Achieving Fairness in the Tangle through an Adaptive Rate Control Algorithm

Luigi Vigneri
IOTA Foundation
10405 Berlin, Germany
luigi.vigneri@iota.org

Wolfgang Welz
IOTA Foundation
10405 Berlin, Germany
wolfgang.welz@iota.org

Alon Gal
IOTA Foundation
10405 Berlin, Germany
alon.gal@iota.org

Vassil Dimitrov
IOTA Foundation
10405 Berlin, Germany
vassil@iota.org

***Abstract*—Throughput is a key property for any distributed ledger technology. However, limited resources, such as bandwidth or node computational power, can lead to network congestion when nodes try to issue more transactions than the network can handle. Consequently, priority criteria are necessary to determine whether a transaction should be accepted or not. In this paper, we present a novel adaptive rate control algorithm for the Tangle, a new-generation distributed ledger allowing large throughput. Our approach combines various concepts, such as resource tests and Proof-of-Work with dynamic difficulty. Our algorithm not only serves as an anti-spam mechanism, but also achieves fair representation. This is to be contrasted with pure Proof-of-Work blockchains, which lead to wasteful mining races.**

I. Introduction

Nakamoto's seminal work on Bitcoin proposed a solution, called *blockchain*, based on Proof-of-Work (PoW) [1] to solve the double spending problem in a decentralized payment system [2]. The growing interest in cryptocurrencies has soon revealed blockchain scaling limitations (e.g., according to [3], Bitcoin only achieves 7 transactions/sec maximum throughput). To improve scalability, researchers have been proposing several alternatives to the Bitcoin blockchain [4]–[6]. In this work, we focus our attention on one of these alternatives, namely the *Tangle* [7], a data structure for storing transactions developed by IOTA [8]. However, network resources (e.g., bandwidth, computational power, disk space) are limited: while PoW-based blockchains come with a built-in rate limitation enforced by the mining difficulty adjustment [2], [4] (but leading to undesirable side effects such as mining races), in the Tangle an explicit rate control mechanism is necessary in order to ensure the network does not exceed its maximum capacity.

In this preliminary paper, we propose a novel adaptive rate control algorithm based on the following ideas:

(i) issuing a transaction requires some computational effort;
(ii) every node in the network is allowed to issue transactions independently of its computational power;
(iii) the difficulty to issue multiple transactions in a short time interval progressively increases.

A mechanism based on (i) acts as an anti-spam technique since malicious nodes are prevented from harming the network through spam attacks. What is more, properties (ii)-(iii) guarantee that the Tangle can achieve a certain degree of fairness, which means that any node (even with low hashing power) has a non-negligible probability to have its transactions approved (see Section II for further information about the approval mechanism). To satisfy the above requirements, nodes are asked to solve a PoW[1] such that its difficulty depends on the number of transactions issued in the past time window and on the collateral owned by the node (Section III-B). This mechanism requires node accountability where each transaction can be associated with the global identity of its issuing node (Section III-A).

978-1-7281-1328-9/19/$31.00 ©2019 IEEE

II. Tangle model

Approval mechanism. In the Tangle, when a node wants to issue a new transaction, it is asked to approve the correctness of two other transactions in the network. This mechanism generates a directed acyclic graph where vertices represent transactions, and edges represent the approval relation. Specifically, edge (x, y) implies transaction x approves transaction y. The details of this procedure are out of the scope of the paper, and we refer the interested reader to [7] for further information.

Maximum throughput. The maximum acceptable throughput in the Tangle depends on the limited resources available. This threshold can be deduced from stress testing the network and seeing its practical maximum throughput, i.e., the maximum rate at which most nodes still stay in sync. We consider this value as an input parameter known by the system manager[2].

Partial synchrony. We consider a bound h on the network latency which means that, if a message is sent at time t, then all nodes will receive the same message within time $t + h$. Note that our definition of latency does not include the time spent for PoW.

III. Rate control algorithm

A. Proof-of-Identity

In order to implement a rate control mechanism in a distributed system, certain transactions issued by specific nodes

[1]We highlight that consensus in the Tangle is reached independently of the PoW which only serves as an anti-spam mechanism.

[2]Although we are dealing with a distributed architecture for transactions, we assume the presence of a central entity, called the system manager, that initializes the network and sets the system parameters.

must be ignored and not further processed by the network. As such, it is necessary to introduce global node identities. In our infrastructure, we envision a node authentication mechanism which does not require any centralized data structure, as this would obviously break the distributed nature of the framework. A potential solution to guarantee global identities is to use common public key cryptography to sign a transaction and to link it to its issuing node in a tamper proof way. Specifically, we require that the issuing node add its public key to every transaction. This way, every node can verify the authenticity of the issuing node.

When each node has an identity, a distributed system becomes vulnerable to Sybil attacks [9], where a malicious entity masquerades many counterfeit identities and uses them to overcome the rate control mechanism to launch a coordinated assault or spam the network. One way to make such an attack harder is the so-called resource testing, where each identity has to prove the ownership of certain difficult-to-obtain resources. Since in the cryptocurrency world users own a certain amount of tokens (*collateral*), we propose a Sybil protection mechanism inspired by Proof-of-Stake [10], where any node with a minimum amount of such a collateral is allowed to issue transactions. However, unlike the original Proof-of-Stake, in our proposal collateral does not leave the users' possession and, thus, cannot be forfeited. As the collateral is linked to its node's identity, everyone in the network can verify its amount. We refer to the above mechanism as Proof-of-Identity.

B. Adaptive rate control algorithm

In a pure PoW-based architecture, a high difficulty value would prevent low-power nodes from issuing transactions, which is not desirable especially in the context of Internet-of-Things; on the other hand, low difficulty can quickly lead to network congestion. We propose an adaptive PoW algorithm to allow every node to issue transactions while penalizing spamming actions.

In our algorithm, when a node decides to issue a transaction, it must solve a cryptographic puzzle where the difficulty is a function of the collateral owned and of the number of transactions issued recently. Assume node i generates n_i^T transactions in the previous T time units where $T \gg h$. The same node i has to set the difficulty of the PoW to d_i defined by

$$d_i = d_0 + w(s_i, n_i^T),$$

where d_0 is a basic difficulty, and $w : \mathbb{R}^+ \to \mathbb{R}^+$ is a function that depends on the collateral s_i and on n_i^T. The time window T, the difficulty d_0 and the weight function $w(\cdot)$ are parameters chosen depending on the fairness level that the system manager wants to achieve: for instance, we expect that powerful nodes are penalized when the time window T becomes large as they need more work to issue several transactions in a short time.

As an additional security measure, we require that the total number of transactions issued by a user is limited, i.e.,

$$n_i^T \leq z(s_i), \qquad \forall i, \tag{1}$$

where $z : \mathbb{R}^+ \to \mathbb{R}^+$ is a function that depends on the collateral s_i such that the larger the stake owned by a node, the higher the number of transactions the same node can issue. The threshold of Eq. (1) brings a twofold benefit: first, it ensures that even a user with infinite computational power cannot arbitrarily spam the network; second, when $z(\cdot)$ is a superlinear function on the collateral s_i, multiple low collateral identities have a lower total threshold (the sum of thresholds of all identities) than a single node with a large collateral, hence discouraging Sybil attacks.

C. Implementation details

For the sake of simplicity, we assume incoming transactions are checked in the same order as they are issued by the sending node. As the expected time needed to perform the PoW is typically much larger than the network latency h, this is a reasonable assumption.

When a transaction is seen for the first time, the node stores the id of the node issuing the transaction, the timestamp t_0 and the PoW difficulty. The identity id of the issuing node as well as its collateral s_{id} can be determined using the methods described in Section III-A. Based on this information, it can then be checked that the number of transactions issued in the recent T time units by the same node does not exceed the allowed maximum $z(s_{id})$ based on its collateral s_{id} and that the difficulty of the most recent transaction is indeed sufficient. This idea is more formally described in Algorithm 1.

Algorithm 1: Rate control algorithm

Input: incoming transaction t, set of known transactions $\mathcal{X}$, time window T, basic difficulty d_0, weight function $w(\cdot)$.

Output: forward or ignore t.

$t_0 \leftarrow$ `time(t)`;
$id \leftarrow$ `nodeId(t)`;
$\mathcal{T} \leftarrow t' \in \mathcal{X}$ such that `time(t')` $\in (t_0 - T, t_0]$ and `nodeId(t')` $= id$;
if $|\mathcal{T}| < z(s_{id})$ **then**
  if `difficulty(t)` $\geq d_0 + w(s_{id}, |\mathcal{T}|)$ **then**
    return forward t;
return discard t;

IV. Conclusion and future research

The discrepancy between smaller general purpose devices and optimized hardware with respect to the PoW performance is several orders of magnitude. Hence, any rate control based on PoW (as in the current IOTA implementation of the Tangle [8]) would eventually leave smaller devices behind. Conversely, a purely stake-based system would lead to a centralization where only the "rich" parties can participate. In this preliminary work, we have proposed a rate control algorithm which combines the above two approaches achieving a reasonable compromise: slow nodes or users with low

collateral can issue (a few) transactions at inexpensive prices, while at the same time faster users cannot spam the network due to a limitation on burst of transactions.

The proposed framework naturally paves the way to interesting future research directions. For example, the need for a minimum collateral necessary to issue transactions (see Section III-A) introduces the following dilemma: on the one hand, a low threshold allows more users to issue transactions but it does not protect sufficiently against the creation of counterfeit identities; on the other hand, a large threshold would drastically reduce the number of potential nodes participating in the network at the cost of a larger security. Another potential research direction is to analyze how to properly tune the parameters d_0, $w(\cdot)$ and $z(\cdot)$ to achieve a certain efficiency against spam or Sybil attacks. However, since the choice of these parameters also affects the fairness metric, such trade-offs are non-trivial and require specific analytical models.

REFERENCES

[1] C. Dwork and M. Naor, "Pricing via processing or combatting junk mail," in *Annual International Cryptology Conference*. Springer, 1992, pp. 139–147.

[2] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[3] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba *et al.*, "On scaling decentralized blockchains," in *Financial Cryptography and Data Security*, J. Clark, S. Meiklejohn, P. Y. Ryan, D. Wallach, M. Brenner, and K. Rohloff, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 106–125.

[4] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, pp. 1–32, 2014.

[5] O. Smith, "Bitcoin price rival: Cryptocurrency 'faster than bitcoin' will challenge market leaders," 2018. [Online]. Available: https://www.express.co.uk/finance/city/907536/cryptocurrency-bitcoin-market-leader-cash-ripple-ethereum

[6] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. D. Caro *et al.*, "Hyperledger fabric: A distributed operating system for permissioned blockchains," *CoRR*, vol. abs/1801.10228, 2018. [Online]. Available: http://arxiv.org/abs/1801.10228

[7] S. Popov, "The tangle," *cit. on*, p. 131, 2016.

[8] "The next generation of distributed ledger technology." [Online]. Available: https://www.iota.org/

[9] J. R. Douceur, "The sybil attack," in *International workshop on peer-to-peer systems*. Springer, 2002, pp. 251–260.

[10] S. King and S. Nadal, "Ppcoin: Peer-to-peer crypto-currency with proof-of-stake," *self-published paper, August*, vol. 19, 2012.