Evaluation des performances du consensus IOTA sous des hypothèses d'implémentation réalistes

Hamed Nazim Mamache, Gabin Mazué, Osama Rashid, Gewu Bu, Maria Potop-Butucaru

Evaluation des performances du consensus IOTA sous des hypothèses d'implémentation réalistes[†]

Hamed Nazim Mamache[1] Gabin Mazué[1] Osama Rashid[1] Gewu Bu[2] Maria Potop-Butucaru[1]

[1]*LIP6, Sorbonne University, Paris France*
[2]*LIMOS, Université Clermont Auvergne, Clermont-Ferrand, France*

Les registres distribués sont des technologies attractives ayant des applications variées allant du monde financier au monde des télécommunications. Dans le paysage des blockchains IOTA est le seul registre partagé dédié à l'internet des objets. Afin d'assurer des propriétés fortes de cohérence la version initiale d'IOTA proposait l'utilisation d'un contrôleur permettant d'assurer un ordre total sur les transactions insérées dans le système. Plus récemment, IOTA a proposé de remplacer le contrôleur par un algorithme de consensus. Pour cela, deux algorithmes de consensus ont été proposés par la fondation IOTA, décrits dans le framework Coordicide : Fast Probabilistic Consensus et Cellular Consensus. Au moment de leur publication, ces algorithmes étaient publicités comme étant la nouvelle brique de consensus de IOTA. Nous avons évalué les performances de ces algorithmes en utilisant des hypothèses d'implémentation réalistes. De plus, nous avons évalué la convergence des ces algorithmes en variant les topologies du réseau sous-jacent. Nos simulations montrent des taux de convergence faible, même sous des adversaires de faible puissance. De plus, nous avons observé de mauvaises performances de passage à l'échelle sauf lors des tests avec des topologies Watts Strogatz. Nos résultats indiquent que la conception de registres distribués dédiés aux IoT reste un problème ouvert et proposent des directions de recherche potentielle. Depuis l'apparition de nos résultats, la fondation IOTA a annoncé la publication imminente d'une nouvelle version de sa brique de consensus.

Mots-clefs : Consensus, Byzantine fault, IOTA

1 Introduction

Distributed Ledger Technologies (DLT) such as blockchains provide a secure way to share information between a high number of independent nodes operating under different authorities, while ensuring high availability and immutability. The use of Distributed Ledger Technologies (DLT) can respond to both security and decentralization needs in the management of IoT devices.

However, Bitcoin [NB08] and similar proposals (e.g Ethereum [W+14]) came with several drawbacks, like mandatory transaction fees and high computational requirements, that prevent them from being used as standards for IoT industry. Therefore, alternative solutions have been opened by IOTA [Pop18]. IOTA's data structure is a *Directed Acyclic Graph* (DAG) based distributed ledger, also known as the *Tangle*, aimed to overcome limitations of Bitcoin-like distributed ledgers when used in IoT environment while preserving equivalent security levels. Similar approaches have been proposed by Spectre or Phantom [SLZ16, SZ18]. Yet, IOTA and similar approaches still have flaws: 1) lack of strong consistency guarantees and 2) unclear resistance to attacks. In order to respond to these criticism IOTA proposed recently in [PMC+20] attacks resilient consensus mechanisms that plugged into the IOTA Tangle will offer strong consistency guarantees. Two consensus algorithms are proposed: Fast Probabilistic Consensus (FPC) and Cellular Consensus (CC). These two proposals have been partially evaluated in [PB21, CMP19].

[†]La version complète du papier est acceptée par IEEE ICC 2022 (HAL: https://hal.archives-ouvertes.fr/hal-03427543)

In this work we investigate the performances of IOTA consensus in several aspects. First, we run Fast Probabilistic Consensus (FPC) and Cellular Consensus (CC) on top of various topologies, from theoretical to practical. (2D Grid, Torus and Watts-Strogatz model [WS98]) then we evaluate their resilience to adversarial behavior. Our evaluation is conducted with OMNET++ simulator enriched with three adversarial models introduced in [PMC+20]. Even though most of the results reported in our study are negative they contain hints in order to design an efficient IoT dedicated blockchain.

2 Consensus, Topology & Byzantines Adversaries

The basic idea of Consensus algorithm is that all honest nodes in the system should agree dynamically with a common opinion in a distributed way, which cannot be changed easily. We follow the setting proposed in [PB21, CMP19], assuming that the time is discrete and divided into *rounds*. Let N be the set of nodes in the system. Each node $i \in N$ has an **opinion**, $O_i(r) \in \{0, 1\}$ at the round r. At the round 0 of a simulation run, all nodes will be given randomly an initial opinion according to the probability $P0$. When $P0 = 0.5$ every node has the same chance to get 0 or 1 as its initial opinion. **Consensus** is achieved, if $\forall i, j \in N, O_i(r_{end}) = O_j(r_{end})$, where r_{end} is the final round of a simulation run. An opinion held by most of the nodes among a group de nodes is a *major opinion*. The **convergence rate** for a consensus algorithm represents by the percentage of simulation runs leading to an achieved consensus.

Two consensus algorithms from IOTA have been investigated: **Fast Probabilistic Consensus** (FPC) bases on the query/reply of opinion among nodes in the network. In a non-initial round r, node i requires randomly a number of other nodes in the network. Node i calculates the mean of received opinions, $\overline{O_i}(r)$. Then i changes its opinion to 1, if $\overline{O_i}(r)$ is higher than a threshold $U_r \in [0, 1]$ chosen randomly each round by i; it changes to 0 if $\overline{O_i}(r)$ is smaller than U_r. In the case where $\overline{O_i}(r) = U_r$, node keeps its original opinion from the last round. An initial threshold τ is fixed for all nodes instead of chosen randomly. We make adjustments from original FPC [PB21] to adapt realistic environment : 1) Remove the assumption that nodes share a common random values sequence. 2) Nodes know only their neighbors, rather than all nodes in the system. Instead, they use random walk to query random nodes far away from them. **Cellular Consensus** (CC) allows nodes only to query opinions from their neighbors and change their opinions to adapt to the major opinion among their neighbors. In addition, CC allows node to verify if neighbors give their opinion honestly, once a liar is detected, it will be banned by its neighbors.

We apply these two consensus on top of three network topologies for our experimentation: **2D Grid** (where central nodes have 4 neighbors, then edge nodes have 2 to 3 neighbors only) and **Torus** (where all nodes have 4 neighbors), two theoretical topologies often used as reference topologies in algorithm analyzing, and **Watts-Strogatz** [WS98], a more realistic network model, that simulates a network with relatively high connection density and relatively low average distance between any two nodes.

Three Byzantines Adversaries are introduced: **Cautious adversaries** are able to lie on every round of a simulation run with a probability P_{lying}. However, the opinion sent during the same round is always the same even though the queries come from different nodes. **Semi-Cautious adversaries** will not lie, however, they may not respond to a query, with a probability $P_{silence}$. Thus delaying the process of convergence and reducing the number of accessible nodes in the network. **Berserk adversaries** behave similar to Cautious, except that they are able to provide different responses to different queries received in the same round. Thus, during the same round, it can send his true opinion, then lie and respond with a wrong opinion.

3 Simulation Results and Conclusion

In this section, we conclude our results with three interesting observations of FPC and CC, and hints to design an efficient IoT dedicated blockchain. Full simulation results are shown in the full version of this article [‡]. We fixed the number of rounds for each simulation run at 30, which is large enough according to our results. P_{lying} and $P_{silence}$ are both set as 50% for Byzantines adversaries. In addition for FPC, we fix the distance of random walk at 4, which means a node can query nodes with a distance of 4, to represent a limited view of each node. The number of queries of each round is 10, which means a node can

‡. HAL: https://hal.archives-ouvertes.fr/hal-03427543

ask maximum of 10 other nodes for opinion. Each results point is the average of 100 simulation runs. Our results show in a similar format where the vertical axis shows the convergence rate, while the horizontal axis represents the probability of the initial opinion distribution, $P0$. We summarize the following observations and hints by showing 4 groups of figures:

(1) Both consensus mechanisms are sensitive to the initial opinion distribution $P0$: FPC will immediately decrease the convergence rate once the initial distribution exceeds a threshold. Conversely, the convergence rate of CC gradually decreases when the initial distribution attempt to 50%. We can clearly see in Figure 1 (a) and (b) and Figure 2 (a) and (b), that there is a different downward trend of the convergence rate in the two consensus. However in Watts-Strogatz topology, the impact of $P0$ is negligible in for both FPC and CC. That could be explained as that in a relatively dense network like Watts-Strogatz, the exchange of information could be easier. Nodes are more likely to query different nodes in each round, which increases the opinion exchange, thus offsetting the effect of uneven initial opinion distribution. Note that the curves of FPC are not symmetrical at $P0 = 0.5$. This small offset is due to the choice of the node's initial threshold τ. That makes nodes more inclined to change their opinion to 0 in the first round.

(2) Both consensus mechanisms are sensitive to the total number of nodes and the average distance among the nodes in the network, which means, the leak of scalability. For example in Figure 1 with the increase of the network size from 49 to 1024, the convergence rate of FPC drops from50% and 100% to 0% around $P0 = 0.5$ in 2D Grid and Torus, respectively. CC has a lower convergence rate compared with FPC in the same condition, it seems less sensitive to the increase of nodes. Although in Watts-Strogatz topology, where a relatively low average distance among nodes is guaranteed, fluctuations caused by network size are minimal for both consensus.

(3) Semi-Cautious has no effect on FPC, while the remaining two, Cautious and Berserk adversaries, have a great impact on FPC. Convergence rate goes down lower than 25% even with less than 10% Cautious adversaries in Figure 3 (a) and (b), and same for Berserk adversaries. That situation is better in Watts-Strogatz topology: The convergence rate goes down below 25% with 20% adversaries. In CC, all adversary methods are effective, but not as serious as in FPC. We tested up to 33% adversaries present in the network shown in Figure 4. Furthermore, if the topology is compact and dense, like in Watts-Strogatz, CC is quite resistant to all three adversaries with small fluctuation shown in Figure 4 (c) for example.

We see that FPC and CC perform poorly in low connectivity topologies. In reality, we cannot ensure the density of network nodes, especially in the IoT situation with varying applications. A potentially feasible solution is to build a high-connectivity logical network topology under the FPC and CC layers, by designing appropriate network layer protocols. Our future work will focus on formal proofs of our conclusions and the design of suitable logical network layer protocols.

References

[CMP19] Angelo Capossele, Sebastian Müller, and Andreas Penzkofer. Robustness and efficiency of leaderless probabilistic consensus protocols within byzantine infrastructures. *arXiv preprint arXiv:1911.08787*, 2019.

[NB08] Satoshi Nakamoto and A Bitcoin. A peer-to-peer electronic cash system. *Bitcoin.–URL: https://bitcoin. org/bitcoin. pdf*, 4, 2008.

[PB21] Serguei Popov and William J Buchanan. Fpc-bi: Fast probabilistic consensus within byzantine infrastructures. *Journal of Parallel and Distributed Computing*, 147:77–86, 2021.

[PMC+20] Serguei Popov, Hans Moog, Darcy Camargo, Angelo Capossele, Vassil Dimitrov, Alon Gal, Andrew Greve, Bartosz Kusmierz, Sebastian Mueller, Andreas Penzkofer, et al. The coordicide. *Accessed Jan*, pages 1–30, 2020.

[Pop18] Serguei Popov. The tangle, 2018. Accessed 30 Juin 2019.

[SLZ16] Yonatan Sompolinsky, Yoad Lewenberg, and Aviv Zohar. SPECTRE: A fast and scalable cryptocurrency protocol. *IACR Cryptol. ePrint Arch.*, 2016:1159, 2016.

[SZ18] Yonatan Sompolinsky and Aviv Zohar. PHANTOM: A scalable blockdag protocol. *IACR Cryptol. ePrint Arch.*, 2018:104, 2018.

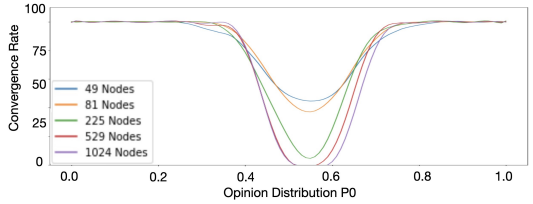

(a) 2D Grid

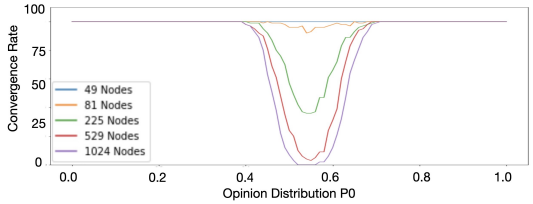

(b) Torus

FIGURE 1: FPC convergence rate according to the initial division probability P_0 for different network sizes, without malicious nodes

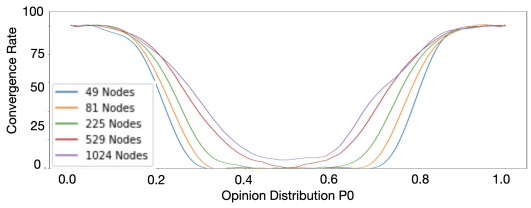

(a) 2D Grid

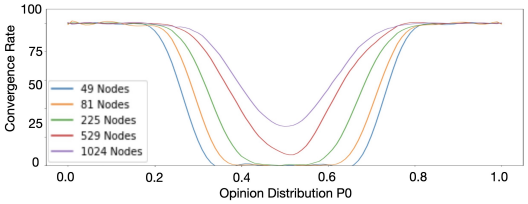

(b) Torus

FIGURE 2: CC convergence rate according to the initial division probability P_0 for different network sizes N, without malicious nodes

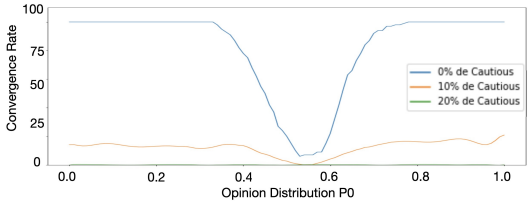

(a) 2D Grid

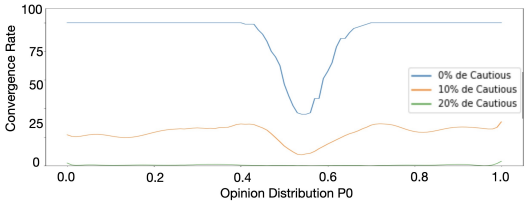

(b) Torus

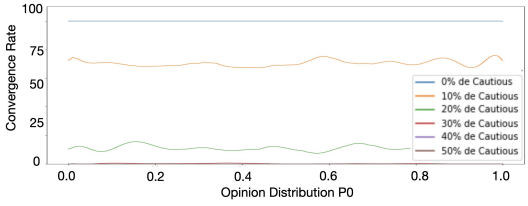

(c) Watts-Strogatz

FIGURE 3: FPC convergence rate according to the initial division probability P_0 for different percentage of Cautious adversaries $P_{malicious}$

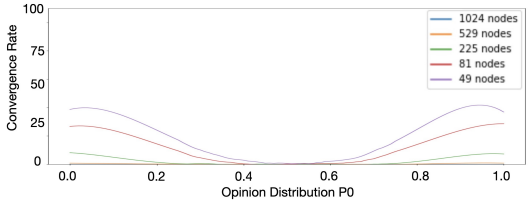

(a) 2D Grid

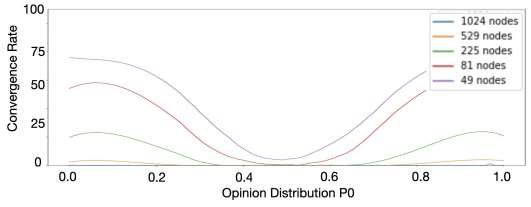

(b) Torus

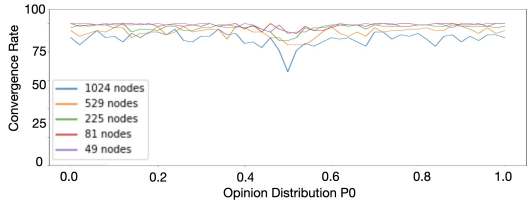

(c) Watts-Strogatz

FIGURE 4: CC convergence rate according to the initial division probability P_0 for different network sizes N, with 33% Cautious Adversaries

[W+14] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.

[WS98] Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440–442, 1998.