

Equilibria in the Tangle

Serguei Popov^{1,*} Olivia Saa² Paulo Finardi³

March 3, 2018

¹Department of Statistics, Institute of Mathematics, Statistics and Scientific Computation, University of Campinas – UNICAMP, rua Sérgio Buarque de Holanda 651, 13083–859, Campinas SP, Brazil

e-mails: popov@ime.unicamp.br and serguei.popov@iota.org

²Department of Applied Mathematics, Institute of Mathematics and Statistics, University of São Paulo – USP, rua do Matão 1010, 05508–090, São Paulo SP, Brazil

e-mail: olivia@ime.usp.br and olivia.saa@iota.org

³Institute of Computing – IC, University of Campinas – UNICAMP, av. Albert Einstein 1251, 13083–852, Campinas SP, Brazil

e-mails: ra144809@ic.unicamp.br and paulo.finardi@iota.org

Abstract

We analyse the Tangle — a DAG-valued stochastic process where new vertices get attached to the graph at Poissonian times, and the attachment’s locations are chosen by means of random walks on that graph. These new vertices (also thought of as “transactions”) are issued by many players (which are the nodes of the network), independently. We prove existence of (“almost symmetric”) Nash equilibria for the system where a part of players tries to optimize their attachment strategies. Then, we also present simulations that show that the “selfish” players will nevertheless cooperate with the network by choosing attachment strategies that are similar to the “recommended” one.

Keywords: random walk, Nash equilibrium, directed acyclic graph, cryptocurrency, tip selection

AMS 2010 subject classifications: Primary 91A15. Secondary 60J20, 68M14.

*corresponding author

1 Introduction

In this paper we study *the Tangle*, a stochastic process on the space of (rooted) Directed Acyclic Graphs (DAGs). This process “grows” in time, in the sense that new vertices are attached to the graph according to a Poissonian clock, but no vertices/edges are ever deleted. When that clock rings, a new vertex appears and attaches itself to locations that are chosen with the help of certain random walks on the state of the process in the *recent past* (this is to model the network propagation delays); these random walks therefore play the key role in the model.

Random walks on random graphs can be thought of as a particular case of Random Walks in Random Environments: here, the transition probabilities are functions of the graph only, i.e., there are no additional variables (such as conductances¹ etc.) attached to the vertices and/or edges of the graph. Still, this subject is very broad, and one can find many related works in the literature. One can mention the internal DLA models (e.g. [13] and references therein), random walks on Erdős-Rényi graphs [5, 13], or random walks on the preferential attachment graphs [4] (which most closely resembles the model of this paper).

The motivation for studying the particular model presented in this paper stems from the fact that it is applied in the IOTA cryptocurrency [1, 19], which uses (non-trivial) DAGs as the primary ledger for the transactions’ data. This is different from “traditional” cryptocurrencies such as the Bitcoin, where that data is stored in a sequence of blocks², also known as *blockchain*. An important observation, which motivates the use of more general DAGs instead of blockchains is that the latter *scale* poorly: when the network is large, it is difficult for it to achieve consensus on which blocks are “valid” in the situations when the new blocks come too frequently. We also cite [2, 3, 16, 20] which deal with other approaches to using DAGs as distributed ledgers.

The main results of the present paper deal with the following question: what if some participants of the network are trying to minimize their *costs* by adopting a behavior different from the “default” one? How will the system behave in such circumstances? To address these kinds of questions, we first provide general arguments to prove existence of (“almost symmetric”) Nash equilibria for the system, see Section 2. Although one can hardly access the explicit form of these equilibria in a purely analytical way, simulations presented in Section 3 show that the “selfish” players will typically still choose attachment strategies that are similar to the default

¹this refers to the well-known relation between reversible Markov chains and electric networks, see e.g. the classical book [7]

²that is, the underlying graph is essentially \mathbb{Z}_+ (after discarding finite forks)

one, meaning that they would prefer *cooperating* with the network rather than simply *using* it).

Let us stress also that, in this paper, we consider only “selfish” players (those who only care about their own costs but still want to use the network in a legitimate way³); we do not consider the case when there are “malicious” ones (those who want to disrupt the network even at a cost to themselves). We are going to treat several types of attacks against the network in the subsequent papers.

1.1 Description of the model

In the following we introduce the mathematical model describing the Tangle [19].

Let $\text{card}(A)$ stand for the cardinality of (multi)set A . For an oriented multigraph $\mathcal{T} = (V, E)$, where V is the set of vertices and E is the multiset of edges, and $v \in V$, we denote by

$$\begin{aligned}\text{deg}_{\text{in}}(v) &= \text{card}\{e = (u_1, u_2) \in E : u_2 = v\}, \\ \text{deg}_{\text{out}}(v) &= \text{card}\{e = (u_1, u_2) \in E : u_1 = v\}\end{aligned}$$

the “incoming” and “outgoing” degrees of the vertex v (counting the multiple edges). In the following, we refer to multigraphs simply as graphs. For $u, v \in V$, we say that u *approves* v , if $(u, v) \in E$. We use the notation $\mathcal{A}(u)$ for the set of the vertices approved by u . We say that $u \in V$ *references* $v \in V$ if there is a sequence of sites $u = x_0, x_1, \dots, x_k = v$ such that $x_j \in \mathcal{A}(x_{j-1})$ for all $j = 1, \dots, k$, i.e., there is a directed path from u to v . If $\text{deg}_{\text{in}}(w) = 0$ (i.e., there are no edges pointing to w), then we say that $w \in V$ is a *tip*.

Let \mathcal{G} be the set of all directed acyclic graphs (also known as DAGs, that is, oriented graphs without cycles) $G = (V, E)$ with the following properties:

- the graph G is finite and the multiplicity of any edge is at most two (i.e., there are at most two edges linking the same vertices);
- there is a distinguished vertex $\varphi \in V$ such that $\text{deg}_{\text{out}}(v) = 2$ for all $v \in V \setminus \{\varphi\}$, and $\text{deg}_{\text{out}}(\varphi) = 0$ (this vertex φ is called *the genesis*);
- any $v \in V$ such that $v \neq \varphi$ references φ ; that is, there is an oriented path⁴ from v to φ (one can say that the graph is *connected towards* φ).

³i.e., want to issue valid transactions and have them confirmed by the rest of the network

⁴not necessarily unique

We now describe the tangle as a continuous-time Markov process on the space \mathcal{G} . The state of the tangle at time $t \geq 0$ is a DAG $\mathcal{T}(t) = (V_{\mathcal{T}}(t), E_{\mathcal{T}}(t))$, where $V_{\mathcal{T}}(t)$ is the set of vertices and $E_{\mathcal{T}}(t)$ is the multiset of directed edges at time t . The process's dynamics are described in the following way:

- The initial state of the process is defined by $V_{\mathcal{T}}(0) = \varnothing$, $E_{\mathcal{T}}(0) = \emptyset$.
- The tangle *grows with time*, that is, $V_{\mathcal{T}}(t_1) \subset V_{\mathcal{T}}(t_2)$ and $E_{\mathcal{T}}(t_1) \subset E_{\mathcal{T}}(t_2)$ whenever $0 \leq t_1 < t_2$.
- For a fixed parameter $\lambda > 0$, there is a Poisson process of incoming *transactions*; these transactions then become the vertices of the tangle.
- Each incoming transaction chooses⁵ two vertices v' and v'' (which, in general, may coincide), and we add the edges (v, v') and (v, v'') . We say in this case that this new transaction was *attached* to v' and v'' (equivalently, v *approves* v' and v'').
- Specifically, if a new transaction v arrived at time t' , then $V_{\mathcal{T}}(t'+) = V_{\mathcal{T}}(t') \cup \{v\}$, and $E_{\mathcal{T}}(t'+) = E_{\mathcal{T}}(t') \cup \{(v, v'), (v, v'')\}$.

Let us write

$$\begin{aligned} \mathcal{P}^{(t)}(x) &= \{y \in \mathcal{T}(t) : y \text{ is referenced by } x\}, \\ \mathcal{F}^{(t)}(x) &= \{z \in \mathcal{T}(t) : z \text{ references } x\} \end{aligned}$$

for the “past” and the “future” with respect to x (at time t). Note that these introduce a *partial order* structure on the tangle. Observe that, if t_0 is the time moment when x was attached to the tangle, then $\mathcal{P}^{(t)}(x) = \mathcal{P}^{(t_0)}(x)$ for all $t \geq t_0$. We also define the *cumulative weight* $\mathcal{H}_x^{(t)}$ of the vertex x at time t by

$$\mathcal{H}_x^{(t)} = 1 + \text{card}(\mathcal{F}^{(t)}(x)); \tag{1}$$

that is, the cumulative weight of x is one (its “own weight”) plus the number of vertices that reference it. Observe that, for any $t > 0$, if y approves x then $\mathcal{H}_x^{(t)} - \mathcal{H}_y^{(t)} \geq 1$, and the inequality is strict if and only if there are vertices different from y which also approve x . Also note that the cumulative weight of any tip is equal to 1.

There is some data associated to each vertex (transaction), created at the moment when that transaction was attached to the tangle. The precise nature of that data is

⁵the precise selection mechanism will be described below

not relevant for the purposes of this paper, so we assume that it is an element of some (unspecified, but finite) set \mathcal{D} ; what is important, however, is that there is a natural way to say if the set of vertices is *consistent* with respect to the data they contain⁶. When it is necessary to emphasize that the vertices of $G \in \mathcal{G}$ contain some data, we consider the *marked* DAG $G^{[\mathfrak{d}]}$ to be $(G, \mathfrak{d}) = (V, E, \mathfrak{d})$, where \mathfrak{d} is a function $V \rightarrow \mathcal{D}$. We define $\mathcal{G}^{[\mathfrak{d}]}$ to be the set of all marked DAGs (G, \mathfrak{d}) , where $G \in \mathcal{G}$.

1.2 Attachment strategies

There is one very important detail that has not been explained, namely: how does a newly arrived transaction choose which two vertices in the tangle it will approve, i.e., what is the *attachment strategy*? Notice that, in principle, it would be good for the whole system if the new transactions always prefer to select tips as attachment places, since this way more transactions would be “confirmed”⁷. In any case, it is quite clear that the appropriate choice of the attachment strategy is essential for the correct functioning (whatever this could mean) of the system.

It is also important to comment that the attachment strategy of a network node is something “internal” to it; what others can see, are the *attachment choices* of the node, but the mechanism behind them need not be publicly known. For this reason, an attachment strategy cannot be *imposed* in the protocol.

We now describe a possible choice of the attachment strategy, used to determine where the incoming transaction will be attached. It is also known as the *recommended tip selection algorithm*, since, due to reasons described above, the recommended nodes’ behavior is always to try to approve tips. We stress again, however, that approving only tips is not imposed in the protocol, since there is usually no way to know if a node “knew” if the transaction it approved was already approved by someone else before (also, there is no way to know which approving transaction was the first).

Let us denote by $\mathcal{L}(t)$ the set of all vertices that are tips at time t , and let $L(t) = \text{card}(\mathcal{L}(t))$. To model the network propagation delays, we introduce a parameter $h > 0$, and assume that at time t only $\mathcal{T}(t - h)$ is known to the entity that issued the incoming transaction. We then define the *tip-selecting random walk*, in the following way. It depends on a parameter q (the backtracking probability) and on a function f .

⁶one may think that the data refers to value transactions between accounts, and consistency means that no account has negative balance as a result, and/or the total balance has not increased

⁷we discuss the exact meaning of this later; for now, think that “confirmed” means “referenced by many other transactions”

The initial state of the random walk is the genesis \wp^8 , and it is stopped upon hitting the set $\mathcal{L}(t - h)$. It is important to observe that $v \in \mathcal{L}(t - h)$ does not necessarily mean that v is still a tip at time t . Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a monotone non-increasing function. The transition probabilities of the walkers are defined in the following way: the walk *backtracks* (i.e., jumps to a randomly chosen site it approves) with probability $q \in [0, 1/2)$; if y approves $x \neq \wp$, then the transition probability $P_{xy}^{(f)}$ is proportional to $f(\mathcal{H}_x - \mathcal{H}_y)$, that is,

$$P_{xy}^{(f)} = \begin{cases} \frac{q}{2}, & \text{if } y \in \mathcal{A}(x), \\ \frac{(1 - q)f(\mathcal{H}_x^{(t-h)} - \mathcal{H}_y^{(t-h)})}{\sum_{z: x \in \mathcal{A}(z)} f(\mathcal{H}_x^{(t-h)} - \mathcal{H}_z^{(t-h)})}, & \text{if } x \in \mathcal{A}(y), \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

(for $x = \wp$ we define the transition probabilities as above, but with $q = 0$). In what follows, we will mostly assume that $f(s) = \exp(-\alpha s)$ for some $\alpha \geq 0$. We use the notation $P^{(\alpha)}$ for the transition probabilities in this case. Intuitively, the smaller is the value of α , the more *random* the walk is⁹. It is worth observing that the case $q = 0$ and $\alpha \rightarrow \infty$ corresponds to the GHOST protocol of [21] (more precisely, to the obvious generalization of the GHOST protocol for the case when a tree is substituted by a DAG).

Now, to select two tips w_1 and w_2 where our transaction will be attached, just run two independent random walks as above, and stop when you first hit $\mathcal{L}(t - h)$. One can also require that w_1 should be different from w_2 ; for that, one may re-run the second random walk in the case its exit point happened to be the same as that of the first random walk. Observe that $(\mathcal{T}(t), t \geq 0)$ is a continuous-time transient Markov process on \mathcal{G} ; since the state space is quite large, it is difficult to analyse this process. In particular, for a fixed time t , it is not easy to study the above random walk since it takes place on a *random* graph, e.g., can be viewed as a random walk in a random environment; it is common knowledge that random walks in random environments are notoriously hard to deal with.

We say that a transaction is *confirmed with confidence* γ_0 (where γ_0 is some pre-defined number, close to 1), if, with probability at least γ_0 , the large- α random walk¹⁰ ends in a tip which references that transaction. It may happen that a transaction

⁸although in practical implementations one may start it in some place closer to the tips

⁹physicists would call the case of small α *high temperature regime*, and the case of large α *low temperature regime* (that is, α stands for the inverse temperature)

¹⁰recall that the large- α random walk is “more deterministic”

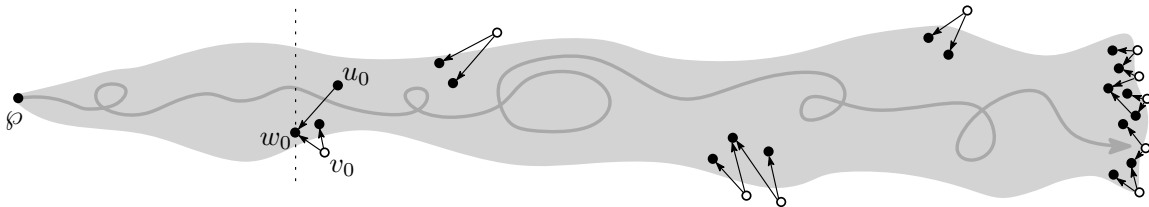


Figure 1: The walk on the tangle and tip selection. Tips are circles, and transactions which were approved at least once are disks.

does not get confirmed (even, possible, does not get approved a single time), and becomes orphaned forever. Let us define the event

$$\mathcal{U} = \{\text{every transaction eventually gets approved}\}.$$

We believe that the following statement holds true; however, we have only a heuristical argument in its favor, not a rigorous proof. In any case, it is only of theoretical interest, since, as explained below, in practice we will find ourselves in the situation where $\mathbb{P}[\mathcal{U}] = 0$. We therefore state it as

Conjecture 1.1. *It holds that*

$$\mathbb{P}[\mathcal{U}] = \begin{cases} 0, & \text{if } \int_0^{+\infty} f(s) ds < \infty, \\ 1, & \text{if } \int_0^{+\infty} f(s) ds = \infty. \end{cases} \quad (3)$$

Explanation. First of all, it should be true that $\mathbb{P}[\mathcal{U}] \in \{0, 1\}$ since \mathcal{U} is a *tail event* with respect to the natural filtration; however, it does not seem to be very easy to prove the 0–1 law in this context (recall that we are dealing with a transient Markov process on an infinite state space). Next, consider a tip v_0 which got attached to the tangle at time t_0 , and assume that it is still a tip at time $t \gg t_0$; also, assume that, among all tips, v_0 is “closest” (in some suitable sense) to the genesis. Let us now think of the following question: what is the probability that v_0 will still be a tip at time $t + 1$?

Look at Figure 1: during the time interval $[t, t + 1]$, $O(1)$ new particles will arrive, and the corresponding walks will travel from the genesis φ looking for tips. Each of these walks will have to cross the dotted vertical segment on the picture, and with positive probability at least one of them will pass through w_0 , one of the vertices

approved by v_0 . Assume that w_0 was already confirmed (i.e., connected to the right end of the tangle via some other transaction u_0 that approves w_0). Then, it is clear (but not easy to prove!) that the cumulative weight of both u_0 and w_0 should be $O(t)$, and so, when in w_0 , the walk will jump to the tip v_0 with probability $f(O(t))$.

This suggests that the probability that $v_0 \in \mathcal{L}(t+1)$ (i.e., that v_0 still is tip at time $t+1$) is $f(O(t))$, and the Borel-Cantelli lemma¹¹ gives that the probability that v_0 will be eventually approved is less than 1 or equal to 1 depending on whether $\sum_n f(n)$ converges or diverges; the convergence (divergence) of the sum is equivalent to convergence (divergence) of the integral in (3) due to the monotonicity of the function f . A standard probabilistic argument¹² would then imply that if the probability that *a given* tip remains orphaned forever is uniformly positive, then the probability that *at least one* tip remains orphaned forever is equal to 1. \square

One may naturally think that it would be better to choose the function f in such a way that, almost surely, every tip eventually gets confirmed. However, as explained in Section 4.1 of [19], there is a good reason to choose a rapidly decreasing function f , because this defends the system against nodes' misbehavior and attacks. The idea is then to assume that a transaction which did not get confirmed during a sufficiently long period of time is “unlucky”, and needs to be reattached¹³ to the tangle. Let us fix some $K > 0$: it stands for the time when an unlucky transaction is reissued (because there is already very little hope that it would be confirmed “naturally”). We call a transaction issued less than K time units ago “unconfirmed”, and if a transaction was issued more than K time units ago and was not confirmed, we call it “orphaned”. In the following, we assume that the system is *stable*, in the sense that the “recent” unconfirmed transactions do not accumulate and the time until a transaction is confirmed (roughly) does not depend on the moment when it appeared in the system¹⁴.

In that stable regime, let p be the probability that a transaction is confirmed K time units after it was issued for the first time; the number of times a transaction should be issued to achieve confirmation is then a Geometric random variable with parameter p (and, therefore, with expected value p^{-1}); so, the mean time until the

¹¹to be precise, a bit more refined argument is needed since the corresponding events are not independent

¹²which is also not so easy to formalize in these circumstances

¹³in fact, the nodes of the network may adopt a rule that instructs to delete the transactions that are older than K and still are tips from their databases

¹⁴simulations indicate that this is indeed the case when α is small; however, it is not guaranteed to happen for large values of α

transaction is confirmed is K/p . Let us then recall the following remarkable fact belonging to the queuing theory, known as the Little’s formula (sometimes also referred to as the Little’s theorem or the Little’s identity):

Proposition 1.2. *Suppose that λ_a is the arrival rate, μ is the mean number of customers in the system, and T is the mean time a customer spends in the system. Then $T = \mu/\lambda_a$.*

Proof. See e.g. Section 5.2 of [6]. To understand intuitively why this fact holds true, one may reason in the following way: assume that, while in the system, each customer pays money to the system with rate 1. Then, at large time t , the total amount of money earned by the system would be (approximately) μt on one hand, and $T\lambda_a t$ on the other hand. Dividing by t and then sending t to infinity, we obtain $\mu = T\lambda_a$. \square

Little’s formula then implies¹⁵ the following

Proposition 1.3. *The average number of unconfirmed transactions in the system is equal to $p^{-1}\lambda K$.*

Proof. Indeed, apply Proposition 1.2 with $\lambda_a = \lambda$ (think of a transaction which was reattached as a customer which returns to the server after an unsuccessful service attempt; this way, the incoming flow of customers still has rate λ). As observed before, the mean time spent by a customer in the system is equal to K/p . \square

In case when the tangle contains data (which, in principle, can make transactions incompatible between each other), one may choose more sophisticated methods of tip selection. As we already mentioned, selecting tips with larger values of α provides better defense against attacks and misbehavior; however, smaller values of α make the system more stable with respect to the transactions’ confirmation times. An example of “mixed- α ” strategy is the following. Define the “model tip” w_0 as a result of the random walk with large α , then select two tips w_1 and w_2 with random walks with small α , but check that

$$\mathcal{P}^{(t-h)}(w_0) \cup \mathcal{P}^{(t-h)}(w_1) \cup \mathcal{P}^{(t-h)}(w_2)$$

is consistent.

¹⁵in the language of queuing systems, a reissued transaction is a customer which goes back to the server after an unsuccessful service attempt

2 Selfish nodes and Nash equilibria

Now, we are going to study the situation when some participants of the network are “selfish” and want to use a customized attachment strategy, in order to improve the confirmation time of their transactions (possibly at the expense of the others).

For a finite set A let us denote by $\mathcal{M}(A)$ the set of all probability measures on A , that is

$$\mathcal{M}(A) = \left\{ \mu : A \rightarrow \mathbb{R} \text{ such that } \mu(a) \geq 0 \text{ for all } a \in A \text{ and } \sum_{a \in A} \mu(a) = 1 \right\}.$$

Let

$$\mathfrak{M} = \bigcup_{G=(V,E) \in \mathcal{G}} \mathcal{M}(V \times V)$$

be the union of the sets of all probability measures on the pairs of (not necessarily distinct) vertices of DAGs belonging to \mathcal{G} . Then, an *attachment strategy* \mathcal{S} is a map

$$\mathcal{S} : \mathcal{G}^{[\mathfrak{d}]} \rightarrow \mathfrak{M}$$

with the property $\mathcal{S}(V, E, \mathfrak{d}) \in \mathcal{M}(V \times V)$ for any $G^{[\mathfrak{d}]} = (V, E, \mathfrak{d}) \in \mathcal{G}^{[\mathfrak{d}]}$; that is, for any $G \in \mathcal{G}$ with data attached to the vertices (which corresponds to the state of the tangle at a given time) there is a corresponding probability measure on the set of pairs of the vertices. Note also that in the above we considered *ordered* pairs of vertices, which, of course, does not restrict the generality.

Let $\kappa > 0$ be a fixed number. We now assume that, for a (very) large N , there are κN nodes that follow the default tip selection algorithm, and N “selfish” nodes that try to minimize their “cost”, whatever this could mean¹⁶. Assume that all nodes issue transactions with the same rate $\frac{\lambda}{(\kappa+1)N}$, independently. The overall rate of “honest” transactions in the system is then equal to $\frac{\lambda\kappa}{\kappa+1}$, and the overall rate of transactions issued by selfish nodes equals $\frac{\lambda}{\kappa+1}$.

Let $\mathcal{S}_1, \dots, \mathcal{S}_N$ be the attachment strategies used by the selfish nodes. To evaluate the “goodness” of a strategy, one has to choose and then optimize some suitable observable (that stands for the “cost”); as usual, there are several “reasonable” ways to do this. We decided to choose the following one, for definiteness and also for technical reasons (to guarantee the continuity of some function used below); one can

¹⁶for example, the cost may be the expected confirmation time of a transaction (conditioned that it is eventually confirmed), the probability that it was not approved during certain (fixed) time interval, etc.

probably extend our arguments to other reasonable cost functions. Assume that a transaction v was attached to the tangle at time t_v , so $v \in \mathcal{T}(t)$ for all $t \geq t_v$. Fix some (typically large) $M_0 \in \mathbb{N}$. Let $t_1^{(v)}, \dots, t_{M_0}^{(v)}$ be the moments when the subsequent (after v) transactions were attached to the tangle. For $k = 1 \dots, M_0$ let $R_k^{(v)}$ be the event that the *default* tip-selecting walk¹⁷ on $\mathcal{T}(t_k^{(v)})$ stops in a tip that *does not* reference v . We then define

$$W(v) = \mathbf{1}_{R_1^{(v)}} + \dots + \mathbf{1}_{R_{M_0}^{(v)}} \quad (4)$$

to be the number of times that the M_0 “subsequent” tip selection random walks do not reference v (in the above, $\mathbf{1}_A$ is the indicator function of an event A). Intuitively, the smaller is the value of $W(v)/M_0$, the bigger is the chance that v is quickly confirmed.

Next, assume that $(v_j^{(k)}, j \geq 1)$ are the transactions issued by the k th (selfish) node. We define

$$\mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_N) = M_0^{-1} \lim_{n \rightarrow \infty} \frac{W(v_1^{(k)}) + \dots + W(v_n^{(k)})}{n}, \quad (5)$$

to be the *mean cost* of the k th node given that $\mathcal{S}_1, \dots, \mathcal{S}_N$ are the attachment strategies of the selfish nodes.

Definition 2.1. *We say that a set of strategies $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ is a Nash equilibrium if*

$$\mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_{k-1}, \mathcal{S}_k, \mathcal{S}_{k+1}, \dots, \mathcal{S}_N) \leq \mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_{k-1}, \mathcal{S}', \mathcal{S}_{k+1}, \dots, \mathcal{S}_N)$$

for any k and any $\mathcal{S}' \neq \mathcal{S}_k$.

Observe that, since the nodes are indistinguishable, the fact that $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ is a Nash equilibrium implies that so is $(\mathcal{S}_{\sigma_1}, \dots, \mathcal{S}_{\sigma(N)})$ for any permutation σ .

Naturally, we would like to prove that Nash equilibria exist. Unfortunately, we could not obtain the proof of this fact in the general case, since the space of all possible strategies is huge. Therefore, we consider the following simplifying assumption (which is, by the way, also quite reasonable since, in practice, one would hardly use the genesis as the starting vertex for the random walks due to runtime issues):

Assumption L. There is $n_1 > 0$ such that the attachment strategies of all nodes (including those that use the default attachment strategy) only depend on the restriction of the tangle to the last n_1 transactions that they see.

¹⁷i.e., the one used by nodes following the default attachment strategy

Observe that, under the above assumption, the set of all such strategies can be thought of as a compact convex subset of \mathbb{R}^d , where $d = d(n_1)$ is sufficiently large. Additionally, we assume that the *actual* set of strategies that can be used by the nodes may be further restricted to some compact convex subset of the above subset. Also, observe that the set of all possible restrictions of elements of \mathcal{G} on a subset of n_1 vertices is finite; we denote that set by \mathcal{G}_{n_1} . The set of all such attachments strategies will be then denoted by \mathfrak{M}_{n_1} .

In this section we use a different approach to model the network propagation delays: instead of assuming that an incoming transaction does not have information about the state of the tangle during last h units of time, we rather assume that it does not have information about the last n_0 transactions attached to the tangle, where $n_0 < n_1$ is some fixed positive number (so, effectively, the strategies would depend on subgraphs induced by $n_1 - n_0$ transactions, although the results of this section do not rely on this assumption). Clearly, these two approaches are quite similar in spirit; however, the second one permits us to avoid certain technical difficulties related to randomness of the number of unseen transactions in the first case¹⁸.

From now on, we assume that vertices contain no data, i.e., the set \mathcal{D} is empty; this is not absolutely necessary because, with the data, the proof will be essentially the same; however, the notations would become much more cumbersome. Also, there will be no reattachments; again, this would unnecessarily complicate the proofs (one would have to work with *decorated* Poisson processes). In fact, we are dealing with a so-called *random-turn game* here (see e.g. Chapter 9 of [15] for other examples).

To proceed, we need the following

Lemma 2.2. *Let P be the transition matrix of an irreducible and aperiodic discrete-time Markov chain on a finite state space E . Let \hat{P} be a continuous map from a compact set $F \subset \mathbb{R}^d$ to the set of all stochastic matrices on E (equipped by the distance inherited from the usual matrix norm on the space of all matrices on E). Fix $\theta \in (0, 1)$, denote $\tilde{P}(s) = \theta P + (1 - \theta)\hat{P}(s)$, and let π_s be the (unique) stationary measure of $\tilde{P}(s)$. Then, the map $s \mapsto \pi_s$ is also continuous.*

Proof. In the following we give a (rather) probabilistic proof of this fact via the Kac's lemma, although, of course, a purely analytic proof is also possible. Irreducibility and aperiodicity of P imply that, for some $m_0 \in \mathbb{N}$ and $\varepsilon_0 > 0$

$$P_{xy}^{m_0} \geq \varepsilon_0 \tag{6}$$

¹⁸also, it will be more natural and convenient to pass from continuous to discrete time

for all $x, y \in E$. Now, (6) implies that

$$\tilde{P}_{xy}^{m_0}(s) \geq \theta^{m_0} \varepsilon_0 \quad (7)$$

for all $x, y \in E$ and all $s \in F$.

Being $(X_n, n \geq 0)$ a stochastic process on E , let us define

$$\tau(x) = \min\{k \geq 1 : X_k = x\}$$

(with the convention $\min \emptyset = \infty$) to be the *hitting time* of the site $x \in E$ by the stochastic process X . Now, let $\mathbb{P}_x^{(s)}$ and $\mathbb{E}_x^{(s)}$ be the probability and the expectation with respect to the Markov chain with transition matrix $\tilde{P}(s)$ starting from $x \in E$. We now recall the Kac's lemma (cf. e.g. Theorem 1.22 of [8]): for all $x \in E$ it holds that

$$\pi_s(x) = \frac{1}{\mathbb{E}_x^{(s)} \tau(x)}. \quad (8)$$

Now, (7) readily implies that, for all $x \in E$ and $n \in \mathbb{N}$,

$$\mathbb{P}_x^{(s)}[\tau(x) \geq n] \leq c_1 e^{-c_2 n} \quad (9)$$

for some positive constants $c_{1,2}$ which do not depend on s . This in its turn implies that the series

$$\mathbb{E}_x^{(s)} \tau(x) = \sum_{n=1}^{\infty} \mathbb{P}_x^{(s)}[\tau(x) \geq n]$$

converges uniformly in s and so $\mathbb{E}_x^{(s)} \tau(x)$ is uniformly bounded from above¹⁹; also, the Uniform Limit Theorem implies that $\mathbb{E}_x^{(s)} \tau(x)$ is continuous in s . Therefore, for any $x \in E$, (8) implies that $\pi_s(x)$ is also a continuous function of s . \square

Consider, for the moment, the situation when all nodes use the same (default) attachment strategy (i.e., there are no selfish nodes). The restriction of the tangle on the last n_1 transactions then becomes a Markov chain on the state space \mathcal{G}_{n_1} . We now make the following technical assumption on that Markov chain:

Assumption D. The above Markov chain is irreducible and aperiodic.

It is important to observe that Assumption D is *not* guaranteed to hold for *every* natural attachment strategy; however, still, this is not a very restrictive assumption in practice because every finite Markov chain may be turned into an irreducible and aperiodic one by an arbitrarily small perturbation of the transition matrix.

Then, we are able to prove the following

¹⁹and, of course, it is also bounded from below by 1

Theorem 2.3. *Under Assumptions L and D, the system has at least one Nash equilibrium.*

Proof. The authors were unable to find a result available in the literature that implies Theorem 2.3 directly; nevertheless, its proof is quite standard and essentially follows Nash’s original paper [17] (see also [10]). There is only one technical difficulty, which we intend to address via the above preparatory steps: one needs to prove the continuity of the cost function.

Denote by $\pi_{\mathcal{S}}$ the invariant measure of the Markov chain given that the (selfish) nodes use the “strategy vector” $\mathbf{s} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$. Then, the idea is to use Lemma 2.2 with $\theta = \frac{\kappa}{\kappa+1}$, P the transition matrix obtained from the default attachment strategy, and $\widehat{P}(s)$ is the transition matrix obtained from the strategy $\mathcal{S}' = N^{-1} \sum_{k=1}^N \mathcal{S}_k$ (observe that N nodes using the strategies $\mathcal{S}_1, \dots, \mathcal{S}_N$, is the same as one node with strategy \mathcal{S}' issuing transactions N times faster). Assumption D together with Lemma 2.2 then imply that $\pi_{\mathbf{s}} := \pi_{\mathcal{S}'}$ is a continuous function of \mathbf{s} .

Let $\mathbb{E}_{\pi_{\mathcal{S}'}}^{\mathcal{S}, \widehat{\mathcal{S}}}$ be the expectation with respect to the following procedure: take the “starting” graph according to $\pi_{\mathcal{S}'}$, then attach to it a transaction according to the strategy \mathcal{S} , and then keep attaching subsequent transactions according to the strategy $\widehat{\mathcal{S}}$ (instead of \mathcal{S}' and $\widehat{\mathcal{S}}$ we may also use the strategy vectors; \mathcal{S}' and $\widehat{\mathcal{S}}$ would be then their averages). Let also $W^{(k)}$ be the random variable defined as in (4) for a transaction issued by the k th node. Then, the Ergodic Theorem for Markov chains (see e.g. Theorem 1.23 of [8]) implies that

$$\mathfrak{c}^{(k)}(\mathcal{S}) = \mathbb{E}_{\pi_{\mathcal{S}'}}^{\mathcal{S}_k, \mathcal{S}'} W^{(k)}. \quad (10)$$

It is not difficult to see that the above expression is a polynomial of the \mathcal{S} ’s coefficients (i.e., the corresponding probabilities) and $\pi_{\mathcal{S}'}$ -values, and hence it is a continuous function on the space of strategies \mathfrak{M}_{n_1} . Using this, the rest of the proof is standard, it is obtained as a consequence of the Kakutani’s fixed point theorem [14] (also with the help of the Berge’s Maximum Theorem, see e.g. Chapter E.3 of [18]). \square

Symmetric games do not always have symmetric Nash equilibria, as shown in [9]. Also, even when such equilibria exist in the class of mixed strategies, they may be “inferior” to asymmetric pure equilibria; for example, this happens in the classical “Battle of the sexes” game (see e.g. Section 7.2 of [15]).

Now, the goal is to prove that, if the number of selfish nodes N is large, then for *any* equilibrium state the costs of distinct nodes cannot be much different. Namely, we have the following

Theorem 2.4. For any $\varepsilon > 0$ there exists N_0 (depending on the default attachment strategy) such that, for all $N \geq N_0$ and any Nash equilibrium $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ it holds that

$$|\mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_N) - \mathfrak{C}^{(j)}(\mathcal{S}_1, \dots, \mathcal{S}_N)| < \varepsilon \quad (11)$$

for all $k, j \in \{1, \dots, N\}$.

Proof. Without restricting generality we may assume that

$$\begin{aligned} \mathfrak{C}^{(1)}(\mathcal{S}_1, \dots, \mathcal{S}_N) &= \max_{k=1, \dots, N} \mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_N), \\ \mathfrak{C}^{(2)}(\mathcal{S}_1, \dots, \mathcal{S}_N) &= \min_{k=1, \dots, N} \mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_N), \end{aligned}$$

so we then need to proof that $\mathfrak{C}^{(1)}(\mathbf{s}) - \mathfrak{C}^{(2)}(\mathbf{s}) < \varepsilon$, where $\mathbf{s} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$. Now, the main idea of the proof is the following: if $\mathfrak{C}^{(1)}(\mathbf{s})$ is considerably larger than $\mathfrak{C}^{(2)}(\mathbf{s})$, then the owner of the first node may decide to adopt the strategy used by the second one. This would not necessarily decrease his costs to the (former) costs of the second node since a change in an individual strategy leads to changes in *all* costs; however, when N is large, the effects of changing the strategy of only one node would be small, and (if the difference of $\mathfrak{C}^{(1)}(\mathbf{s})$ and $\mathfrak{C}^{(2)}(\mathbf{s})$ were not small) this would lead to a contradiction to the assumption that \mathbf{s} was a Nash equilibrium.

So, let us denote $\mathbf{s}' = (\mathcal{S}_2, \mathcal{S}_2, \mathcal{S}_3, \dots, \mathcal{S}_N)$, the strategy vector after the first node adopted the strategy of its “more successful” friend. Let

$$\mathcal{S} = \frac{1}{N}(\mathcal{S}_1 + \dots + \mathcal{S}_N) \text{ and } \mathcal{S}' = \frac{1}{N}(2\mathcal{S}_2 + \mathcal{S}_3 + \dots + \mathcal{S}_N)$$

be the two “averaged” strategies. In the following, we are going to compare $\mathfrak{C}^{(2)}(\mathbf{s}) = \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_2, \mathcal{S})} W^{(2)}$ (the “old” cost of the second node) with $\mathfrak{C}^{(1)}(\mathbf{s}') = \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S}')} W^{(1)}$ (the “new” cost of the first node, after it adopted the second node’s strategy). We need the following

Lemma 2.5. For any measure π on \mathcal{G}_{n_1} and any strategy vectors $\mathbf{s} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$ and $\mathbf{s}' = (\mathcal{S}'_1, \dots, \mathcal{S}'_N)$ such that $\mathcal{S}_k = \mathcal{S}'_k$ for all $k = 2, \dots, N$, we have

$$|\mathbb{E}_{\pi}^{(\mathcal{S}_j, \mathcal{S})} W^{(j)} - \mathbb{E}_{\pi}^{(\mathcal{S}'_j, \mathcal{S}')} W^{(j)}| \leq \frac{M_0}{N} \quad (12)$$

for all $j = 2, \dots, N$.

Proof. Let us define the event

$$A = \{\text{among the } M_0 \text{ transactions there is at least one issued by the first node}\},$$

and observe that, by the union bound, the probability that it occurs is at most M_0/N . Then, using the fact that $\mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S})}(W^{(j)} \mathbf{1}_{A^c}) = \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S}')} (W^{(j)} \mathbf{1}_{A^c})$ (since, on A^c , the first node does not “contribute” to $W^{(j)}$), write

$$\begin{aligned} & \left| \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S})} W^{(j)} - \mathbb{E}_\pi^{(\mathcal{S}'_j, \mathcal{S}')} W^{(j)} \right| \\ &= \left| \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S})} (W^{(j)} \mathbf{1}_A) + \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S})} (W^{(j)} \mathbf{1}_{A^c}) - \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S}')} (W^{(j)} \mathbf{1}_A) - \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S}')} (W^{(j)} \mathbf{1}_{A^c}) \right| \\ &= \left| \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S})} (W^{(j)} \mathbf{1}_A) - \mathbb{E}_\pi^{(\mathcal{S}_j, \mathcal{S}')} (W^{(j)} \mathbf{1}_A) \right| \\ &\leq \frac{M_0}{N}, \end{aligned}$$

where we also used that $W^{(j)} \leq 1$. This concludes the proof of Lemma 2.5. \square

We continue proving Theorem 2.4. First, by symmetry, we have

$$\mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S}')} W^{(1)} = \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S})} W^{(2)}. \quad (13)$$

Also, it holds that

$$\left| \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S}')} W^{(2)} - \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S})} W^{(2)} \right| \leq \frac{M_0}{N} \quad (14)$$

by Lemma 2.5. Then, similarly to the proof of Theorem 2.3, we can obtain that the function

$$(\mathcal{S}, \mathcal{S}', \mathcal{S}'') \mapsto \mathbb{E}_{\pi_{\mathcal{S}''}}^{(\mathcal{S}, \mathcal{S}')} W^{(2)}$$

is continuous; since it is defined on a compact, it is also uniformly continuous. That is, for any $\varepsilon' > 0$ there exist $\delta' > 0$ such that if $\|(\mathcal{S}, \mathcal{S}', \mathcal{S}'') - (\tilde{\mathcal{S}}, \tilde{\mathcal{S}}', \tilde{\mathcal{S}}'')\| < \delta'$, then

$$\left| \mathbb{E}_{\pi_{\mathcal{S}''}}^{(\mathcal{S}, \mathcal{S}')} W^{(2)} - \mathbb{E}_{\pi_{\tilde{\mathcal{S}}''}}^{(\tilde{\mathcal{S}}, \tilde{\mathcal{S}}')} W^{(2)} \right| < \varepsilon'.$$

Choose $N_0 = \lceil 1/\delta' \rceil$. We then obtain from the above that

$$\left| \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S})} W^{(2)} - \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S}')} W^{(2)} \right| < \varepsilon'. \quad (15)$$

The relations (13), (14), and (15) imply that

$$\left| \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S}')} W^{(1)} - \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S})} W^{(2)} \right| \leq \varepsilon' + \frac{M_0}{N}.$$

On the other hand, since we assumed that \mathbf{s} is a Nash equilibrium, it holds that

$$\mathbb{E}_{\pi_{\mathcal{S}'}}^{(\mathcal{S}_2, \mathcal{S}')} W^{(1)} = \mathfrak{C}^{(1)}(\mathbf{s}') \geq \mathfrak{C}^{(1)}(\mathbf{s}) = \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_1, \mathcal{S})} W^{(1)}, \quad (16)$$

which implies that

$$\mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_1, \mathcal{S})} W^{(1)} - \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_2, \mathcal{S})} W^{(2)} \leq \varepsilon' + \frac{M_0}{N}.$$

This concludes the proof of Theorem 2.4. \square

Now, let us define the notion of *approximate* Nash equilibrium:

Definition 2.6. For a fixed $\varepsilon > 0$, we say that a set of strategies $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ is an ε -equilibrium if

$$\mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_{k-1}, \mathcal{S}_k, \mathcal{S}_{k+1}, \dots, \mathcal{S}_N) \leq \mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_{k-1}, \mathcal{S}', \mathcal{S}_{k+1}, \dots, \mathcal{S}_N) + \varepsilon$$

for any k and any $\mathcal{S}' \neq \mathcal{S}_k$.

The motivation for introducing this notion is that, if ε is very small, then, in practice, ε -equilibria are essentially indistinguishable from the “true” Nash equilibria.

Theorem 2.7. For any $\varepsilon > 0$ there exists N_0 (depending on the default attachment strategy) such that, for all $N \geq N_0$ and any Nash equilibrium $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ it holds that $(\mathcal{S}, \dots, \mathcal{S})$ is an ε -equilibrium, where

$$\mathcal{S} = \frac{1}{N} \sum_{k=1}^N \mathcal{S}^{(k)} \quad (17)$$

(that is, all selfish nodes use the same “averaged” strategy defined above). The costs of all selfish nodes are then equal to

$$\frac{1}{N} \sum_{k=1}^N \mathfrak{C}^{(k)}(\mathcal{S}_1, \dots, \mathcal{S}_N),$$

that is, the average cost in the Nash equilibrium.

In other words, for large N one can essentially assume that all selfish nodes follow the same attachment strategy.

Proof. To begin, we observe that the proof of the second part is immediate, since, as already noted before, for an external observer, the situation where there are N nodes with strategies $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ is indistinguishable from the situation with one node with averaged strategy.

Now, we need to prove that, for any fixed $\varepsilon' > 0$ it holds that

$$\mathfrak{C}^{(1)}(\mathcal{S}, \dots, \mathcal{S}) \leq \mathfrak{C}^{(1)}(\tilde{\mathcal{S}}, \mathcal{S}, \dots, \mathcal{S}) + \varepsilon' \quad (18)$$

for all large enough N (the claim would then follow by symmetry). Recall that we have

$$\mathfrak{C}^{(1)}(\mathcal{S}, \dots, \mathcal{S}) = \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}, \mathcal{S})} W^{(1)}, \quad (19)$$

$$\mathfrak{C}^{(1)}(\mathcal{S}_1, \dots, \mathcal{S}_N) = \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_1, \mathcal{S})} W^{(1)}, \quad (20)$$

and

$$\mathfrak{C}^{(1)}(\tilde{\mathcal{S}}, \mathcal{S}, \dots, \mathcal{S}) = \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\tilde{\mathcal{S}}, \mathcal{S}')} W^{(1)}, \quad (21)$$

where

$$\mathcal{S}' = \frac{1}{N}(\tilde{\mathcal{S}} + (N-1)\mathcal{S}) = \frac{1}{N}\left(\tilde{\mathcal{S}} + \frac{N-1}{N}(\mathcal{S}_1 + \dots + \mathcal{S}_N)\right).$$

Now, the second part of this theorem together with Theorem 2.4 imply²⁰ that, for any fixed $\varepsilon > 0$

$$\left| \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}, \mathcal{S})} W^{(1)} - \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_1, \mathcal{S})} W^{(1)} \right| < \varepsilon \quad (22)$$

for all large enough N .

Next, let us denote

$$\mathcal{S}'' = \frac{1}{N}(\tilde{\mathcal{S}} + \mathcal{S}_2 + \dots + \mathcal{S}_N).$$

Then, again using the uniform continuity argument (as in the proof of Theorem 2.4), we obtain that, for any $\varepsilon'' > 0$

$$\left| \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\tilde{\mathcal{S}}, \mathcal{S}')} W^{(1)} - \mathbb{E}_{\pi_{\mathcal{S}''}}^{(\tilde{\mathcal{S}}, \mathcal{S}'')} W^{(1)} \right| < \varepsilon'' \quad (23)$$

for all large enough N . However,

$$\mathbb{E}_{\pi_{\mathcal{S}''}}^{(\tilde{\mathcal{S}}, \mathcal{S}'')} W^{(1)} = \mathfrak{C}^{(1)}(\tilde{\mathcal{S}}, \mathcal{S}_2, \dots, \mathcal{S}_N) \geq \mathfrak{C}^{(1)}(\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_N) = \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}_1, \mathcal{S})} W^{(1)},$$

since $(\mathcal{S}_1, \dots, \mathcal{S}_N)$ is a Nash equilibrium. Then, (22)–(23) imply that

$$\left| \mathbb{E}_{\pi_{\mathcal{S}}}^{(\mathcal{S}, \mathcal{S})} W^{(1)} - \mathbb{E}_{\pi_{\mathcal{S}'}}^{(\tilde{\mathcal{S}}, \mathcal{S}')} W^{(1)} \right| < \varepsilon + \varepsilon'',$$

and, recalling (19) and (21), we conclude the proof of Theorem 2.7. \square

²⁰note that Theorem 2.4 implies that, when N is large, the nodes already have “almost” the same cost in the Nash equilibrium $(\mathcal{S}_1, \dots, \mathcal{S}_N)$

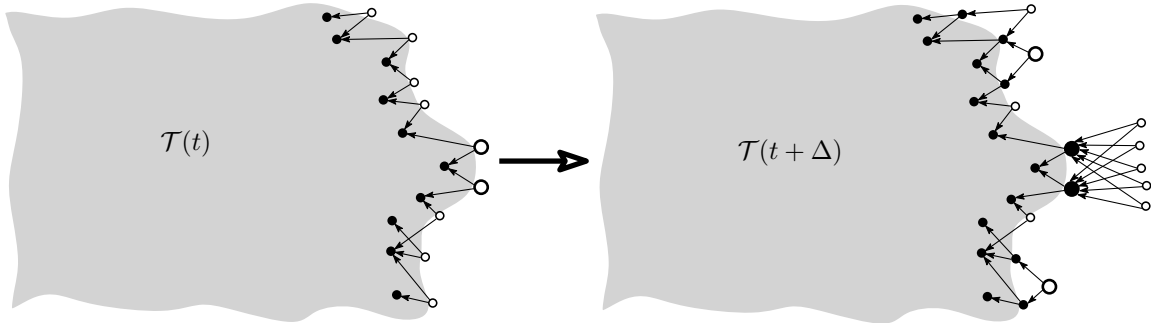


Figure 2: Why the “greedy” tip selection strategy will not work (the two “best” tips are shown as larger circles).

3 Simulations

In this section we investigate Nash equilibria between selfish nodes via simulations. This is motivated by the following important question: since the choice of an attachment strategy is not enforced, there may indeed be nodes which would prefer to “optimise” their strategies in order to decrease the mean confirmation time of their transactions. So, can this lead to a situation where the corresponding Nash equilibrium is “bad for everybody”, effectively leading to the system’s malfunctioning (again, we do not specify the exact meaning of that)?

Due to Theorem 2.7 we may assume that all selfish nodes use the same attachment strategy. Even then, it is probably unfeasible to calculate that strategy exactly; instead, we resort to simulations, which indeed will show that the equilibrium strategy of the selfish nodes will not be much different from the (suitably chosen) default strategy. But, before doing that, let us explain the intuition behind this fact. Naively, a natural strategy for a selfish node would be the following:

- (1) Calculate the exit distribution of the tip-selecting random walk.
- (2) Find the two tips where this distribution attains its “best”²¹ values.
- (3) Approve these two tips.

However, this strategy fails when other selfish nodes are present. To understand this, look at Figure 2: *many* selfish nodes attach their transactions to the two “best” tips. As a result, the “neighborhood” of these two tips becomes “overcrowded”: there

²¹i.e., the maximum and the second-to-maximum

is so much competition between the transactions issued by the selfish nodes, that the chances of them being approved soon actually decrease²².

To illustrate this fact, several simulations have been done. All the results depicted here were generated using (2) as the transition probabilities, with $q = 1/3$, and a network delay of $h = 1$. Also, a transaction will be reattached if the two following criteria are met:

- (1) the transaction is older than 20 seconds²³;
- (2) the transaction is not referenced by the tip selected by a random walk with $\alpha = \infty$ ²⁴.

This way, we guarantee not only that the unconfirmed transactions will be eventually confirmed, but also that all transactions that were never reattached are referenced by most of the tips. Note that when the reattachment is allowed in the simulations, if a new transaction references an old, already reattached transaction together with its newly reissued counterpart, there will be a double spending. Even though the odds of that are low (since when a transaction is re-emitted, it will be old enough to be almost never chosen by the random walk algorithm), a specific procedure was included in the simulations in order to not allow double spendings.

Figure 3 depicts the typical cumulative distribution of the time of the first approval, for a low α and $\lambda h = 50$. Note that roughly 95% of the transactions will be approved before $t = 5$ s, and almost the totality of transactions will be approved before $t = 10$ s. That behaviour will be similar for all studied parameters. The average cost defined in equations (5) and (4) will have a certain meaning, depending on the choice of M_0 . This average cost will be related to the average time of approval of a transaction (indeed, the average time will be approximately W/λ). So, in both cases ($\lambda = 25$ and $\lambda = 50$), the mean cost was calculated over the transactions attached at a interval of time of approximately 10s ($M_0 = 500$ for $\lambda = 50$ and $M_0 = 250$ for $\lambda = 25$), what makes this cost something reasonable to optimise.

²²the “new” best tips are not among them, as shown on Figure 2 on the right

²³even though this is the first mention to a time variable in this paper, the simulation compares actual times in this criteria

²⁴Here, when the random walk must choose among n transactions with the same weight, it will choose randomly, with equal probabilities

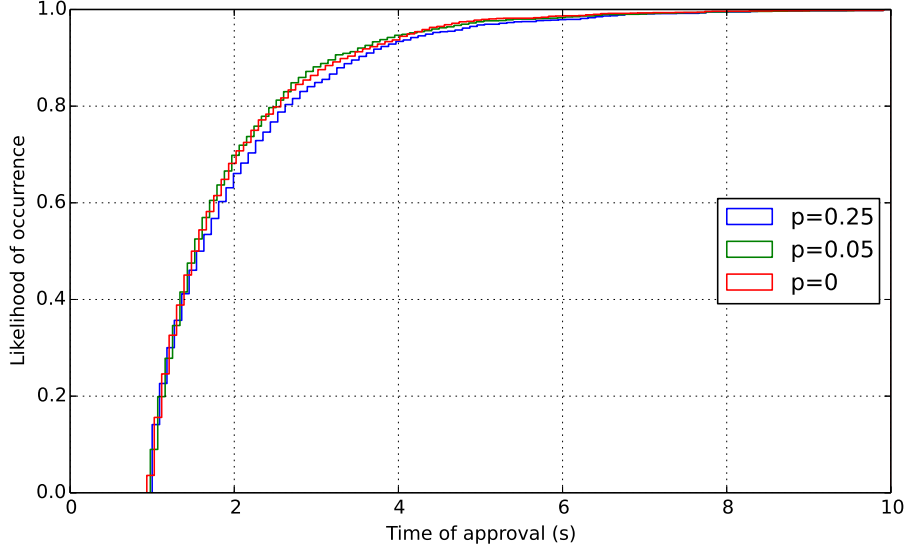


Figure 3: Cumulative distribution of time of approvals for some values of p (that will be defined later)

3.1 One dimensional Nash equilibria

In this section, we will study the Nash equilibria of the tangle problem, considering the following strategy subspace:

$$\mathcal{S}_i = \mathcal{S} = (1 - \theta)\mathcal{S}^0 + \theta\mathcal{S}^1 \quad \text{for each } i = 1, \dots, N,$$

where \mathcal{S}^0 is the default tip selection strategy, \mathcal{S}^1 is the selfish strategy defined in the beginning of this section and $\theta \in [0, 1]$. The goal is to find the Nash equilibria relative to the costs defined in the last section (equations (5) and (4)). The selfish nodes will try to optimise their transaction cost with respect to θ .

Now, suppose that we have a fixed fraction γ of selfish nodes, that chooses a strategy among the possible \mathcal{S} . The non-selfish nodes will not be able to choose their strategy, so they will be restricted, as expected, to \mathcal{S}^0 . Note that, since they can not choose their strategy, they will not “play” the game. Since the costs are linear over \mathcal{S} , such mixed strategy game will be equivalent to a second game where only a fraction $p = \gamma\theta \leq \gamma$ of the nodes chooses \mathcal{S}^1 over \mathcal{S}^0 , and the rest of the nodes chooses \mathcal{S}^0 over \mathcal{S}^1 .

Figure 4(a) represents a typical graph of average costs of transactions issued under \mathcal{S}^0 and \mathcal{S}^1 , as a function of the fraction p , for a low α and two different values of λ . As already demonstrated, when in equilibrium, the selfish nodes should issue transactions with the same average cost. That means that the system should reach equilibrium in one of the following states:

- (1) some selfish nodes choose \mathcal{S}^0 and the rest choose \mathcal{S}^1 ;
- (2) all selfish nodes choose \mathcal{S}^1 ;
- (3) all selfish nodes choose \mathcal{S}^0 .

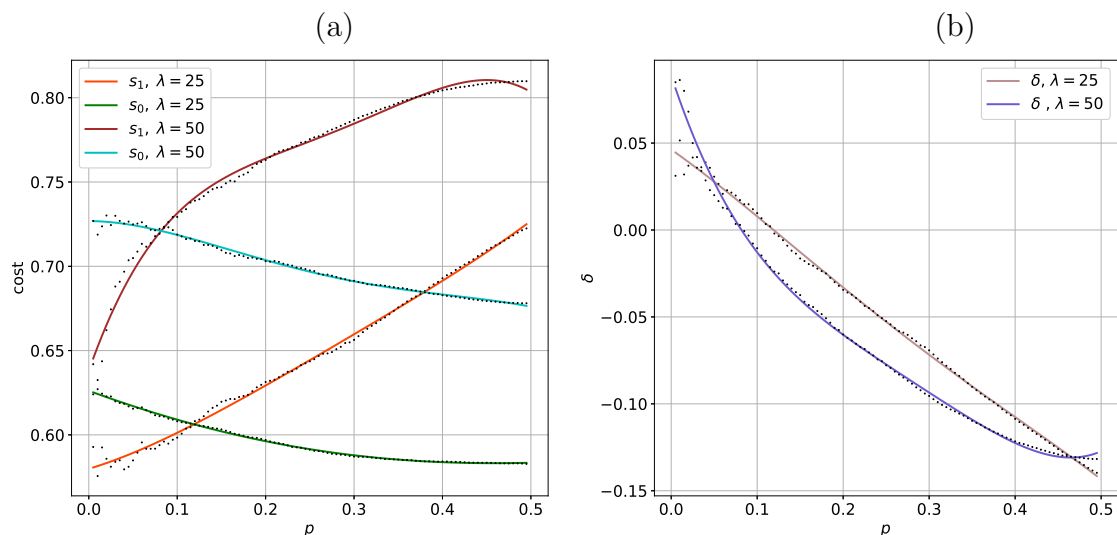


Figure 4: Costs (a) and gain of the strategy \mathcal{S}^1 over \mathcal{S}^0 ; (b) for $\alpha = 0.01$.

If the two curves on the graphs do not intersect, the equilibrium should be clearly at state (2) or (3), depending on which of the average costs is larger. If the two curves on the graphs intercept each other, we will also have the intersection point as a Nash equilibrium candidate. We call \bar{s} the vector of strategies on equilibrium and \bar{p} the fraction of nodes that will issue transactions under \mathcal{S}^1 when the system is in \bar{s} . We define $p^- = \bar{p} - \frac{\gamma}{N}$ and $p^+ = \bar{p} + \frac{\gamma}{N}$, meaning that p^- and p^+ will be deviations from \bar{p} , that result from one node switching strategies, from \mathcal{S}^0 to \mathcal{S}^1 and from \mathcal{S}^1 to \mathcal{S}^0 , respectively. We also define \bar{s}_1 and \bar{s}_2 as strategy vectors related to p^- and p^+ . Note on Figure 3.1 that this kind of Nash equilibrium candidate may not be a real equilibrium. In the first example (3.1(a)), when the system is at point \bar{p} and a

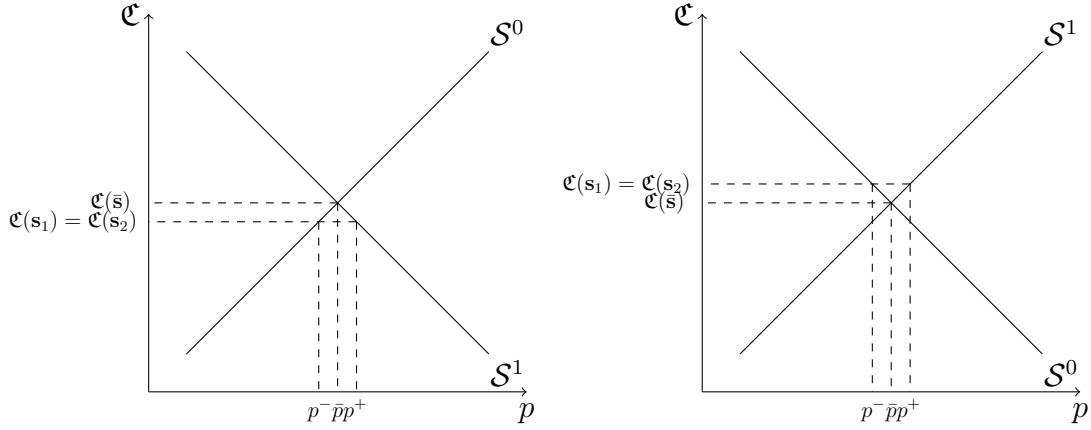


Figure 5: Different Nash equilibrium points in systems with similar curves

node switches strategies, from \mathcal{S}^0 to \mathcal{S}^1 , the cost actually decreases, so \bar{p} cannot be a Nash equilibrium. On the other hand, the second example (3.1(b)) shows a Nash equilibrium at point \bar{p} , since deviations to p^- and p^+ will increase costs.

Now, let us re-examine Figure 4(a). Here, the Nash equilibrium will occur at the point \bar{p} , since we have a situation as on Figure 3.1(b). That point is easily found at Figure 4(b), when $\delta = 0$. Note that the Nash equilibrium for a larger λ will be at a smaller θ_0 than the Nash equilibrium for a smaller λ . This was already expected, since, for a larger λh , the tips will be naturally more “overcrowded”, so the effect depicted at Figure 2 will be amplified. Thus, the Nash equilibrium for the higher λh cases must occur with a smaller proportion of transactions issued with the pure strategy \mathcal{S}^1 .

Reconsider now the mixed strategy game. In the case when all the nodes are allowed to choose between the two pure strategies (\mathcal{S}^0 and \mathcal{S}^1), the Nash equilibrium will be indeed at $\theta_0 = \bar{p}$ (as expected, since in this case $\gamma = 1$). If just a fraction $\gamma = p/\theta > \bar{p}$ of the nodes is selfish, then the Nash equilibrium will occur when $\theta_0 = \bar{p}/\gamma$. Now, if $\gamma \leq \bar{p}$, the costs of the nodes will not coincide²⁵. In that case, the average cost of transactions under \mathcal{S}^1 will always be smaller than the average cost of transactions under \mathcal{S}^0 , meaning that the Nash equilibrium will be met at $\theta_0 = 1$. Summing up, the Nash equilibrium θ_0 , in these cases, will be met at:

$$\theta_0 = \min\{\bar{p}/\gamma, 1\}.$$

²⁵That is the case for the range of studied parameters. That is not a theoretical claim, and it will be not necessarily true for all strategy spaces and parameters

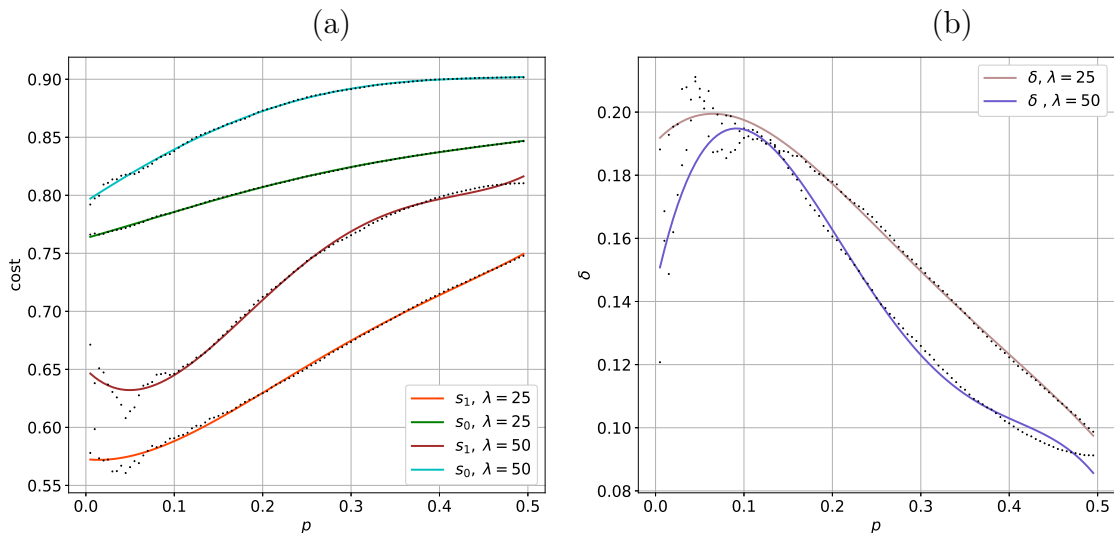


Figure 6: Costs (a) and gain (b) of the strategy \mathcal{S}^1 over \mathcal{S}^0 ; for $\alpha = 0.5$.

Figure 6(a) represents a typical graph of average costs of transactions under \mathcal{S}^0 and transactions under \mathcal{S}^1 as a function of fraction p , for a higher α . In that case, even though the average costs of transactions under \mathcal{S}^0 and transactions under \mathcal{S}^1 do not coincide for any reasonable p (meaning that, here, the Nash equilibrium will be met at $\theta = 1$), the typical difference between the possible pure strategies (that, from now on, we will call absolute gains) will be low, as depicted on Figure 6(b).

Figure 7 shows the average cost increase imposed on the nodes following the default strategy by the nodes issuing transactions under \mathcal{S}^1 . Let $W(p)$ be the non-greedy nodes costs depicted in Figure 6(a). The cost increase is calculated as $(W(p) - W(0))/W(0)$, so it will be the percentual difference of the cost of a non-selfish node in the presence of a percentage p of selfish transactions and the cost of a non-selfish node when there are no selfish transactions at all. This difference is low, meaning that the presence of selfish nodes do not harm the efficiency of the non-selfish nodes. Note that this difference is small for all reasonable values of p , but even for the larger available p , the difference is less than 25%. An interesting phenomenon, as shown in the same graph, is that the average cost increase imposed on the non-greedy nodes may actually be negative. For low values of α , just a small fraction of the transactions under \mathcal{S}^0 will share the approved tips with the transactions under \mathcal{S}^1 . This fraction of transactions will approve overcrowded tips, and will have their costs increased. All the other transactions under \mathcal{S}^0 will have their sites less crowded, since an increase

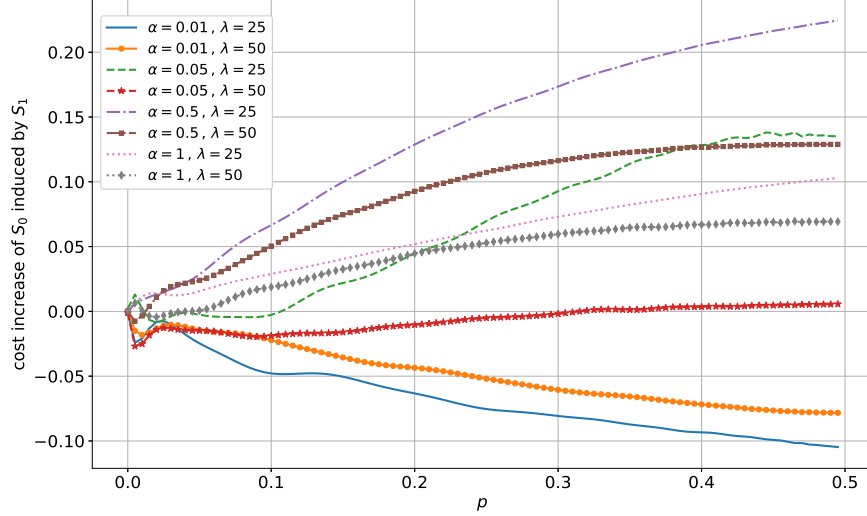


Figure 7: Cost increase of the transactions issued by the strategy \mathcal{S}^0 induced by the presence of transactions emitted by the strategy \mathcal{S}^1 , in percentage.

in \mathcal{S}^1 will mean a decrease in competition over these transactions. Finally, on average, the honest nodes will have their costs decreased.

Figures 8 and 9 are analogous to the first figures, for other values of α and λh ; part (a) of each figure represents average costs and part (b) absolute gains.

3.2 Multidimensional Nash equilibria

In the same way as the last section, a game with a mixed strategy space:

$$\mathcal{S} = \sum_{i=0}^N \theta_i \mathcal{S}^i,$$

where $\sum_{i=0}^N \theta_i = 1$ and $\{\mathcal{S}^i\}_{i=0,\dots,N}$ a set of strategies is equivalent to a game where the nodes must choose among the simple strategies $\{\mathcal{S}^i\}$. That means all the possible strategies must have the same cost in order to the game reach some equilibrium. In this case, in order to simplify the data, the studied object was the probability of a given tip to be chosen by the selfish nodes and the non-selfish nodes. These studied tips were ordered by the random walk exit probability, from the most probable to the least probable. Figure 10 represents the typical probability profile for the selfish and non-selfish nodes. The typical gain of the selfish nodes and the increase of the

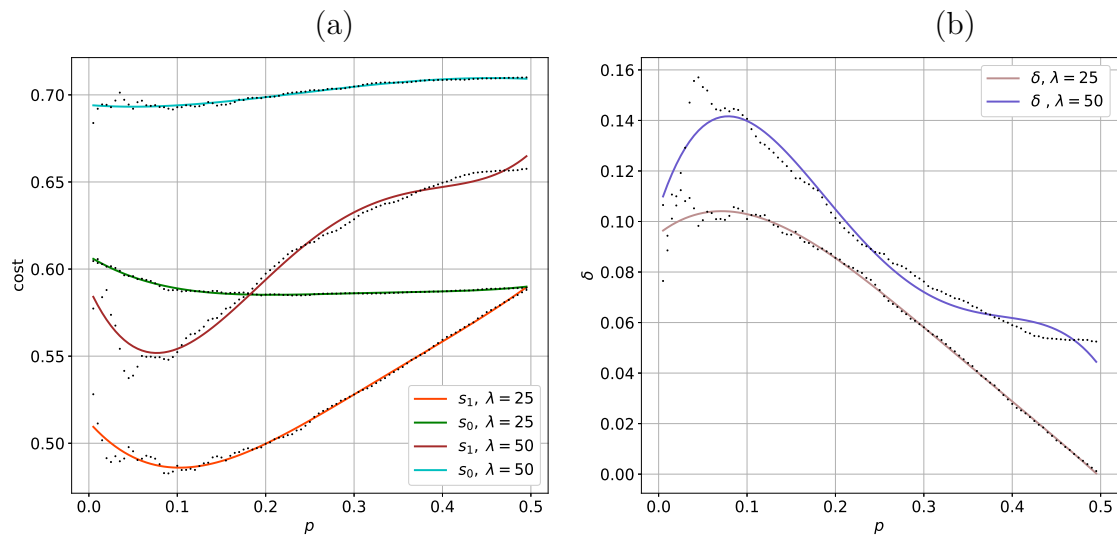


Figure 8: Costs (a) and gain (b) of the strategy \mathcal{S}^1 over \mathcal{S}^0 ; for $\alpha = 0.05$.

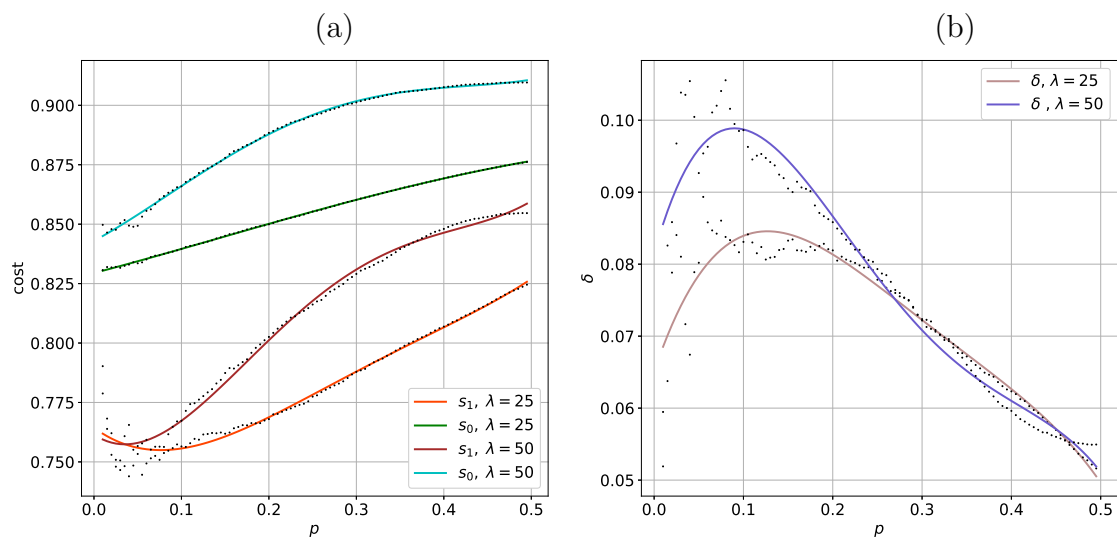


Figure 9: Costs (a) and gain (b) of the strategy \mathcal{S}^1 over \mathcal{S}^0 ; for $\alpha = 1$.

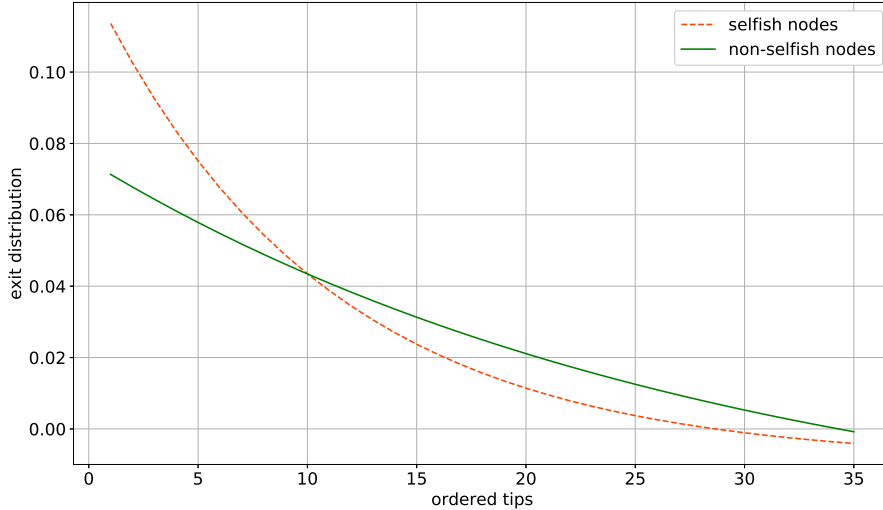


Figure 10: Exit probabilities in the equilibrium for the selfish and the non-selfish nodes.

non-selfish nodes average cost due to the presence of selfish nodes are both small; in our simulations they were always less than 10% (3.615% for the maximum gain and 7.0% for the maximum of the cost increase of the non-selfish nodes induced by the selfish ones, comparing all parameters of λ and α).

4 Conclusions and future work

In the first part of this paper, we prove the existence of (“almost symmetric”) Nash equilibria for a game in the tangle where a part of players tries to optimise their attachment strategies. In the second part of the paper, we numerically determine, for a simple space strategy and some range of parameters, where these equilibria are located.

Our results show that the studied selfish strategy outperform the non-selfish ones by a reasonable order of magnitude. The data show a 25% (in the most extreme scenario) difference in the nodes gains, which in some situations, may be large enough. Nevertheless, the computational cost of a selfish strategy is intrinsically larger than the computational cost of the non-selfish strategies, since the selfish strategy uses the probability distribution of the tips, which is costly to calculate for a random walk with backtracking. They will also have to monitor the tangle, to know its parameters

(like λ , h etc) and act accordingly. Also, even a extreme scenario, where almost half of the transactions were issued by a selfish node, is not enough to harm the non-selfish ones in a meaningful way.

On the other hand, our results raise further questions. The obtained data exhibit a deep qualitative dependence on the parameter α of the simulation. This parameter is related to the randomness of the random walk (a low α implies a high randomness; a higher α implies a low randomness, meaning that the walk will be almost deterministic). Further simulations will be done in order to study the effect of that variable in the equilibria. Also, we only studied equilibria for a given cost, relative to the probability of confirmation of the transactions in a certain interval of time. Since this probability depends heavily on the interval of time chosen (because the probability distribution of the confirmations is far from uniform), another time intervals (that will have another practical meaning) must be analysed.

Finally, the equilibrium in the multidimensional strategy space should be studied in a more quantitative and analytic way, since it should depend strongly on α and p ; and until now it was studied in just a narrow range of parameters. Further research will also be done in order to optimise the default tip selection strategy in a way that minimises this cost imposed by the selfish strategies. Through implementing research methods and techniques from the cross-reactive fields of measure theory, game theory, and graph theory, progress towards resolving the tangle-related open problems has been well under way and will continue to be under investigation.

Acknowledgements

The authors thank Alon Gal, Gur Huberman, Bartosz Kumierz, John Licciardello, Samuel Reid, and Clara Shikhelman for valuable comments and suggestions.

References

- [1] <http://www.iota.org/>
- [2] L. BAIRD (2016) The Swirls hashgraph consensus algorithm: fair, fast, Byzantine fault tolerance. <http://www.swirls.com/downloads/SWIRLDS-TR-2016-01.pdf>
- [3] A. CHURYUMOV (2016) Byteball: a decentralized system for storage and transfer of value. <https://byteball.org/Byteball.pdf>

- [4] C. COOPER, A. FRIEZE (2007) The cover time of the preferential attachment graph. *J. Comb. Theory B* **97** (2), 269–290.
- [5] C. COOPER, A. FRIEZE, S. PETT (2017) The covertime of a biased random walk on $G_{n,p}$. arXiv:1708.04908
- [6] R.B. COOPER (1981) *Introduction to Queueing Theory* (2nd ed.). North Holland.
- [7] P.G. DOYLE, J.L. SNELL (1984) *Random walks and electric networks*. Carus Mathematical Monographs **22**, Mathematical Association of America, Washington.
- [8] R. DURRETT (2012) *Essentials of Stochastic Processes* (2nd. ed.) Springer.
- [9] M. FEY (2012) Symmetric games with only asymmetric equilibria. *Games Econ. Behavior* **75** (1), 424–427.
- [10] A.M. FINK (1964) Equilibrium in a stochastic n -person game. *J. Sci. Hiroshima Univ. Ser. A-I Math.* **28** (1), 89–93.
- [11] A. FRIEZE, M. KRIVELEVICH, P. MICHAELI, R. PELED (2017) On the trace of random walks on random graphs. arXiv:1508.07355
- [12] G. IACOBELLI, D.R. FIGUEIREDO, G. NEGLIA (2017) Transient and slim versus recurrent and fat: random walks and the trees they grow. arXiv:1711.02913
- [13] D. JERISON, L. LEVINE, S. SHEFFIELD (2014) Internal DLA and the Gaussian free field. *Duke Math. J.* **163** (2), 267–308.
- [14] S. KAKUTANI (1941) A generalization of Brouwer’s fixed point theorem. *Duke Math. J.* **8** (3), 457–459.
- [15] A.R. KARLIN, Y. PERES (2017) *Game Theory, Alive*. American Mathematical Society.
- [16] S.D. LERNER (2015) DagCoin: a cryptocurrency without blocks. <https://bitslog.wordpress.com/2015/09/11/dagcoin/>
- [17] J.F. NASH, JR. (1950) Equilibrium points in n -person games. *Proc. Natl. Acad. Sci.* **36** (1), 48–49.
- [18] E.A. OK (2007) *Real Analysis with Economics Applications*. Princeton University Press.

- [19] S. POPOV (2015) The tangle. https://iota.org/IOTA_Whitepaper.pdf
- [20] Y. SOMPOLINSKY, Y. LEWENBERG, A. ZOHAR (2016) SPECTRE: Serialization of proof-of-work events: confirming transactions via recursive elections. <https://eprint.iacr.org/2016/1159.pdf>
- [21] Y. SOMPOLINSKY, A. ZOHAR (2013) Secure high-rate transaction processing in Bitcoin. <https://eprint.iacr.org/2013/881.pdf>