On the timestamps in the tangle

Serguei Popov*

February 25, 2018

Description of the problem and proposed algorithms

In the following we consider the tangle [1] at a fixed moment of time; that is, the state of the tangle is fixed, and so we denote it simply by $\mathbb{T}$. Some notations: if v is a site (transaction) on the tangle, we denote by $\mathcal{A}(v)$ the set of the two transactions approved by v. We say that v *references* (indirectly approves) w if there is a sequence of sites $v = v_0, v_1, \ldots, v_k = w$ such that $v_j \in \mathcal{A}(v_{j-1})$ for all $j = 1, \ldots, k$. Let us write

$$\mathcal{P}(v) = \big\{w \in \mathbb{T} : w \text{ is referenced by } v\big\},$$
$$\mathcal{F}(v) = \big\{z \in \mathbb{T} : z \text{ references } v\big\}$$

for the "past" and the "future" with respect to v (see Figure 1). In other words, the above introduces a *partial order* structure on the tangle. Also we denote by $\mathrm{Ind}(v) = \mathbb{T} \setminus \big(\mathcal{P}(v) \cup \mathcal{F}(v)\big)$ the set of transactions which neither reference nor are referenced by v (i.e., which are, in a way, *independent* from v, hence the abbreviation used). Observe that, by definition, $v \in \mathrm{Ind}(v)$.

Next, we assume that for each transaction v there is a triple $(\mathbf{t}^-(v), \mathbf{t}(v), \mathbf{t}^+(v))$, such that $\mathbf{t}^-(v) \le \mathbf{t}(v) \le \mathbf{t}^+(v)$. In the above, $\mathbf{t}(v)$ is the *timestamp* of v; that is, in principle, $\mathbf{t}(v)$ corresponds to the time when the transaction was attached to the tangle. The quantities $\mathbf{t}^-(v)$ and $\mathbf{t}^+(v))$ are the lower and the upper limits for the timestamp; they are necessary since the nodes' clocks may be imprecise.

It is required that $\mathbf{t}^+(v) > \mathbf{t}^-(w_{1,2})$ where $\mathcal{A}(v) = \{w_1, w_2\}$, for all $v \in \mathbb{T}$; in other words, a transaction cannot approve another transaction whose timestamp interval is "from the future". Our standing assumption will be that the *large majority* (in

*author's contact information: serguei.popov@iota.org

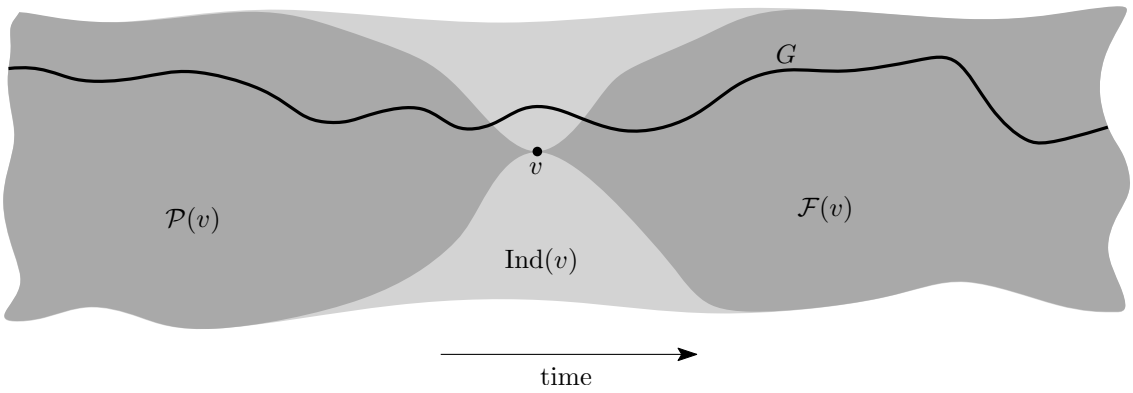

Figure 1: The tangle seen from a vertex v.

some reasonable sense) of the transactions were issued by honest nodes equipped with (more-or-less) reliable clocks, so they report the timestamps as well as the upper and lower limits with a reasonable precision.

It should be observed, though, that it would be good if there are incentives for the honest nodes to honestly report their $\mathbf{t}^-$'s and $\mathbf{t}^+$'s. There incentives may be provided using the ideas described in [3].

Now, consider some transaction v (which is, typically, already "deep inside" the tangle), it has the timestamp triple $(\mathbf{t}^-(v), \mathbf{t}(v), \mathbf{t}^+(v))$. The problem is that we do not know if it was issued by a honest or malicious node (or, maybe, it could be issued by a honest node whose clock is wrong for some reason), therefore we cannot be sure if $\mathbf{t}(v)$ is (even approximately) the real time T_v when v was attached to the tangle. Our goal is to construct the *confidence interval* $[a_v, b_v]$ for T_v; that is, we want (deterministic) $a_v \le b_v$ such that the event $\{T_v \in [a_v, b_v]\}$ occurs with high probability.

In the following, we consider two procedures for constructing such an interval.

Procedure 1. Fix $\beta \in (0, \frac{1}{2})$, and consider two data collections (rigorously speaking, two multisets)

$$\mathcal{D}^- = \big(\mathbf{t}^-(w) : w \in \mathrm{Ind}(v)\big),$$
$$\mathcal{D}^+ = \big(\mathbf{t}^+(w) : w \in \mathrm{Ind}(v)\big).$$

Then, define a_v to be the β-quantile of $\mathcal{D}^-$ and b_v to be the $(1-\beta)$-quantile of $\mathcal{D}^+$.

We have to explain why it is not a good idea to have $\beta = 0$ (i.e., take a_v and b_v the minimal and the maximal values in the above data collection). The reason is

that a malicious entity can mess with the procedure, by pushing a_v to 0 (by issuing new transactions that approve other transactions that are deep in the past, and have small $\mathbf{t}$-values) and b_v to infitity (by issuing new transactions that do not reference v and have very large $\mathbf{t}$-values). Those "disrupting" transactions will be cut off if β is sufficiently away from 0.

In principle, by increasing β we also increase our defences against the malicious behaviour described above; on the other hand, if β is "too close" to $\frac{1}{2}$, the result would be "too random" (observe that, for β close to $\frac{1}{2}$, $\mathbf{t}(v)$ itself may not make it to the confidence interval, even in the case when v was issued by a honest node!). It is unclear, for now, what would be the "optimal" value of β; in fact, one has to make assumptions on the proportion of the malicious nodes in the network and their modus operandi to perform such an analysis. In any case, as a rule of thumb, a value of β in $[0.2, 0.3]$ would probably work.

As a drawback of the above procedure, note that one has to do the calculations for each v separately (observe that, in general, $\mathrm{Ind}(v) \neq \mathrm{Ind}(w)$ when $v \neq w$); this may pose computational difficulties in case when the tangle is very large. On the other hand, the outcome of the procedure is deterministic, provided of course that the nodes use the same β and see the same state of the tangle.

Let us now describe another procedure, which is computationally easier and works for many transactions at once; on the other hand, it produces a *random*[1] result.

Procedure 2. Run the *model-tip selecting* random walk described in Section 4.1 of [1] or in Section 1.1 of [2] (i.e., a random walk with a large α or maybe with $\alpha = \infty$ which corresponds to the GHOST protocol) starting from some site deep inside the tangle, and let G be the (random) set of sites visited by this random walk (again, see Figure 1). Then, define

$$a_v = \max\{\mathbf{t}^-(w) : w \in G \cap \mathcal{P}(v)\},$$
$$b_v = \min\{\mathbf{t}^+(w) : w \in G \cap \mathcal{F}(v)\}.$$

Observe that, if a malicious node tries to forge the timestamp interval of a fixed transaction, this transaction is unlikely to be picked by the random walk. Indeed, in case the malicious node puts the timestamp "from the past" (i.e., $\mathbf{t}^+(v)$ is much less than the real time T_v when v was issued), this corresponds to the "lazy tip" case of Section 4.1 of [1]; if the malicious node puts the timestamp "from the future"

[1] that is, different nodes may arrive to different confidence intervals for T_v even if they see the same tangle

(i.e., $\mathbf{t}^-(v)$ is much greater than the real time when v was issued), such a transaction would not be referenced by anyone for long time, which again makes it unlikely that the random walk eventually passes through it.

Conclusion

The tangle is a graph with only a partial order structure, which makes it difficult (in fact, generally impossible) to establish the correct *time order* of transactions. Even if all transactions have timestamps on them, we cannot be sure that all these timestamps are accurate (there can be some malicious nodes that want to fool the network about the true time when their transactions appear, and/or some nodes with a wrong clock). Nevertheless, one can determine the *confidence intervals* for timestamps with reasonable accuracy. In the above text we described two possible algorithms for doing that; the first one is computationally more difficult, but produces a deterministic result (two nodes that use the same β and see the same state of the tangle will get the same confidence interval). The other algorithm is simpler and works for determining the timestamps' confidence intervals for several transactions at once, but produces a random result (it depends on the path that the random walk have actually chosen).

References

[1] S. Popov (2015) The tangle. `https://iota.org/IOTA_Whitepaper.pdf`

[2] S. Popov, O. Saa, P. Finardi (2017) Equilibria in the Tangle. arXiv:1712.05385

[3] S. Popov (2018) Local modifiers in the Tangle. Work in progress.