



Distributed Privacy Preserving Platform for Ridesharing Services

Yevhenii Semenko, Damien Saucez

► **To cite this version:**

Yevhenii Semenko, Damien Saucez. Distributed Privacy Preserving Platform for Ridesharing Services. [Research Report] Inria - Sophia Antipolis. 2019. <hal-01968399>

HAL Id: hal-01968399

<https://hal.inria.fr/hal-01968399>

Submitted on 2 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed Privacy Preserving Platform for Ridesharing Services

Yevhenii Semenko
Amadeus, France
semenkoe@gmail.com

Damien Saucez
Université Côte d'Azur, Inria, France
damien.saucez@inria.fr

Abstract—The sharing economy fundamentally changed business and social interactions. Interestingly, while in essence this form of collaborative economy allows people to directly interact with each other, it is also at the source of the advent of eminently centralized platforms and marketplaces, such as Uber and Airbnb. One may be concerned with the risk of giving the control of a market to a handful of actors that may unilaterally fix their own rules and threaten privacy. In this paper, we propose a decentralized ridesharing architecture which gives the opportunity to shift from centralized platforms to decentralized ones. Digital communications in our proposition are specifically designed to preserve data privacy and avoid any form of centralization. We integrate a blockchain in our proposition to guarantee the essential roles of a marketplace, but in a decentralized way. Our numerical evaluation quantifies the advantages and limits of decentralization.

Index Terms—data privacy, decentralized, sharing economy.

I. INTRODUCTION

Mass digitalization of our daily life and deep involvement of user data boosted the sharing economy by allowing the advent of massively used platforms, such as Airbnb and Uber, that make our day-to-day tasks much easier. However, such platforms raise concerns about the concentration of power and information in a handful of mercantile companies with the risk of inappropriate personal data usage [1], personal information leakage [2], or market control [3].

This paper proposes to tackle these issues for the case of ridesharing services by proposing a multilayered distributed architecture that removes the need of a centralized platform. The shift to a fully decentralized paradigm is not straightforward. First, with a centralized approach, personal data are concentrated and, unless security breaches or inadequate policies, the data can only be manipulated by well identified entities (i.e., the platform operator). In a fully decentralized entity data are scattered among the constituents of the decentralized entity platform with inherently more risks of information leakage. Second, decisions to authorize or not a given client or producer to be part of the platform or to define a price are more complex in a decentralized platform as the system as a whole must answer these questions. Finally, responsiveness of the system can become a challenge as it is impossible to tune the resources as well as with centralized platforms.

Ridesharing services are particularly challenging as they are by nature highly dynamic (clients come and leave frequently), they involve users willing to minimize data communication as

they are roaming, they require to exchange sensitive personal information, and they are subject to rules that can change from region to region while users can move between these regions.

Privacy preserving solutions for ridesharing have been proposed ([4]–[10]) but we are the first to take a holistic approach where we consider simultaneously privacy, accountability, business, communications, and scalability.

In Sec. II we define a decentralize privacy preserving platform for ridesharing where peers forms an overlay communication network and rely on a blockchain to organize their marketplace. In Sec. III we evaluate its scalability based on analytical bounds and on a trace-driven numerical evaluation. Finally, in Sec. IV we conclude.

II. DECENTRALIZED PLATFORM

Ridesharing platforms such as Uber or Lyft rely on centralized infrastructures involving third-parties, which raises privacy concerns. In this paper, we remove this dependency to avoid any lock-in with a third-party and to preserve privacy. To that aim, we conceive the platform as a fully decentralized peer-to-peer network [11] where each node is equally important and interchangeable and rely on a multilayer architecture.

A. Multilayered architecture

First, we build an overlay network including all the business actors to constitute the *network layer*. As ridesharing services can span over multiple regions with different constraints and can have several millions of users, the underlying peer-to-peer networking technology must provide a way to segregate traffic in arbitrary regions and it must be able to scale. For these reasons we must use a protocol where node *IDentifiers* (ID) can be structured and scoped, such as in Chord [12] or in Kademia [13]. It is therefore straightforward to limit the scope of communications to particular regions. For example, if communications must remain within peers of a country, it is enough to partition the ID space with one partition per country and impose that every peer of a country uses an ID within the partition of the country. Fig. 1 shows an example of how to scope communications per region with Chord.

Second, as the platform does not involve a data storage third party, all actors involved in the platform must take part in the storage effort. Keeping in mind that data privacy is essential, the best approach is to build a Distributed Hash Table (DHT) to implement the *Data storage layer*. Even though the

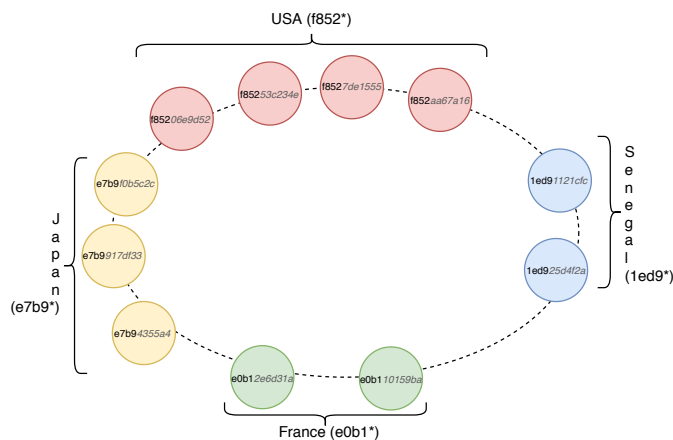


Fig. 1. Example of device grouping with Chord.

data storage layer by itself is not sufficient to guarantee data privacy, when it is combined with the other layers it enables data privacy as detailed in Sec. II-B.

User Interface (UI) and user experience can't be altered by the decentralization even though the *User application layer* where the UI is implemented fundamentally changes. In centralized platforms, user applications interact with the rest of the platform via REST APIs. In a decentralized solution, the application must include the other layers of the platform and implement, at least partially, the logic of the platform.

Finally, the *Management layer* that is simple by nature in a centralized approach becomes particularly complex to implement in a fully distributed system with no trusted parties. Particularly complex issue to solve in such an environment are how to control business processes, manage data access, and provide accountability to solve potential disputable situations. These issues arise as there is no central to solve debatable issues, to control or to set the required system parameters.

To implement the *Management layer* we use a blockchain [14] which plays the role of an intermediary for controlling data access, guarantee that parties act in conformance with the rules of the platform, and log events. The business logic of the platform is then implemented as a smart contract and data privacy is ensured.

By itself a blockchain does not provide data privacy protection, actually it would tend to reduce privacy as it spreads the data to open locations. However, thanks to smart contracts data owners can specify who and how their data can be accessed. In Sec. II-B we detail how to guarantee data privacy with our multilayered distributed platform.

B. Data Privacy

The first condition to preserve privacy is to minimize the amount of data that is spread in the system. That is, for each data, the designer of the service must determine whether or not the data must be shared, who can access it, and under which conditions. Consumers of the data are either identified at the time of the publication or they are not.

The first case is the simplest, the data is encrypted with the public key of the target such that only the target can read the data. In this case, the data can be stored anywhere in the DHT, multiple targets means multiple copies of the data. As our architecture does not rely on trusted third parties, there is no trusted *Public Key Infrastructure* (PKI) on which we can rely to obtain the public key. However, if like in Bitcoin the ID is made such that it is unambiguously linked with the public key of the party then it is possible to safely retrieve the public key for that node [15]: 1) the producer of data determines the ID of its target, 2) it requests the system (or the target directly) to provide its public key, and 3) it verifies that the ID and the public key are linked (e.g., the ID contains a hash of the public key).

In the second case where the target of the data is not known while publishing the data, it is not possible to use encryption directly. Instead, the producer uses the Shamir's Threshold Scheme [16]. The producer generates as many parts of its secret as needed and determines a DHT key for each part such that they can be published in the DHT according to its requirements. Here the choice of the key for each part to be published in the DHT is important as it determines the nodes that will store the data. In parallel to publishing the data in the DHT, the producer also publishes a smart contract in the blockchain that specifies the policy of accessing the data. As result, when a party wants to consume a data, it sends a request to the overlay and if the smart contract is fulfilled then the storing nodes provide their parts to the requester that will be able to re-construct the data. The threshold value and where to store the parts (and their redundancy level) depends on the service's threat model.

C. Protocol

Road transportation clients and drivers are by nature on the move and thus have poor, expensive, and unreliable connectivity. For that reason they do not directly take place in the peer-to-peer network. Instead the peer-to-peer network is run by arbitrary machines willing to share their resources for the community. The nodes taking part in the peer-to-peer network then act as service nodes to access the network for the clients and drivers by the means of an HTTPS REST API.

Client and driver applications select a node from the peer-to-peer network as service node according to their own preferences (e.g., randomly or topologically close to them, or a trusted service node).

Any ride can be decomposed in three phases. First, the client hails for a ride. During that phase, the client and a driver mutually agree on making the ride together. Second, the ride starts and the client is dropped-off at its desired location. Finally, the client and driver finish the transaction by paying and they comment the service if they want.

Below we detail the protocol that we designed to implement these three phases in our distributed platform. Fig. 2 and Fig. 3 depict the exchanges and the messages, respectively.

1) *Requesting phase*: When a client wants to ride from point *A* to point *B*, it sends a *Request* message to its

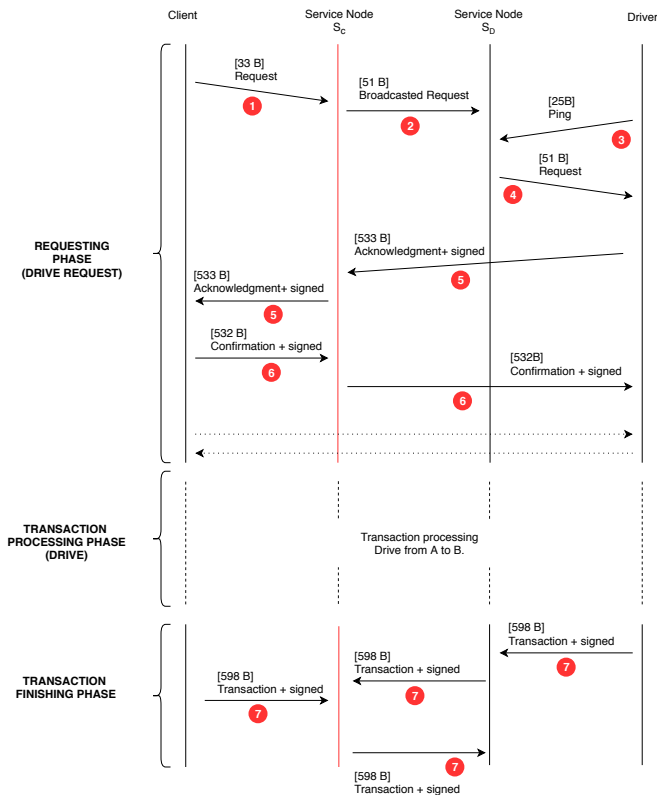


Fig. 2. Ridesharing service protocol workflow.

service node. The message contains a nonce, a timestamp and information on about the desired type of car. The nonce is made to avoid leaking the identity of the client in the network and is the HMAC of a random number generated by the client. The information about the ride are the anonymized origin and destination location of the ride, the exact distance for the ride as computed with a traffic route planner API, the price per kilometer that the client is ready to pay for the ride, and the categories of cars the client is willing to use. The client does not provide its actual origin and destination coordinates to avoid leaking sensitive information to the system. Instead it provides approximate locations. The price per kilometer is computed according to an auction mechanism. The location anonymization and price computation mechanisms are out of the scope of this paper.

When a service node receives a `Request`, it broadcasts it in the peer-to-peer network and tags the message with its own IP address and port number. In the meanwhile, the service node locally stores the request in its *pending ride requests queue*. When a node in the peer-to-peer network receives a broadcasted `Request` message, it locally stores it in its pending ride requests queue.

Every driver periodically polls its service node with a `Ping` message to know the pending ride requests.¹ The `Ping` message contains the anonymized location of the driver, its type of car, and a timestamp. If the driver is in a location

close to a request in the pending request queue of the service node, the service node replies to the driver by forwarding it the oldest compatible “close enough” `Request` message in its pending ride requests queue. If the driver doesn’t want to accept this ride, it silently ignores it. Otherwise, it sends an `Acknowledgement` message directly to the service node of the client, without going through the peer-to-peer network as it knows the IP and port that can be used for direct communications with the service node. It is worth to notice that at this stage the driver leaked potential private information as it publicly reveals its ID, current approximate location, and will to achieve a ride that is uniquely identified by a nonce. Unfortunately, the only solution to avoid this would be to have a mechanism to encrypt the acknowledgment such that only the client can decrypt it. But as we are in an untrusted environment, that would imply to leak the client ID as drivers would have to verify that the request was not forged by a service node. In the meanwhile, only the ID and an approximate location is leaked, which is not a severe threat for drivers as they are anyway clearly publicly identified in real life for legal reasons (e.g., with a specific sticker on their cab), while it could be a sever threat for clients. This is the reason while we designed our protocol in that way.

The `Acknowledgement` message contains the nonce taken from the `Request`, the driver’s ID and public key, it’s car type, its rank, a price per kilometer, a random number, and ancillary data that depend on the local regulations and habits (e.g., driver name, car make and model. . .). In addition, it provides the IP address and port number to use to have a direct communication with it.² The message is signed by the driver. When the client service node receives this message, it forwards the acknowledgment to the client and removes the request from its pending requests queue such that any further `Acknowledgement` message will be silently dropped.

When a client receives an `Acknowledgement` message, it decides whether or not it agrees to make the ride with that particular driver (e.g., based on the driver’s rank). If it doesn’t accept the driver proposition, it silently ignores the proposition in order to avoid leaking personal information. On the contrary, it is accepts the proposition, it sends a `Confirmation` message directly to the driver’s device (for which it knows the IP address and port number to use). The message contains the nonce and the random number used by the client to generate the nonce, the client’s ID, public key, and rank, the exact origin and destination points of the ride, the random number sent by the driver, and ancillary data that depend on the local regulations and habits (e.g., client’s surname). The message is signed by the client and is fully encrypted with the public key of the driver. With the nonce and its associated random number, the client’s public key, and the signature, the driver can verify that the client acknowledgment is generated by the owner of the ID that initiated the `Request` message and with the echo of the random number that it provided in the

²We assume that drivers can set in place hole-punching mechanisms, e.g., via their home box, to allow direct connection to their device while working.

¹Time between `Pings` must be randomized to avoid race conditions.

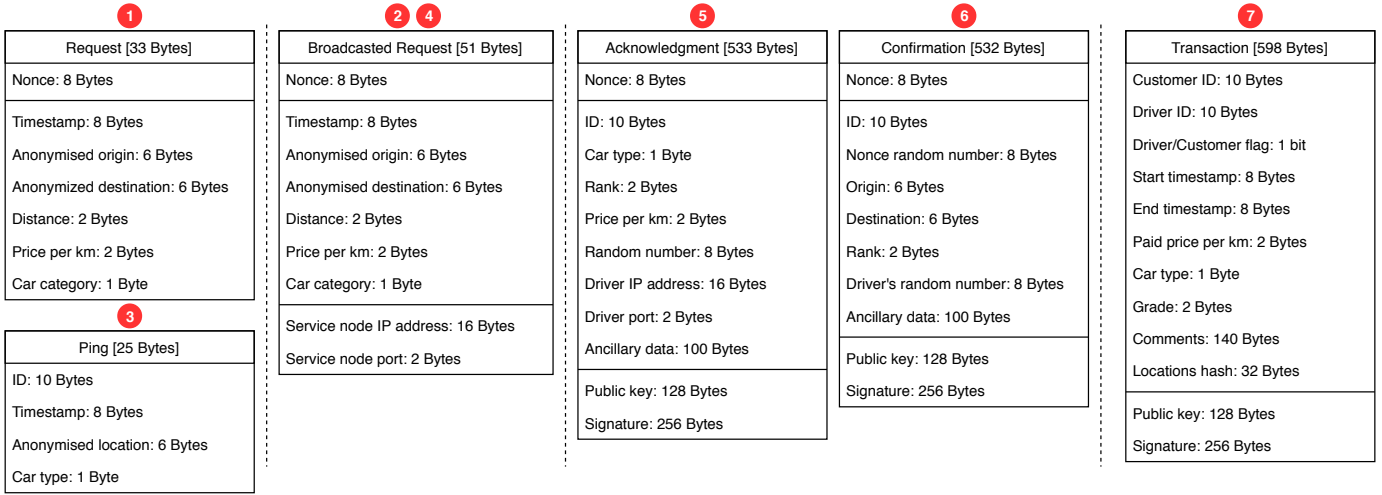


Fig. 3. Ridesharing service messages.

Acknowledgment message, it prevents replay attacks.

At this stage, both the client and the driver have fully disclosed their personal information, but this information is known only by themselves. The client and driver start exchanging their actual GPS coordinates until the client is picked up. To minimize the number of message exchanged they rely on a traffic route planner API and only exchange information if they deviate from the plan. These communications are encrypted and directly sent between the client and the driver.

It is worth to mention that even though the client's service node doesn't know the ID or exact coordinates of the client, it knows its IP address (as it has direct communication with it) and knows the ID of the driver. There are no practical solution to avoid such leakage of information but the client can tackle this issue by using a trusted service node (e.g., by deploying one of its own or using a service node that respects privacy).

2) *Driving phase*: The drive starts when the client is picked up by the driver. At this time, devices can collect GPS exchange format (GPX) data independently from each other. During the drive there are no messages exchanged between the devices or with the rest of the network. Due to this, a third party will never be able to determine the time of making the drive or intercept any data concerning the transaction participants GPS positions. This information is not published anywhere and is for strict personal usage. It can, for example, be used during a litigation where parties can compare their traces to justify their disagreement.

3) *Cloture phase*: When the applications determine that the ride is finished (e.g., once the destination is reached or after a user intervention), they ask their participant to give a grade (e.g., stars) and a comment on the drive, like it is done in most ride sharing services. The driver and the client independently advertise their ride to the network with a Transaction message. This message contains the ID of the client and of the driver, a start and an end timestamp, a flag indicating whether the message is from a client or a driver, the proposed grade, a comment, the price paid per kilometer, and a hash of the

concatenation of the origin and destination of the ride. The message is signed by the emitter of the message (i.e., the driver or the client) and contains its public key.

The only personal information that can be inferred from a Transaction message is that a given client did a ride with a given driver at a given time. Other information, e.g., location, are never disclosed, unless they appear in the comments.

Transaction Data are eventually validated in the blockchain. Each node in the blockchain checks the values of both transaction data. If both are similar and correct, i.e., timestamps are reasonably close in both transaction messages, the hash of the locations is the same, and the price per kilometer is the same, then they are considered as validated transactions and added to the new next block. In the case transactions cannot be validated (e.g., only one of the two has been received, or the transactions are incompatible), they are still eventually added to the blockchain but they are marked as being in conflict.

Periodically, a *summary block* is published in the blockchain. The new summary block contains the list of observed IDs seen since the last summary block and their new rank. The rank is computed based on the last published rank for this ID and the grades that appear in the transactions with this ID. The function used to re-compute the rank must be robust to malicious users.

III. EVALUATION

In the following we evaluate the scalability of our proposition when it is implemented with Chord. We first perform an analytical study that we comment with data coming from on a real ride trace.

A. Platform dimensioning

1) *Traffic load*: With the protocol described above, we can precisely determine the data traffic generated by each ride for a network composed of n service nodes. To compute the maximum amount of traffic, we consider the case where messages are exchanged in the peer-to-peer network using

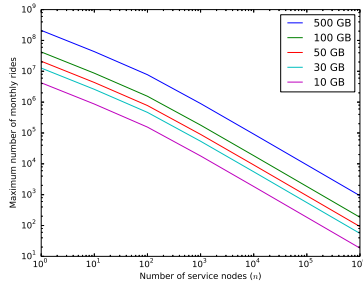


Fig. 4. Total maximum number of monthly rides supported by the platform, given a maximum monthly data volume budget per service node.

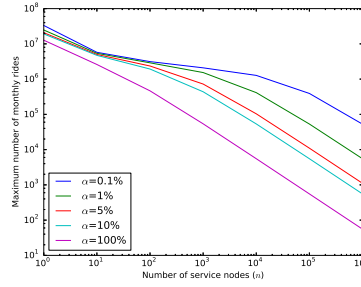


Fig. 5. Total maximum number of monthly rides supported by the platform when service nodes are grouped per zone for a maximum monthly data volume budget of 30 GB per service node.

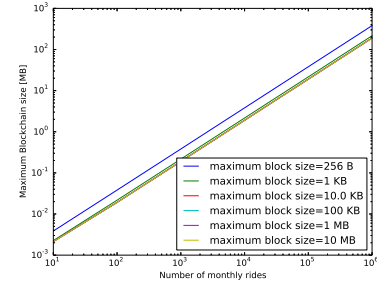


Fig. 6. Maximum increase of blockchain size per month as a function of the number of monthly rides.

UDP segments in IPv6 packets. To determine the amount of traffic supported by service nodes for one ride, we have to identify the three different types of service nodes.

(i) The *client's service node* receives the client's Request that it broadcasts to its neighbors (i.e., $\log_2(n)$ service nodes). It also receives the driver's Acknowledgment messages from the d drivers willing to take the ride and forwards the first one to the client. At the end of a ride, it receives the Transaction message from the its client and forwards it to its neighbors (i.e., $\log_2(n)$ service nodes) but also the one from the driver that it also broadcasts.

(ii) The *driver's service node* receives one broadcasted Request message and forwards it to its neighbors (i.e., $\log_2(n)$ service nodes) and the driver's Ping message to which it answers by forwarding the broadcasted Request message. Similar to the client's node, this node will receive and broadcast transactions.

(iii) The *relay service nodes* have neither the client nor a driver. Such relay nodes receive one broadcasted Request message and two Transaction messages that they have to broadcast it to their neighbors (i.e., $\log_2(n)$ nodes).

The worst case is if the client requests are always issued from clients connected to the same client's service node and when one driver of each service node (including the client's service node) answers to the request.

Fig. 4 shows the evolution of the maximum number of rides that the whole platform can handle every month for a given monthly data volume quota (from 10 GB to 500 GB). The x -axis gives the number of service nodes and the y -axis gives the maximum number of rides per month. Both axes are in log-scale. As shown in Fig. 4, the maximum number of rides that can be handled by the system exponentially decreases with the number of service nodes because every request has to be broadcasted in the network. For example, for 10,000 service nodes having a budget of 30 GB per month, the platform will handle at least 5,524 rides per month. This is three orders of magnitude lower than the number of requests currently seen by Uber or Lyft in a city like San Francisco [17]. Nevertheless, it is important to remember that this value represents the minimum number of rides that the platform can support. Based

on a trace study in Sec. III-B we will see that in practice our platform can handle much more rides, and be used for current taxi services.

To alleviate this scalability issue, we can leverage the possibility to limit the scope of the messages in our Chord network thanks to a clever construction of IDs and identify two types of communications. On the one hand the Transaction messages must be broadcasted to all the service nodes in order to be treated safely by the blockchain. All the other messages can remain in the scope of the region where the number of service nodes can be small (e.g., a dozen) without impairing the security of the system if taken randomly in a large pool of supposedly independent nodes. If α is the fraction of service nodes that compose a region the traffic volume at a service node is in $\mathcal{O}(\alpha \cdot n + \log(n))$ instead of $\mathcal{O}(n + \log(n))$ which substantially decreases the service node load for small values of α . Similarly to Fig. 4, Fig. 5 shows the maximum number of rides that the platform can support for various α and a monthly budget of 30 GB. We can see an important improvement as regions decrease in size. For example, with 10,000 nodes, if each region is limited to 10 service nodes (i.e., $\alpha = 0.1\%$), the platform will support at least 1,275,197 rides per month, which is of the same order of magnitude as what ride sharing services are experiencing today in large cities like San Francisco [17].

2) *Blockchain bloat*: As blockchains are strictly growing structures, it is important to estimate the growth of the structure with time. In our case, we can compute the worst case increase of the blockchain size as it is directly linked to the number of rides supported by the platform: for every ride, at most two transactions will be published in a block of the blockchain. Each transactions in a block contains the IDs of the driver and client, a flag to indicate if the emitter was the driver or the client, a rank, the price per kilometer and the pointer to the comment stored in the DHT. A block contains a list of such transactions, a timestamp, the hash of the block, and the hash of its previous block. Summary blocks are also created with the list of IDs and their rank for IDs with a modified rank since the emission of the last summary block. In the worst case, if the IDs are observed only once between two summary blocks and if their rank changed, the summary

block will then contain all the IDs.

Fig. 6 shows the maximum monthly increase of the blockchain size as a function of the number of monthly rides. Both axes are in log scale. As long as the maximum block size allows to store a large enough number of transactions, the size of the block has no particular impact on the blockchain size. On the contrary, when the block size is too small, it cannot contain many transactions and the block overhead (i.e., its timestamp, hash value, and predecessor pointer) is not negligible anymore. Nevertheless, for reasonable block size, increasing the maximum block size only marginally influences the block chain size but it makes it less reactive as the time span between two block creations increases, a linear function of the transaction rate.

Fig. 6 shows that for realistic parameters of the system, the monthly increase of blockchain size is reasonable. For example, for a monthly budget of 30 GB, $\alpha = 0.1\%$, and 10,000 service nodes, the number of rides would be 1,275,197. With a maximum block size of 100 KB, a new block would be created about every 20.57 minutes and a new summary block about every 2.48 hours for a total of 241.10 MB. If the system can be less reactive, the maximum block size can be increased to 1 MB. In this case, blocks are created every 3.50 hours and summary blocks are created every 1.03 days and the blockchain size monthly increase is 240.84 MB.

B. San-Francisco Taxi trace evaluation

In Sec. III-A we show the scalability limits of our distributed platform in the worst possible case. In this section, we re-evaluate the scalability of our platform in a realistic scenario. To that aim, we use the San Francisco taxi cabs mobility trace from Piorkowski et al. [18]. This dataset tracked GPS coordinates and status of 536 taxi cabs for one month in the Bay area, between May 17th 2008 and June 9th 2008. In total, this data set logs 437,377 rides over a month.

If the platform is composed of 10,000 service nodes, in the worst case where all available drivers acknowledge ride requests and $\alpha = 1$, the maximum monthly traffic at a service node would only be 45.96 GB while the theoretical worst case is higher than 500 GB for the same amount of rides, as depicted by Fig. 4. These results show that despite in theory our distributed platform does not scale as well as a centralized platform, in practice the distributed scheme is perfectly usable. The reason is that in practice the offer and the demand are aligned and only few resources remain unused. In practice, $d \ll n$ for large enough number of service nodes.

Finally, the blockchain size as computed from the trace, shows that in practice the blockchain bloat is limited. For instance, with maximum block size values between 1 KB and 10 MB, the total blockchain size to store informations from the trace stays between 78 MB and 79 MB.

IV. CONCLUSION

The sharing economy mostly relies on centralized platforms, which causes serious threats on security, privacy, and concentration of power. To tackle these issues we present a fully

decentralized and privacy preserving solution. Communication between clients and providers is ensured by a peer-to-peer network and distributed storage. A blockchain plays the role of the marketplace to compute the prices and provide proven trust between the clients and the providers. We carefully designed our communication protocol and data manipulations to ensure data privacy.

We have specifically design our solution for ridesharing business and our analytical study shows that scalability of a privacy-preserving distributed platform remains a challenge in theory. However, in practice, when applied to real taxi services, we can see that it scales enough to enable new respectful sharing economy businesses.

REFERENCES

- [1] C. Cadwalladr and E. Graham-Harrison, "Revealed: 50 million facebook profiles harvested for cambridge analytica in major data breach," *The Guardian*, vol. 17, 2018.
- [2] J. M. Robbins and A. M. Sechooler, "Once more unto the breach: What the equifax and uber data breaches reveal about the intersection of information security and the enforcement of securities laws," *Criminal Justice*, vol. 33, no. 1, pp. 4–7, 2018.
- [3] D. Fletcher, "How facebook is redefining privacy," 2010.
- [4] "Hailing Rides Down Crypto Lane: The Future Of Ridesharing," <https://www.forbes.com/sites/andrewrossow/2018/07/18/hailing-rides-down-crypto-lane-the-future-of-ridesharing/>, accessed: 2018-12-11.
- [5] T. V. A. Pham, I. I. Dacosta Petrocelli, G. F. M. Endignoux, J. R. Troncoso-Pastoriza, K. Huguenin, and J.-P. Hubaux, "Oride: A privacy-preserving yet accountable ride-hailing service," *Proceedings of the 26th USENIX Security Symposium*, 2017. [Online]. Available: <http://infoscience.epfl.ch/record/228219>
- [6] "Arcade City: The Future of Ridesharing Is Decentralized," <https://fee.org/articles/arcade-city-the-future-of-ridesharing-is-decentralized/>.
- [7] "LaZooz," <http://lazooz.org>, accessed: 2018-12-11.
- [8] U. M. Aivodji, S. Gams, M.-J. Huguette, and M.-O. Killijian, "Meeting points in ridesharing: A privacy-preserving approach," *Transportation Research Part C: Emerging Technologies*, vol. 72, pp. 239–253, 2016.
- [9] C. Dai, X. Yuan, and C. Wang, "Privacy-preserving ridesharing recommendation in geosocial networks," in *International Conference on Computational Social Networks*. Springer, 2016, pp. 193–205.
- [10] P. Goel, L. Kulik, and K. Ramamohanarao, "Privacy-aware dynamic ride sharing," *ACM Transactions on Spatial Algorithms and Systems (TSAS)*, vol. 2, no. 1, p. 4, 2016.
- [11] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A survey and comparison of peer-to-peer overlay network schemes," *IEEE Communications Surveys & Tutorials*, vol. 7, no. 2, pp. 72–93, 2005.
- [12] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 4, pp. 149–160, 2001.
- [13] P. Maymounkov and D. Mazieres, "Kademlia: A peer-to-peer information system based on the xor metric," in *International Workshop on Peer-to-Peer Systems*. Springer, 2002, pp. 53–65.
- [14] D. Tapscott and A. Tapscott, *Blockchain revolution: how the technology behind bitcoin is changing money, business, and the world*. Penguin, 2016.
- [15] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [16] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.
- [17] "TNCs TODAY," <http://tncstoday.sfcta.org><http://tncstoday.sfcta.org>, accessed: 2018-08-17.
- [18] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "A Parsimonious Model of Mobile Partitioned Networks with Clustering," in *The First International Conference on COMMunication Systems and NETWORKS (COMSNETS)*, January 2009. [Online]. Available: <http://www.comsnets.org>