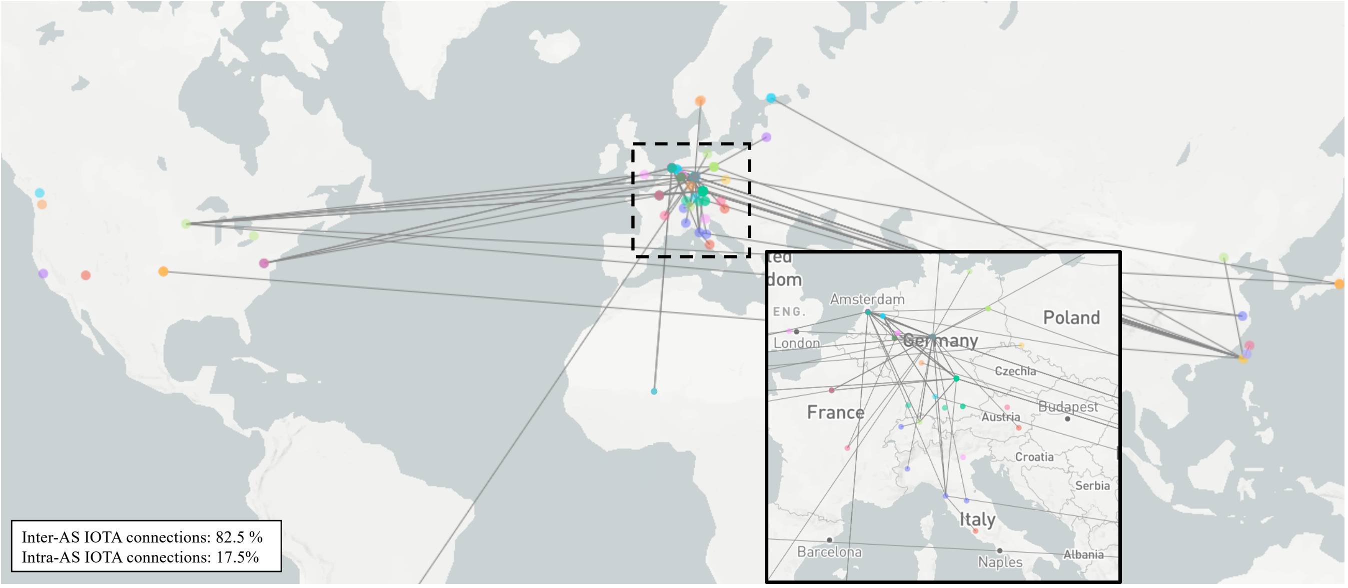

Figure 2. Global map of IOTA full nodes with neighboring relationships.

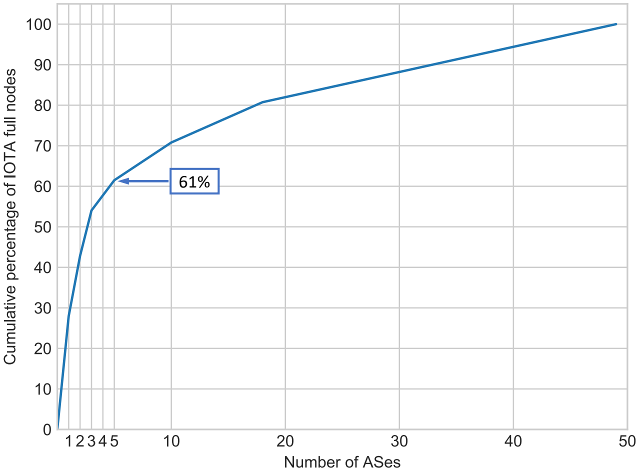

Figure 3. Percentage of IOTA full nodes vs number of their hosting ASes.

shows also a zoom of the global map in Europe. From the zoomed map we can observe that most of the IOTA full nodes are located in Germany, i.e., where the IOTA Foundation is located.

Figure 3 shows the cumulative fraction of IOTA full nodes as a function of the number of their hosting ASes. This metric has been iteratively computed by choosing at every iteration the AS that contains the greatest number of IOTA full nodes in the reconstructed IOTA topology. We see that just 5 ASes host 61% of all IOTA full nodes. As a consequence, a group of few malicious ASes could naturally intercept the connections of the majority of the IOTA full nodes. Due to this high concentration, the IOTA full nodes network traffic is very susceptible to attacks by malicious ASes.

Starting from the reconstructed topology of the IOTA full nodes, we further investigated the impact of malicious ASes against IOTA by mapping all the IOTA full node connections in the *AS-level* topology. An AS-level topology is a graph in which a node represents an AS and a link between two nodes represents the neighboring relationship between two ASes. Moreover, every link is labelled with the *business relationship* between the two ASes, which may be customer-to-provider, provider-to-customer, or peer-to-peer. We extracted the AS-level topology using the overall AS business relationships provided by CAIDA [20]. We then inferred the routes between any two ASes that host a full node by means of the algorithm described in [14], which takes into consideration the business relationships between ASes.

Figure 4 depicts the cumulative percentage of IOTA full node connections that can be intercepted by an increasing number of ASes, e.g., by colluding with each other. To obtain this metric, we iteratively picked the AS that appears in greatest number of paths related to the IOTA full nodes connections in the reconstructed IOTA topology. We see that only 5 ASes can together intercept 66% of IOTA connections. Note that, respect to Figure 3 in which 50 ASes host all the IOTA nodes, it is possible to intercept *all* IOTA full nodes connections with just 29 ASes. This can be explained by considering that large transit providers naturally intercept the majority of the Internet traffic, and thus also of the IOTA traffic.

Finally, Figure 5 depicts the percentage of IOTA full nodes having a given percentage of their connections (on

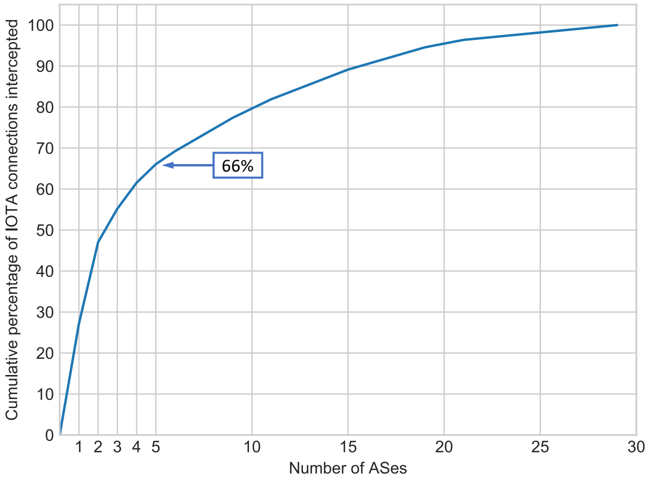

Figure 4. Cumulative percentage of full nodes connections intercepted by ASes.

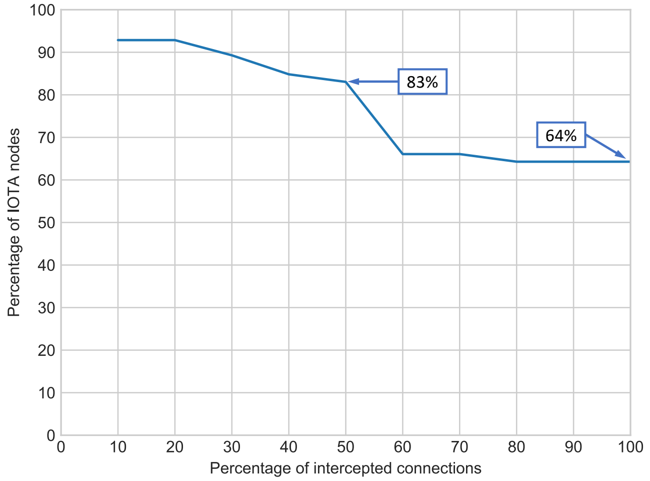

Figure 5. Percentage of IOTA full nodes having a given percentage of their connections intercepted by at least one AS *in addition to their direct provider*.

the abscissas) intercepted by at least one AS *in addition to their direct provider*. This metric has been computed by taking into account the number of connections that a IOTA full node has, i.e., we counted how many full nodes have a given percentage of their connection intercepted by at least one AS that is not their direct provider. In other words, we evaluated how many IOTA connections belonging to a IOTA full node transit trough at least one, same, AS. We found out that, for 83% of the IOTA full nodes, there is at least one AS (other than their provider) that can intercept at least 50% of their connections. Also, for 64% of IOTA full nodes, there is at least one AS (other than their provider) that can intercept *all* their connections, and thus it has the complete control over which transactions the full node receives and which not. In short, 64% of IOTA full nodes can be isolated by an AS that is not their provider.

V. Routing-Level Attacks Against Post-Coordicide IOTA

IOTA is currently a currency with a centralized consensus due to the role of the coordinator. However, the IOTA Foundation is planning to move to a fully distributed consensus, which will work without a coordinator (*coordicide*) [7]. Generally speaking, the transition from a centralized cryptocurrency to a decentralized one typically influences both the technical feasibility and the impact of the attacks from malicious ASes. In particular, the feasibility should become harder, because multiple full nodes must be intercepted in order the attack to have an appreciable effect, but the effect should become more devastating, since it changes from denials of service to more dangerous consensus faults (e.g., double spending, forking, etc.). In the following, we will analyze these techniques and their resistance against malicious ASes.

A. Auto-Peering and Node Discovery

The post-coordicide IOTA full nodes will have an autopeering and an automatic node discovery features. The aim is to avoid that some malicious entity can either force any two full nodes to become peers, or foresee that any two full nodes will become peers in the future. In other words, no one should be able to either “drive” the IOTA network topology to assume specific configurations, or to foresee the future IOTA network topologies. This is achieved by making the node peering relationships driven by verifiably random variables. The main objective is to avoid *eclipse attacks*, in which a group of colluding full nodes manages to become all the neighbors of a victim full node, in such a way all the contacts between the victim node and the rest of the network is intercepted by the colluding nodes. Unfortunately these features do not help in preventing attacks by malicious ASes able to perform BGP hijacking to attract a full node’s traffic, whatever its neighbors are. Note also that the autopeering mechanism of the post-coordicide IOTA network will choose neighbors in a *verifiable* random manner, so once a full node has randomly chosen its neighbors, such neighbors will be necessarily public. Given that the adversary knows the IP prefix of the victim node and of its neighbors, then all the outgoing and incoming traffic are hijackable in principle. An exception is in case the victim full node and the neighbor belong to the same AS, so that they use intra-AS routes that are not affected by BGP hijacking. Note however that we cannot force full nodes to prefer neighbors in the same AS in order to increase resistance against BGP hijacking, because this would reduce the IOTA network’s global connectivity.

B. Fast Probabilistic Consensus

At the time of writing, it is not clear from the last coordicide white paper [7] which consensus algorithm will be adopted by the IOTA Foundation for the post-coordicide IOTA. Possible options are *Fast Probabilistic consensus* (FPC) and *Cellular Consensus*. Since FPC is treated in much more detail than Cellular Consensus in the white paper, we believe that FPC is the most probable choice. So we focus on FPC in the present analysis. The FPC is a Byzantine Fault-Tolerance protocol that promises to be more lightweight than traditional BFT protocols at the cost of reaching consensus only with high probability, and not with probability one. In the presence of two conflicting transactions (e.g., two transactions spending the same money), FPC will lead full nodes to reach (with high probability) a consensus on which transaction is "good". Each full node starts with an initial opinion regarding the "goodness" of the transaction. The initial state is computed from the time of arrival of the transaction at the full node. If the node received the transaction at time t, no received transaction is conflicting with it, and no transaction conflicting with it is received until time $t + \Delta$, then the node assumes the initial opinion: "transaction is good". Otherwise, it assumes the initial opinion: "transaction is bad". Then the node divides the FPC algorithm in rounds, in each of which it retrieves the value of a distributed computed random threshold ratio X. Successively, the node asks to k randomly chosen full nodes their opinion on the "goodness" of the transaction. If a node does not respond, the asking node will assume that the other node is down or faulty. As a consequence, it will randomly chose another node to ask to. If more than $X \cdot k$ of the queried nodes has the opinion "transaction is good", then the node assume the opinion: "transaction is good". Otherwise, it assumes the opinion: "transaction is bad". After enough rounds in which the node has not changed its opinion, such an opinion is finalized into a consensus value.

In the presence of an adversary capable of performing BGP hijacking and establishing a man in the middle between full nodes, both the initial opinion and the finalized opinion can be manipulated. The initial opinion can be manipulated by simply delaying a transaction A which conflicts with another transaction B, in such a way to induce the nodes to assume the initial opinions: "transaction A is bad" and "transaction B is good". The finalized opinion can be manipulated by dropping those opinions that the adversary dislikes, and forwarding the ones that the adversary likes. Note that if the message carrying an opinion is dropped by a man in the middle, the asking node will randomly chose another node to ask to. If there exist at least k full nodes whose opinion is liked by the adversary, then the victim node will eventually ask to these nodes. As a consequence, the victim node will convince itself that the vast majority of the network has such an opinion, whereas the converse may be true. With the above techniques, a BGP hijacker cn in principle break the Fast Probabilistic Consensus and induce at least a couple of full nodes to finalize on a different consensus value, thus causing double spending.

VI. Conclusions

In this paper, we presented a first security analysis of IOTA with respect to Internet routing attacks. In particular we identified three notable attacks that a malicious AS can mount against IOTA, namely *address freeze*, *targeted denial of consensus* and *general denial of consensus*. Such attacks can lead, respectively, to a money loss, an interruption of the consensus mechanisms for a given victim node, and an interruption of consensus for the whole IOTA cryptocurrency. We also reconstructed the current IOTA network topology and we study how much vulnerable such a network is with respect to possible malicious ASes. Finally, basing on the last published IOTA coordicide white paper [7], we also analyzed the security of *post-coordicide* IOTA with respect to malicious ASes, and we showed some possible vulnerabilities that may eventually lead to double spending. Based on the above analyses, we conclude that IOTA could solve many of its routing-level vulnerabilities by supporting secure channels between its nodes. Secure channels do not avoid traffic interception, but they impede malicious ASes to selectively drop messages, which is the key to mount the attacks presented herein.

Acknowledgment

This work was partially supported by the Italian Ministry of Education and Research (MIUR) in the framework of the CrossLab project (Departments of Excellence) and by the projects PRA_2018_81 ”Wearable sensor systems: personalized analysis and data security in healthcare”, PRA_2020_92 ”Quantum computing, technologies and applications” funded by the University of Pisa.

References

[1] S. Popov, "The tangle," 2016.

[2] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system."

[3] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, no. 2014, pp. 1–32, 2014.

[4] B. Shabandri and P. Maheshwari, "Enhancing IoT security and privacy using distributed ledgers with iota and the tangle," in *2019 6th International Conference on Signal Processing and Integrated Networks (SPIN)*, 2019, pp. 1069–1075.

[5] J. Brogan, I. Baskaran, and N. Ramachandran, "Authenticating health activity data using distributed ledger technologies," *Computational and Structural Biotechnology Journal*, vol. 16, pp. 257 – 266, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S2001037018300345

[6] A. Arena, P. Perazzo, and G. Dini, “Virtual private ledgers: Embedding private distributed ledgers over a public blockchain by cryptography,” in *Proceedings of the 23rd International Database Applications Engineering Symposium*, ser. IDEAS ’19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: https://doi.org/10.1145/3331076.3331083

[7] S. Popov, H. Moog, D. Camargo, A. Capossele, V. Dimitrov, A. Gal, A. Greve, B. Kusmierz, S. Mueller, A. Penzkofer *et al.*, “The coordicide,” 2020.

[8] G. De Roode, I. Ullah, and P. J. M. Havinga, “How to break IOTA heart by replaying?” in *2018 IEEE Globecom Workshops (GC Wkshps)*, Dec 2018, pp. 1–7.

[9] M. Colavita and G. Tanzer, “A cryptanalysis of IOTA’s curl hash function,” 2018.

[10] E. Heilman, N. Narula, G. Tanzer, J. Lovejoy, M. Colavita, M. Virza, and T. Dryja, “Cryptanalysis of Curl-P and other attacks on the IOTA cryptocurrency,” 2019.

[11] S. Shafeeq, S. Zeadally, M. Alam, and A. Khan, “Curbing address reuse in the IOTA distributed ledger: A cuckoo-filter-based approach,” *IEEE Transactions on Engineering Management*, pp. 1–12, 2019.

[12] A. Cullen, P. Ferraro, C. K. King, and R. Shorten, “Distributed ledger technology for iot: Parasite chain attacks,” *CoRR*, vol. abs/1904.00996, 2019. [Online]. Available: http://arxiv.org/abs/1904.00996

[13] Winston (pseudonym), “XY attack vector (IOTA’s version of the 34% attack). [forum post],” https://forum.helloiota.com/469/XY-Attack-Vector-IOTAs-version-of-the-34-attack, 2017.

[14] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, “How secure are secure interdomain routing protocols,” *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 4, p. 87–98, Aug. 2010. [Online]. Available: https://doi.org/10.1145/1851275.1851195

[15] M. Apostolaki, A. Zohar, and L. Vanbever, “Hijacking bitcoin: Routing attacks on cryptocurrencies,” in *2017 IEEE Symposium on Security and Privacy (SP)*, May 2017, pp. 375–392.

[16] Y. Rekhter, T. Li, and S. Hares, “A border gateway protocol 4 (bgp-4),” Internet Requests for Comments, RFC Editor, RFC 4271, January 2006.

[17] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, “BGP hijacking classification,” in *2019 Network Traffic Measurement and Analysis Conference (TMA)*, June 2019, pp. 25–32.

[18] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, “ARTEMIS: Neutralizing BGP hijacking within a minute,” *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2471–2486, Dec 2018.

[19] J. Jiang, S. Wen, S. Yu, Y. Xiang, and W. Zhou, “Identifying propagation sources in networks: State-of-the-art and comparative studies,” *IEEE Communications Surveys Tutorials*, vol. 19, no. 1, pp. 465–481, 2017.

[20] “The caida as relationships dataset,” https://www.caida.org/data/as-relationships/.