

Where are you Bob? Privacy-Preserving Proximity Testing with a Napping Party

Ivan Oleynikov¹, Elena Pagnin², and Andrei Sabelfeld¹

¹ Chalmers University, Gothenburg, Sweden

² Lund University, Lund, Sweden

{ivanol, andrei}@chalmers.se, elena.pagnin@eit.lth.se

Abstract Location based services (LBS) extensively utilize proximity testing to help people discover nearby friends, devices, and services. Current practices rely on full trust to the service providers: users share their locations with the providers who perform proximity testing on behalf of the users. Unfortunately, location data has been often breached by LBS providers, raising privacy concerns over the current practices. To address these concerns previous research has suggested cryptographic protocols for privacy-preserving location proximity testing. Yet general and precise location proximity testing has been out of reach for the current research. A major roadblock has been the requirement by much of the previous work that for proximity testing between Alice and Bob both must be present online. This requirement is not problematic for one-to-one proximity testing but it does not generalize to one-to-many testing. Indeed, in settings like ridesharing, it is desirable to match against ride preferences of all users, not necessarily ones that are currently online.

This paper proposes a novel privacy-preserving proximity testing protocol where, after providing some data about its location, one party can go offline (nap) during the proximity testing execution, without undermining user privacy. We thus break away from the limitation of much of the previous work where the parties must be online and interact directly to each other to retain user privacy. Our basic protocol achieves privacy against semi-honest parties and can be upgraded to full security (against malicious parties) in a straight forward way using advanced cryptographic tools. Finally, we reduce the responding client overhead from quadratic (in the proximity radius parameter) to constant, compared to the previous research. Analysis and performance experiments with an implementation confirm our findings.

Keywords: Secure proximity testing · privacy-preserving location based services · MPC.

1 Introduction

As we digitalize our lives and increasingly rely on smart devices and services, *location based services (LBS)* are among the most widely employed ones. These

range from simple automation like “switch my phone to silent mode if my location is office”, to more advanced services that involve interaction with other parties, as in “find nearby coffee shops”, “find nearby friends”, or “find a ride”.

Preserving Privacy for Location Based Services. Current practices rely on full trust to the LBS providers: users share their locations with the providers who manipulate location data on behalf of the users. For example, social apps Facebook and Tinder require access to user location in order to check if other users are nearby. Unfortunately, location and user data has been often breached by the LBS providers [27]. The ridesharing app Uber has been reported to violate location privacy of users by stalking journalists, VIPs, and ex-partners [22], as well as ex-filtrating user location information from its competitors [40]. This raises privacy concerns over the current practices.

Privacy-Preserving Proximity Testing. To address these concerns previous research has suggested cryptographic protocols for privacy-preserving location services. The focus of this paper is on the problem of *proximity testing*, the problem of determining if two parties are nearby without revealing any other information about their location. Proximity testing is a useful ingredient for many LBS. For example, ridesharing services are often based on determining the proximity of ride endpoints [18]. There is extensive literature (discussed in Section 6) on the problem of proximity testing [44,42,41,11,29,30,36,38,20,21,35,24].

Generality and Precision of Proximity Testing. Yet general and precise location proximity testing has been out of reach for the current research. A major roadblock has been the requirement that proximity testing between Alice and Bob is only possible in a pairwise fashion and both must be present online. As a consequence, Alice cannot have a single query answered with respect to multiple Bobs, and nor can she be able to check proximity with respect to Bob’s preferences unless Bob is online.

The popular ridesharing service BlaBlaCar [4] (currently implemented as a full-trust service) is an excellent fit to illustrate our goals. This service targets intercity rides which users plan in advance. It is an important requirement that users might go off-line after submitting their preferences. The goal is to find rides that start and end at about the same location. Bob (there can be many Bobs) submits the endpoints of a desired ride to the service and goes off-line (napping). At a later point Alice queries the service for proximity testing of her endpoints with Bob’s. A key requirement is that Alice should be able to perform a one-to-many proximity query, against all Bobs, and compute answer even if Bob is offline. Unfortunately, the vast majority of the previous work [44,42,41,11,30,36,38,20,21,35,24] fall short of addressing this requirement.

Another key requirement for our work is precision. A large body of prior approaches [44,42,41,11,29,30,31,26] resort to grid-based approximations where the proximity problem is reduced to the problem of checking whether the parties are located in the same cell on the grid. Unfortunately, grid-based proximity

suffers from both false positives and negatives and can be exploited when crossing cell boundaries [9]. In contrast, our work targets precise proximity testing.

This paper addresses privacy-preserving proximity testing with respect to napping parties. Beyond the described offline functionality and precision we summarize the requirements for our solution as follows: (1) security, in the sense that Alice may not learn anything else about Bob’s location other than the proximity; Bob should not learn anything about Alice’s location; and the service provider should not learn anything about Alice’s or Bob’s locations; (2) generality, in the sense that the protocol should allow for one-to-many matching without demanding all users to be online; (3) precision, in the sense of a reliable matching method, not an approximate one; (2) lightweight client computation, in the sense of offloading the bulk of work to intermediate servers. We further articulate on these goals in Section 2.

Contributions. This paper proposes OLIC (OffLine Inner-Circle), a novel protocol for proximity testing (Section 4). We break away from the limitation of much of the previous work where the parties must be online. Drawing on Hallgren et al.’s two-party protocol InnerCircle [20] we propose a novel protocol for proximity testing that utilizes two non-colluding servers. One server is used to blind Bob’s location in such a way that the other server can unblind it for any Alice. Once they have uploaded their locations users in our protocol can go offline and retrieve the match outcome the next time they are online.

In line with our goals, we guarantee security with respect to semi-honest parties, proving that the only location information leaked by the protocol is the result of the proximity test revealed to Alice (Section 4.2). We then show how to generically mitigate malicious (yet non-colluding) servers by means of zero knowledge proofs and multi-key homomorphic signatures (Section 4.3). Generality in the number of users follows from the fact that users do not need to be online in a pairwise fashion, as a single user can query proximity against the encrypted preferences of the other users. We leverage InnerCircle to preserve the precision, avoiding to approximate proximity information by grids or introducing noise. Finally, OLIC offloads the bulk of work from Bob to the servers, thus reducing Bob’s computation and communication costs from quadratic (in the proximity radius parameter) to constant. We note, that while InnerCircle can also be trivially augmented with an extra server to offload Bob’s computations to, this will add extra security assumptions and make InnerCircle less applicable in practice. OLIC, on the other hand, already requires the servers for Bob to submit his data to, and we get offloading for free. On Alice’s side, the computation and communication costs stay unchanged. We develop a proof of concept implementation of OLIC and compare it with InnerCircle. Our performance experiments confirm the aforementioned gains (Section 5).

On the 2 Non-Colluding Servers Assumption. We consider this assumption to be realistic, in the sense that it significantly improves on the current practices of a single full-trust server as in BlaBlaCar, while at the same time being compatible with common assumptions in practical cryptographic privacy-preserving

systems. For example, Sharemind [39] requires three non-colluding servers for its multi-party computation system based on 3-party additive secret sharing. To the best of our knowledge, OLIC represents the first 2-server solution to perform proximity testing against napping users in ridesharing scenarios, where privacy, precision and efficiency are all mandatory goals. Notably, achieving privacy using a single server is known to be impossible [17]. Indeed, if Bob was to submit his data to only one server, the latter could query itself on this data and learn Bob’s location via trilateration attack.

2 Modeling Private Proximity Testing Using Two Servers

The goal of private proximity testing is to enable one entity, that we will call Alice, to find out whether another entity, Bob, is within a certain distance of her. Note that the functionality is asymmetric: only Alice learns whether Bob lies in her proximity. The proximity test should be performed without revealing any additional information regarding the precise locations to one another, or to any third party.

Our Setting. We consider a multi party computation protocol for four parties (Alice, Bob, Server₁, and Server₂) that computes the proximity of Alice’s and Bob’s inputs in a private way and satisfies the following three constraints:

- (C-1) Alice does not need to know Bob’s identity before starting the test, nor the parties need to share any common secret;
- (C-2) Bob needs to be online only to update his location data. In particular, Bob can ‘nap’ during the actual protocol execution.
- (C-3) The protocol is executed with the help of two servers.

In detail, constraint (C-1) ensures that Alice can look for a match in the database without necessarily targeting a specific user. This may be relevant in situations where one wants to check the availability of a ride ‘near by’, instead of searching if a specific cab is at reach and aligns with our generality goal towards one-to-many matching. Constraint (C-2) is rarely considered in the literature. The two most common settings in the literature are either to have Alice and Bob communicate directly to one another (which implies that either the two parties need to be online at the same time) [20], or to rely on a single server, which may lead either to ‘hiccup’ executions (lagging until the other party rejoins online) [30] or to the need for a trusted server. In order to ensure a smooth executions even with a napping Bob and to reduce the trust in one single server, we make use of two servers to store Bob’s contribution to the proximity test, that is constraint (C-3). This aligns with our goal of lightweight client computation. We remark that, for a napping Bob, privacy is lost if we use a single server [17].

Finally, unlike [30] we do not require the existence of a social network among system users, to determine who trusts whom, nor do we rely on shared secret keys among users.

Formalizing ‘Napping’. We formalize the requirement that ‘Bob may be napping’ during the protocol execution in the following way. It is possible to split the protocol in two sequential phases. In the first phase, Bob is required to be online and to upload data to the two servers. In the second phase Alice comes online and perform her proximity test query to one server, that we call Server_1 . The servers communicate with one another to run the protocol execution, and finally Server_2 returns the result to Alice.

Ideal Functionality. We adopt an ideal functionality that is very close to the one in [20]: if Alice and Bob are within a distance r^2 of each other, the protocol outputs 1 (to Alice), otherwise it outputs 0 (to Alice). Figure 1 depicts this behavior. Alice and Bob are the only parties giving inputs to the protocol. Server_1 and Server_2 do not give any input nor receive any output. This approach aligns with our goal towards precision and a reliable matching method, and brakes away from approximate approaches. For simplicity, we assume the threshold value r^2 to be set a priori, but our protocol accommodates for Alice (or Bob) to choose this value.




Figure 1: Figurative representation of the functionality implemented by privacy-preserving location proximity.

Practical efficiency. Finally, we are interested in solutions that are efficient and run in reasonable time on commodity devices (e.g., laptop computers, smartphones). Concretely, we aim at reducing the computational burden of the clients—Alice and Bob—so that their algorithms can run in just a few seconds, and can in the worst case match the performance of *InnerCircle*.

Attacker Model. We make the following assumptions:

- (A-1) $\text{Server}_1, \text{Server}_2$ are not colluding;
- (A-2) All parties (Alice, Bob, Server_1 , and Server_2) are honest but curious, i.e., they meticulously follow the protocol but may log all messages and attempt to infer some further knowledge on the data they see.

In our model any party could be an attacker that tries to extract un-authorized information from the protocol execution. However, we mainly consider attackers that do so without deviating from their expected execution. In Section 4.3, we show that it is possible to relax assumption (A-2) and guarantee users’ privacy against 2 malicious servers. In this case, the servers may misbehave and output a different result than what expected from the protocol execution. However, we do not let the two server collude and intentionally share information. While we can handle malicious servers, we cannot tolerate a malicious Alice or Bob. Detecting attacks such as Alice or Bob providing fake coordinates, or Alice trilaterating Bob is outside our scope and should be addressed by different means. We regard such attacks as orthogonal to our contribution and suggest to mitigate them by employing tamper-resistant location devices or location tags techniques [30].

Privacy. Our definition of privacy goes along the lines of [30]. Briefly, a protocol is private if it reveals no information other than the output of the computation, i.e., it has the same behavior as the ideal functionality. To show the privacy of a protocol we adopt the standard definition of simulation based security for deterministic functionalities [28]. Concretely, we will argue the indistinguishability between a real world and an ideal world execution of our protocol, assuming that the building blocks are secure, there is no collusion among any set of parties and all parties are honest but curious.

General limitations. Proximity testing can be done in a private way only for a limited amount of requests [30,20] in order to avoid trilateration attacks [43]. Practical techniques to limit the number of requests to proximity services are available [34]. Asymmetric proximity testing is suitable for checking for nearby people, devices and services. It is known [44] that the asymmetric setting not directly generalize to mutual location proximity testing.

3 Recap of the InnerCircle Protocol

We begin by giving a high-level overview of the InnerCircle protocol by Hallgren et al. [20] and then describe its relation to our work.

The InnerCircle protocol in [20] allows two parties, Alice and Bob, to check whether the Euclidean distance between them is no greater than a chosen value r . The protocol is privacy preserving, i.e., Alice only learns a bit stating whether the Bob lies closer than r or not, while it does not reveal anything to Bob or to any external party.

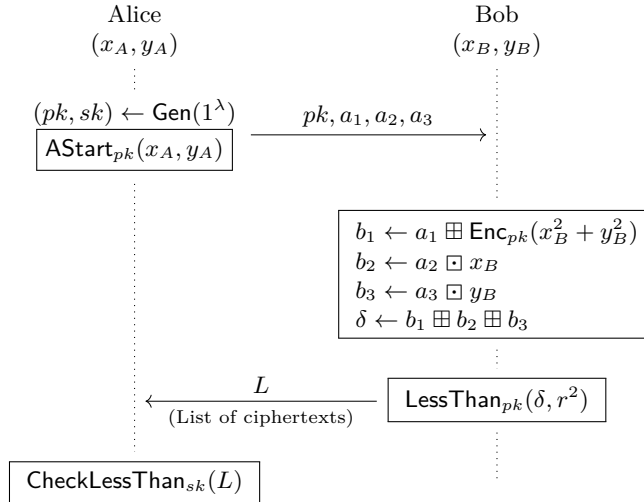


Figure 2: Diagram of the InnerCircle protocol.

More formally, let (x_A, y_A) and (x_B, y_B) respectively denote Alice's and Bob's locations. Throughout the paper, we will use the shorthand $D = (x_A - x_B)^2 +$

$(y_A - y_B)^2$ to denote the squared Euclidean distance between Alice and Bob and δ is an encryption of D . For any fixed non-negative integer r , the functionality implemented by the `InnerCircle` protocol is defined by:

$$F_{\text{InnerCircle}}((x_A, y_A), (x_B, y_B)) = (z, \varepsilon), \text{ where}$$

$$z = \begin{cases} 1, & \text{if } D \leq r^2 \text{ (} D = (x_A - x_B)^2 + (y_A - y_B)^2 \text{)} \\ 0, & \text{otherwise.} \end{cases}$$

`InnerCircle` is a single round protocol where the first message is from Alice to Bob and the second one is from Bob to Alice (see Figure 2).

`InnerCircle` has three major steps (detailed in Figure 3):

(1) Alice sends to Bob encryptions of her coordinates together with her public key pk ; (2) Bob computes δ (an encryption of D) using Alice's data and his own coordinates, after that he runs `LessThan` to produce a list L of ciphertexts which encode whether $D \leq r^2$ (without knowing D himself) and sends L to Alice; (3) Alice decrypts the list using `CheckLessThan` and learns whether $D \leq r^2$ or not.

Figure 3 contains the definitions of the three procedures used in `InnerCircle` that will carry on to our OLIC protocol. Below we describe the major ones (`LessThan` and `CheckLessThan`) in more detail.

The procedure `LessThan` (Figure 3b) produces a set of ciphertexts from which Alice can deduce whether $D \leq r^2$ without disclosing the actual value D . The main idea behind this algorithm is that given $\delta = \text{Enc}_{pk}(D)$ Bob can “encode” whether $D = i$ in the value $x \leftarrow \delta \oplus \text{Enc}_{pk}(-i)$, and then “mask” it by multiplying it by a random element $l \leftarrow x \boxtimes r$. Observe that if $D - i = 0$ then $\text{Dec}_{sk}(l) = 0$; otherwise if $D - i$ is some invertible element of \mathcal{M} then $\text{Dec}_{sk}(l)$ is uniformly distributed on $\mathcal{M} \setminus \{0\}$. (When we instantiate our OLIC protocol for specific cryptosystem, we will ensure that $D - i$ is either 0 or an invertible element, hence we do not consider here the third case when neither of these holds.) The `LessThan` procedure terminates with a shuffling of the list, i.e., the elements in L are reorganized in a random order so that the entry index no longer correspond to the value i .

<code>AStart</code> _{$pk(x_A, y_A)$}	
1 :	$a_1 \leftarrow \text{Enc}_{pk}(x_A^2 + y_A^2)$
2 :	$a_2 \leftarrow \text{Enc}_{pk}(2x_A)$
3 :	$a_3 \leftarrow \text{Enc}_{pk}(2y_A)$
4 :	return (a_1, a_2, a_3)

(a) The `AStart` algorithm.

<code>LessThan</code> _{$pk(\delta, r^2)$}	
1 :	for $i \in \{0 \dots r^2 - 1\}$
2 :	$x_i \leftarrow \delta \boxplus \text{Enc}_{pk}(-i)$
3 :	$t_i \leftarrow \mathcal{M} \setminus \{0\}$
4 :	$l_i \leftarrow x_i \boxtimes t_i$
5 :	$L \leftarrow [l_0, \dots, l_{r^2-1}]$
6 :	return <code>Shuffle</code> (L)

(b) The `LessThan` algorithm.

<code>CheckLessThan</code> _{$sk(L)$}	
1 :	$[l_0, l_2 \dots l_{r^2-1}] \leftarrow L$
2 :	for $i \in \{0 \dots r^2 - 1\}$
3 :	$v \leftarrow \text{Dec}_{sk}(l_i)$
4 :	if $v = 0$
5 :	return 1
6 :	return 0

(c) The `CheckLessThan` algorithm.

Figure 3: The core algorithms in the `InnerCircle` protocol.

The procedure `CheckLessThan` (called `inProx` in [20] depicted in Figure 3c. This procedure takes as input the list L output by `LessThan` and decrypts one by one all of its components. The algorithm returns 1 if and only if there is a list element that decrypts to 0. This tells Alice whether $D = i$ for some $i < r^2$. We remark that Alice cannot infer any additional information, in particular she cannot extract the exact value i for which the equality holds. In the `LessThan` procedure Bob computes such “encodings” l for all $i \in \{0 \dots r^2\}$ and accumulates them in a list. So if $D \in \{0 \dots r^2 - 1\}$ is the case, then one of the list elements will decrypt to zero and others will decrypt to uniformly random $\mathcal{M} \setminus \{0\}$ elements, otherwise all of the list elements will decrypt to random nonzero elements. If the cryptosystem allowed multiplying encrypted values, Bob could compute the product of the “encodings” instead of collecting them into a list and send only a single plaintext, but the cryptosystem used here is only additively homomorphic and does not allow multiplications.

4 Private Location Proximity Testing with Napping Bob

Notation. We denote an empty string by ε , and a list of values by $[\cdot]$ and the set of plaintexts (of an encryption scheme) by \mathcal{M} . We order parties alphabetically, protocols’ inputs and outputs follow the order of the parties. Finally, we denote the computationally indistinguishability of two distributions D_1, D_2 by $D_1 \stackrel{c}{\equiv} D_2$.

4.1 OLIC: Description of the Protocol

In what follows, we describe our protocol for privacy-preserving location proximity testing with a napping Bob. We name this protocol OLIC (for OffLine InnerCircle) to stress its connection with the `InnerCircle` protocol and the fact that it can run while Bob is offline. More specifically, instead of exchanging messages directly with one another (as in `InnerCircle`), in OLIC Alice and Bob communicate with (and through) two servers.

The Ideal Functionality of OLIC. At a high level, OLIC takes as input two locations (x_A, y_A) and (x_B, y_B) , and a radius value r ; and returns 1 if the Euclidean distance between the locations is less than or equal to r , and 0 otherwise. We perform this test with inputs provided by two parties, Alice and Bob, and the help of two servers, `Server1` and `Server2`. Formally, our ideal functionality has the same input and output behavior as `InnerCircle` for Alice and Bob, but it additionally has two servers whose inputs and outputs are empty strings ε .

$$F_{\text{OLIC}}((x_A, y_A), (x_B, y_B), \varepsilon, \varepsilon) = (\mathbf{res}, \varepsilon, \varepsilon, \varepsilon),$$

$$\text{where } \mathbf{res} = \begin{cases} 1, & \text{if } (x_A - x_B)^2 + (y_A - y_B)^2 \leq r^2 \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

In our protocol we restrict the input/output spaces of the ideal functionality of Equation 4.1 to values that are meaningful to our primitives. In other words, we

require that the values x_A, y_A, x_B, y_B are admissible plaintext (for the encryption scheme employed in OLIC), and that the following values are invertible (in the ring of plaintext) $x_B, y_B, D - i$ for $i \in \{0, \dots, r^2 - 1\}$ and $D = (x_A - x_B)^2 + (y_A - y_B)^2$. We will denote the set of all suitable inputs as $\mathcal{S}_\lambda \subseteq \mathcal{M}$.

Figure 4 depicts the flow of the OLIC protocol. Figure 5 contains the detailed description of the procedures called in Figure 4.

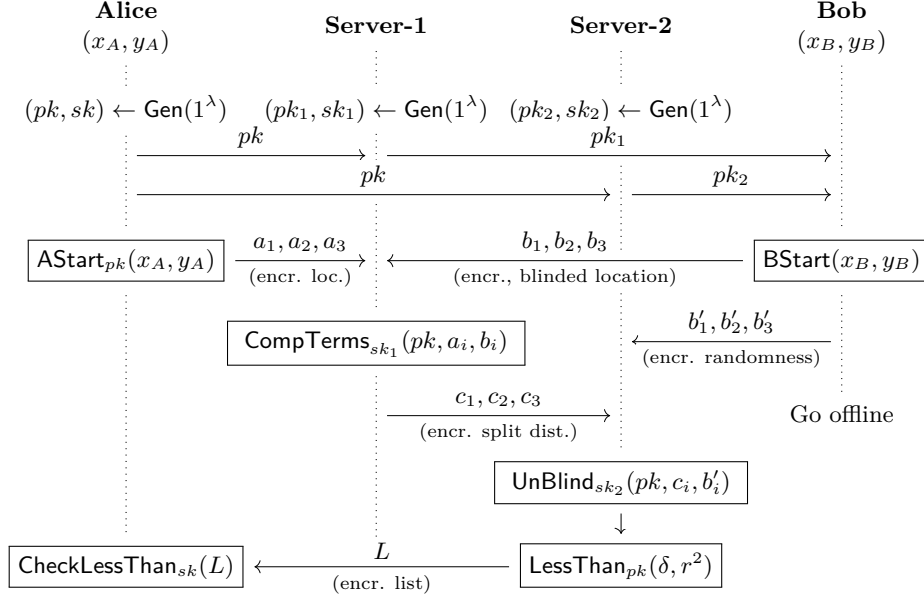


Figure 4: Overview of the message flow of our OLIC protocol. The message exchanges are grouped to reduce vertical space; in a real execution of the protocol Bob may submit at any time before **CompTerms** is run.

Step-0: Alice, Server₁, and Server₂ independently generate their keys. Alice sends her pk to Server₁ and Server₂; Server₁ and Server₂ send their respective public keys pk_1, pk_2 to Bob.

Step-1: At any point in time, Bob encodes his location using the **BStart** algorithm (Figure 5a), and sends its blinded coordinates to Server₁, and the corresponding blinding factors to Server₂.

Step-2: At any point in time Alice runs **AStart** (Figure 3a) and sends her ciphertexts to Server₁.

Step-3: once Server₁ collects Alice's and Bob's data it can proceed with computing 3 addsends useful to obtain the (squared) Euclidean distance between

their locations. To do so, Server_1 runs **CompTerms** (Figure 5b), and obtains:

$$\begin{aligned} c_1 &= \text{Enc}_{pk}(x_A^2 + y_A^2 + (x_B^2 + y_B^2 + r_1)) \\ c_2 &= \text{Enc}_{pk}(2x_A + (x_B + r_2)) \\ c_3 &= \text{Enc}_{pk}(2y_A + (y_B + r_3)) \end{aligned}$$

Server_1 sends the above three ciphertexts to Server_2 .

Step-4: Server_2 runs **UnBlind** (Figure 5c) to remove Bob’s blinding factors from c_1, c_2, c_3 and obtains the encrypted (squared) Euclidean distance between Alice and Bob:

$$\begin{aligned} \delta &= c_1 \boxplus \text{Enc}_{pk}(-r_1) \boxplus c_2 \boxminus \text{Enc}_{pk}(r_2^{-1}) \boxplus c_3 \boxminus \text{Enc}_{pk}(r_3^{-1}) \\ &= \text{Enc}_{pk}((x_A - x_B)^2 + (y_A - y_B)^2) \end{aligned}$$

Then Server_2 uses δ and the radius value r to run **LessThan** (as in **InnerCircle**, Figure 3b), obtains the list of ciphertexts L and returns L to Alice.

Step-5: Alice runs **CheckLessThan** (Figure 3c) to learn whether $D \leq r^2$ in the same way as done in **InnerCircle**.

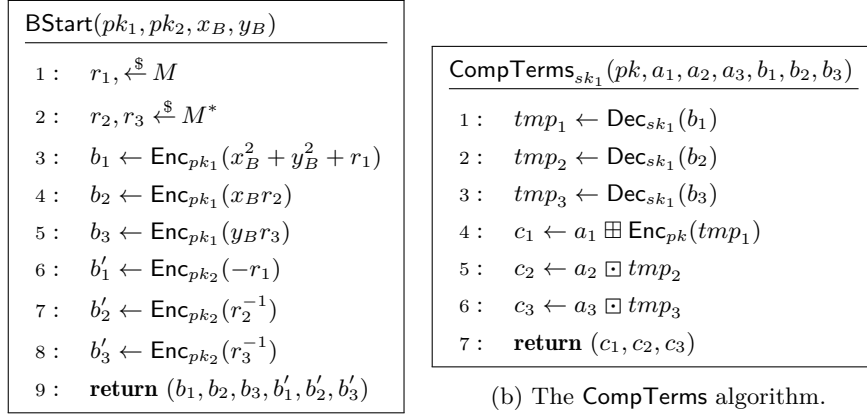
The correctness of OLIC follows from the correctness of **InnerCircle** and a straightforward computation of the input-output of **BStart**, **CompTerms**, and **UnBlind**.

4.2 OLIC: Privacy of the Protocol

We prove that our OLIC provides privacy against semi-honest adversaries. We do so by showing an efficient simulator for each party involved in the protocol and then arguing that each simulator’s output is indistinguishable from the view of respective party. More details on the cryptographic primitives and their security model can be found in [13,28].

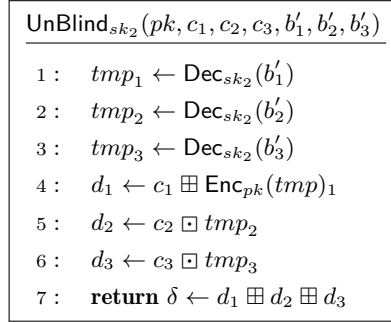
Theorem 1. *If the homomorphic encryption scheme $\text{HE} = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ used in OLIC is IND-CPA secure and function private then OLIC securely realizes the privacy-preserving location proximity testing functionality (in Equation 4.1) against semi-honest adversaries.*

At a high level, to prove Theorem 1 we construct one simulator per party and prove that each simulator’s output is computationally indistinguishable (by nonuniform algorithms) from the views of each single party in a real protocol execution. The simulators for Bob, Server_1 , and Server_2 are trivial — all the messages in their views are either uniformly distributed or encrypted, both of which is easy to simulate (any encrypted value can be simulated as an encryption of 0). The simulator for Alice is slightly more involved, it needs to generate a list of ciphertexts L whose distribution is defined by the output bit **res**: when **res** = 0 the L is a list of encryptions of random nonzero elements of \mathcal{M} , otherwise one of its elements (the position of which is chosen uniformly at random) contains an encryption of zero. Given **res** this can also be done efficiently. Due to space restriction the detailed proof of Theorem 1 can be found in the extended version of this paper [32].



(a) The BStart algorithm.

(b) The CompTerms algorithm.



(c) The UnBlind algorithm.

Figure 5: The new subroutines in OLIC.

Privacy Remark on Bob. In case the two servers collude, Bob loses privacy, in the sense that Server₁ and Server₂ together can recover Bob's location (by unblinding the tmp_i in CompTerms, Figure 5b). However Alice retains her privacy even in case of colluding servers (thanks to the security of the encryption scheme).

4.3 Security Against Malicious Servers

We describe a generic way to turn OLIC into a protocol secure against malicious, but non-colluding, servers. To do so we employ a suitable non-interactive zero knowledge proof system (NIZK) and a fairly novel cryptographic primitive called multi-key homomorphic signatures (MKHS) [10,1]. The proposed maliciously secure version of OLIC is currently more of a feasibility result rather than a concrete instantiation: to the best of our knowledge there is no combination of MKHS and NIZK that would fit our needs.

Homomorphic signatures [5,14] enable a signer to authenticate their data in such a way that any third party can homomorphically compute on it and obtain (1) the result of the computation, and (2) a signature vouching for the correctness of the latter result. In addition to what we just described, MKHS make it possible to authenticate computation on data signed by multiple sources. Notably, homomorphic signatures and MKHS can be used to efficiently verify that a computation has been carried out in the correct way on the desired data without need to access the original data [10]. This property is what we leverage to make OLIC secure against malicious servers.

At a high level, our proposed solution to mitigate malicious servers works as follows. Alice and Bob hold distinct secret signing keys, \mathbf{sk}_A and \mathbf{sk}_B respectively, and sign their ciphertexts before uploading them to the servers. In detail, using the notation in Figure 4, Alice sends to Server_1 the three messages (ciphertexts) a_1, a_2, a_3 along with their respective signatures $\sigma_i^A \leftarrow \text{MKHS.Sign}(\mathbf{sk}_A, a_i, \ell_i)$, where ℓ_i is an opportune label.³ Bob acts similarly. Server_1 computes f_i (the circuit corresponding to the computation in `CompTerms` on the i -th input) on each ciphertext and each signature, i.e., $c_i \leftarrow f(a_i, b_i)$ and $\sigma'_i \leftarrow \text{MKHS.Eval}(f, \{\mathbf{pk}_A, \mathbf{pk}_B\}, \sigma_i^A, \sigma_i^B)$. The unforgeability of MKHS ensures that each multi-key signature σ'_i acts as a proof that the respective ciphertext c'_i has been computed correctly (i.e., using f on the desired inputs⁴). Server_2 can be convinced of this fact by checking whether $\text{MKHS.Verif}(f, \{\ell_j\}, \{\mathbf{pk}_A, \mathbf{pk}_B\}, c'_i, \sigma'_i)$ returns 1. If so, Server_2 proceeds and computes the value δ (and its multi-key signature σ) by evaluating the circuit g corresponding to the function `UnBlind`. As remarked in Section 4.2, privacy demands that the random coefficients, x_i :s, involved in the `LessThan` procedure are not leaked to Alice. However, without the x_i :s Alice cannot run the MKHS verification (as this randomness should be hardwired in the circuit h corresponding to the function `LessThan`). To overcome this obstacle we propose to employ a zero knowledge proof system. In this way Server_2 can state that it knows undisclosed values x_i :s such that the output data (the list $L \leftarrow [l_1, \dots, l_{r^2-1}]$) passes the MKHS verification on the computations dependent on x_i :s. This can be achieved by interpreting the MKHS verification algorithm as a circuit v with inputs x_i (and l_i).

Security Considerations. To guarantee security we need the MKHS to be context hiding (e.g., [37], to prevent leakage of information between Server_1 and Server_2); unforgeable (in the homomorphic sense [10]); and the final proof to be zero knowledge. Implementing this solution would equip OLIC with a quite advanced and heavy machinery that enables Alice to detect malicious behaviors from the servers.

A Caveat on Labels. There is final caveat on data labels [10] needed for the MKHS schemes. We propose to set labels as a string containing the public information: day, time, identity of the user producing the location data, and

³ More details on labels at the end of the section.

⁴ The ‘desired’ inputs are indicated by the labels, as we discuss momentarily.

data type identifier (e.g., $(1, 3)$) to identify the ciphertext b_3 , sent by Bob to Server_1 , $(2, 1)$ to identify the ciphertext b'_1 , sent by Bob to Server_2 , and $(0, 2)$ to identify Alice’s ciphertext a_2 —Alice only sends data to Server_1). We remark that such label information would be retrievable by `InnerCircle` as well, as in that case Alice knows the identity of her interlocutor (Bob), and the moment (day and time) in which the protocol is run.

5 Evaluation

In this section we evaluate the performance of our proposal in three different ways: first we provide asymptotic bounds on time complexity of the algorithms in `OLIC`; second, we provide bounds on the total communication cost of the protocol; finally we develop a proof-of-concept implementation of `OLIC` to test its performance (running time and communication cost) and compare it against the `InnerCircle` protocol. Recall that `InnerCircle` and `OLIC` implement the same functionality (privacy-preserving proximity testing), however the former requires Bob to be online during the whole protocol execution while the latter does not.

Parameters. As described in Section 4.1, `OLIC` requires Alice to use an additive homomorphic encryption scheme. However, no special property is needed by the ciphertexts from Bob to the servers and between servers. Our implementation employs the ElGamal cryptosystem over a safe-prime order group for ciphertexts to and from Alice, while for the other messages it uses the Paillier cryptosystem. We refer to this implementation as **(non-EC)**, as it does not rely on any elliptic curve cryptography (ECC). In order to provide a fair comparison with its predecessor (`InnerCircle`), we additionally instantiate `OLIC` using only ECC cryptosystems (**EC**), namely Elliptic Curve ElGamal.

We note that additive homomorphic ElGamal relies on the fact that a plaintext value m is encoded into a group element as g^m (where g denotes the group generator). In this way, the multiplication of $g^{m_1} \cdot g^{m_2}$ returns an encoding of the addition of the corresponding plaintext values $m_1 + m_2$. In order to ‘fully’ decrypt a ciphertext, one would have to solve the discrete logarithm problem and recover m from g^m , which should be unfeasible, as this corresponds to the security assumption of the encryption scheme. Fortunately, this limitation is not crucial for our protocol. In `OLIC`, indeed, Alice only needs to check whether a ciphertext is an encryption of zero or not (in the `CheckLessThan` procedure), and this can be done efficiently without fully decrypting the ciphertext. However, we need to keep it into account when implementing `OLIC` with ECC. Indeed, in `OLIC` the servers need to actually fully decrypt Bob’s ciphertexts in order to work with the corresponding (blinded) plaintexts. We do so by employing a non-homomorphic version of ElGamal based on the curve M383 in (**EC**) and Paillier cryptosystem in (**non-EC**).

In our evaluation, we ignore the computational cost of initial key-generation as, in real-life scenarios, this process is a one-time set up and is not needed at every run of the protocol.

5.1 Asymptotic complexity

Table 1 shows both the concrete numbers of cryptographic operations and our asymptotic bounds on the time complexity of the algorithms executed by each party involved in OLIC. We assume that ElGamal and Paillier encryption, decryption and arithmetic operations, are performed in time $\mathcal{O}(\lambda m)$ (assuming binary exponentiation), where λ here is the security parameter and m is the cost of λ -bit interger multiplication — depends on the specific multiplication algorithm used, e.g., $m = \mathcal{O}(n \log n \log \log n)$ for Schönhage–Strassen algorithm. This applies to both **(EC)** and **(non-EC)**, although **(EC)** in practice is used with a lot smaller values of λ .

Table 1: Concrete number of operations and asymptotic time complexity for each party in OLIC (m is the cost of modular multiplication, λ the security parameter and r the radius). In our implementation Alice decryption is actually an (efficient) zero testing.

Party	Cryptosystem operations	Time bound
Alice	3 encryptions r^2 decryptions	$\mathcal{O}(r^2 \lambda m)$
Bob	6 encryptions	$\mathcal{O}(\lambda m)$
Server ₁	3 decryptions, 3 homomorphic operations,	$\mathcal{O}(\lambda m)$
Server ₂	3 decryptions, $2r^2 + 5$ arithmetic operations, r^2 encryptions	$\mathcal{O}(r^2 \lambda m)$

Communication cost. The data transmitted among parties during a whole protocol execution amounts to $r^2 + 6$ ciphertexts between Alice and the servers and 6 ciphertexts between Bob and servers. Each ciphertext consists of 4λ bits in case of **(EC)**, and 2λ bits in case of **(non-EC)** (for both ElGamal and Paillier). Asymptotically, both cases require $\mathcal{O}(\lambda r^2)$ bit of communication—although **(EC)** is used with a lot smaller λ value in implementation.

5.2 Implementation

We developed a prototype implementation of OLIC in Python (available at [32]). The implementation contains all of the procedures shown in pseudocode in Figures 3 and 5. To ensure a fair compare between OLIC and InnerCircle, we implemented latter in the same environment (following the nomenclature in Figure 2).

Our benchmarks for the total communication cost of the OLIC protocol are reported in Figure 6. We measured the total execution time of each procedure of both InnerCircle and OLIC. Figure 7 shows the outcome of our measurements for values of the proximity radius parameter r ranging from 0 to 100. Detailed values can be found in extended version of this paper [32].

Setup. We used cryptosystems from the cryptographic library `bettertimes` [19], which was used for benchmarking of the original `InnerCircle` protocol. The benchmarks were run on Intel Core i7-8700 CPU running at a frequency of 4.4 GHz. For **(non-EC)** we used 2048-bit keys for ElGamal (the same as in `InnerCircle` [20]), and 2148-bit keys for Paillier to accommodate a modulus larger than the ElGamal modulus (so that any possible message of Bob fits into the Paillier message space). For **(EC)** we used curves `Curve25519` for additive homomorphic encryption and `M383` for the ciphertexts exchanged among Bob and the servers from the `ecc-pycrypto` library [7]. We picked these curves because they were available in the library and also because `M383` uses a larger modulus and allows us to encrypt the big values from the plaintexts field of `Curve25519`.

The plaintext ring for **(non-EC)** ElGamal-2048 has at least $|\mathcal{M}| \geq 2^{2047}$ elements, which allows Alice and Bob’s points to lie on the grid $\{1, 2 \dots 2^{1023}\}^2$ ensuring that the squared distance between them never exceeds 2^{2047} . The corresponding plaintext ring size for **(EC)** is $|\mathcal{M}| \geq 2^{251}$ (the group size of `Curve25519`), and the grid is $\{1, 2 \dots 2^{125}\}^2$. Since Earth equator is $\approx 2^{26}$ meters long, either of the two grids is more than enough to cover any location on Earth with 1 meter precision.

The plaintext ring for **(non-EC)** ElGamal-2048 has at least $|\mathcal{M}| \geq 2^{2047}$ elements, which allows Alice and Bob’s points to lie on the grid $\{1, 2 \dots 2^{1023}\}^2$ ensuring that the squared distance between them never exceeds 2^{2047} . The corresponding plaintext ring size for **(EC)** is $|\mathcal{M}| \geq 2^{251}$ (the group size of `Curve25519`), and the grid is $\{1, 2 \dots 2^{125}\}^2$. Since Earth equator is $\approx 2^{26}$ meters long, either of the two grids is more than enough to cover any location on Earth with 1 meter precision.

Optimizations. In `InnerCircle` [20], the authors perform three types of optimizations on the `LessThan` procedure:

- (O-1)** Iterating only through those values of $i \in \{0, \dots, r^2 - 1\}$ which can be represented as a sum of two squares, i.e., such i that $\exists a, b : i = a^2 + b^2$.
- (O-2)** Precomputing the ElGamal ciphertexts $\text{Enc}_{pk}(-i)$, and only for those i described in **(O-1)**.
- (O-3)** Running the procedure in parallel, using 8 threads.

We adopt the optimizations **(O-1)** and **(O-2)** in our implementation as well. Note that **(O-1)** reduces the length of list L (Figure 4), as well as the total communication cost. We disregard optimization **(O-3)** since we are interested in the total amount of computations a party needs to do (thus this optimization is not present in our implementations of `OLIC` and `InnerCircle`).

5.3 Performance Evaluation

Figure 7 shows a comparison of the total running time of each party in `OLIC` versus `InnerCircle`. One significant advantage of `OLIC` is that it offloads the execution of the `LessThan` procedure from Bob to the servers. This is reflected in




Figure 6: Total communication cost of `OLIC`.




Figure 7: Running times of each party in InnerCircle and OLIC for both (**non-EC**) and (**EC**) instantiations. (Reported times are obtained as average of 40 executions.)

Figure 7, where the running time of Bob is flat in OLIC, in contrast to quadratic as in InnerCircle. Indeed, the combined running time of the two servers in OLIC almost matches the running time of Bob in InnerCircle. Note that the servers have to spend total 10-12 seconds on computations when $r = 100$, which is quite a reasonable amount of time, given that servers are usually not as resource-constrained as clients, and that the most time-consuming procedure `LessThan` consists of a single loop which can easily be executed in parallel to achieve even better speed. Finally, we remark that the amount of data being sent (Figure 6) between the parties in OLIC is quite moderate.

In case Alice wants to be matched with multiple Bobs, say, with k of them, the amount of computations that she and the servers perform will grow linearly with k . The same applies to the communication cost. Therefore, one can obtain the counterparts of Figures 7 and 6 for multiple Bobs by simply multiplying all the plots (except Bob’s computations) by k .

6 Related Work

Zhong et al. [44] present the Louis, Lester and Pierre protocols for location proximity. The Louis protocol uses additively homomorphic encryption to compute the distance between Alice and Bob while it relies on a third party to perform the proximity test. Bob needs to be present online to perform the protocol. The Lester protocol does not use a third party but rather than performing proximity testing computes the actual distance between Alice and Bob. The Pierre protocol resorts to grids and leaks Bob’s grid cell distance to Alice.

Hide&Crypt by Freni et al. [11] splits proximity in two steps. Filtering is done between a third party and the initiating principal. The two principals then execute computation to achieve finer granularity. In both steps, the granule which a principal is located is sent to the other party. C-Hide&Hash by Mascetti et al. [29] is a centralized protocol, where the principals do not need to communicate pairwise but otherwise share many aspects with Hide&Crypt. FriendLocator by Šikšnys et al. [42] presents a centralized protocol where clients map their position to different granularities, similarly to Hide&Crypt, but instead of refining via the second principal each iteration is done via the third party. VicinityLocator also by Šikšnys et al. [41] is an extension of FriendLocator, which allows the proximity of a principal to be represented not only in terms of any shape.

Narayanan et al. [30] present protocols for proximity testing. They cast the proximity testing problem as equality testing on a grid system of hexagons. One of the protocol utilizes an oblivious server. Parties in this protocol use symmetric encryption, which leads to better performance. However, this requires to have preshared keys among parties, which is less amenable to one-to-many proximity testing. Saldamli et al. [36] build on the protocol with the oblivious server and suggest optimizations based on properties from geometry and linear algebra. Nielsen et al. [31] and Kotzanikolaou et al. [26] also propose grid-based solutions.

Šeděnka and Gasti [38] homomorphically compute distances using the UTM projection, ECEF (Earth-Centered Earth-Fixed) coordinates, and the Haversine formula that make it possible to consider the curvature of the Earth. Hallgren et al. [20] introduce InnerCircle for parallizable decentralized proximity testing, using additively homomorphic encryption between two parties that must be online. The MaxPace [21] protocol builds on speed constraints of an InnerCircle-style protocol as to limit the effects of trilateration attacks. Polakis [34] study different distance and proximity disclosure strategies employed in the wild and experiment with practical effects of trilateration.

Sakib and Huang [35] explore proximity testing using elliptic curves. They require Alice and Bob to be online to be able to run the protocol. Järvinen et al. [24] design efficient schemes for Euclidean distance-based privacy-preserving location proximity. They demonstrate performance improvements over InnerCircle. Yet the requirement of the two parties being online applies to their setting as well. Hallgren et al. [18] how to leverage proximity testing for endpoint-based ridesharing, building on the InnerCircle protocol, and compare this method with a method of matching trajectories.

The computational bottle neck of privacy-preserving proximity testing is the input comparison process. Similarly to [20,33], we rely on homomorphic encryption to compare a private input (the distance between the submitted locations) with a public value (the threshold). Other possible approaches require the use of the arithmetic black box model [6], garbled circuits [25], generic two party computations [12], or oblivious transfer extensions [8].

To summarize, the vast majority [44,42,41,11,30,36,38,20,21,35,24] of the existing approaches to proximity testings require both parties to be online,

thus not being suitable for one-to-many matching. A notable exception to the work above is the C-Hide&Hash protocol by Mascetti et al. [29], which allows one-to-many testing, yet at the price of not computing the precise proximity result but its grid-based approximation. Generally, a large number of approaches [44,42,41,11,29,30,31,26] resort to grid-based approximations, thus losing precision of proximity tests.

There is a number of existing works, which consider the problem of computing generic functions in the setting, where clients are not online during the whole execution. Hallevi et al. [17] consider a one-server scenario and show that the notion of security against semi-honest adversary (which we prove for our protocol) is impossible to achieve with one server. Additionally, the model from [17] lets all the parties know each other’s public keys, i.e., the clients know all the other clients who supply inputs for the protocol—this does not allow one-to-many matching, which we achieve in our work. Further works [23,15,16,2,3] also consider one-server scenarios.

7 Conclusions

We have presented OLIC, a protocol for privacy-preserving proximity testing with a mapping party. In line with our goals, (1) we achieve privacy with respect to semi-honest parties; (2) we enable matching against offline users which is needed in scenarios like ridesharing; (3) we retain precision, not resorting to grid-based approximations, and (4) we reduce the responding client overhead from quadratic (in the proximity radius parameter) to constant.

Future work avenues include developing a fully-fledged ridesharing system based on our approach, experimenting with scalability, and examining the security and performance in the light of practical security risks for LBS services.

Acknowledgments. This work was partially supported by the Swedish Foundation for Strategic Research (SSF), the Swedish Research Council (VR) and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under grant agreement No 669255 (MPCPRO). The authors would also like to thank Carlo Brunetta for insightful discussion during the initial development of the work.

References

1. Aranha, D.F., Pagnin, E.: The simplest multi-key linearly homomorphic signature scheme. In: LATINCRYPT. pp. 280–300. Springer (2019)
2. Beimel, A., Gabizon, A., Ishai, Y., Kushilevitz, E., Meldgaard, S., Paskin-Cherniavsky, A.: Non-interactive secure multiparty computation. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO. LNCS, vol. 8617, pp. 387–404. Springer (2014)
3. Benhamouda, F., Krawczyk, H., Rabin, T.: Robust non-interactive multiparty computation against constant-size collusion. In: Katz, J., Shacham, H. (eds.) CRYPTO. LNCS, vol. 10401, pp. 391–419. Springer (2017)

4. BlaBlaCar - Trusted carpooling. <https://www.blablacar.com/>
5. Boneh, D., Freeman, D.M.: Homomorphic signatures for polynomial functions. In: EUROCRYPT. pp. 149–168. Springer (2011)
6. Catrina, O., De Hoogh, S.: Improved primitives for secure multiparty integer computation. In: SCN. pp. 182–199. Springer (2010)
7. ChangLiu: Ecc-pycrypto Library (2019 (accessed April 14, 2020)), <https://github.com/lc6chang/ecc-pycrypto>
8. Couteau, G.: New protocols for secure equality test and comparison. In: ACNS. pp. 303–320. Springer (2018)
9. Cuéllar, J., Ochoa, M., Rios, R.: Indistinguishable regions in geographic privacy. In: SAC. pp. 1463–1469 (2012)
10. Fiore, D., Mitrokotsa, A., Nizzardo, L., Pagnin, E.: Multi-key homomorphic authenticators. In: ASIACRYPT. pp. 499–530. Springer (2016)
11. Freni, D., Vicente, C.R., Mascetti, S., Bettini, C., Jensen, C.S.: Preserving location and absence privacy in geo-social networks. In: CIKM. pp. 309–318 (2010)
12. Garay, J., Schoenmakers, B., Villegas, J.: Practical and secure solutions for integer comparison. In: PKC. pp. 330–342. Springer (2007)
13. Gentry, C., Halevi, S., Vaikuntanathan, V.: i-hop homomorphic encryption and rerandomizable yao circuits. In: Annual Cryptology Conference. pp. 155–172. Springer (2010)
14. Gorbunov, S., Vaikuntanathan, V., Wichs, D.: Leveled fully homomorphic signatures from standard lattices. In: STOC. pp. 469–477. ACM (2015)
15. Gordon, S.D., Malkin, T., Rosulek, M., Wee, H.: Multi-party computation of polynomials and branching programs without simultaneous interaction. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT. LNCS, Springer (2013)
16. Halevi, S., Ishai, Y., Jain, A., Komargodski, I., Sahai, A., Yorgev, E.: Non-interactive multiparty computation without correlated randomness. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT. LNCS, vol. 10626, pp. 181–211. Springer (2017)
17. Halevi, S., Lindell, Y., Pinkas, B.: Secure computation on the web: Computing without simultaneous interaction. In: CRYPTO. pp. 132–150 (2011)
18. Hallgren, P., Orlandi, C., Sabelfeld, A.: PrivatePool: Privacy-Preserving Ridesharing. In: CSF. pp. 276–291 (Aug 2017)
19. Hallgren, P.: BetterTimes Python Library (2017 (accessed January 22, 2020)), <https://bitbucket.org/hallgrep/bettertimes/>
20. Hallgren, P.A., Ochoa, M., Sabelfeld, A.: InnerCircle: A parallelizable decentralized privacy-preserving location proximity protocol. In: PST. pp. 1–6 (2015)
21. Hallgren, P.A., Ochoa, M., Sabelfeld, A.: MaxPace: Speed-Constrained Location Queries. In: CNS (2016)
22. Hern, A.: Uber employees 'spied on ex-partners, politicians and Beyoncé' (2016), <https://www.theguardian.com/technology/2016/dec/13/uber-employees-spying-ex-partners-politicians-beyonce>
23. Jarrous, A., Pinkas, B.: Canon-mpc, a system for casual non-interactive secure multi-party computation using native client. In: Sadeghi, A., Foresti, S. (eds.) WPES. pp. 155–166. ACM (2013)
24. Järvinen, K., Kiss, Á., Schneider, T., Tkachenko, O., Yang, Z.: Faster privacy-preserving location proximity schemes. In: CANS. LNCS, vol. 11124, pp. 3–22. Springer (2018)
25. Kolesnikov, V., Sadeghi, A.R., Schneider, T.: Improved garbled circuit building blocks and applications to auctions and computing minima. In: CANS. pp. 1–20. Springer (2009)

26. Kotzanikolaou, P., Patsakis, C., Magkos, E., Korakakis, M.: Lightweight private proximity testing for geospatial social networks. *Computer Communications* **73**, 263–270 (2016)
27. Lee, D.: Uber concealed huge data breach (2017), <http://www.bbc.com/news/technology-42075306>, <http://www.bbc.com/news/technology-42075306>
28. Lindell, Y.: How to simulate it—a tutorial on the simulation proof technique. In: *Tutorials on the Foundations of Cryptography*, pp. 277–346. Springer (2017)
29. Mascetti, S., Freni, D., Bettini, C., Wang, X.S., Jajodia, S.: Privacy in geo-social networks: proximity notification with untrusted service providers and curious buddies. *VLDB J.* **20**(4), 541–566 (2011)
30. Narayanan, A., Thiagarajan, N., Lakhani, M., Hamburg, M., Boneh, D.: Location privacy via private proximity testing. In: *NDSS* (2011)
31. Nielsen, J.D., Pagter, J.I., Stausholm, M.B.: Location privacy via actively secure private proximity testing. In: *PerCom Workshops*. pp. 381–386. IEEE CS (2012)
32. Oleynikov, I., Pagnin, E., Sabelfeld, A.: Where are you Bob? Privacy-preserving proximity testing with a napping party (2020), <https://www.cse.chalmers.se/research/group/security/olic/>
33. Pagnin, E., Gunnarsson, G., Talebi, P., Orlandi, C., Sabelfeld, A.: TOPPool: Time-aware Optimized Privacy-Preserving Ridesharing. *PoPETs* **2019**(4), 93–111 (2019)
34. Polakis, I., Argyros, G., Petsios, T., Sivakorn, S., Keromytis, A.D.: Where’s wally?: Precise user discovery attacks in location proximity services. In: *CCS* (2015)
35. Sakib, M.N., Huang, C.: Privacy preserving proximity testing using elliptic curves. In: *ITNAC*. pp. 121–126. IEEE Computer Society (2016)
36. Saldamli, G., Chow, R., Jin, H., Knijnenburg, B.P.: Private proximity testing with an untrusted server. In: *WISEC*. pp. 113–118. ACM (2013)
37. Schabhüser, L., Butin, D., Buchmann, J.: Context hiding multi-key linearly homomorphic authenticators. In: *CT-RSA*. pp. 493–513. Springer (2019)
38. Sedenka, J., Gasti, P.: Privacy-preserving distance computation and proximity testing on earth, done right. In: *AsiaCCS*. pp. 99–110 (2014)
39. Sharemind MPC Platform. <https://sharemind.cyber.ee/sharemind-mpc/multi-party-computation/>
40. Shu, C.: Uber reportedly tracked Lyft drivers using a secret software program named ‘Hell’. <https://techcrunch.com/2017/04/12/hell-o-uber/> (2017)
41. Siksnyš, L., Thomsen, J.R., Saltenis, S., Yiu, M.L.: Private and flexible proximity detection in mobile social networks. In: *MDM*. pp. 75–84 (2010)
42. Siksnyš, L., Thomsen, J.R., Saltenis, S., Yiu, M.L., Andersen, O.: A location privacy aware friend locator. In: *SSTD*. pp. 405–410 (2009)
43. Veytsman, M.: How I was able to track the location of any tinder user (February 2014), <http://blog.includesecurity.com/>
44. Zhong, G., Goldberg, I., Hengartner, U.: Louis, lester and pierre: Three protocols for location privacy. In: *PET*. pp. 62–76 (2007)