

Feedback control for distributed ledgers: An attack mitigation policy for DAG-based DLTs

P. Ferraro, A. Penzkofer, C. King, R. Shorten

Abstract—In this paper we present a feedback approach to the design of an attack mitigation policy for DAG-based Distributed Ledgers. We develop a model to analyse the behaviour of the ledger under the so called *Tips Inflation Attack*, which endangers the liveness of transactions, and we design a control strategy to counteract this attack strategy. The efficacy of this approach is showcased through a theoretical analysis, in the form of two theorems about the stability properties of the ledger with and without the controller, and extensive Monte Carlo simulations of an agent-based model of the distributed ledger.

I. INTRODUCTION

Bitcoin, and the technology that underpins it, *Blockchain*, have become, in the past years, a source of great debate and controversy in both business and scientific communities. The objective of Distributed Ledger Technologies (DLTs) (the agnostic term for Blockchain and related technologies) is to create a protocol to add entries to a spreadsheet where a record of transactions and other account information is transcribed, accessible and (potentially) owned by every node of a Peer to Peer (P2P) network, with an intrinsic mechanism to enforce consensus among its users. While conceptually simple in its formulation, due to their trust-less and decentralised nature Distributed Ledger Technologies (DLTs) have to operate in highly adversarial environments that are characterised by large numbers of untrustworthy agents [1]. This makes the design of DLTs quite complex as design choices need to take into account threats coming from a very large space of possibilities. Indeed much of the recent literature on this topic is concerned with developing DLTs that are robust in the presence of such actors [2][3]. In addition to a large body of work on identifying attack scenarios and vulnerabilities in ledger architectures work on attack mitigation has developed along two main directions. One approach, based on game theory is to design incentive mechanisms for ledgers that make certain attacks prohibitively expensive for attackers, so that the resulting protocol is characterised by a Nash equilibrium [4][5]. A second approach is based on developing techniques to identify specific attacks and to make ad-hoc modifications to the protocol [6][7]. Often these approaches affect the basic operation of the ledger structure in order to eliminate attack vulnerabilities.

We propose a new direction in this work in which concepts from control theory can be used to develop policies that protect ledgers against certain adversarial attacks. The rationale for developing such strategies is that they preserve the operational efficiencies of the basic ledger in honest environments, but also provide protection when needed (potentially at the cost of some operational efficiency during periods of attack). The protection mechanism draws from the capability of nodes to agree, with high accuracy, on some parameters of the ledger. This knowledge is then utilised by the nodes to act in a coordinated way to counter the effects of the attack.

The attack vector that motivates this work is the so called *Tips Inflation Attack* that is of concern in certain Directed Acyclic Graph (DAG) based DLTs. DAG based ledgers have emerged in recent years as a generalisation and alternative to Blockchains, in which the ledger structure evolves not as a sequence of blocks, but rather graph like [8]. The utilisation of a DAG structure allows for several interesting features, such as faster block writing times and partial ordering, all of which enable interesting use cases of DLTs [9], [10]. DAGs, however, raise several concerns with regards to the vulnerability of the data structure that are not present in classical Blockchains. For example in chain-based systems, the growth of the data structure is limited by the block issuance rate and blocks that do not extend the longest chain are removed or “orphaned”. In contrast, in DAG-based systems the extension of the data structure is less clear and depends significantly on the protocol that is adopted for the validation of new blocks¹ This enables new attack vectors to the data structure itself that must be studied and, where possible, averted. This paper describes an attack on a well known DAG-based DLT, namely IOTA [11]. In this attack the adversary attempts to modify the block validation process in the block-DAG, which is also called *Tangle*, such that the number of *orphaned* blocks grows unbounded. The basic idea is that at some time a malicious agent that obtains a proportion $q \in [0, 1]$ of the writing access to the DAG, starts publishing blocks that approve only own and old blocks. This, as we will see, results in an increase to the number of unconfirmed blocks or “*tips*” in the Tangle, potentially leading to an instability of the system.

Specifically, the contributions of this paper are as follows.

- We perform a theoretical analysis of the Tips Inflation Attack. Sufficient conditions for its success are provided in the form of a stability theorem for the Tangle.

Pietro Ferraro and Robert Shorten are with Dyson School of Design Engineering, Imperial College London, London.

Andreas Penzkofer and Pietro Ferraro are with the IOTA Foundation, Berlin. Christopher King is with Department of Mathematics, Northeastern University, Boston.

All the code used to generate the figures in this work can be found at this [link](#).

¹For example, in [11] the author proposes only to extend the data structure at places that “support” the preferred state of the ledger.

- A modification of the IOTA protocol, inspired by control theory. The efficacy of this approach is shown through theoretical results and extensive simulations.

We note that this work builds on top of previous research [12], [6], in which the authors derived a fluid model for the Tangle [12] and proposed an alternative tip selection algorithm [6] to ensure the stability of the network. The aim of [12] and [6] is to study the behaviour of the Tangle under nominal conditions whereas this work focuses on the Tangle under adversarial conditions.

II. STRUCTURE OF THE PAPER

The remainder of this paper is organised as follows. Sections III and IV present the original version of the IOTA protocol and the Tips Inflation Attack. In Section V we present the mathematical model used to analyse the attack. In Section VI we introduce a modification of the tip selection algorithm that offers a solution to the attack. In Section VII we validate the theoretical results through extensive Monte Carlo simulations and finally, in Section VIII we summarise our findings and present future lines of research.

III. THE TANGLE

We are interested in a particular DLT architecture that makes use of a Directed Acyclic Graph (DAG) $D = (V, E)$, with vertex set V and directed edges set E , to achieve consensus on a shared ledger. A particular instance of a DAG based DLT is introduced in [11]. In this DAG, called *Tangle*, the vertices are blocks issued by network participants, called *nodes*, and each block contains at most one transaction.² The edges are formed by blocks approving previous blocks that we call *parents*. Since the approval can be done only toward past blocks, the graph is acyclic, starting with a first block, called the *genesis*. All yet unapproved blocks are called *tips* and the set of all unapproved blocks is called the *tips set*. A node selects tips to approve through a *tip selection algorithm*. To issue a block a node selects k blocks from the tips set.³ This process, called *approval*, is represented by an edge in the graph. If there exists a directed edge from vertex i to j , we say j is *directly approved* by i . If there is a directed path from i to j we say that j is *indirectly approved* by i .

Due to various factors, there is a delay between the selection of the parents of a block and their removal from the tip pool. First, there may be delays at the issuing node due to Proof of Work (PoW). The role of PoW is to prevent malicious users from spamming the network, thus, the required PoW is less computationally intense than in Blockchain counterparts [13], and can be easily carried out by common IoT devices, e.g. smartphones and smart appliances. Second, the time until the parents are removed in the tip pools of other nodes is delayed by propagation times and processing. As a consequence it is possible that a block is approved several times by multiple other blocks. This also leads to each node having an own local

version of the Tangle. We call the period between selection of the parents and the appearance in the tip pool, the *delay period*.

As an example of this process, refer to figure, Figure 1.

A. Tip selection algorithm

In the IOTA protocol the participants in the network, i.e. nodes, are the creators of blocks. In order to extend the Tangle and advance the confirmation of blocks, nodes are required to select previous blocks and reference these [14].

In [15] new nodes select tips as on the basis of k weighted random walks, from the interior of the DAG to the tips set. It can be proven that this protocol represents a Nash equilibrium for the Tangle, which each node is incentivised to follow, as any deviation from this tip selection algorithm results into orphanage of the own blocks. We note the algorithm protects against an adversary in the presence of an honest majority of issuers. However, the computational cost of running the random walk algorithm in the graph is prohibitive in an IoT scenario.

A more computationally efficient algorithm for selecting parent blocks is the *uniform random tip selection* (URTS), in which each tip has the same probability of being selected by a new block. However, in order to prevent penalisation of block issuers for how they append, it must be ensured that the attachment location has no or little impact on the liveness of the blocks. This can be, for example, achieved by operating a consensus mechanism prior to adding blocks to the Tangle, e.g. [16], which ensures that no conflicts enter the tip set.

IV. TIPS INFLATION ATTACK

In the previous section we introduced the Tangle and how it can be constructed using a URTS algorithm. Under honest conditions this leads to a stable tip pool size. We now describe an attack that attempts to increase the tip pool size, potentially unbounded.

Typically blocks in the Tangle reference more than one parent in order to keep the tip pool size bounded. In the attack an adversary modifies the tip selection algorithm, such that the adversary's blocks only approve old blocks (which are not tips). This may be, for example, the genesis or alternatively "old" blocks. Given that the writing power of the honest users is $p \in [0, 1]$, if an attacker manages to gain a proportion $q = 1 - p$ of the writing power of the network, then for sufficiently large q , the attacker can artificially increase the size of the tips set arbitrarily, provided that the attack is maintained.⁴ Figure 2 shows an example of this process. This would lead to liveness problems, where honest blocks are not approved for an extensive amount of time or, in the worst case not approved at all.

One may consider, as a basic protection, to only allow nodes to approve blocks with a certain maximum age, i.e. blocks that are issued within a certain recent time Δ . However,

²This limitation is not a necessity and the number of transactions per block can, in principle, be increased.

³Nodes are not necessarily obliged to select a given number of parents. But it is reasonable to expect that most of them will follow the suggested number.

⁴The writing power q depends on the specific protocol that the network is following: it can be the computing power of the attacker with respect to the rest of the network or some weight [17] that is associated with a specific node. For the purpose of this work, the means through which an attacker gets writing power q is irrelevant.




Fig. 1: Sequence to issue a new block: an instance of the Tangle with three new incoming blocks (left panel). Blue blocks have already been approved, red blocks represent the current tips of the Tangle and grey blocks are incoming, i.e. are in the delay period. The newly arriving tips remain unconfirmed (dashed lines) until this process is over. Once the delay period is over, the selected tips become confirmed vertices and the grey blocks are added to the tips set (right panel).




Fig. 2: Example of the Tips Inflation Attack, in the extreme case in which all the new blocks do not approve tips. The amount of tips keeps increasing over time.

this would not mitigate the problem, as the attacker would be effectively decreasing the chance of honest block to be selected, regardless. What is more, since this introduces an *expiration time* this may lead to orphanage of blocks, including the honest ones. Accordingly, in this paper we consider a version of the protocol with no constraints on the maximum allowed age Δ .

Remark: the reader might notice some similarities between a Denial of Service (DoS) attack [18] and the Tips Inflation attack. While DoS and tips inflation attacks might technically look similar, the difference between the two is that in the former a user would be denied access to a service provided by a server due to the congestion of the server. In this instance the tips inflation attack does not prevent a user to access the Tangle but it increases the number of tips to the point where honest blocks are no longer approved.

V. MODEL

As mentioned in the previous section, there is a delay period between the selection of the parents for a block and the appearance of the block in the tip pool. Generally, the time of appearance for a given block is different for each node, i.e. each node has an own local version of the Tangle. However, for simplicity, we assume that each node has instantaneous reading access to a “global” Tangle, while the writing access is delayed by h . Thus, during a time period h the approvals of the selected tips are pending and tips may still be available for selection by other new blocks, although the block has already been selected by a node. After the delay time h the new directed edges are added to the global graph, directed from the new site to its parent sites. After this point the parent sites are no longer tips, and thus are no longer available for selection by other

new blocks.⁵ In an honest environment, the URTS leads to a stabilisation of the global Tangle tip pool size at the value $L = \frac{k}{k-1} \lambda h$, where λ is the block issuance rate [14].

Our model involves the following variables:

- 1) $L(t)$ is the number of tips at time t
- 2) $W(t)$ is the number of ‘pending’ tips at time t which are being considered for approval by some new block(s)
- 3) $X(t) = L(t) - W(t)$ is the number of ‘free’ tips at time t
- 4) T_a is the time when block a is created
- 5) $N(t)$ is the number of blocks created up to time t
- 6) $U(T_a) \in \{0, 1, 2\}$ is the number of free tips selected for approval by block a at time T_a

We have the relations

$$N(t) = \sum_{a: T_a \leq t} 1 \quad (1)$$

$$W(t) = \sum_{a: t-h < T_a \leq t} U(T_a) \quad (2)$$

$$X(t) = N(t-h) - \sum_{a: T_a \leq t} U(T_a) \quad (3)$$

$$L(t) = N(t-h) - \sum_{a: T_a \leq t-h} U(T_a) \quad (4)$$

Assuming the uniform random tip selection algorithm, $U(T_a)$ is a random variable whose distribution depends only on the values of X , W and L at time T_a . (For convenience we assume that block a is created at time $T_a + 0$, so it sees the state of the Tangle at time T_a). The expected value of $U(T_a)$ is

$$\mathbb{E}[U(T_a)] = 2 \frac{X(T_a)}{L(T_a)} - \frac{X(T_a)}{L(T_a)^2} \quad (5)$$

⁵If a block z is referenced by two blocks x and y , at times t_x and t_y respectively, where $t_x < t_y$, then block z ceases to be a tip at $t_x + h$.

A. Fluid Model for two selections

In order to gain some understanding of the system of equations (1)-(4) we consider the asymptotic regime of large arrival rate, where the time between consecutive blocks is very small. In this regime it should be reasonable to approximate the system (1)-(4) by a fluid model.

We introduce a scaling parameter λ so that the arrival rate is proportional to λ , and let $\lambda \rightarrow \infty$ to reach the fluid model. The rescaled variables $\{\lambda^{-1}L(t), \lambda^{-1}X(t), \lambda^{-1}W(t)\}$ are assumed to converge to deterministic limits as $\lambda \rightarrow \infty$, and the limits are represented in the fluid model by real-valued functions $\{l(t), x(t), w(t)\}$. The creation of new blocks in the fluid model is described by an arrival rate $a(t)$ so that $\int_0^t a(s) ds$ corresponds to the limit of $\lambda^{-1}N(t)$. Furthermore the variable $U(T_a)$ is replaced by its time average (over a short time interval), which by the law of large numbers is equivalent to the ensemble average (5). By rescaling variables and letting $\lambda \rightarrow \infty$, this expected value converges to $u(t) = 2x(t)/l(t)$. Referring to the equation (3) the change of X over a small time increment γ can be approximated as

$$\begin{aligned} X(t+\gamma) - X(t) &= N(t-h+\gamma) - N(t-h) - \sum_{a:t < T_a \leq t+\gamma} U(T_a) \\ &\simeq \gamma \lambda a(t-h) - \gamma \lambda a(t) \frac{2x(t)}{l(t)} \\ &= \gamma \lambda \left[a(t-h) - a(t) \frac{2x(t)}{l(t)} \right] \end{aligned}$$

Applying similar reasoning to the other equations we get the following set of delay differential equations (DDE) for the fluid model:

$$\frac{dx}{dt}(t) = a(t-h) - a(t)u(t) \quad (6)$$

$$\frac{dl}{dt}(t) = a(t-h) - a(t-h)u(t-h) \quad (7)$$

$$w(t) = l(t) - x(t) = \int_{t-h}^t a(s)u(s)ds \quad (8)$$

where

$$u(t) = \frac{2x(t)}{l(t)} \quad (9)$$

The solution of these equations $\{x(t), l(t), w(t)\}$ can be interpreted as a fluid model which describes the dynamics of the Tangle with very high arrival rate, using the random tip selection algorithm. Note that the DDE system (6) - (8) must be supplemented by initial conditions in the interval $[-h, 0]$.

B. Fluid Model for k selections and Tips Inflation Attack

In the presence of a Tips Inflation Attack, a malicious actor will attach a proportion q of the new tips to blocks that have already been referenced, effectively reducing the number of tips that are removed from the tips set. Accordingly, using the same derivation outlined in the previous section, we can derive a fluid model that describes the dynamics of the Tangle with very high arrival rate, using a random tip selection algorithm but now with k parents:

$$\frac{dx}{dt}(t) = a(t-h) - p a(t) u(t) \quad (10)$$

$$\frac{dl}{dt}(t) = a(t-h) - p a(t-h) u(t-h) \quad (11)$$

$$w(t) = l(t) - x(t) = \int_{t-h}^t p a(s) u(s) ds \quad (12)$$

where

$$u(t) = \frac{k x(t)}{l(t)} \quad (13)$$

Considering this new model, for a constant arrival rate $a(t) = a(t-h) = \beta$, it is straightforward to derive the equilibrium $(\hat{x}, \hat{l})$ of the system at steady state. From relations 10-11 we obtain:

$$0 = 1 - p k \frac{\hat{x}}{\hat{l}} \quad (14)$$

while from relation (12) we obtain:

$$\hat{l} - \hat{x} = \int_{t-h}^t p \beta k \frac{\hat{x}}{\hat{l}} ds = p \beta k \frac{\hat{x}}{\hat{l}} (t - t + h) = \beta p k h \frac{\hat{x}}{\hat{l}} \quad (15)$$

These lead to

$$\hat{x} = \frac{h\beta}{pk-1} \quad (16)$$

$$\hat{l} = \frac{pkh}{pk-1} \beta \quad (17)$$

$$\hat{w} = \frac{pkh-h}{pk-1} \beta \quad (18)$$

Notice that for $p = 1$ (corresponding to no attack) and $k = 2$ the equilibrium for the tips is $\hat{l} = 2h\beta$.

Notice also, that for $p \rightarrow \frac{1}{k}$, $\hat{l} \rightarrow \infty$ and for all (p, k) such that $pk - 1 < 0$, $\hat{l} < 0$.

Remark: Notice that we made the assumption that the arrival rate β is constant. While this might appear as a simplification, the Tangle is designed to work close to capacity all the time (i.e. the amount of blocks that are added to the Tangle is the maximum allowed by the bandwidth of the system). Indeed, if we assume that any free bandwidth can be occupied at no or little cost, the adversary can (and would) increase the throughput up to the maximum allowed value. Therefore, while in a usual setting, the arrival rate might be variable, the only way that a tips inflation attack would occur is indeed at the maximum allowed throughput. For this reason, in the remainder of the paper we make the assumption that the arrival rate β is constant.

The previous observations lead to the following result:

Theorem 1. Consider the system (10 - 13) with constant arrival rate $a(t) = \beta$. For each pair (p, k) , $p \in [0, 1]$, $k > 0$, the system has a locally stable solution $\{\hat{x}, \hat{l}, \hat{w}\}$, if $pk - 1 > 0$.

Proof. Given a time-independent solution $\{\hat{x}, \hat{l}, \hat{w}\}$ of the system, the linearized model is constructed by letting

$$\begin{aligned} l(t) &= \hat{l} + \theta(t), \\ x(t) &= \hat{x} + \zeta(t), \\ w(t) &= \hat{w} + \eta(t) \end{aligned} \quad (19)$$

and keeping only those terms in the equations which are linear in $\theta(t), \zeta(t), \eta(t)$. Set $pk = c \neq 1$. Derivation of the DDEs to linear order follows the methods outlined in chapter 4.6 of [19]. In this specific instance, they are

$$\dot{\zeta} = -(c-1)\zeta(t) + \frac{c-1}{c}\theta(t) \quad (20)$$

$$\dot{\theta} = -(c-1)\zeta(t-h) + \frac{c-1}{c}\theta(t-h) \quad (21)$$

The resulting linear system of delay differential equations is denoted $\mathcal{L}^{(0)}(\theta, \zeta, \eta) = 0$. The solution $\{\hat{x}, \hat{l}, \hat{w}\}$ is locally stable if

$$\max\{|\theta(t)|, |\zeta(t)|, |\eta(t)|\} \rightarrow 0 \quad \text{as } t \rightarrow \infty \quad (22)$$

for all solutions of the linear system. The solution is locally unstable if there is some solution satisfying $\mathcal{L}^{(0)}(\theta, \zeta, \eta) = 0$ such that

$$\max\{|\theta(t)|, |\zeta(t)|, |\eta(t)|\} \rightarrow \infty \quad \text{as } t \rightarrow \infty \quad (23)$$

The spectrum of the linear system $\mathcal{L}^{(0)}(\theta, \zeta, \eta) = 0$ is the set of complex values z for which the system has a nonzero solution of the form

$$(\theta(t), \zeta(t), \eta(t)) = e^{zt} (\theta, \zeta, \eta) \quad (24)$$

for some constants (θ, ζ, η) . The solution is stable if the spectrum is contained in the open left half of the complex plane, and is unstable if there is some element of the spectrum in the open right half of the complex plane [19][20].

Accordingly we investigate the spectrum $\mathcal{L}^{(0)}(\theta, \zeta, \eta) = 0$ by considering solutions of the form

$$\zeta(t) = \zeta e^{zt}, \quad \theta(t) = \theta e^{zt} \quad (25)$$

Substituting (25) into (20) and (21) we obtain

$$(hz + c + 1)\zeta = \frac{c-1}{c}\theta \quad (26)$$

$$\theta e^{hz} = -\frac{c-1}{hz}\zeta + \frac{c-1}{hzc}\theta \quad (27)$$

Algebraic manipulation of (26 - 27) shows that z needs to satisfy the equation:

$$c + \frac{c}{c-1}hz = e^{-hz} \quad (28)$$

It can be verified, by proposition 4.9 in [19] (a proof can be found in [21]), that this equation has no roots in the closed right half plane for all $c > 1$. This implies that $\max_{c+c/(c-1)hz=e^{-hz}} \Re(z) < 0$. By theorem 4.8 of [19] it follows that the equilibrium $\{\hat{x}, \hat{l}, \hat{w}\}$ of system (10 - 13) are locally locally stable for $c = pk > 1$. This concludes the proof.




Fig. 3: Simulations of the Tangle for different values of p and $k = 2$. Each simulation shows 100 realizations of the model with block times, T_a , generated according to a Poisson distribution with $\beta = 100$ and $h = 1$. For $p = 0.5$ (black line) the system becomes unstable. For all other values of p the average steady state value of the system reaches the equilibrium predicted by the fluid model (dashed lines, computed using equations (16)-(18)).

C. Validation of the previous results

To validate the model and the result presented in the previous paragraph we compare its behaviour with an agent based version of the Tangle: at each time step a random number of blocks arrive, according to a Poisson distribution, and for each one of these blocks the tip selection algorithm is performed on the current tips set in order to generate graph structures equivalent to the ones presented in detail in Section III. In other words, this agent based model simulates the behaviour of each block, therefore providing an accurate replica of the mechanism described in Section III. To simulate a Tips Inflation Attack, we assume that each new block will select two tips with proportion p (representing the honest issuers) or no tips with proportion $q = 1-p$ (representing the adversary).⁶ The variables used for the comparison are the number of leaves $L(t)$, which are obtained by enumerating the number of tips at the end of each time step. Due to the stochastic nature of the Tangle, 100 Monte Carlo simulations are performed in order to obtain statistically meaningful results. Figure 3 shows agent based simulation for 5 different values of p and the corresponding equilibria that would be obtained according to equations 16)-(18) (dashed lines). It is easy to see, by visual inspection, that the values and the stability behaviour predicted by the fluid model are also exhibited by the average steady state of the agent-based system.⁷ Furthermore, the interested reader can refer to [22] for a detailed analysis of the convergence properties of the stochastic model (1)-(4) to the fluid limit (10 - 13).

⁶Technically the latter is the case, where all referenced blocks are no longer tips.

⁷As a side note, we want to point out that the same conclusions appear to hold also for scenarios with low arrival rates in which the fluid limit approximation would not hold. This has been observed through a large number of Monte Carlo simulations that can be found at this link.

VI. CONTROL ALGORITHM

Given the previous result, we propose a control algorithm solution to tackle the Tips Inflation Attack. The approach is to add a feedback loop on the system that regulates the number of tips k that each new block selects.

$$k(l(t)) = 2 + \frac{m}{\beta h} (l(t) - 2\beta h)^+ \quad (29)$$

Combining (29) and (10)-(13), under the assumption of a constant arrival rate $a(t) = \beta$, we obtain the following system of equations:

$$\frac{dx}{dt}(t) = \beta(1 - p k(l(t)) \frac{x(t)}{l(t)}) \quad (30)$$

$$\frac{dl}{dt}(t) = \beta(1 - p k(t-h) \frac{x(t-h)}{l(t-h)}) \quad (31)$$

$$w(t) = l(t) - x(t) = \int_{t-h}^t p \beta k(s) \frac{x(s)}{l(s)} ds \quad (32)$$

$$k(l(t)) = \frac{m}{\beta h} (l(t) - 2\beta h) \quad (33)$$

For $p = 1$ the equilibrium occurs at $k = 2$, with $\tilde{l} = 2\beta h$ and $\tilde{x} = \beta h$. For $p < 1$ the equilibrium occurs at $\tilde{k} > 2$, with

$$\tilde{l} = K \frac{\beta h}{2pm}, \quad \tilde{x} = K \frac{\beta h}{pm(4p - 4pm + K)} \quad (34)$$

and

$$\tilde{k} = \frac{K}{2p} + 2 - 2m \quad (35)$$

where

$$K = 3pm + 1 - 2p + \sqrt{(2p - 1 - pm)^2 + 4pm} \quad (36)$$

A. Stability of the proposed control law

Without loss of generality we take $\beta = 1$ and consider the following more general control law:

$$k(l(t)) = 2 + r (l(t) - C)^+ \quad (37)$$

Notice that (29) is a special case of (37). Under these assumptions, the model depends on four parameters r, C, h, p , where $r, C, h > 0$ and $0 < p \leq 1$. The equations are

$$\frac{d}{dt} x(t) = 1 - p k(t) \frac{x(t)}{l(t)} \quad (38)$$

$$\frac{d}{dt} l(t) = 1 - p k(t-h) \frac{x(t-h)}{l(t-h)} \quad (39)$$

$$w(t) = p \int_{t-h}^t k(s) \frac{x(s)}{l(s)} ds \quad (40)$$

From (38-40) we get

$$w(t) = h - x(t) + x(t-h), \quad \frac{d}{dt} (l(t) - x(t) - w(t)) = 0$$

We will assume continuous initial conditions for x, w, l on the interval $[-h, 0]$, satisfying

$$x(t), w(t), l(t) \geq 0, \quad l(t) = x(t) + w(t) \forall t \in [-h, 0] \quad (41)$$

With these conditions, it follows that for all $t \geq 0$ the values of $l(t), w(t)$ are fully determined by the values $(x(t), x(t-h))$, that $l(t) = x(t-h) + h$, and that the function $x(t)$ satisfies the following delay differential equation:

$$\frac{d}{dt} x(t) = 1 - p \frac{[2 + r(x(t-h) + h - C)^+]}{x(t-h) + h} x(t) \quad (42)$$

The method of steps can be used to show that the equation (42), together with initial conditions satisfying (41), has a unique positive continuous solution for all $t \geq 0$. The equations $w(t) = h - x(t) + x(t-h)$ and $l(t) = x(t) + w(t)$ then provide the unique solution of the system (38-40) for all $t \geq 0$.

Theorem 2. Suppose that x, w, l satisfy the system (38-40) with initial conditions on $[-h, 0]$ satisfying (41). Suppose also that $l(t) \leq C$ for all $t \in [-h, 0]$. Then for all $T \geq 0$,

$$l(T) \leq h + \max \left[C, \frac{C}{2p}, \frac{1}{rp} \right] \quad (43)$$

Proof: If $l(T) \leq C + h$ then the bound (43) holds. Suppose $l(T) > C + h$, and define

$$S = \sup \{s \leq T : l(s) \leq C\}$$

Since $l(0) \leq C$, it follows that $S \geq 0$, and continuity of l implies that $l(S) = C$. Clearly

$$l(t) \geq C \quad \text{for all } t \in [S, T] \quad (44)$$

Equation (39) implies that

$$l(t+h) = l(t) + \int_t^{t+h} l'(s) ds \leq l(t) + h \quad (45)$$

It follows that for all $t \in [S, S+h]$ we have $l(t) \leq l(S) + h \leq C + h$, and therefore $T \geq S + h$. Define

$$K_1 = \min \left(r, \frac{2}{C} \right)$$

$$K_2 = \max \left(C, \frac{1}{pK_1} \right)$$

For all $t \in [S, T]$,

$$\begin{aligned} \frac{k(t)}{l(t)} &= \frac{2 - rC}{l(t)} + r \\ &\geq \begin{cases} r & \text{if } rC \leq 2 \\ 2C^{-1} & \text{if } rC \geq 2 \end{cases} \quad (\text{using the bound (44)}) \\ &\geq K_1 \end{aligned}$$

This is because, if $rC \leq 2$ then $\frac{2-rC}{l(t)} + r = F + r \geq r$, since $F \geq 0$ and if $rC > 2$, due to bound (44), $\frac{2-rC}{l(t)} + r \geq \frac{2-rC}{C} + r \geq r = 2C^{-1}$ (as $\frac{2-rC}{C} \leq 0$).

This implies

$$\frac{d}{dt} x(t) \leq 1 - p K_1 x(t) \quad \text{for all } t \in [S, T]$$

and hence for all $t \in [S, T]$

$$x(t) \leq \left(x(S) - \frac{1}{pK_1} \right) e^{-pK_1(t-S)} + \frac{1}{pK_1}$$

Since $x(S) \leq l(S) = C$,

$$\begin{aligned} x(t) &\leq \left(C - \frac{1}{pK_1}\right) e^{-pK_1(t-S)} + \frac{1}{pK_1} \\ &\leq \begin{cases} \frac{1}{pK_1} & \text{if } C \leq \frac{1}{pK_1} \\ C & \text{if } C \geq \frac{1}{pK_1} \end{cases} \end{aligned}$$

Therefore for all $t \in [S, T]$ we have

$$x(t) \leq \max\left(C, \frac{1}{pK_1}\right) = K_2 \quad (46)$$

Using $l(T) = x(T) + w(T) = x(T-h) + h$, and $T-h \geq S$, we deduce that

$$\begin{aligned} l(T) \leq K_2 + h &= \max\left(C, \frac{1}{p} \max\left(\frac{1}{r}, \frac{C}{2}\right)\right) + h \\ &= \max\left[C, \frac{1}{rp}, \frac{C}{2p}\right] + h \end{aligned}$$

which completes the proof.

Algorithm 1 Tips Management Control

- 1: For each new block a
 - 2: Compute $k(t)$
 - 3: Compute $\bar{k} = \lfloor k(t) \rfloor, s = k(t) - \bar{k}$;
 - 4: Sample a random number x from the uniform distribution in $[0, 1]$
 - 5: **if** $x \geq s$ **then**
 - 6: Block a selects $\bar{k}$ tips
 - 7: **else**
 - 8: Block a selects $\bar{k} + 1$ tips
 - 9: **end if**
-

VII. IMPLEMENTATION DETAILS AND SIMULATIONS

The control law (29) is not implementable directly on the Tangle as each block needs to select an integer number of blocks. To circumvent this issue, consider Algorithm 1. It is easy to show that this stochastic behaviour converges in the fluid limit to control law (29). In fact, call $K(t)$ the random variable associated with the number of selections from a block arriving at time t . The total amount of selections over a time interval δt will be

$$P(t, \delta t) = \sum_{T_a \in [t, t+\delta t]} K(T_a) \quad (47)$$

Assuming that the arrival rate is proportional to λ (as we did in Section V.a) and letting $\lambda \rightarrow \infty$ and $\delta t \rightarrow 0$ then, due to the independence of the random number sampling in step 4 of Algorithm 1, by the law of large numbers, $\lambda^{-1}P(t, \delta t)$ will converge to the expected value of the random variable $K(t)$:

$$\begin{aligned} \mathbb{E}[K(t)] &= (\bar{k} + 1)s + \bar{k}(1 - s) \\ &= (\bar{k} + 1)(k(t) - \bar{k}) + \bar{k}(1 - k(t)) \\ &= k(t). \end{aligned} \quad (48)$$

To showcase the performance, we consider the same agent based model of the Tangle, used in Section V-C, with $\beta = 100$,

$h = 1$ being subject to a Tips Inflation Attack with $p = 0.4$ in two scenarios: a continued attack from the beginning and for different values of m ; a discontinuous attack, starting at $t = 20$ and ending at $t = 40$ with $m = 4$. Figures 4 and 5 show the average evolution of 100 different realizations of the Tangle for the two scenarios. Notice that, for all choices of m , the system remains stable and that the equilibria, predicted by (34), match the average steady state of the controlled agent based system. In the second scenario, the system reacts to the attack, rapidly adjusting the amount of tips selected by the honest nodes and manages to keep the system stable for the whole duration of the attack.

Two elements that we believe are worth stressing is that the controller does not reduce the effect of the attack to zero, nor it would be possible for new block to approve as many tips as possible. While it might be possible to design a more complex controller (for instance including an integral action) to achieve this goal and in theory, every new block might approve as many tips as possible (e.g., new honest blocks approving the whole tips set), this would translate on requiring large amounts of computational power on the side of the honest nodes to calculate the control action. While this is generally a trivial requirement under normal circumstances, in a DLT setting it is important to keep the amount of computations that honest nodes have to perform to as low as possible (which also means, that it is important to keep the amount of approvals, as low as possible). Therefore, while the design of more complex controllers will be explored in a future work, we believe that the obtained ratio between performance and simplicity represents a good trade off, given the current setting.

VIII. CONCLUSIONS AND FUTURE LINES OF RESEARCH

In this paper we presented a fluid model of a class of DAG-based DLTs and designed a feedback control to reduce the effects of a Tips Inflation Attack. Our claims are backed by a thorough theoretical analysis, in the form of two stability theorems, and by extensive Monte Carlo simulations. The development of more complex control algorithms will be the focus of future research.

IX. ACKNOWLEDGEMENTS

The authors wish to thank Daria Dziubałtowska for her help in the realization of this manuscript.

REFERENCES

- [1] M. Rauchs, A. Glidden, B. Gordon, G. C. Pieters, M. Recanatini, F. Rostand, K. Vagneur, and B. Z. Zhang, "Distributed ledger technology systems: A conceptual framework," *Available at SSRN 3230013*, 2018.
- [2] E.-E. Gojka, N. Kannengießer, B. Sturm, J. Bartsch, and A. Sunyaev, "Security in distributed ledger technology: An analysis of vulnerabilities and attack vectors," in *Intelligent Computing: Proceedings of the 2021 Computing Conference, Volume 3*. Springer, 2021, pp. 722–742.
- [3] M. Saad, J. Spaulding, L. Njilla, C. Kamhoua, S. Shetty, D. Nyang, and D. Mohaisen, "Exploring the attack surface of blockchain: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 1977–2008, 2020.
- [4] W. Li, M. Cao, Y. Wang, C. Tang, and F. Lin, "Mining pool game model and nash equilibrium analysis for pow-based blockchain networks," *IEEE Access*, vol. 8, pp. 101 049–101 060, 2020.




Fig. 4: Simulation results for the Tangle with Algorithm 1 for different values of m and $p = 0.4$. Each simulation shows the average of 100 realizations of the model. The top panel shows the average number of tips $L(t)$ over the 100 realization for different values of m , while the bottom panel shows the average value of the controller $k(t)$ for the corresponding values of m . Notice that the predicted equilibria of the fluid model (dashed lines, computed using equations (34)) correspond to average steady state value of the agent based system for each value of m .




Fig. 5: Simulations of the Tangle with Algorithm 1 for $m = 4$ and a Tips Inflation Attack with $p = 0.4$ starting at $t = 20$ and ending at $t = 40$. The top panel shows the average number of tips $L(t)$ over the 100 realizations while the bottom panel shows the average value of the controller $k(t)$. The controller increases the amount of tips that are selected (on average) by new blocks during the attack and returns to its nominal value when that stops being the case.

- [5] S. Kim and S.-G. Hahn, "Mining pool manipulation in blockchain network over evolutionary block withholding attack," *IEEE Access*, vol. 7, pp. 144 230–144 244, 2019.
- [6] P. Ferraro, C. King, and R. Shorten, "On the stability of unverified transactions in a dag-based distributed ledger," *IEEE Transactions on Automatic Control*, vol. 65, no. 9, pp. 3772–3783, 2019.
- [7] A. Cullen, P. Ferraro, C. King, and R. Shorten, "On the resilience of dag-based distributed ledgers in iot applications," *IEEE Internet of Things Journal*, vol. 7, no. 8, pp. 7112–7122, 2020.
- [8] Q. Wang, J. Yu, S. Chen, and Y. Xiang, "Sok: Diving into dag-based blockchain systems," *ArXiv*, vol. abs/2012.06128, 2020.
- [9] R. Overko, R. Ordóñez-Hurtado, S. Zhuk, P. Ferraro, A. Cullen, and R. Shorten, "Spatial positioning token (sptoken) for smart mobility," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [10] P. Ferraro, L. Zhao, C. King, and R. Shorten, "Personalised feedback control, social contracts, and compliance strategies for ensembles," *arXiv preprint arXiv:2103.07261*, 2021.
- [11] S. Müller, A. Penzkofer, N. Polyanskii, J. Theis, W. Sanders, and H. Moog, "Tangle 2.0 leaderless nakamoto consensus on the heaviest dag," *IEEE Access*, vol. 10, pp. 105 807–105 842, 2022.
- [12] P. Ferraro, C. King, and R. Shorten, "Distributed ledger technology for smart cities, the sharing economy, and social compliance," *Ieee Access*, vol. 6, pp. 62 728–62 746, 2018.
- [13] G. O. Karame and E. Androulaki, *Bitcoin and blockchain security*. Artech House, 2016.
- [14] B. Kusmierz, W. Sanders, A. Penzkofer, A. Caposelle, and A. Gal, "Properties of the Tangle for Uniform Random and Random Walk Tip Selection," in *2019 IEEE International Conference on Blockchain (Blockchain)*, 2019, pp. 228–236.
- [15] S. Popov, O. Saa, and P. Finardi, "Equilibria in the tangle," *Computers & Industrial Engineering*, vol. 136, pp. 160–172, 2019.
- [16] S. Müller, A. Penzkofer, B. Kuśmierz, D. Camargo, and W. J. Buchanan, "Fast probabilistic consensus with weighted votes," in *Proceedings of the Future Technologies Conference (FTC) 2020, Volume 2*, K. Arai, S. Kapoor, and R. Bhatia, Eds. Cham: Springer International Publishing, 2021, pp. 360–378.
- [17] A. Cullen, P. Ferraro, W. Sanders, L. Vigneri, and R. Shorten, "Access Control for Distributed Ledgers in the Internet of Things: A Networking Approach," *IEEE Internet of Things Journal*, pp. 1–1, 2021.
- [18] G. Carl, G. Kesidis, R. R. Brooks, and S. Rai, "Denial-of-service attack-detection techniques," *IEEE Internet computing*, vol. 10, no. 1, pp. 82–89, 2006.
- [19] H. Smith, *An introduction to delay differential equations with applications to the life sciences*. springer, 2011.
- [20] J. K. Hale and S. M. V. Lunel, *Introduction to functional differential equations*. Springer Science & Business Media, 2013, vol. 99.
- [21] F. Brauer, "Absolute stability in delay equations," *Journal of differential equations*, vol. 69, no. 2, pp. 185–191, 1987.
- [22] C. King, "The fluid limit of a random graph model for a shared ledger," *Advances in Applied Probability*, vol. 53, no. 1, pp. 81–106, 2021.