

ORide: A Privacy-Preserving yet Accountable Ride-Hailing Service

Anh Pham
EPFL

Italo Dacosta
EPFL

Guillaume Endignoux
EPFL

Juan Ramon Troncoso-Pastoriza
EPFL

Kévin Huguenin
UNIL

Jean-Pierre Hubaux
EPFL

Abstract

In recent years, ride-hailing services (RHSs) have become increasingly popular, serving millions of users per day. Such systems, however, raise significant privacy concerns, because service providers are able to track the precise mobility patterns of all riders and drivers. In this paper, we propose ORide (Oblivious Ride), a privacy-preserving RHS based on somewhat-homomorphic encryption with optimizations such as ciphertext packing and transformed processing. With ORide, a service provider can support the matching of riders and drivers without learning their identities or location information. ORide offers riders with fairly large anonymity sets (e.g., several thousands), even in sparsely-populated areas. In addition, ORide supports key RHSs features such as easy payment, reputation scores, accountability, and retrieval of lost items. Using real data-sets consisting of millions of rides, we show that the computational and network overhead introduced by ORide is acceptable. For example, ORide only adds several milliseconds to ride-hailing operations and the extra driving distance for a driver is less than 0.5 km in more than 75% of the cases evaluated. In short, we show that a RHS can offer strong privacy guarantees to both riders and drivers while maintaining the convenience of its services.

1 Introduction

Ride-hailing services (RHSs), such as Uber and Lyft, enable millions of riders and drivers worldwide to set up rides via their smartphones. Their popularity over traditional taxi services is due to the convenience of their services, e.g., ride requests at the touch of a button, fare estimation, automatic payments, and reputation ratings. Moreover, accountability provided by RHSs is a key feature to riders and drivers, as it makes them feel safer [11, 15]. For instance, in case of a criminal investigation, the RHS provider can offer law-enforcement agencies with the location trace of a particular ride and the identities of the participants.

To offer such services, however, RHSs collect a vast amount of sensitive information that puts at risk the privacy of riders and drivers. First, for each ride, the location traces

and rider’s and driver’s identities are known to the service provider (SP). As a result, the SP, or any entity with access to this data, can infer sensitive information about riders’ activities (such as one-night stands [35]), monitor the locations of riders in real-time for entertainment [19], track the whereabouts of their ex-lovers [40], look up trip information of celebrities [26], and even mount revenge attacks against journalists critical of such services [44]. The SP can charge riders differently based on their profiles [1]. In the case of drivers, there are reports of SPs that track drivers to find if the drivers attended protests [2]. Second, due to the release of drivers’ Personal Identifiable Information (PII) early in the ride set-up procedure, an outsider adversary can massively collect drivers’ PII [37]. Third, there is evidence that RHS drivers and riders are discriminated based on the racial and/or gender information specified in their profiles [20]. Hence, there is a strong need to provide privacy and anonymity for *both* riders and drivers w.r.t. the SP and w.r.t. each other.

To the best of our knowledge, the only privacy-friendly alternative to current RHSs is PrivateRide, recently proposed by Pham et al. [37]. The work, however, has weaknesses, i.e., it does not provide strong privacy guarantees for riders, and offers less accountability and usability, compared to the current RHSs (see Section 2). Therefore, a mechanism with more robust privacy and accountability guarantees is needed.

We present ORide, a privacy-preserving RHS inspired by PrivateRide; it reuses only one operation from PrivateRide, i.e., the proximity check to prevent drivers’ PII from being harvested (see Section 5.4). ORide enables the SP to efficiently match riders and drivers without leaking either their identities or their locations, while providing accountability to deter misbehavior. ORide provides strong privacy for both riders and drivers, i.e., all users in the system are part of large anonymity sets, even if they are in sparsely-populated residential areas. Even in the extreme case of targeted attacks (i.e., a strong curious SP wants to know the destination of a specific rider given the time and location of her pick-up event), the location privacy of the rider’s destination is still guaranteed. For this purpose, ORide

relies on state-of-the-art somewhat-homomorphic encryption system [17] (SHE), to which we apply optimizations for ciphertext packing and transformed processing [36] hence enabling a notable boost in performance and a reduction in overhead w.r.t. naive cryptographic solutions.

Often considered as important as privacy in RHSs are accountability and usability [11, 15]. This introduces challenges in resolving the uneasy tension between privacy, accountability and usability. To achieve accountable privacy, ORide enables the SP to revoke the anonymity of misbehaving riders or drivers when needed. However, the SP does not have full control over this re-identification operation, i.e., it is able to do it only with the support from the attacked party. In addition, to preserve the convenience of the service, ORide supports automatic payment through credit cards and enables riders to contact drivers for lost items. ORide also preserves the reputation-rating operations of the current RHSs.

The evaluation of ORide by using real data-sets from NYC taxi cabs [42] shows that, even with strong bit-security of more than 112 bits, ORide introduces acceptable computational and bandwidth costs for riders, drivers and the SP. For example, for each ride request, a rider needs to download only one ciphertext of size 186 KB with a computational overhead of milliseconds. ORide also provides large anonymity sets for riders at the cost of acceptable bandwidth requirements for the drivers: e.g., for rides in boroughs of Queens and Bronx, a ride would have an anonymity set of about 26,000, and the drivers are only required to have a data-connection speed of less than 2 Mbps. In addition, our results show that ORide is scalable, as we considered a request load that is significantly higher than the one in current RHSs, e.g., Uber accounts for only 15% of the ride pick-up requests in NYC [41].

In summary, we make the following contributions:

- *A novel, oblivious, and efficient matching mechanism.* ORide includes a novel protocol based on quantum-resistant SHE to match riders and drivers, without revealing their identities and locations to the SP. We optimize our SHE-based protocol to considerably reduce the bandwidth requirements and the processing overhead, compared to a vanilla SHE-based protocol; and we propose an efficient extension to deal with malicious drivers.
- *The design and prototype of ORide.* ORide supports the matching of riders and drivers, and different accountability mechanisms, and it reduces the amount of sensitive information revealed to the SP. In particular, ORide supports functionalities that are often considered as important as privacy, such as credit-card payment, reputation rating, and contacting drivers in case of lost items and traceability in case of criminal activity during a ride.
- *Thorough performance and usability evaluation.* Using real data-sets and robust security parameters (i.e., 112 bits security), we show that ORide provides strong privacy guarantees for riders and drivers. In addition, the computational

and network overhead introduced by ORide is practical for riders, drivers and SP. We also show that ORide has a negligible impact on the accuracy of matching riders and drivers compared with current RHSs.

2 Related Work

Researchers have proposed different privacy-enhancing solutions for ride sharing (i.e., car pooling) services [6, 14, 21, 22, 38] and public transportation ticketing systems [8, 27, 32]. In addition, along the line of privacy-preserving solutions for location-based services, many solutions, including solutions based on location cloaking have emerged e.g., [16, 25]. However, little work exists in the area of privacy and security for RHSs, probably due to their relative novelty. According to our literature review, the most significant work in this area is PrivateRide, recently proposed by Pham et al. [37]. PrivateRide offers a practical solution that enhances location privacy for riders and protect drivers' information from harvesting attacks while maintaining the convenience of the service.

Although PrivateRide offers a first solution to the privacy risks in RHSs, it has several important weaknesses that are addressed by our work. First, PrivateRide cannot guarantee the same level of privacy to all the riders because the size of the anonymity set in a particular cloaked area depends on the density of riders in that area. For instance, the anonymity set is smaller for ride requests in areas outside a city center. Also, the tradeoff between the size of a cloaked area and the accuracy of the ride-matching results prevents the use of larger cloaking areas (i.e., to achieve larger anonymity sets). Second, PrivateRide does not protect drivers's privacy, which is also important [2]. Third, PrivateRide provides limited accountability features to deal with relatively common scenarios such as drivers and riders physically attacking each other (i.e., safety concerns) or items being lost during a ride; for many users, such features can be as important as their privacy. Fourth, PrivateRide usability is reduced w.r.t. current RHSs because the supported payment mechanism is less convenient (i.e., PrivateRide requires payments with e-cash bought in advance before a ride). Moreover, ride-matching is suboptimal because the distance between rider and drivers is estimated using the centers of the cloaked areas, instead of exact locations, resulting in additional waiting time for riders.

3 System Model

Our goal is to design a RHS that provides stronger privacy guarantees to both riders and drivers, as well as better or equivalent usability and accountability compared with PrivateRide [37] and current RHSs (e.g., Uber, Lyft, and Easy Taxi). For this purpose, we assume a system consisting of three parties: riders, drivers and the service provider (SP). We now describe our adversarial and system assumptions.

3.1 Adversarial Assumptions

In our model, riders and drivers are active adversaries. The SP is a passive adversary (i.e., honest-but-curious). We assume that most riders and drivers do not collude with the SP, as drivers are independent contractors rather than SP’s employees. The case of covertly active SP is discussed in Section 7.2. However, it is worth mentioning here that, even if the SP is covertly active, it does not provide riders and drivers with malicious apps. This is a valid assumption, because unlike server code, an app can be reverse-engineered, and the risk of public exposure and reputation loss is a strong deterrent against this attack from the SP.

Given that they were observed in current RHSs (i.e., higher chance), we focus on the following attacks:

- (A1) The riders and drivers might attempt to assault each other [46]; in extreme cases, a driver might attempt to kidnap or kill the rider, or vice versa.
- (A2) The SP uses its knowledge about side information about riders and drivers, including their home/work addresses, together with protocol transcripts, to do *large-scale* inference attacks to profile riders’ and drivers’ activities [35].
- (A3) The SP might attempt to carry out *targeted attacks* on specific riders. That is, besides their home/work addresses, the SP knows the precise pick-up location and time of a specific rider and wants to know the drop-off location and time of this ride, or vice versa [26, 40, 44].

3.2 Design Goals

The goal of ORide is to defend against the attacks listed in Section 3.1, and to offer the same level of accountability and usability as current RHSs, as follows.

- Riders and drivers are held accountable for their behaviors during their rides, i.e., the SP is able to identify misbehaving riders or drivers when needed, e.g., if one party physically attacks the other. However, the SP is only able to identify the misbehaving party with support from the attacked party.
- The system preserves the convenience and usability properties offered by current RHSs, such as payment through credit cards and reputation rating. In addition, once a rider is matched with a driver, she can track the location of the driver approaching the pick-up location, and they can contact each other for pick-up coordination. The system also enables riders to contact drivers of their past rides to find lost items.

3.3 System Assumptions

We assume that the metadata of the network and lower communication layers cannot be used to identify riders and drivers or to link their activities. Such an assumption is

Notation	Description
k_s	Ephemeral private key
k_p	Ephemeral public key
$cert_X$	Public-key certificate of X
loc_X	Planar coordinates of X , $loc_X = (x_X, y_X)$
n	Number of drivers
d	Degree of the polynomial
dt	Deposit token
r_{dt}	A random number to create a deposit token
z	A geographical zone
$sig_X\{m\}$	Message m and signature of X on m
$Bsig_{SP}(m)$	Blind signature of the SP on message m

Table 1: Table of notations

practical because, in most cases, the smartphones of drivers and riders do not have fixed public IP addresses; they access the Internet via a NAT gateway offered by their cellular provider. If needed, a VPN proxy or Tor could be used to hide network identifiers.

In addition, we assume that, besides localization capabilities, the rider’s and driver’s smartphones support peer-to-peer wireless communication e.g., Bluetooth and WiFi Direct. Also, for all location-based computations, the apps use a coordinate system such that the Euclidean distances correspond to the great-circle distances, e.g., by using map-projection systems for local areas such as UTM [45] to convert a pair of (latitude, longitude) to planar coordinates (x, y) . Moreover, drivers use a navigation app that does not leak their locations to the SP. This can be done by using a third-party navigation app (e.g., TomTom) or prefetching the map of their operating areas (e.g., a city) and using the navigation app in off-line mode.

3.4 Notation

Throughout the rest of this work, we denote polynomials and scalar values with lowercase letters, variables and rings with uppercase letters, and vectors with boldface letters. $\lceil \cdot \rceil$ denotes rounding to the nearest integer. Polynomials of degree $(d-1)$ will be interchangeably denoted as $a = \sum_{i=0}^{d-1} a_i X^i$ or in their vector form \mathbf{a} when there is no ambiguity. The used symbols and terms are summarized in Table 1.

4 Oblivious Ride-Matching Protocol

One of the challenges in privacy-preserving RHSs is how to efficiently match ride requests to ride offers without revealing the riders’ and drivers’ locations to each other and to the SP. For this, ORide relies on somewhat-homomorphic encryption (see Section 4.1) where the riders and drivers send their encrypted locations to the SP, from which the SP can compute the encrypted squared Euclidean distances between them. We detail this in the following sections. For other cryptographic primitives used in ORide, we refer the readers to Appendix A.3.

4.1 Somewhat-Homomorphic Encryption

Somewhat-Homomorphic Encryption (SHE) is a special kind of malleable encryption that allows a certain number of operations (additions and multiplications) over ciphertexts, without the need to decrypt them first. All SHE cryptosystems present semantic security, i.e., it is not (computationally) possible to distinguish whether two different encryptions conceal the same plaintext. Therefore, it is possible for a party without the private key (in our case, the SP), to operate on the ciphertexts produced by riders and drivers, without obtaining any information about the plaintext values. Additionally, we choose one of the most recent and efficient SHE schemes based on ideal lattices, the FV scheme [17]. This scheme relies on the hardness of the Ring Learning with Errors (RLWE) problem [30]. Note that whenever working with cryptosystems based on finite rings, we must work with integer numbers, thus, from here on, we will assume that all inputs are adequately quantized as integers. Here, we briefly describe the main functions of the FV scheme.

For plaintext elements in a polynomial quotient ring $m \in R_t = \mathbb{Z}_t[X]/(X^d + 1)$ and ciphertext elements in $R_q = \mathbb{Z}_q[X]/(X^d + 1)$, where q and t are positive integers $q > t$ defining the upperbound of the ciphertext and plaintext coefficients, respectively. Let $\Delta = \lfloor q/t \rfloor$ and χ_k, χ_n be two *short* noise random distributions in R_q , the FV encryption of a message $m \in R_t$ with secret key $k_s = s \sim \chi_k$ and public key $\mathbf{k}_p = [p_0, p_1] = [(-a \cdot s + e), a] \in R_q^2$, with e drawn from χ_n and a randomly chosen in R_q , generated by FV.GenKeys , results in a vector expressed as

$$\mathbf{c} = \text{FV.Enc}(\mathbf{k}_p, m) = [p_0 \cdot u + e_1 + \Delta \cdot m, p_1 \cdot u + e_2], \quad (1)$$

where u is drawn from χ_k , and e_1, e_2 are short random polynomials from the error distribution χ_n . All operations are $\text{mod } q$ in R_q .

Decryption of a ciphertext $\mathbf{c} = [c_0, c_1]$ works as

$$m = \text{FV.Dec}(k_s, \mathbf{c}) = (\lfloor t \cdot [c_0 + c_1 \cdot s \text{ mod } q] / q \rfloor) \text{ mod } t.$$

The scheme allows to seamlessly add (FV.Add), subtract (FV.Sub) and multiply (FV.Mul) two encryptions, obtain the encryption of the added, subtracted, and multiplied plaintexts respectively; multiplications consider the encryptions as polynomials in s : $[c_0, c_1] \rightarrow c_0 + c_1 \cdot s$, such that the product between \mathbf{c} and \mathbf{c}' is evaluated as: $[c_0, c_1] \cdot [c'_0, c'_1] \rightarrow c_0 \cdot c'_0 + (c_0 \cdot c'_1 + c_1 \cdot c'_0)s + c_1 \cdot c'_1 \cdot s^2 \rightarrow [c''_0, c''_1, c''_2]$, which results in a ciphertext in R_q^3 , with one extra polynomial; it is possible to recover a fresh-like encryption with two polynomials by employing a relinearization primitive, which requires the usage of a matrix (relinearization key) composed of encrypted pieces of the secret key (we refer the reader to [17] for further details).

4.2 Naive Approach

SHE can be applied to the ride-matching problem in RHS as follows. When a rider wants to make a ride request, she




Figure 1: Our optimized ride-matching approach enables the SP to send to the rider a single ciphertext containing all the squared distances ($\text{dist}_{D_i}^2$) between the rider and available drivers as opposed to one ciphertext per driver (naive approach).

generates an ephemeral FV public/private key-pair together with a relinearization key. She uses the public key to encrypt her planar coordinates and obtains their encrypted forms. She then informs the SP the zone where her pick-up location is, the public and relinearization keys and her encrypted planar coordinates. When this information arrives at the SP, the SP broadcasts the public key to all drivers available in that zone. Each driver uses the public key to encrypt their planar coordinates and sends them to the SP. The SP computes, based on their encrypted coordinates, the encrypted distances between the rider and the drivers, and it returns the encrypted distances to the rider, from which the rider can decrypt and select the best match, e.g., the driver who is the closest to her pick-up location.

However, naive applications of SHE would incur unfeasible computational and bandwidth costs for the riders and the SP, due to the high ciphertext expansion. Furthermore, for each rider request the SP would need to separately compute the ciphertext distances between the rider and each of the drivers: For n drivers, this would mean n distance calculations between encrypted polynomials of d coefficients each, and n ciphertext distances returned to the rider. This would impose an unfeasible overhead in terms of computations for the SP, therefore delaying the ride-matching for the rider and a considerable bandwidth overhead at the rider-SP link (hundreds of MBs if the system has several thousand drivers) (see Section 9.3).

4.3 Optimized Approach

We propose three optimizations that take advantage of the polynomial structure and enables us to operate on d elements of \mathbb{Z}_t packed as a polynomial in R_t in a *single* ciphertext, such that each encrypted operation affects all the coefficients in parallel. When the rider decrypts this ciphertext, she can recover these d values by looking at all the coefficients. From here on, we assume that $d \geq n$, which will be usually the case due to the security bounds on d (see Section 8); in other cases, $\lceil n/d \rceil$ encryptions can be used to pack the whole set of distances analogously.

Packing itself reduces the bandwidth overhead, but it is not

enough for our goal. As we show in Section 5.4, we use all the d packed values independently of each other to calculate all the distances homomorphically in the same encrypted operation, so we need coefficient-wise homomorphic operations. While polynomial additions and subtractions are naturally coefficient-wise, polynomial multiplication in R_t (and its homomorphic counterpart in R_q) is a *convolution product* of the coefficients. A well-known method for transforming convolution products into coefficient-wise products (and vice-versa) in polynomial rings is the *number-theoretic transform* (NTT) [36], a Fourier transform specialized for finite fields. This transform is commonly used in the ciphertext space to speed up polynomial multiplications that are then implemented as coefficient-wise products.

Particularly, in our case, we apply an inverse-NTT to plaintexts before encryption and an NTT after decryption, such that products in the encrypted domain are translated into coefficient-wise products in the plaintext domain. The NTT being linear, it does not affect additions and subtractions. We note that the NTT exists only for certain values of d and t , in particular when t is a prime and d divides $t - 1$. In our case, in order that operations in \mathbb{Z}_t simulate operations in \mathbb{N} on our values, we choose $d = 2^l$ as a power of two and t as a sufficiently large Proth prime (of the form $k2^l + 1$, see [36]) such that all squared-Euclidean distances are less than t . As a result, we improve on both the bandwidth and the computation overhead.

For a final optimization, and due to the low degree of the evaluated operations (squared Euclidean distances), we avoid the use of relinearizations at the SP, which (a) reduces the need to generate and send the relinearization key from the rider to the SP, (b) reduces the noise inside the encryptions, and (c) enables more efficient operations at the SP, at the cost of one extra polynomial to represent the encrypted distance returned to the rider, as we will see in Section 9.

5 ORide

In this section, we present our ORide constructions. We start with an overview description of the system, and then detail ORide operations.

5.1 ORide Overview

ORide provides strong location privacy and anonymity for riders and drivers while still guaranteeing service accountability, secure payment and reputation rating operations. For this purpose, the riders and the drivers have to possess *ride prerequisites* (Section 5.2), including anonymous credentials (ACs), deposit tokens, and digital certificates issued by the SP. To participate in the system, both riders and drivers create anonymous sessions by *logging in to the service* (Section 5.3) with their respective ACs. Drivers periodically report to the SP the geographical zone where they are located. These zones are defined by the SP to balance the network load in the system and the size of the anonymity set of each ride (Section 9.4). Note that, differently from PrivateRide, in ORide, expand-

ing the size of a zone *does not* affect the optimality of the ride-matching and fare-calculation operations (Section 5.4).

When a rider initiates a ride request, the SP, the rider and drivers are involved in a *ride set-up* procedure (Section 5.4) that matches the rider to a driver. In addition, during this stage, as in current RHSs, the rider and the driver agree on the fare based on the estimated distance and duration of the ride. Once the rider and the driver are matched together, they terminate their anonymous sessions. When the ride is completed, the driver creates a new anonymous session and notifies the SP that she is available again. Note that drop-off times and locations are not reported to the SP. In addition, some time after the ride finishes, i.e., at the end of the day, the rider and driver perform *ride-payment and reputation-rating* operations (Section 5.5).

5.2 Ride Prerequisites

Digital certificates. We assume each rider and driver has a digital certificate denoted as $cert_R$ or $cert_D$, issued by the SP at registration time. Each certificate contains a public key and a randomly generated ID. The SP can use this random ID to find the real identity of the certificate holder. Note that the digital certificates are not used by the riders and drivers to log in to the service, and they are not revealed to the SP during a ride. They are used by the riders and drivers to identify each other during the ride as part of ORide’s accountability mechanism (Section 6).

Anonymous credentials. ORide relies on Anonymous Credentials Light (ACL) [9], a linkable anonymous credential system, i.e., a user should use an AC only once to avoid her transactions from being linkable. To use the service anonymously, each user (rider or driver) requests ACs in advance from the SP, using their digital certificate. Hereafter, we denote the anonymous credential for a user X as AC_X , where X is R for riders or D for drivers. Each AC_X contains the average reputation score rep_X , an expiration date exp_X , and the secret key sk_X associated with the public key pub_X in the digital certificate of the AC holder. To differentiate between riders and drivers in the system, an AC also contains a role attribute $role_X$, i.e., $role_X = 1$ if $X = D$, and $role_X = 0$ if $X = R$.

Note that to prevent the SP from de-anonymizing users by correlating the time an AC is issued with the time it is used, or by relying on the AC’s expiration date, the user app could automatically request ACs from the SP at a certain time (e.g., at mid-night), and the expiration date is coarse-grained, e.g., all ACs issued in a day expire at the end of that day). Note that users do not show their reputation scores to the SP during login operations (Section 5.3), hence the reputation scores are the precise reputation of the AC holder. In addition, to prevent DoS attacks against the system, the SP defines a threshold on the number of ACs a rider or driver can acquire per day.

Deposit token. To deter riders’ and drivers’ misbehavior during the ride, each rider is required to possess a *deposit token* before the ride. A deposit token, denoted as dt , is a

random number generated by the rider, blindly signed by the SP (by using blind-signature schemes e.g., [13]) such that the SP is not able to link a token it issued and a token spent by a rider. A deposit token is worth a fixed amount of money defined by the SP. A rider deposits a token to the SP in the beginning of the ride, and she is issued a new token (of the same monetary value) by the SP after the ride payment is successfully completed. Note that the driver is not required to make a deposit, because during the ride set-up operation, the rider and driver exchange their digital certificates with each other. Consequently, if the driver misbehaves, the SP can identify the driver by collaborating with the rider. We discuss this in more detail in Section 5.6.

5.3 Log in to the Service

To use the service, the rider and the driver need to create anonymous sessions to the SP: to do so, they use their anonymous credentials AC_R and AC_D respectively.

Rider. The rider sends to the SP the rider-role $role_R$ and the expiry date exp_R stated in her AC_R . In addition, she proves to the SP that the claimed values are correct, and that, in a zero-knowledge fashion, she knows the secret key sk_R tied to the AC_R .

Driver. Similarly to the rider, the driver follows the same aforementioned procedure by using her AC_D to anonymously log in to the service.

The SP assigns a one-time session ID to each anonymous session, to keep track of that session for coordination. For the sake of simple exposition, in the following, we exclude this one-time session ID from messages exchanged between the rider/driver and the SP.

5.4 Ride Set-up

When a rider requests a ride, the operations performed by the rider, the drivers and the SP are as follows.

1. The rider generates an ephemeral FV public/private key pair, denoted as (\mathbf{k}_p, k_s) . She first computes $p_{x_R} = \sum_{i=0}^{d-1} x_R X^i$ and $p_{y_R} = \sum_{i=0}^{d-1} y_R X^i$. She then applies the inverse-NTT and uses \mathbf{k}_p to encrypt these values: $\mathbf{c}_{x_R} = \text{FV.Enc}(\mathbf{k}_p, \text{NTT}^{-1}(p_{x_R}))$ and similarly for \mathbf{c}_{y_R} . She then sends the zone of her pick-up location (denoted as z), deposit token dt , \mathbf{k}_p , \mathbf{c}_{x_R} and \mathbf{c}_{y_R} to the SP. In order to obtain a new deposit token (of the same monetary value) at the end of the ride, the rider generates a random number r_{dt} , blinds it to r'_{dt} , and sends r'_{dt} in the request. Note that the SP might try to de-anonymize the rider based on the pick-up zone reported by the rider; this is discussed in Section 9.4.
2. The SP checks the validity of the deposit token, i.e., it was not used before. If the token is valid, the SP marks it as used. It then sends to each driver in zone z a different randomly permuted index $0 \leq i < n$ and the public key \mathbf{k}_p .
3. The i -th driver encodes her coordinates in the i -th coefficient: $q'_{x_D} = x_{D_i} X^i$ and $q'_{y_D} = y_{D_i} X^i$. Similarly to the

rider, she applies the inverse-NTT, encrypts these values and sends them to the SP: $\mathbf{c}_{x_D}^i = \text{FV.Enc}(\mathbf{k}_p, \text{NTT}^{-1}(q'_{x_D}))$ and analogously for $\mathbf{c}_{y_D}^i$.

4. The SP sums all drivers' ciphertexts by using the homomorphic property of the cryptosystem to pack them together: $\mathbf{c}_{x_D} = \sum_{i=0}^{n-1} \mathbf{c}_{x_D}^i$ and similarly for \mathbf{c}_{y_D} . It then homomorphically computes the n packed squared values of the Euclidean distances between the n drivers and the rider in parallel, due to the packing $\mathbf{c}_{dist} = (\mathbf{c}_{x_R} - \mathbf{c}_{x_D})^2 + (\mathbf{c}_{y_R} - \mathbf{c}_{y_D})^2$, and it sends the result to the rider (see Fig. 1).
5. The rider decrypts the ciphertext and applies the NTT to obtain a squared distance in each coefficient: $\mathbf{dist} = \text{NTT}(\text{FV.Dec}(\mathbf{k}_s, \mathbf{c}_{dist}))$. Then, she selects the driver with the smallest squared distance.
6. The SP notifies the selected driver. If she declines the offer, the SP asks the rider to select a different driver; it repeats this operation, until one driver accepts. The SP confirms with the rider and the driver that they have been assigned to each other.
- 7a. The rider and the driver establish a secure channel via the SP, e.g., using the unauthenticated Diffie-Hellman protocol, to exchange data that should not be observed by the SP.¹ From the information used to derive the secret key of the secure channel, they compute a shared secret *pairing PIN*. This *pairing PIN* will be used for the proximity-check operation in Step 8.

With this secure channel, the rider and the driver reveal their reputation scores to each other. The trustworthiness of the revealed values are proved by showing that they are indeed the values in the rider's and driver's ACs. If the rider's reputation is too low, the driver can abort the protocol at this step. Likewise, the rider can select another driver, by using the list of cleartext squared Euclidean distances she obtained in step 5.

- 7b. Via the secure channel, the rider and the driver exchange their precise locations (i.e., loc_R and loc_D , respectively). In addition, they exchange their digital certificates ($cert_R$ and $cert_D$) with each other. This provides accountability for the rider and driver, i.e., if the rider is attacked by the driver, she can report the driver's certificate to the SP, or vice versa.

The driver drives from her current location loc_D to the pick-up location loc_R , using an off-line navigation app or a third-party navigation app (such as Google Maps or TomTom). She sends, in real time, her precise locations to the rider, via the secure channel, thus the rider can track the movements of the car. Also, at this point, the rider and the driver can call or message each other through their ride-hailing apps, if needed. They can also terminate their anonymous sessions with the SP.

¹Detection of man-in-the-middle attacks by the SP is done in step 8.




Figure 2: ORide ride setup protocol. The dashed arrows represent the secure channel (via the SP), and the dotted arrows represent the proximity channel.

8. When the rider and the driver are in proximity, the driver performs a proximity check to verify the *physical presence* of the rider before releasing her identifying information. That is, they use a short-range wireless technology (e.g., Bluetooth or WiFi Direct) to set up a proximity channel using the *pairing PIN*. If the channel is successfully established, the driver can verify that the rider is in her proximity. This is similar to the approach proposed in [37] to prevent drivers' PII from being harvested. If this step fails, the driver can decide to abort the protocol. Also, via the proximity channel, the rider and the driver can check that the secure channel (established at step 7a) was not tampered with by the SP.
9. The driver releases her identifying information, including her vehicle's license plate number and her profile picture, to the rider. This information helps the rider to identify the driver and her car and to prevent certain threats, e.g., fake drivers [43]. Therefore, the required communication distance between the phones of the rider and the driver is small (e.g., several meters).
10. The rider and the driver create a *fare report*. A fare report

is a token generated by the rider and, at the end of the day, the driver deposits it to the SP to be paid (Section 5.5). A fare report is created as follows. The rider sends her drop-off location to the driver via the secure channel, they agree on the path, and based on the estimated path, they compute the fare. The rider and driver then sign a message consisting of the fare and their certificates, i.e., $fare\ report = sig_{R-D}\{fare, cert_R, cert_D\}$, using the private key associated with their $cert_R$ and $cert_D$. Note that this upfront-fare method has been implemented in current RHSs, such as in Uber [34] and in Lyft [39]. Once the driver receives the fare report from the rider, the ride begins. The rider's and driver's app do not report any information to the SP at this stage and during the ride. Also, to prevent the SP from inferring the starting time of the ride based on the interactions between the rider and the driver over the secure channel, the rider and driver can randomly send bogus information to each other through the secure channel.

Intuitively, with the aforementioned ride-matching algorithm where the distances between the rider and drivers

are computed based on their (encrypted) *precise* locations, the expansion of a zone, unlike in PrivateRide, will not result in negative effects on the optimality of the ride-matching and fare-calculation operations. In addition, with ciphertext packing, we reduce by a factor of n the communication between the SP and the rider. However, this comes at the cost of an added vulnerability against potentially malicious drivers. Furthermore, note that in step 1 of the protocol, any valid rider can generate an ephemeral public/private key pair. Consequently, if the SP was active, it could track the locations of the drivers, thus indirectly track the locations of the riders. We discuss solutions to these potential issues in Section 7.

5.5 Ride Payment and Reputation Rating

When the car arrives at the drop-off location, the driver creates a new anonymous session to the SP. This enables her to receive ride-request broadcasts from the SP. Note that the driver *does not* report to the SP that the ride is completed.

The driver sends to the SP the fare report $sig_{R-D}\{fare, cert_R, cert_D\}$ she received during the ride set-up operation (step 10, Section 5.4). The SP checks the correctness of the rider certificate $cert_R$ in the fare report and the correctness of the signature. If they are valid, the SP charges the rider according to her payment method. It also issues a new *deposit token*, denoted as dt' , to the rider, by blindly signing the blinded random number r'_{dt} it received in the rider's ride request (i.e., $dt' = sig_{SP}\{r'_{dt}\}$, step 1 in Section 5.4), and it sends this blind signature to the rider's account. It then subtracts the service fee, and deposits the remainder to the driver.

Once the payment is successfully completed, the rider and driver can vote for the reputation of each other, similarly to current RHSs. They can log in to the service with their real credentials and provide the reputation score for the party whom they rode with.

Note that ORide preserves the payment and reputation-rating operations of the current RHSs. That is, unlike PrivateRide, it does not require the rider to purchase e-cash in advance, and it does not require the rider and the driver to generate and keep extra cryptographic tokens for the reputation-rating operation. In addition, ORide does not require the rider and the driver to hide their identifying information to the SP during the payment and reputation-rating operations, because the rider and driver are both anonymous during the ride. However, it is important to note that, in order to prevent the SP from de-anonymizing the rider and the driver by correlating the time a fare report is deposited with the drop-off event of the ride, the payment operation should *not* occur immediately after the ride, e.g., the drivers deposit the fare reports to the SP at the end of the day.

5.6 Ride Cancellation

As in current RHSs, a rider or a driver can cancel a ride any time before or during the ride. This, however, is discouraged by the SP, because it can lead to malicious behaviors: For example, once a rider and a driver are assigned to each other

by the SP, they meet at the pick-up location and start the ride as normal; but, to avoid the service fee, the rider or the driver can send a fake cancellation notification to the SP. Therefore, similarly to current RHSs, if a rider or a driver cancels a ride a certain amount of time after the ride request, they should be punished by the SP, e.g., their reputation scores are lowered or fees are charged [12].

In ORide, when a rider cancels a ride, the SP can offer her two options: lose her *deposit token* (i.e., pay a penalty) or reveal her $cert_R$ and have her reputation score lowered. If a driver cancels a ride, the SP can ask the rider to reveal the $cert_D$, from which the SP can identify and punish the driver according to its policy.

6 Accountability

In this section, we discuss the accountability goals (mentioned in Section 3.2) of ORide. This includes defense mechanisms against the attack (A1) in Section 3.1 and additional features such as retrieval of lost items, ensurance of payment, and integrity of the reputation-rating operation. Attacks (A2) and (A3) are discussed in Section 8.

(A1) Accountability. ORide enables the rider and the driver to exchange, during the ride set-up procedure, their digital certificates, i.e., $cert_R$ and $cert_D$, respectively, and the fare report. This provides accountability for riders and drivers, i.e., an attacked party can report to the SP the digital certificate of the attacker and the fare report, from which the SP can identify the attack to charge her a fee, lower her reputation and/or support legal action. However, the SP is only able to identify the attacker with support from the attacked party. Likewise, the attacked party cannot obtain the real identity of the attacker without support from the SP.

To provide accountability in the extreme cases where the rider's phone is not available, e.g., kidnapping or the phone is broken, promptly after a rider receives $cert_D$ and fare report from the driver, she can share them, via out-of-band channels such as messaging apps, with her friends. This is similar to the approach used in personal safety apps, such as Google Trusted Contacts [23]. Similarly, during the ride, via out-of-band channels, she can share her GPS trace with her friends using (k, l) threshold secret sharing [18], i.e., each GPS location point is split into l parts so that any k out of l parts reconstruct the original coordinate.

Another risk in RHSs is ensurance of payment. A rider cannot avoid paying the fare of a ride, because she is identified with her digital certificate $cert_R$ in the fare report. As the rider and driver agree on the fare and sign it before the ride, they cannot subsequently increase or decrease this fare. However, they might collude to underpay the service fee to the SP, by agreeing on a small fare and paying the difference in cash. Yet in this case, ORide offers the same guarantees as current RHSs, because riders can already request a small ride through the application and then pay in cash for a longer ride once they have met the driver. In future work, we will

explore mechanisms to protect against this attack.

Moreover, the bilateral rating system enables the SP to ban abusive riders and drivers from the service. A rider or driver cannot claim a better reputation for herself, because the proof for attributes in her AC will not be correct w.r.t. her falsely claimed reputation. They also cannot arbitrarily vote for the reputation of each other, because a payment record is needed (the deposit of a *fare report*). In addition, as discussed in Section 5.6, similarly to current RHSs, ORide enables the SP to hold riders and drivers accountable for ride cancellations.

SP incentives. From an economic perspective, ride-hailing service SPs would have incentives to deploy ORide as it provides privacy and security for the riders and still preserves their business models (i.e., the SP can still charge a commission for each ride). In order to monetize ride data, the SPs can provide a discount for riders if they reveal (part of) their GPS traces. In addition, privacy and security for RHSs might be enforced by law and legislation, and ORide shows that it is technically possible to achieve a strong level of protection. As such, this work lays the foundation for the design of a privacy-preserving and secure RHSs.

Additional features. Similarly to current RHSs, ORide enables the riders to retrieve lost items (i.e., items forgotten in the car), as drivers’ certificates $cert_D$ and car information are provided during the ride set-up procedure. As discussed earlier, the riders can share $cert_D$ with their friends, thus, even if the riders lose their phones, they can still be able to retrieve the $cert_D$ from their friends and contact the SP. In addition, due to the secure channel established between the rider and the driver, the rider can still track the driver trajectory while waiting at her pick-up location, or they can contact each other (e.g., messaging or calling) when needed.

7 Protecting against Malicious Behaviors

In this section, we describe how the protocol presented in Section 5 can be extended to defend against malicious drivers and a covertly active SP.

7.1 Malicious Drivers: Masking

As mentioned in Section 5.4, if a driver behaves maliciously, she could encrypt non-zero values in the slots other than her allotted one, thus corrupting the inputs from other drivers. Our protocol can cope with this malicious behavior by adding one extra step in which the SP homomorphically multiplies each driver ciphertext by a mask $m_i = \text{NTT}^{-1}(X^i)$ for the driver’s index i (see notations from Section 5.4), which preserves only the contents in the allocated slot. However, the mask does not hold any sensitive information and it is known by the SP, hence consequently, a naive homomorphic multiplication with an encrypted mask would impose an unjustified overhead. Therefore, we propose, instead, a more efficient multiplication operation, denoted \star , as follows.

Given a ciphertext $\mathbf{c} = [c_0, c_1] \in R_q^2$ corresponding to a plaintext $m \in R_t$, and a mask $m_i \in R_t$, we want to obtain a

ciphertext $\mathbf{c}' = \mathbf{c} \star m_i$ corresponding to the masked plaintext $\text{FV.Dec}(k_s, \mathbf{c}) \cdot m_i$. m_i can be thought of as its own noiseless and unscaled encryption (Equation (1) on page 4, evaluated for $u, e_1, e_2 = 0$, and no scale Δ), being a vector in R_q^2 with only one non-zero component ($[m_i, 0] \in R_q^2$). Therefore, the product results

$$\mathbf{c} \star m_i = [c_0 \cdot m_i, c_1 \cdot m_i].$$

The \star operation consists of two polynomial multiplications, it avoids encryption of m_i , halves the number of products with respect to an encrypted homomorphic multiplication, and keeps the cipher size from growing after the product, hence considerably improving the performance of this operation.

In any case, this precaution is only needed in case the drivers are malicious, and random checks on their locations can be implemented instead if the drivers are just covertly active (i.e., they refrain from cheating if there is a negligible chance of being caught in the act).

7.2 Covertly Active SP

As mentioned in Section 5.4, any valid rider can generate an ephemeral key to make a ride request. Since the SP issues credentials for riders and drivers, it can impersonate a rider or a driver in its own system.

If the SP *continuously* impersonates a rider, it could learn the drivers’ locations, from which it could learn the coarse-grained pick-up locations of the riders. In other words, if a rider chooses the driver who is the closest to her pick-up location, the SP would know that she is in the Voronoi cell of her selected driver. Below, we present a mechanism to deter this attack. We note that the attack is not trivial, due to the high dynamics of the system, i.e., drivers can arbitrarily go on-line and off-line anytime. The SP would not have strong incentives to perform this attack, because it would add computational and bandwidth overhead to the service, thus negatively affecting the productivity of the service itself. To deter this attack, we introduce the notion of *Proof-of-Ride* (PoR), defined and used as explained below. An illustration of the protocol with PoR is shown in Appendix A.1.

A PoR is a random number $rand$ generated by the rider, signed by the driver using the secret key associated with her $cert_D$, and then *blindly* signed by the SP by using blind-signature schemes such as [13], i.e., $\text{PoR} = \text{Bsig}_{SP}\{\text{sig}_D\{rand\}\}$. It is used to prove to the drivers that the rider is real, i.e., she did a ride in the past. When a rider makes a ride request, she has to provide in her ride request a PoR, the $cert_D$ and the random number $rand$ used in the PoR. A PoR can be used only once. For the first ride, $\text{PoR} = cert_R$.

To prevent the SP from creating its own $cert_R$ and $cert_D$ to create its own fake PoR, the SP has to provide certificate transparency [28]: the SP maintains and publishes a publicly auditable and append-only log of all rider and driver certificates it has issued and revoked. Whenever a driver receives

	Identities	Pick-up loc.	Pick-up time	Drop-off loc.	Drop-off time	Loc. trace	Fare
Current RHSs	Rider, Driver	Precise	Precise	Precise	Precise	Full	Yes
PrivateRide	Driver	Zone	Obfuscated	Zone	Obfuscated	Partial	Yes
ORide	N/A	Zone	Obfuscated	N/A	N/A	N/A	N/A

Table 2: Information observed by the SP during ride set-up procedure w.r.t. different RHS designs. Note that the zone in ORide is larger than the zone in PrivateRide without affecting the ride-matching optimality (see Section 9.4). Also note that, the payment operation in ORide reveals some information about the riders, but it cannot be used to break the anonymity of the rides (explained in Section 8).

a PoR, she can check whether the rider’s certificate $cert_R$ (in the case of the first ride), or the driver’s certificate indicated in the PoR, is in the list of certificates published by the SP. This way, if the SP internally creates fake accounts, it can be detected by auditing authorities, similarly to the cases of companies opening fake user accounts [10]. Similarly, to prevent a rider from double-spending a PoR, and the SP from reusing a valid PoR to perform the aforementioned active attack, the SP maintains and publishes an append-only logs of PoRs that have been spent, or cancelled (due to ride cancellation).

Note that PoR could create a point of linkability, i.e., the SP is able to know that the rider is in the set of identities indicated in the fare reports deposited by a specific driver. This can be easily prevented by using anonymous-reputation and anonymous-payment systems (e.g., e-cash), as used in the PrivateRide system [37].

8 Privacy and Security Analysis

In this section, we present an analysis of ORide to show that it effectively addresses against the privacy attacks described in Section 3.1.

The SP cannot de-anonymize a rider or driver through their anonymous logins by using their ACs. These are guaranteed due to the anonymity and unlinkability properties of the ACL anonymous credential system [9]. Additionally, the SP cannot obtain extra information from the riders’ and drivers’ encrypted locations and their encrypted distances; this is due to the semantic security of the FV encryption scheme [17]. Table 2 compares the information observed by the SP during ride set-up procedure w.r.t. different RHS designs.

Specifically, in ORide, the information observed by the SP from ORide operations can be put in two databases, as follows.

- Ride DB, in which each entry contains the role and expiration date of the AC, the pick-up zone and obfuscated pick-up time. The role and expiration date are coarse-grained, i.e., all riders or drivers have the same role, and all ACs issued on the same day expire at the end of that day.
- Payment DB, in which each entry contains a rider’s ID, a fare, and the day the fare report is deposited to the SP. Note that this database does not exist if payment is done through e-cash, as used in the PrivateRide system.

(A2) Large-scale inference attacks by the SP. To profile riders’ and drivers’ activities, the SP needs to learn the identities, the locations, and the times associated with their rides.

By using the Payment DB, the SP would know which specific rider took a ride with a specific driver on which day and what its fare was. Since drivers are often licensed to operate in a city or state, knowing that a rider took a ride with a specific driver, the SP might be able to know the city where the rider took a ride, but it does not know the specific location in the city. In addition, knowing the home/work addresses and the fares of the rides, the SP might be able to know if a rider went from home to work. Such rides are not sensitive, however, compared to others, such as one-night stands, going to abortion clinics or political-party meetings. Otherwise, anonymous-payment methods, such as e-cash, could be used to decouple the riders’ identities from the fares, thus preventing the SP from knowing rides between home and work of the riders.

By using the Ride DB, the SP might be able to guess the identities of the riders, only if the pick-up zone has a limited number of ride activities *and* riders, e.g., a zone where only one rider lives. This case, however, will never happen in ORide, because the zones are defined by the SP in such a way that each zone has *at least* a large minimum number of ride requests per day, while balancing the bandwidth requirements for the drivers. We illustrate this in Section 9. Note that the SP would be detected if it lied about the activity densities in the zones, because these densities are public knowledge [41], and the drivers would notice if they received very few ride requests from a certain zone.

Assuming a stronger adversarial SP that, besides home/work addresses of the riders, knows that a rider took a ride from a given zone and it wants to know the time of that ride. In this case, the anonymity set of a ride is the number of rides that occurred on the same day in that zone. Similarly to the aforementioned case, the anonymity set of a ride depends on the ride density in its pick-up zone. However, note that, as this requires more side information, the anonymity-set size in this case is the *lower-bound* estimation of the anonymity set for the aforementioned case. This lower bound is used in the evaluation of the anonymity set achieved by ORide, presented in Section 9.

(A3) Targeted attacks by the SP. In the case where the SP knows the precise pick-up location and time of a ride, it still cannot know the drop-off location and time of the ride, because, in ORide, the drop-off event is not reported to the SP. Knowing the fare from the Payment DB, the SP might be able to check whether the target went home or to work, but it could not know other destinations. However, as mentioned earlier, such rides are not very sensitive. Otherwise, anonymous-payment methods, such as e-cash could be used to prevent these attacks.

PII- and location-harvesting attacks by outsiders. ORide reuses the proximity-check mechanism proposed by

PrivateRide, hence it provides the same guarantees for harvesting-attacks against drivers’ PII. However, a malicious outsider might attempt to triangulate drivers, to obtain a snapshot of the locations of all drivers in a zone: It could make three ride requests at the same time to obtain the distances, and cancels these requests immediately. ORide mitigates this attack by applying two measures: (1) requiring a deposit token from each rider per request, thus enabling the SP to identify riders who make many requests and cancel (as discussed in Section 5.6), and (2) permuting the list of drivers’ indices for each ride request (step 2 in Section 5.4). Also, the SP can further prevent this problem by defining a threshold on the number of ACs each rider account can obtain per day, thus limiting the number of ride requests a rider can make.

9 Evaluation

In this section, we evaluate our protocols using a real taxi-trace data-set. We first evaluate the performance of the ride-matching operation in terms of computational and bandwidth requirements for the riders and drivers. We then evaluate the effect of Euclidean distances on the optimality of ride-matching operations.

9.1 Data-Sets

Our data-set consists of over 1.1 billion taxi rides in New York from January 2009 to June 2015 [42]. We extracted data for the month of November in 2013, one of the busiest months in the data-set, which resulted in a subset of 15,004,556 rides. In this subset, the average duration of the rides is 13 minutes, respectively. The GPS traces of the rides are not given; however, the precise pick-up and drop-off locations and times, and pseudo-IDs of the taxi drivers associated with the rides are provided.

We make the following assumptions. First, the drop-off location of a driver is her waiting location. Second, a ride-request event is a pick-up event (i.e., consisting of a pick-up location and pick-up time) in our data-set. Third, for each ride-request event, the set of drivers available for that request consists of drivers who have at least one drop-off event in the last 30 minutes since the ride-request timestamp. The 30-minute interval is chosen, because the data-set shows that 99th percentile of the time gap between the drop-off event of a driver and her next pick-up event is approximately 30 minutes.

9.2 Implementation Details

Our ORide prototype features the main cryptographic operations for the ride matching in the ride set-up procedure (Section 5.4). Other cryptographic operations needed for requesting a ride, i.e., AC operations and blind signatures, and for setting up the proximity channel between the rider’s app and the driver’s app, can be found in the evaluation of the PrivateRide system [37].

To measure the cryptographic overhead of ride-matching operations, we implemented a proof-of-concept ORide

Setting	Rider		Driver	
	Upload (KB)	Download (KB)	Download (KB)	Upload (KB)
S1	372	761856	124	248
S2	372	186	124	248
S3	372	186	124	248

Table 3: Per-ride bandwidth requirements of ORide, with $d = 4096$, $\log_2(q) = 124$, and there are 4096 drivers available for a ride request ($n = 4096$).

in C++, by relying on the NFFlib library [31]. In our experiments, the SP, the rider, and the driver are located on the same computer, hence network delays are not considered. However, the network delay would not impose a considerable overhead, because a ride-matching operation requires only one round-trip message between the rider and the SP, and one round-trip message between the SP and each driver. Also, the data exchanged between the rider and the SP, and the SP and the drivers, is small, as discussed in Section 9.4. Note that, similarly to the current RHSs, the SP can define a timeout for responses from the drivers such that the latency is reasonable for the services. Due to the dependency requirements of the NFFlib, it is not trivial to port the implementation to mobile platforms. However, to make our experiments close to the performance of smartphones, in all of our implementations, we *do not* use SSE or AVX optimizations for Intel processors. The ORide proof-of-concept implementation on smartphones is work in progress.

9.3 Per-Ride Overhead

In this section, we describe our experimental setup, and presents the bandwidth and computational overhead for a rider and a driver per ride request.

We used ORide’s prototype to estimate the overhead added for ride-matching operations in three settings: (**S1**) when naive SHE approach (Section 4.2) without using relinearizations at the SP is used, (**S2**) when ciphertext-packing optimizations are used and the drivers correctly write their coordinates to their assigned allocated slots (Section 5.4), and (**S3**) when ciphertext-packing optimizations are used and the drivers write arbitrary values to slots of other drivers, hence masking (Section 7.1) is needed.

Experimental Setup. To measure the performance of our system, we used a computer (Intel i5-4200U CPU, 2.6 GHz, 6 GB RAM) with Debian Jessie (Linux kernel 3.16). The security parameters used in our experiments are tuned to achieve an equivalent bit-security of more than 112 bits, therefore exceeding current NIST standards [5] for 2016-2030. With this security target, and a plaintext size of 20 bits the needed polynomial dimension is $d = 4096$, with coefficients of size 124 bits (each polynomial has a size of 62 KB). These parameters guarantee both 112-bits of security and correct homomorphic operations for homomorphically adding up to 4096 encrypted locations in the same ciphertext

Setting	Rider			Driver		SP	
	Gen. keys (ms)	Encrypt (ms)	Decrypt (ms)	Load key (ms)	Encrypt (ms)	Load key (ms)	Compute Dist. (ms)
S1	1.51±0.06	2.6±0.2	7823.4±573.4	0.53±0.01	2.6±0.2	0.53±0.01	113868.8±6553
S2	1.51±0.06	2.6±0.2	2.2±0.1	0.53±0.01	2.6±0.2	0.53±0.01	208.9±4
S3	1.51±0.06	2.6±0.2	2.2±0.1	0.53±0.01	2.6±0.2	0.53±0.01	745.5±24.5

Table 4: Per-ride computational overhead of ORide (without AVX/SSE optimizations), for $d=4096$, $\log_2(q)=124$, and there are 4096 drivers available for a request. Statistics ($avg \pm std$) were computed from 1000 experiments.

and calculating the corresponding Euclidean distances.² We refer the reader to Section A.2 for more details about the possible granularity a geographical area can have.

Assuming a rider makes a ride request to the SP, and there are 4096 drivers available for the request ($n=4096$), with the aforementioned security parameters, the bandwidth requirements and computational overhead per ride request, for a rider and a driver, are shown in Table 3 and Table 4, and explained below.

- *Bandwidth overhead for a rider:* In all three settings, for each ride request, a rider sends to the SP a public key and two ciphertexts for her encrypted planar coordinates. This totals 6 polynomials, a payload size of 372 KB.

Regarding the number of distance ciphertexts a rider receives from the SP, in setting **S1**, it is equal to n , i.e., the number of responding drivers. In settings **S2** and **S3**, it is significantly reduced to $\lceil n/d \rceil$, due to ciphertext packing. A ciphertext distance, when avoiding relinearizations (see Section 5.4), consists of 3 polynomials, thus having a total size of 186 KB. Assuming there are 4096 drivers responding a ride request, setting **S1** would require the SP to send 4096 distance ciphertexts (744 MB) to the rider, while **S2** would require only one distance ciphertext (186 KB).

- *Bandwidth overhead for a driver:* in all three settings, for each request: On the downlink, the SP forwards to each driver a public key, 2 polynomials of size 124 KB. On the uplink, each driver sends back to the SP her encrypted planar coordinates, totaling 4 polynomials of size 248 KB.
- *Computational overhead:* As shown in Table 4, in all three settings, the computational overhead introduced by key generation and encryption operations are small, i.e., 1.5 ms and 2.6 ms, respectively for both riders and drivers. Due to masking, setting **S3** introduces a small computational overhead for the SP in homomorphic squared-Euclidean-distance computation, compared to setting **S2** (745 ms vs. 208.9 ms). However, noticeably, due to ciphertext packing, settings **S2** and **S3** significantly reduce the computational cost for the SP (208.9 and 745 ms compared to 113868.8 ms required by **S1**). It also

²We refer the reader to Section 6 in [17] for more details on the choice of cryptographic parameters for FV. It is worth noting that we have considered pessimistic bounds in order to cope with recently published attacks that reevaluate the security of lattice-based cryptosystems [7].

significantly reduces the decryption overhead for the rider, from 7823 ms in setting **S1** to 2.2 ms in settings **S2** and **S3**.

Note that the results for the rider and driver are optimistic, as we used a laptop instead of a smartphone (however, as stated before, CPU optimizations were not used to reduce the difference). While such comparisons are not straightforward, we can do a rough estimation of the expected performance of ORide in smartphones. For instance, comparing the performance scores of top multicore CPUs in smartphones [3] with top multicore CPUs in desktops [4], we can see that the difference is less than an order of magnitude. Assuming such difference, then we can see that the computational overheads for key generation, encryption and decryption are still acceptable in smartphones, around 15 ms, 26 ms, and 22 ms respectively. The overhead is still acceptable even if we consider two orders of magnitude difference, as the total time to hail a ride is in the order of minutes.

9.4 Riders’ Anonymity and Drivers’ Bandwidth Requirements

In this section, we present the trade-off between the ride-anonymity set vs. bandwidth requirements for the riders and drivers, by using the real data-set presented in Section 9.1.

Due to the high demand of taxi rides in Manhattan w.r.t. lower activity other boroughs in NYC (from our data-set, Manhattan accounts for 90% of ride requests), we define two zone settings as follows.

- Setting one (Z1): Manhattan is divided into census tracts (CTs). Each CT is one zone. The boroughs of Queens and Bronx are merged into one zone, and the boroughs of Brooklyn and Staten Island are merged into one zone.
- Setting two (Z2): Manhattan is divided into neighborhood tabulation areas (NTAs). Each NTA is one zone. Similarly to setting one, the boroughs of Queens and Bronx are merged into one zone, and the boroughs of Brooklyn and Staten Island are merged into one zone.

Estimation of the anonymity set. As explained in Section 8, the number of rides in a day from a zone is a *lower-bound* estimation of the anonymity set for a ride. Fig. 3a shows the experimental cumulative distribution function (CDF) of the lower-bound anonymity-set size. It can be observed that, for Manhattan with the zone granularity of census tracts, 81.7% of the rides have an anonymity set

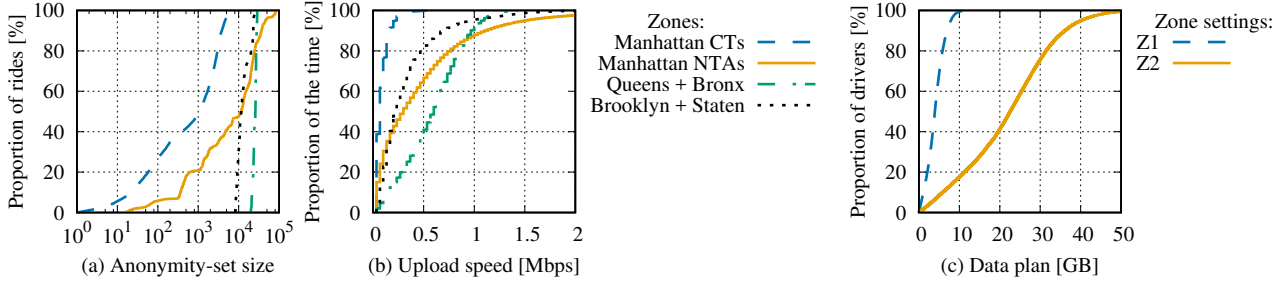


Figure 3: System performance. (a) Anonymity-set size, (b) Upload speed requirement for the drivers, (c) Monthly-data plan requirement for the drivers

of size *at least* 50, and for a zone consisting of Queens and Brooklyn, all of the rides have an anonymity-set size of *at least* approximately 26,000.

Bandwidth requirements for riders. As mentioned earlier, the bandwidth requirements for a rider, per ride request, depends on the number of available drivers. Our experiments show that for both zone settings, for all ride requests, the number of available drivers is less than 3,500. This means, with the security parameter chosen (as presented in Section 9.3) and when proposed optimized packing approaches are used, a rider needs to download only one ciphertext distance, i.e., 186 KB, which is negligible.

Bandwidth requirements for drivers. Fig. 3b shows the CDF of the upload speed required for the drivers; the upload speed is computed by multiplying the number of requests a driver receives per second with the size of the ciphertexts she has to upload per request. Note that the required downlink speed is half of the uplink speed, because the downlink payload is half the size of the uplink payload (Section 9.3). It shows that for Manhattan with the zone granularity of census tracts, the required upload speed is less than 0.5 Mbps, and for other zones, the required upload speed is less than 2 Mbps, which is provided by 3G or 3.5G networks.

Monthly-data plan required for the drivers. Fig. 3c shows the CDF of a data plan required for the drivers for two aforementioned zone settings; this is calculated by multiplying the total number of requests a driver would receive during her waiting time with the uplink- and downlink-payloads per request. The result shows that with the zone setting Z1, a driver needs *at most* 10 GB of data per month, and with the zone setting Z2, 60% of drivers need less than 25 GB of data per month. This requirement is reasonable: For example, in the U. S. , an unlimited data plan typically offers 20-26 GB of high-speed data with less than \$100 [33]. In addition, the drivers can reduce his data-plan consumption by using free WiFi networks, such as LinkNYC [29]. Also, note that the results presented also show that ORide can scale, because current RHSs (e.g., Uber) accounts for only 15% of the ride pick-up requests in NYC [41].

The requirements on bandwidth for the drivers and the anonymity-set sizes for riders enables the SP to define the zones that balance the trade-off between the two aforemen-




Figure 4: Effect of Euclidean distance on the extra distances for the drivers (left) and on the waiting time for the riders (right) w.r.t. different zone settings.

tioned requirements. For example, for areas that have a high density of ride activities such as Manhattan, the SP could discretize the borough into zones of CTs or NTAs, or combinations of CTs and NTAs. Note that, as shown earlier, at the granularity level of CTs (Z1), the anonymity set provided by ORide for the case of a very strong adversarial SP is already large. In the corner cases of punctual events such as concerts and sporting events, the SP can split a crowded zone into sub-zones, in order to balance the aforementioned trade-off. For areas that have less ride activities, such as other boroughs in NYC or other cities, an entire borough or city can be a zone. For example, a zone consisting of the boroughs of Queens and Bronx would guarantee an anonymity set of at least 26,000 for a ride, while requiring the drivers to have an Internet connection of only at most 2 Mbps.

9.5 Effect on Ride Matching

To minimize the extra costs for both the drivers (gas and driving time to pickup) and the riders (waiting times at pickup locations), ideally, the ride-matching algorithm should take into account the road networks and real-time traffic conditions. ORide uses a simpler matching metric, i.e., Euclidean distances between the riders' and drivers' locations due to the limited operations supported by SHE. In addition, due to the bandwidth constraints, the ORide ride-matching algorithm matches a rider to drivers in the same zone, hence suboptimality, e.g., if a rider is close to the border of a zone, the closest driver might be in one of the neighboring zones.

Fig. 4 shows the CDFs of the relative extra costs due to the suboptimality of ORide, compared to the ideal solution,

w.r.t. these different zone settings: Z_1 , Z_2 , and the entire city of New York (NYC). The experiment was done on a set of 1,000 randomly selected ride requests. For the ideal matching, we use the Google Maps Distance Matrix APIs [24] to compute the times and distances between a pick-up request and the available drivers (see assumptions in Section 9.1). To reduce the number of requests made to the Google APIs³, from the set of available drivers, we selected 100 drivers who are closest to the pick-up location as the potential candidates for the ideal matching.

It can be observed that the median extra costs are small: when Z_1 is used, in more than 45% of the cases, the driver selected by ORide and the ideal solution is the same, and, in nearly 80% of the cases, the extra driving distance is less than 0.5 km. In addition, the size of the zone only has negligible effects on the optimality of the matching algorithm: If the set of all the drivers available in NYC was used for the ORide matching algorithm, compared to the ideal solution, 78.7% of the cases would have an extra distance of less than 0.5 km, compared to 76.2 % and 76.8 % of the cases when Z_1 and Z_2 were used, respectively.

10 Conclusion

In this paper, we have proposed ORide, a practical solution that efficiently matches riders and drivers in a privacy-preserving way while still offering key RHSs features such as easy payment, reputation scores, accountability, and retrieval of lost items. ORide enables the SP to choose a balanced trade-off between anonymity sets for riders vs. bandwidth requirements for the drivers. For example, for a lower-bound anonymity-set size of 26,000 for rides from the boroughs of Queens and Bronx, drivers only need to have an Internet connection of at most 2 Mbps. The trade-off enables the SP to define the zones such that all users in the system are guaranteed large anonymity sets, even if they are in sparsely-populated residential areas with sparse ride activities (by expanding the zones). We have also shown that, even in the extreme case of targeted attacks, i.e., a curious SP wants to know the destination of a rider given the time and location of a rider's pick-up event, the location privacy of the rider's destination is still guaranteed.

As part of our future work, we plan to implement a full prototype of the system on mobile platforms, and to evaluate the performance of the system on real RHS data-sets.

Acknowledgements

The authors are grateful to Chris Soghoian for his suggestions about accountability and usability of RHSs in the beginning of the work. The authors also would like to thank Zhicong Huang and Christian Mouchet for their comments on the finalization of the paper.

³The number of requests per day is limited.

References

- [1] <https://www.bloomberg.com/news/articles/2017-05-19/uber-s-future-may-rely-on-predicting-how-much-you-re-willing-to-pay>. Last visited: Jun. 2017.
- [2] <http://www.theverge.com/2015/6/14/8778111/uber-threatens-to-fire-drivers-attending-protests-in-china>. Last visited: Jan. 2017.
- [3] <http://browser.primatelabs.com/android-benchmarks>. Last visited: Jan. 2017.
- [4] <http://browser.primatelabs.com/processor-benchmarks>. Last visited: Jan. 2017.
- [5] Recommendation for Key Management, Part 1: General, SP 800-57 Part 1 Rev. 4. Online: <http://dx.doi.org/10.6028/NIST.SP.800-57pt1r4>, January 2016. Last visited: Feb. 2017.
- [6] AÏVODJI, U. M., GAMBS, S., HUGUET, M.-J., AND KILLIJIAN, M.-O. Meeting points in ridesharing: A privacy-preserving approach. *Transportation Research Part C: Emerging Technologies* (2016).
- [7] ALBRECHT, M. R. On dual lattice attacks against small-secret LWE and parameter choices in HELib and SEAL. Cryptology ePrint Archive, Report 2017/047, 2017. <http://eprint.iacr.org/2017/047>.
- [8] ARFAOUI, G., LALANDE, J.-F., TRAORÉ, J., DESMOULINS, N., BERTHOMÉ, P., AND GHAROUT, S. A Practical Set-Membership Proof for Privacy-Preserving NFC Mobile Ticketing. *Proc. of the 15th Privacy Enhancing Technologies Symposium* (2015).
- [9] BALDIMTSI, F., AND LYSYANSKAYA, A. Anonymous credentials light. In *Proc. of Conference on Computer & communications security* (2013), pp. 1087–1098.
- [10] <https://www.bloomberg.com/view/articles/2016-09-09/wells-fargo-opened-a-couple-million-fake-accounts>. Last visited: Jan 2017.
- [11] <http://uk.businessinsider.com/despite-its-problems-uber-is-still-the-safest-way-to-order-a-taxi-2014-12?r=US&IR=T>. Last visited: Feb. 2017.
- [12] <https://newsroom.uber.com/updated-cancellation-policy/>. Last visited: Jan. 2017.
- [13] CHAUM, D. Blind signatures for untraceable payments. In *Proc. of CRYPTO* (1983).
- [14] DAI, C., YUAN, X., AND WANG, C. Privacy-preserving ridesharing recommendation in geosocial networks. In *Proc. of Conference on Computational Social Networks* (2016).
- [15] <http://www.dailydot.com/via/uber-lyft-safety-background-checks/>. Last visited: Feb. 2017.
- [16] DAMIANI, M. L., SILVESTRI, C., AND BERTINO, E. Fine-grained cloaking of sensitive positions in location-sharing applications. *IEEE Pervasive Computing* 10, 4 (2011), 64–72.
- [17] FAN, J., AND VERCAUTEREN, F. Somewhat Practical Fully Homomorphic Encryption. Cryptology ePrint Archive, Report 2012/144, 2012. <http://eprint.iacr.org/2012/144>.
- [18] FELDMAN, P. A practical scheme for non-interactive verifiable secret sharing. In *Proc. of Symposium on Foundations of Computer Science* (1987).
- [19] <http://www.forbes.com/sites/kashmirhill/2014/10/03/god-view-uber-allegedly-stalked-users-for-party-gods-viewing-pleasure/>. Last visited: Jan. 2017.
- [20] GE, Y., KNITTEL, C. R., MACKENZIE, D., AND ZOEPEF, S. Racial and gender discrimination in transportation network companies. Tech. rep., National Bureau of Economic Research, 2016. <http://www.nber.org/papers/w22776>. Last visited: Jan. 2017.
- [21] GOEL, P., KULIK, L., AND RAMAMOHANARAO, K. Optimal pick up point selection for effective ride sharing. *IEEE Transactions on Big Data* (2016).

- [22] GOEL, P., KULIK, L., AND RAMAMOCHANARAO, K. Privacy-aware dynamic ride sharing. *Trans. on Spatial Algorithms and Systems* (2016).
- [23] <https://support.google.com/trustedcontacts/?hl=en>. Last visited: Jan. 2017.
- [24] <https://developers.google.com/maps/documentation/distance-matrix/>. Last visited: Jan. 2017.
- [25] GRUTESER, M., AND GRUNWALD, D. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proc. of the 1st conference on Mobile systems, applications and services* (2003), ACM, pp. 31–42.
- [26] <http://www.ibtimes.co.uk/former-uber-employee-reveals-drivers-used-tracking-technology-stalk-celebrities-politicians-1596263>. Last visited: Jan 2017.
- [27] ISERN-DEYÀ, A. P., VIVES-GUASCH, A., MUT-PUIGSERVER, M., PAYERAS-CAPELLÀ, M., AND CASTELLÀ-ROCA, J. A secure automatic fare collection system for time-based or distance-based services with revocable anonymity for users. *The Computer Journal* (2013).
- [28] LAURIE, B., LANGLEY, A., AND KASPER, E. Certificate transparency. <https://tools.ietf.org/html/rfc6962>, 2013.
- [29] <https://www.link.nyc>. Last visited: Feb. 2017.
- [30] LYUBASHEVSKY, V., PEIKERT, C., AND REGEV, O. On Ideal Lattices and Learning with Errors Over Rings. Cryptology ePrint Archive, Report 2012/230, 2012. <http://eprint.iacr.org/2012/230>.
- [31] MELCHOR, C. A., BARRIER, J., GUELTON, S., GUINET, A., KILLIJIAN, M., AND LEPOINT, T. NTLlib: NTT-Based Fast Lattice Library. In *Proc. of the RSA conference - The Cryptographers' Track* (2016).
- [32] MILUTINOVIC, M., DECROIX, K., NAESSENS, V., AND DE DECKER, B. Privacy-preserving public transport ticketing system. In *Proc. of Conference on Data and Applications Security and Privacy* (2015).
- [33] <https://www.nerdwallet.com/blog/utilities/best-unlimited-data-plans/>. Last visited: Feb. 2017.
- [34] <https://newsroom.uber.com/philippines/new-upfront-fares-on-uberx/>. Last visited: Jan. 2017.
- [35] http://www.oregonlive.com/today/index.ssf/2014/11/sex_the_single_girl_and_ubers.html. Last visited: Jan. 2017.
- [36] PEDROUZO-ULLOA, A., TRONCOSO-PASTORIZA, J. R., AND PEREZ-GONZALEZ, F. Number theoretic transforms for secure signal processing. *Trans. on Information Forensics and Security* (2017).
- [37] PHAM, T. V. A., DACOSTA PETROCELLI, I. I., JACOT-GUILLARMOD, B., HUGUENIN, K., HAJAR, T., TRAMÈR, F., GLIGOR, V., AND HUBAUX, J.-P. PrivateRide: A Privacy-Enhanced Ride-Hailing Service. In *Proc. of Privacy Enhancing Technologies Symposium* (2017).
- [38] SÁNCHEZ, D., MARTÍNEZ, S., AND DOMINGO-FERRER, J. Co-utile P2P ridesharing via decentralization and reputation management. *Transportation Research Part C: Emerging Technologies* (2016).
- [39] <https://techcrunch.com/2016/11/29/just-like-uber-lyft-launches-upfront-fares/>. Last visited: Jan. 2017.
- [40] <http://thenextweb.com/insider/2016/12/13/uber-tracks-customers-long-after-their-ride-is-over/>. Last visited: Jan. 2017.
- [41] <http://toddschneider.com/posts/analyzing-1-1-billion-nyc-taxi-and-uber-trips-with-a-vengeance/>. Last visited: Feb. 2017.
- [42] <https://uofi.app.box.com/NYctaxidata>. Last visited: Jan. 2017.
- [43] <http://www.usatoday.com/story/tech/columnist/stevenpetrow/2016/10/12/fake-uber-drivers-dont-become-next-victim/91903508/>. Last visited: Jan 2017.
- [44] <http://www.usatoday.com/story/tech/2014/11/19/uber-privacy-tracking/19285481/>. Last visited: Jan. 2017.
- [45] <https://www.uwgb.edu/dutchs/FieldMethods/UTMSystem.htm>. Last visited: Feb. 2017.
- [46] <http://www.wcnc.com/news/crime/uber-driver-attacked-rider-over-politics-man-says/339458660>. Last visited: Jan 2017.

A Appendix

A.1 Covertly Active SP

Fig. 5 illustrates the changes introduced to the original ride set-up procedure (Section 5.4) to handle a *covertly active* SP.

A.2 Plaintext Space

Assume a geographical area of size $s \times s$ and a plaintext space of b bits to represent the squared-Euclidean distances between points in the area. The area can be quantized into a grid with cells of size $s/2^{(b-1)/2} \times s/2^{(b-1)/2}$, with the explanation as follows. Assuming the area is discretized into a grid of $v \times v$ cells, the largest possible squared-Euclidean distance between any two points on the grid is $2 \times v^2$, and this has to be at most 2^b . Therefore, $v \leq 2^{(b-1)/2}$. In other words, each edge of size s can be discretized into v points, and the distance between any pair of two consecutive points is $s/2^{(b-1)/2}$. Therefore, the area can be represented by a grid with cells of size $s/2^{(b-1)/2} \times s/2^{(b-1)/2}$.

For example, with 20-bit plaintext space, a geographical area of size 60 km², such as the borough of Manhattan in NYC, would be quantized into a grid of resolution approximately 10 m \times 10 m.

A.3 Cryptographic Primitives

In this section, we briefly describe the cryptographic building blocks used in ORide.

Blind signatures. A blind-signature scheme [13] is a form of digital-signature schemes in which the signer does not know the content of the message that she is signing. This is achieved by enabling the signature requester to ‘blind’ (i.e., randomize) the message before sending it to the signer. When the signature requester receives the signature on her blinded message, she ‘unblinds’ it to obtain a valid signature for the original message. The signer, when is asked to verify the signature of an unblinded message, is not able to link this message back to the blinded version she signed.

Anonymous credentials. An anonymous credential (AC) is a cryptographic token with which the credential owner can prove to another party that she satisfies certain properties without revealing her real identity. In ORide, a user is identified when she obtains ACs from the SP. However, when she wants to start an anonymous session, she reveals to the SP only the expiration date and the role specified in the AC (i.e., rider or driver), and she proves to the SP,

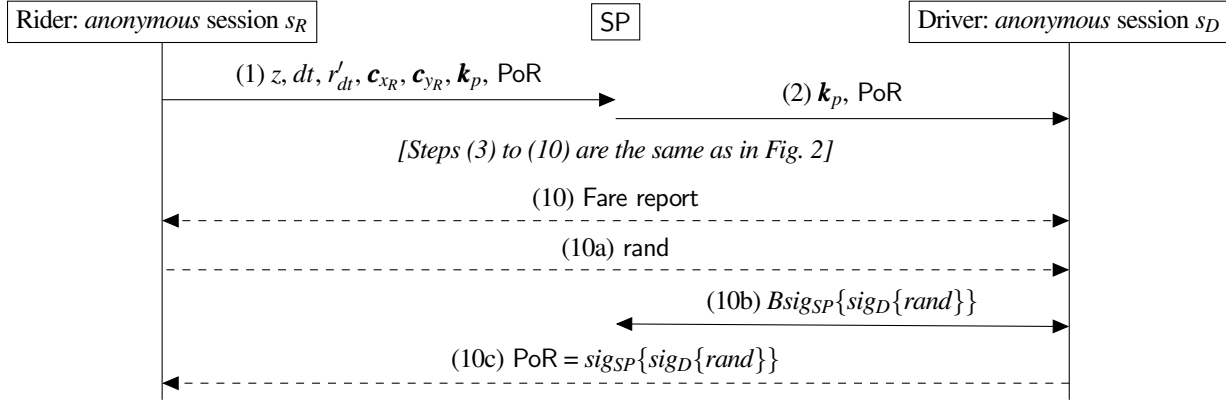


Figure 5: Changes introduced to the original ride set-up protocols (Fig. 2) to handle *covertly active* SP.

in a zero-knowledge fashion, that she knows the private key associated with the AC. To prove her reputation to a driver, a rider reveals to the driver the reputation score specified in her AC together with the proof to show that the revealed value is trustworthy. ORide relies on the Anonymous Credentials Light (ACL). However, note that ACL is a linkable anonymous credential scheme, i.e., a user can only use a credential once to avoid her transactions from being linkable.