

Are Distributed Ledger Technologies Ready for Smart Transportation Systems?

Mirko Zichichi, Stefano Ferretti, Gabriele D'Angelo
Department of Computer Science and Engineering
University of Bologna
Bologna, Italy
{mirko.zichichi2, s.ferretti, g.dangelo}@unibo.it

Abstract—The aim of this paper is to understand if Distributed Ledger Technologies (DLTs) are ready to support complex services, such as those related to smart transportation systems. In smart transportation services, a huge amount of sensed data is generated by a multitude of vehicles. While DLTs provide very interesting features, such as immutability, traceability and verifiability of data, some doubts on the scalability and responsiveness of these technologies appear to be well-founded. We propose an architecture for smart transportation systems that resorts to DLT features. Moreover, we provide experimental results of a real test-bed over IOTA, a promising DLT for IoT. Results clearly show that, while the viability of the proposal cannot be rejected, further work is needed on the responsiveness of DLT infrastructures.

Index Terms—Blockchain, Distributed Ledger Technologies, IOTA, Smart Transportation Systems

I. INTRODUCTION

It is widely recognized that next generation Internet services will massively resort to crowd-sourced and crowd-sensed data, coming from multiple sensors installed on multiple devices. Data aggregation provides the backbone for analyses able to capture some data findings that would not be possible from single sensors. This is true in smart transportation systems as well, where services are built through data sensed by vehicles [1]. Transportation efficiency, travel safety, vehicle security, environment monitoring, are just few examples of types of services that might be offered [2].

While the amount of possible services is countless, a number of issues must be considered, that are basically related to the gathering, storing and level of trust of the data. In fact, in order to share, aggregate and trade data coming from vehicles, some features must be provided by the digital services in use, such as access control, authenticity, verifiability and proof-of-location [3]. This is where a new kind of technology can come to aid. Distributed Ledger Technologies (DLTs) are thought to provide a trusted and decentralized ledger of data. DLTs are a novel keyword, that extends the famous “blockchain” buzzword, to include those technological solutions that do not organize the data ledger as a linked list of blocks. Blockchains gathered momentum when Bitcoin and other crypto-currencies skyrocket. Then, the interest was mainly devoted to the possibility of building decentralized applications based on smart contracts [4], [5]. Currently, DLTs are widely utilized

in scenarios where: i) there are multiple parties that concur in handling some shared data, ii) there is no complete trust among these parties, and often iii) parties compete to the access/ownership of such data. This is a typical scenario of smart transportation services that exploit data sensed from multiple sources (vehicles). Hence, the question now is if DLTs can be efficiently employed in such scenarios.

As a matter of fact, there are DLTs which have been designed with the intent to support the Internet of Things (IoT) [6]–[9]. The main features of these novel technologies are concerned with the attempt to solve some main limitations that are commonly attributed to other blockchains, such as the lack of scalability, sustainability, transaction verification rate (i.e. how fast is the system to add novel data to the ledger). Examples of these novel DLTs for IoT are IOTA [10] and Radix [11]. However, while their design is very interesting, at the time of writing we are aware of just few, and usually simplified, experimental studies on these technologies [3], [12]–[16]; none that demonstrate the viability of these proposed technologies in IoT and smart cities scenarios.

The aim of this work is, first, to propose a novel system architecture that exploits DLTs for the support of smart transportation systems. Second, we present an experimental evaluation on DLTs, based on the use of real data traces to emulate the data generation of a smart city traffic application. We analyze the performance of the IOTA DLT, through tests that measure its degree of scalability and responsiveness in real-time scenarios. Through our tests, we demonstrate how the Masked Authenticated Messaging (MAM) extension module of the IOTA protocol can be used to reliably and securely store and share sensed data in smart mobility applications. However, the latencies for the transactions’ validation result quite high (i.e. 23 sec, on average). Clearly enough, these latencies might not be acceptable in certain real-time application scenarios. Thus, there is still room for improvement.

Moreover, we report on some tests over the Radix alphanet test network. However, due the infancy of the Radix project, we are able to provide only some preliminary outcomes. Still, obtained results seem to be encouraging, but further studies are needed on this DLT.

The remainder of this paper is organized as follows. Section II provides some background on the IOTA DLT. Section III describes the application scenario that has been built

⁰ Working draft paper, submitted for evaluation to an international conference.

to perform the study. Section IV presents all the details of the experimental evaluation, how we conducted the experiments and which metrics have been considered. In Section V, we describe results of the extensive experimental evaluation we conducted over IOTA. Section VI provides a discussion on the obtained results and on possible techniques to improve the DLTs performance. Finally, Section VII provides some concluding remarks.

II. BACKGROUND

A DLT is a software infrastructure maintained by a peer-to-peer network, where the network participants must reach a consensus on the states of transactions submitted to the distributed ledger, to make the transactions valid. Every participant to a DLT contains a local replica of the ledger, which provides data transparency to network participants and ensures high availability of the system. The information recorded to a DLT is append-only, using cryptographic techniques that guarantee that, once a transaction has been added to the ledger, it cannot be modified.

In this work, we mainly focus on IOTA, a specific DLT that is well suited for the IoT and smart transportation systems. This project aims to solve problems about scalability, control centralization, as well as post-quantum security issues, which are present in other blockchain technologies [12]. IOTA is a lightweight, permissionless DLT that enables participants to transfer immutable data and value among each other. From a distributed system point of view, IOTA nodes are organized as a peer-to-peer overlay, where nodes exchange messages containing updates on the decentralized ledger. Nodes that run the entire DLT protocol are commonly referred as full nodes. Being the IOTA architecture still in its infancy, currently a “coordinator node” is present in the system. Its task is to perform a periodic checkpointing of the ledger, with the aim to sustain possible large-scale security attacks. It releases milestone transactions that confirm that all the previous transactions are valid. The purpose of the IOTA foundation is that, after the transient phase, the coordinator will be shut off, hence making IOTA a pure peer-to-peer system [17].

The IOTA decentralized ledger is not structured as a blockchain, but as a Direct Acyclical Graph (DAG) called the Tangle [10]. In the Tangle, graph vertices represent transactions and edges represent approvals. When a new transaction is issued, it must approve two previous transactions and the result is represented by means of directed edges. This process whereby a node selects two random tip transactions from its ledger is termed “tip selection”. In addition to the tip selection, in order to attach a novel transaction to the Tangle, a node must perform a Proof of Work (PoW), i.e. a computation to obtain a piece of data which satisfies certain requirements and which is difficult (costly and time-consuming) to produce but easy for others to verify [10]. The purpose of PoW is to deter denial of service attacks and other service abuses.

The validation approach is thought to address two major pain points that are associated to traditional blockchain-based DLTs, i.e. latency and fees. IOTA has been designed to offer

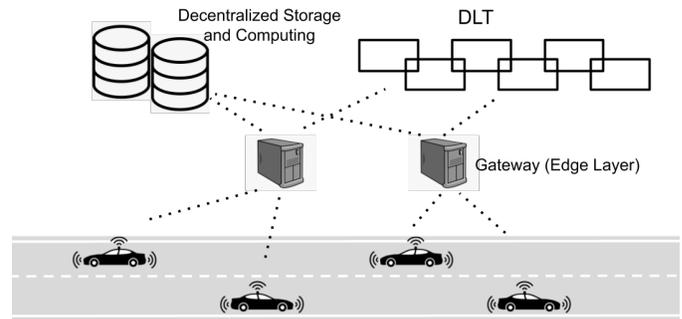


Fig. 1: System model.

fast validation, and no fees are required to add a transaction to the Tangle [13]. This makes IOTA an interesting choice to support smart services built through crowd-sourced data.

An important feature offered by IOTA is the Masked Authenticated Messaging (MAM). MAM is a second layer data communication protocol which adds functionality to emit and access an encrypted data stream over the Tangle. Data streams assume the form of channels, formed by a linked list of transactions in chronological order. Once a channel is created, only the owner can publish encrypted messages. Users that possess the MAM channel encryption key (or set of keys, since each message can be encrypted with a different key) are enabled to decode the message. Messages are pushed on the channel in chronological order, and each message has a link to the next message to be created. Thus, once a user gains access to the MAM channel, he is enabled to see data from that moment on, whilst he cannot look back through the history of the channel before his entrance [13]. In other words, MAM enables users to subscribe and follow a stream of data, generated by some devices. The data access to new data may be revoked simply by using a new encryption key.

III. ON THE USE OF DLTs FOR SMART TRANSPORTATION

We consider a set of vehicles, equipped with sensors that can generate data of some interest (see Figure 1). Such sensed data can be transmitted through a network to an edge computing infrastructure. Thus, each vehicle interacts with a gateway, transmitting sensed data on a periodical basis. The gateway collects and handles the data, based on the specific service being realized. The nature of this specific platform is out of the scope of this work, since it truly depends on the kind of service to be hosted. For instance, it might be organized as a classic cloud system, rather than a distributed file system to store data, e.g. IPFS [18].

In order to provide a level of traceability, verifiability and immutability of the generated data, the data itself, or a related digest (when the data is a large file or when it is a sensitive information), is added to a DLT [3]. We assume the gateway is able to issue messages to a DLT node, thanks to authentication. These messages are converted to transactions added to the ledger. In general, all DLTs provide such kind of functionalities. For instance, in IOTA, Radix and Ethereum (e.g. through the INFURA APIs), there are APIs that allow

entities, external to the DLT, to send a novel transaction. The main point here is that these transactions must be registered in the DLT in a fast way. Second, a good level of scalability must be guaranteed. Third, since a high amount of data is produced, the DLT should offer low fees (or no costs at all). Finally, we need to treat all these transactions as a data-stream, easy to retrieve. By its design, IOTA is recognized as a responsive, scalable, feeless DLT, with MAM channels as the tool to treat data as streams. For this reason, in the evaluation we will focus on IOTA.

IV. EXPERIMENTAL EVALUATION

In this work, we are interested in evaluating the goodness of the adoption of IOTA as the immutable registry for smart transportation systems. Thus, we focused on the transmission of sensed data to IOTA, measuring latencies needed to issue, insert and validate transactions, and also the level of reliability of the full nodes.

A. The Trace-driven Vehicles Simulation

We conducted a trace-driven experimental evaluation. Traces were generated using the RioBuses dataset, a real dataset of mobility traces of buses in Rio de Janeiro (Brasil) [19]. Based on these traces, we simulated a number of buses that, during their path, generate sensed data. (The type and purpose of such data is out of the scope of this evaluation, since we are mainly interested in the behaviour of the DLT; it suffices to assume that they represent typical, small sized sensed data, such as a temperatures, air pollution values, etc.) We assume that the time spent to fetch such data is negligible, with the respect to the time to publish it to the DLT.

These messages were utilized to generate real transactions transmitted to the DLT. Each message was sent to a given DLT node. How this node was selected is discussed in the next subsection. Figure 2 shows the paths of 10 buses, as an example, that were considered during our tests. We varied the number of buses in the range: 60, 120, 240. For each bus, we utilized one hour of trace data. Based on the paths, each bus was set to generate approximately 45 message/hour. Thus, we made one hour long tests, where each bus generated, on average, a message to be issued to the DLT every 80 sec, which is a reasonable time interval to sense data in an urban scenario. For each test configuration, we replicated the experiment 12 times.

For each transaction, we recorded the outcome of the request, i.e. successful or unsuccessful, due to some DLT nodes internal error, as well as the latency between the transmission of the transaction and the confirmation of its insertion in the ledger.

B. IOTA Setup

Each bus was emulated by a single process (issuing messages based on the data trace). Thus, the first task was to find, for each bus, a full node of the IOTA DLT to interact with. In IOTA, network full nodes do not usually allow to get their neighbors in the P2P overlay, through API. This

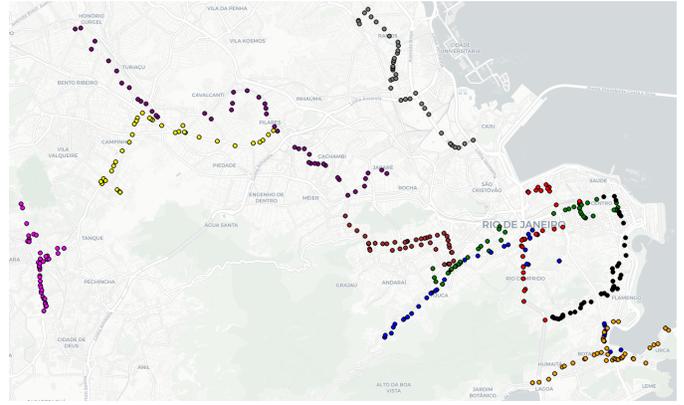


Fig. 2: The 1 hour long path of 10 buses in Rio de Janeiro (Brazil).

hinders the possibility to perform a in-depth graph search on the overlay, in order to retrieve an up-to-date list of active nodes to interact with. Thus, in our tests we were enabled to rely only on services that maintain a public list of active nodes [20]. With this in view, the scheme we designed to select the IOTA nodes to contact, is as follows. Given the list of public nodes, a filter is applied to keep only nodes that are fully synchronized, i.e. the node has solidified all the milestones up to the latest one released by the coordinator, and that allows remote PoW. During testing these nodes were ~ 60 . Then, we designed three heuristics for the selection of a full node to pair to each bus from the public pool:

- 1) **Fixed Random:** Each bus is assigned to a random IOTA full node from the pool, during the setup phase; then, every transaction generated by that bus is handled by this node, for the whole duration of the test.
- 2) **Dynamic Random:** A random node from the pool is selected every time a message has to be published by a bus.
- 3) **Adaptive RTT:** For each bus, its associated node actively changes every time a message has to be published, while the previous one is still pending. Based on results of past interactions, the known IOTA nodes are ranked through the experienced Round Trip Time (RTT) [21]. Then, a new node is chosen by selecting the best known node or, if every known node is in the process of publishing a message, a new node is picked randomly from the pool.

We used a MAM channel associated to each single bus. Every message to be published in the MAM channel requires three transactions to be issued, i.e. one containing the data and two other messages for the signature. The advantage of this approach is that through each MAM channel it is possible to easily retrieve the bus's data stream and that only the channel owner can publish on it. An example of a (private) MAM channel, specifically created for a bus during our tests, can be found by querying the IOTA DLT with the root: JEIJZEVPUHHKEKKDSFFEYYTVSFRXOU YWFH9LZIKKKQEDO9L9MK9LIVOZUIPCML9RCHNDR QYPNGNOUOGO. The entire dataset and the scripts used

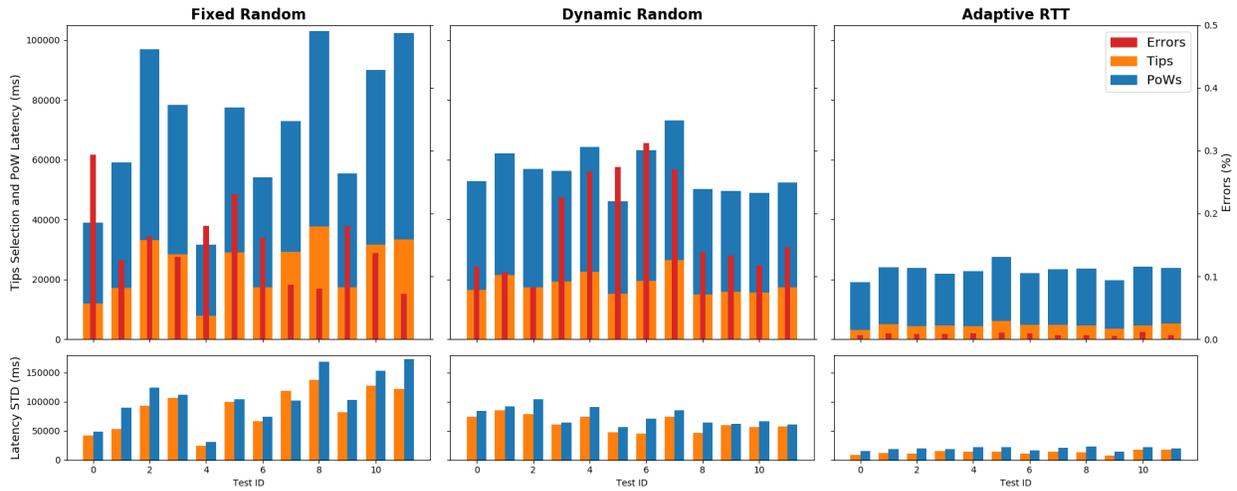


Fig. 3: 60 bus tests: average latencies, standard deviation and errors for the three different schemes

are stored in a github repository [22]. For each transaction, we measured the time required to perform the tip selection as well as the PoW. The tip selection depth parameter, i.e. the number of milestones to go back to start the random walk to select tips, was set to 3, whilst the minimum weight magnitude, i.e. the number of trailing zeros of a transaction hash, was 14 (minimum standard value for the IOTA mainnet).

V. RESULTS

Figure 3 shows results obtained for different test repetitions, when the number of emulated buses was set equal to 60. In particular, we show the results for each scheme we employed for the selection of the nodes. In the upper part, the histograms report the average latencies measured during a single test. The orange (lighter) part of the histogram shows the average latency to perform the tip selection, while the blue (darker) part shows the average latency associated to the PoW. The red (central and smaller) bars refer to the percentage of errors (the related y-axis is shown on the right of the figure), i.e. amount of transactions that failed to be added to the Tangle, due to full nodes' errors. On the lower part of the figure, we show the average standard deviations related the specific tests, both for the tip selection and PoW. From the figure, it is possible to appreciate how in general a random selection of the full node to issue a transaction does not lead to good results. The amount of errors is quite high, as well as the measured latencies. Thus, these tests seem to conclude that, at the time of writing, the IOTA DLT is not fully structured to support smart services for transportation systems. On the other hand, the good news is that if we carefully select the full node to issue a transaction, the performances definitely improve. In fact, our third scheme "Adaptive RTT" has a low amount of errors, on average around 0.8%. Measured latencies are lower than other approaches. Still, the average latency amounts to 23 seconds, which is far from a real-time update of the DLT. The level of acceptability of latency values truly depends on the application scenario.

These first results suggest that some scalability tests might give further insights on the viability of the use of IOTA as the

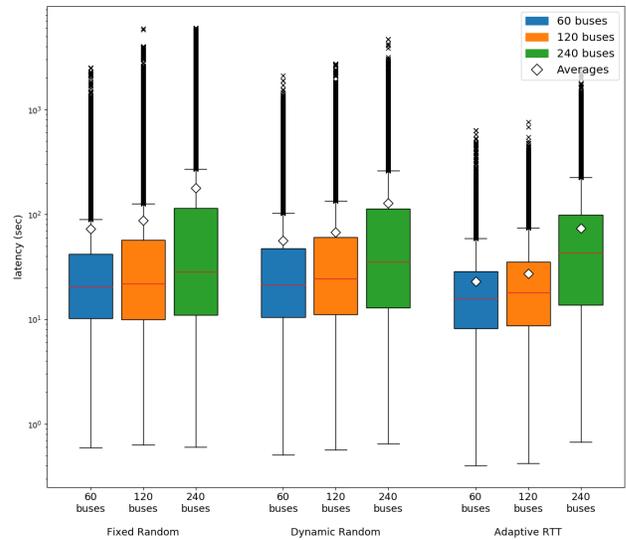


Fig. 4: Boxplots for tests with 60, 120, 240 buses

DLT to support smart transportation system. For this reason, we made some tests with an increasing number of buses.

Figure 4 shows average results obtained using our three considered schemes, when varying the number of buses. Results are reported as box plots. Thus, each box plot corresponds to the average results for a scheme in a given scenario. This allows to assess the scalability of each scheme, by looking at the results for an increasing amount of buses. At the same time, it is possible to compare the three schemes by looking at their performance for each scenario. In the box plot, the diamond represents the mean value of the overall latency (i.e. the time from the transaction transmission to the node to the acknowledgement that it has been added in the Tangle). The rectangle identifies the Inter-Quartile Range (IQR), i.e. values from the 25th to the 75th percentile. The middle box thus represents the middle 50% of values. Hence, the lower part of the box (let denote it Q_1) is the first quartile (25th percentile),

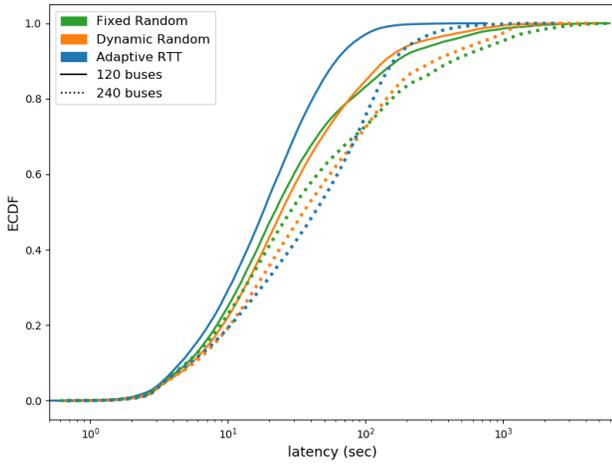


Fig. 5: Empirical Cumulative Distribution Function for tests with 120 and 240 buses

the highest (denote it Q3) is the third quartile (75th percentile). The red line inside the box is the median value. The lower and upper values identified by the vertical line are the whiskers. In box plot, the whiskers are defined as 1.5 times the IQR. Thus, the lower whisker is $Q1 - 1.5 \cdot IQR$, while the upper whisker is $Q3 + 1.5 \cdot IQR$; they represent a common way to describe the dispersion of the data. Finally, the “×” symbols outside the whiskers are the outliers. To better show the obtained results, the y-axis is reported in a log scale.

Results confirm that “Adaptive RTT” provides better results. Average latencies are definitely lower than other schemes. It is worth noticing that, being the y-axis in log scale, the difference on the performance is relevant. In particular, the first two schemes have outliers well over 10^3 sec. In all cases, average latencies increase with the number of buses. This suggests that the number of full nodes devoted to the transaction management should increase proportionally to the number of buses. Indeed, if we assume that 60 full nodes are used, in the 240 buses tests we have 4 buses per node, that receive ~ 0.5 msg/sec, on average. This means that every 2 sec a IOTA node receives a request for a novel transaction, that requires 23 sec (on average, using “Adaptive RTT”). Results confirm an important difference between the 240 buses scenario (i.e. ~ 0.5 msg/sec) and the 120 buses scenario (i.e. ~ 0.25 msg/sec, on average). This means that further improvements are needed to solve scalability issues.

To better emphasize the outcomes, Table I reports some summarized statistics (shown in the box plots) and the error rates. Actually, the main difference on the performance of the approaches is on the amount of errors. While the average error for “Adaptive RTT” is $\sim 1\%$, for the other two schemes we have errors well above 15%. These error rates are clearly unacceptable, meaning that these approaches are unusable.

Finally, Figure 5 shows the empirical cumulative distribution function obtained for the compared schemes in the 120 and 240 bus scenarios. In this case, for the sake of a better visualization, the x-axis is shown in a log-scale. These

charts further confirm the better performance obtained by the “Adaptive RTT” scheme.

VI. DISCUSSION

A. On the Performance of IOTA

Obtained results require some discussion. In fact, on one side, it is shown that through a proper selection of full nodes, it is possible to obtain a reliable ledger update (low errors), thus making viable the use of IOTA to support smart transportation systems. On the other side, however, the measured latencies are relevant. In our tests, we employed public available IOTA full nodes to add transactions. Thus, we refer to such nodes to perform the tip selection and the PoW. The rationale behind this choice was based on the assumption that sensors placed on the buses do not have computation capabilities to behave as full nodes [14].

IOTA offers a view of the status of public full nodes [20]. Thus, it is possible to monitor their computation capabilities and the workload. During our experimental evaluation, all these nodes had usually a low computational load. Nonetheless, results confirm that the selection of the node is quite relevant. As a further confirmation of this claim, in our preliminary tests we tried to exploit a heuristics, alternative to those presented in the previous section. The idea was to select the best N full nodes, in terms of available resources, and use them to issue the transactions. With $N=10$, we measured very poor performances. This was due to the fact that, while apparently well provisioned, certain full nodes were not able to sustain the workload coming from our application (~ 0.77 message/sec). Trying to increase the scalability of the system and better balance the nodes workload, we increased the value of N. However, with $N=20$, we noticed a high variability on the performance of the employed full nodes, with substantial difference between the highly ranked and the lower ranked public full node. For this reason, we found that it was simpler (with similar performances) employing the “Fixed Random” approach.

An alternative approach might be to employ an edge computing system model, where the execution of the PoW is executed locally by the gateway (see Figure 1). (The tip selection must be always accomplished at a full node, that maintains a complete copy of the Tangle.) The rationale would be to relieve the IOTA node from the computational burden of the PoW. However, this would force to equip the gateway with sufficient computational capabilities to perform the PoW for all the transactions generated by the buses it handles.

Finally, it would be possible to ask the gateway to act as a full node for the DLT. This would actually resemble the testbed we considered in this work (due also to the fact that the exploited IOTA public nodes had a low workload overhead, concurrent to our tests). In this case, the difference would be that it would be possible to have a direct control of the full node. Its hardware characteristics might be properly set to tolerate a certain predicted workload, and this node might be reserved to handle transactions from the specific

TABLE I: Results on IOTA, with 60, 120, 240 buses.

# buses	Heuristic	Avg Latency	Conf. Int. (95%)	Errors
60	Fixed Random	72.68 sec	[70.43, 74.94] sec	15.37%
	Dynamic Random	56.0 sec	[54.51, 57.5] sec	18.26%
	Adaptive RTT	22.99 sec	[22.69, 23.29] sec	0.81%
120	Fixed Random	87.75 sec	[85.38, 90.12] sec	29.49%
	Dynamic Random	67.6 sec	[66.29, 68.9] sec	18.99%
	Adaptive RTT	27.35 sec	[27.11, 27.58] sec	1.1%
240	Fixed Random	177.62 sec	[174.25, 181.0] sec	42.81%
	Dynamic Random	128.2 sec	[126.28, 130.12] sec	44.85%
	Adaptive RTT	73.26 sec	[72.68, 73.85] sec	7.55%

smart transportation system application, only. This scenario represents an interesting further work.

B. On the Use of Alternative DLTs

We conducted preliminary tests with other blockchains, such as the well known Ethereum. However, Ethereum was not designed to register a huge amount of transactions containing (typically small sized) sensed data. The costs to issue a transaction in a block are usually quite high. Moreover, the confirmation times and scalability limitations are two other main factors that discourage the adoption of this technology. In fact, because of a hard-coded limit on computation per block, the Ethereum blockchain currently supports roughly 15 transactions per second. All this makes Ethereum an impractical technology to be used in our scenario.

Clearly enough, it might be interesting to evaluate the performance of DLTs, thought to support smart transportation scenarios, that implement novel techniques to improve scalability. An example is sharding, i.e. breaking the ledger into smaller, more manageable chunks, and distributing those chunks across multiple nodes, in order to spread the load and maintain a high throughput.

A novel DLT that implements sharding techniques is Radix [11]. At the time of writing, the Radix technology is still at its infancy and a main net does not exist, yet. Nevertheless, we exploited the alphanet test network to issue transactions on the ledger. This gave us some preliminary results that we report in Table II. In the table, we show the average latency, confidence interval and error percentage to add transactions on this Radix alphanet. Results are averaged over an amount of 12 test repetitions. In this case, we obtained very low latencies (below 1 sec), with a non negligible (but low) error rate. It is worth to point out that these results can be indicative of the functioning of Radix. However, we claim that it is difficult to compare these results with those obtained for IOTA. In fact, in IOTA we exploited the main net, while in Radix we had to employ a preliminary testnet, with few nodes involved to the ledger management (~ 6 nodes) and basically no additional workload, apart from our tests. As a matter of fact, comparable results can be obtained if tests are executed on the IOTA test net, where the PoW is faster (we obtained average latencies around ~ 2 sec).

While there are novel interesting proposals to improve scalability, such as sharding or the Ethereum plasma [23], a main problem refers to the high fees that should be associated

TABLE II: Preliminary results on Radix.

# buses	Avg Latency	Conf. Int. (95%)	Errors
120	777.17 msec	[774.68, 779.65] msec	2.73%

to every transaction. IOTA is designed to be feeless, in order to let billions of devices and sensors to interact with the Tangle without costs. Conversely, Radix and Ethereum use fees. The possible costs may be acceptable only when the transaction fees are negligible with respect to the value of the data. However, we claim that, in general, the need for fees hinders the use of a DLT in smart transportation scenarios.

VII. CONCLUSIONS

In this paper, we proposed an architectural solution resorting to DLTs to support smart transportation systems. The benefits on the use of the distributed ledgers is that they would allow to safely and securely store sensed data, offering authenticity, verifiability and immutability features. Moreover, the use of DLTs can be employed to provide proof-of-location [3].

We analyzed the main characteristics of current DLTs and focused on the DLT that, among the others, promises to be the best solution for smart transportation scenarios, i.e. IOTA. We thus made an extensive experimental evaluation, whose results have been summarized and analyzed. The conclusion is that, probably, work must be done, in order to provide effective distributed ledgers for smart transportation systems. In fact, it is important to be able to select proper nodes to interact with in order to have acceptable error rates. Moreover, measured latencies resulted higher than 20 sec, which is quite high if we think at real-time applications, reasonable for less time demanding services. In any case, this might be a transient problem, that could be solved by improving the IOTA peer-to-peer infrastructure.

Furthermore, in our tests all the work (i.e. tip selection and PoW) was performed by the full nodes. The rationale was to relieve sensors and devices from this task [14]. An alternative solution might be to delegate the PoW to some other entity, such as a gateway in between the vehicle sensors and the full node. Moving the PoW from the full nodes elsewhere might strongly improve the performances of the DLT nodes. The study of this possible improvement is ongoing.

A technique to improve the responsiveness may be based on the use of sharding. Indeed, we studied a novel technology, i.e. Radix, that is specifically based on sharding, obtaining

interesting results. However, an open security question arises, i.e. if we decrease the amount of nodes that validate transactions (as the sharding does), then does the risk of a hack increase?

ACKNOWLEDGMENT

This project has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie ITN EJD grant agreement No 814177.

REFERENCES

- [1] R. W. van der Heijden, F. Engelmann, D. Mödinger, F. Schönig, and F. Kargl, "Blockchain: Scalability for resource-constrained accountable vehicle-to-x communication," in *Proc. of the 1st Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers*, 2017.
- [2] H. Mousannif, I. Khalil, and H. Al Moatassime, "Cooperation as a service in vanets." *J. UCS*, vol. 17, no. 8, pp. 1202–1218, 2011.
- [3] M. Zichichi, S. Ferretti, and G. D'Angelo, "A distributed ledger based infrastructure for smart transportation system and social good," in *IEEE Consumer Communications and Networking Conference (CCNC)*, 2020.
- [4] G. D'Angelo, S. Ferretti, and M. Marzolla, "A blockchain-based flight data recorder for cloud accountability," in *Proc. of the 1st Workshop on Cryptocurrencies and Blockchains for Distributed Systems (CryBlock)*, 2018.
- [5] M. Zichichi, M. Contu, S. Ferretti, and G. D'Angelo, "Likestarter: a smart-contract based social DAO for crowdfunding," in *Proc. of the 2nd Workshop on Cryptocurrencies and Blockchains for Distributed Systems (CryBlock)*, 2019.
- [6] R. Di Pietro, X. Salleras, M. Signorini, and E. Waisbard, "A blockchain-based trust system for the Internet of Things," in *Proc. of the 23th ACM on Symposium on Access Control Models and Technologies (SACMAT)*, 2018.
- [7] G. D'Angelo, S. Ferretti, and V. Ghini, "Simulation of the Internet of Things," in *In proc. of International Conference on High Performance Computing Simulation (HPCS)*, 2016.
- [8] —, "Multi-level simulation of internet of things on smart territories," *Simulation Modelling Practice and Theory (SIMPAT)*, vol. 73, 2017.
- [9] S. Ferretti and G. D'Angelo, "On the Ethereum blockchain structure: A complex networks theory perspective," *Concurrency and Computation: Practice and Experience*, Wiley, in press.
- [10] S. Popov, "The Tangle," 2016. [Online]. Available: https://iota.org/IOTA_Whitepaper.pdf
- [11] "Radix knowledge base," 2019. [Online]. Available: <https://docs.radixdlt.com/kb/>
- [12] P. C. Bartolomeu, E. Vieira, and J. Ferreira, "IOTA feasibility and perspectives for enabling vehicular applications," <https://arxiv.org/abs/1808.10069>, 2018.
- [13] J. Brogan, I. Baskaran, and N. Ramachandran, "Authenticating health activity data using distributed ledger technologies," *Computational and Structural Biotechnology Journal*, vol. 16, 2018.
- [14] A. Elsts, E. Mitskas, and G. Oikonomou, "Distributed ledger technology and the Internet of Things: A feasibility study," in *Proc. of the 1st Workshop on Blockchain-enabled Networked Sensor Systems (BlockSys)*, 2018.
- [15] S. K. Pinjala and K. M. Sivalingam, "Dcaci: A decentralized lightweight capability based access control framework using iota for internet of things," in *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*, April 2019, pp. 13–18.
- [16] S. Suhail, C. S. Hong, and A. Khan, "Orchestrating product provenance story: When IOTA ECOSYSTEM meets the electronics supply chain space," *CoRR*, vol. abs/1902.04314, 2019. [Online]. Available: <http://arxiv.org/abs/1902.04314>
- [17] IOTA Foundation, "The Coordicide," 2019. [Online]. Available: https://files.iota.org/papers/Coordicide_WP.pdf
- [18] J. Benet, "IPFS - content addressed, versioned, P2P file system," *ArXiv*, vol. abs/1407.3561, 2014.
- [19] D. Dias and L. H. M. K. Costa, "CRAWDAD dataset coppe-ufrrj/riobuses (v. 2018-03-19)," Downloaded from <https://crawdad.org/coppe-ufrrj/RioBuses/20180319>, 2018.
- [20] "IOTA nodes," October 2019. [Online]. Available: <https://iota-nodes.net/>
- [21] V. Jacobson, "Congestion avoidance and control," in *ACM SIGCOMM Computer Communication Review*, vol. 18, no. 4. ACM, 1988.
- [22] "Dataset and scripts github repository," October 2019. [Online]. Available: <https://github.com/miker83z/testingIOTA>
- [23] J. Poon and V. Buterin, "Plasma: Scalable autonomous smart contracts," *White paper*, pp. 1–47, 2017.