

Resilient Privacy Protection for Location-based Services through Decentralization

HONGYU JIN, Networked Systems Security Group, KTH Royal Institute of Technology, Sweden

PANOS PAPADIMITRATOS, Networked Systems Security Group, KTH Royal Institute of Technology, and RISE SICS, Sweden

Location-based Services (LBSs) provide valuable services, with convenient features for mobile users. However, the location and other information disclosed through each query to the LBS erodes user privacy. This is a concern especially because LBS providers can be *honest-but-curious*, collecting queries and tracking users' whereabouts and infer sensitive user data. This motivated both *centralized* and *decentralized* location privacy protection schemes for LBSs: anonymizing and obfuscating LBS queries to not disclose exact information, while still getting useful responses. Decentralized schemes overcome disadvantages of centralized schemes, eliminating anonymizers, and enhancing users' control over sensitive information. However, an insecure decentralized system could create serious risks beyond private information leakage. More so, attacking an improperly designed decentralized LBS privacy protection scheme could be an effective and low-cost step to breach user privacy. We address exactly this problem, by proposing security enhancements for mobile data sharing systems. We protect user privacy while preserving accountability of user activities, leveraging pseudonymous authentication with mainstream cryptography. We show our scheme can be deployed with off-the-shelf devices based on an experimental evaluation of an implementation in a static automotive testbed.

Additional Key Words and Phrases: Location privacy, honest-but-curious, pseudonymous authentication

1 INTRODUCTION

A Location-based Service (LBS) query targets a location/region and expresses one or more specific user interests; the LBS server responds with the most up-to-date relevant information, e.g., the latest menu of a restaurant, movies at a cinema, remaining parking slots at a shopping mall, or traffic conditions in the area. During this process, users' current or future whereabouts and interests are disclosed to the LBS server through their queries. All submitted information is deemed necessary to best serve users, and the LBS server is entrusted with rich data. However, many studies [5, 15, 62] reveal that service providers can be honest-but-curious, aggressively collecting information to profile users, e.g., identifying home or working places or inferring interests for commercial purposes.

LBS privacy is studied extensively. Location k -anonymity [16] ensures that at least $k - 1$ other users are involved in an obfuscated region, \mathcal{R} , used as the querier's location. Therefore, even in the presence of a local observer in \mathcal{R} , the query cannot be linked to a certain user; the LBS server only learns the querier is one among k users in \mathcal{R} . Protection can be achieved by centralized schemes [16, 45, 46] that introduce an anonymizer, a proxy between users and the LBS server, which anonymizes user queries before sending them to the LBS server. However, the assumed anonymizer trustworthiness merely "shifts" the trust from the LBS server to the anonymizer, holding rich information the same way that the LBS server would. Simply put, an anonymizer could be itself honest-but-curious.

Decentralized approaches [19, 53] can eliminate the need for an anonymizer and protect user privacy in a collaborative manner: for example, form an obfuscated area with k users within each other's communication range [19]. However, if such k users are too close, e.g., in a church, a shopping mall, or a cinema, such symbolic "addresses" can still be disclosed. Thus, it is hard to define how large k should be to ensure an appropriate level of protection.

An alternative collaborative privacy protection approach is to pass/share LBS-obtained information among users, to decrease exposure to the LBS server [31, 55]. The sharing approach requires

nodes to cache information received from the LBS and pass it to neighbors when requested. Moreover, sharing (and thus reduced exposure to the LBS server) is orthogonal to location obfuscation; the two could complement each other.

Nonetheless, opening up the system functionality is a double-edged sword: sharing reduces user exposure to the curious provider (LBS or anonymizer) but it also exposes her to possibly faulty, misbehaving or curious peers. In fact, studies on peer-to-peer (P2P) systems [32, 40, 63] show that insecure decentralized schemes face serious problems: sensitive information could be exposed to peers and malicious nodes could pass on bogus data.

Signed LBS server responses can be self-verifiable when passed to peers [55] and partially address the second above-mentioned concern regarding malicious nodes. However, queries and cached information from different users could be diverse, making it necessary to share one or multiple complete LBS-obtained responses (each with a signature attached), even though only a subset of each LBS-obtained response might be needed by the querying peer. For example, the query could be *What is the menu of restaurant A?* and *How many parking spaces are available at B?* while another node could query *What are the opening hours of restaurant A?* and *What stores are in shopping mall B?* It is not straightforward to decide whether the peer responses include information as if the response were directly obtained from the LBS server, thus, guaranteeing quality of service. Even though nodes can assess the completeness of gathered peer responses [44] and wait until the required completeness is met, peer responses equivalent to an LBS response can be hard to achieve; and, in any case, incur significant delays and overhead, i.e., redundant cryptographic fields and information. Last but not least, peer queries, openly submitted to nodes' neighborhood, could expose users to other nodes and passive eavesdroppers.

These challenges are addressed here by our security architecture for secure and resilient decentralized/collaborative privacy protection for LBSs. We propose new components orthogonal to the LBS functionality. We leverage pseudonymous authentication to provide privacy-enhancing message authentication and integrity for communication with other users/peers and with infrastructure entities. We also leverage proactive caching of Point of Interest (POI) data (e.g., for a region) by a small fraction of users (termed serving nodes) that serve others, sharing the cached POI data. This ensures that peer responses can provide the same quality as direct LBS responses without extra communication incurred by using an obfuscated area. This also enables cross-checking multiple peer responses to the same query, to detect false responses from malicious nodes. The burden is balanced among users through a periodical randomized role assignment by the infrastructure. While users benefit from the information sharing system, our scheme minimizes their exposure to the LBS server and limits exposure to curious serving nodes (thanks to encrypted querying and serving node communication). Our evaluation shows both effective exposure reduction and highly successful valid POI provision in a realistic intelligent transportation setting, even with a very high (in practice) fraction (e.g., 20 %) of curious or malicious nodes.

The contributions of this paper are: (1) use of pseudonymous authentication for privacy-preserving user authentication, (2) privacy-enhancing randomized serving node assignment and encrypted P2P query-response exchange, (3) cross-checking of P2P responses and privacy-preserving proactive LBS response validation, (4) qualitative and extensive quantitative simulation-based security and privacy analyses, (5) comparison to the state-of-the-art, MobiCrowd [55], in terms of achieved privacy and security, and efficiency (notably, communication overhead), and (6) experimental evaluation leveraging a field-operational-test grade automotive communication testbed, emulating real-world conditions.

The rest of the paper is organized as follows: We discuss related work in Sec. 2. We outline the system and adversarial models, as well as requirements in Sec. 3. Then, we present the proposed scheme in Sec. 4. We provide a qualitative security and privacy analysis in Sec. 5.1, followed by an

extensive quantitative simulation-based evaluation in Sec. 5.2. We show experimental evaluation results with our scheme implemented on an automotive testbed, and discuss design choices for its deployment in real-world scenarios (Sec. 6) before we conclude (Sec. 7).

2 RELATED WORK

User queries, including user locations and interests, are essential information for LBSs to best serve users; this very same information can also be used to breach user privacy. Queries from the same user can be linked and eventually used to reconstruct user traces and infer user activities, e.g., home or work addresses. Anonymization or pseudonymization, i.e., removal of user identities or their replacement with pseudonymous identities, prevents the LBS server from correlating real user identities with queries. The LBS server only sees a pool of user queries without any information that immediately identifies the querying users. However, studies have shown that simply removing identities from frequent queries could not prevent tracking and even profiling (e.g., inferring real names) of specific users [24, 26, 28].

Location obfuscation: Location obfuscation uses cloaked areas or perturbed locations to represent actual user locations. Cloaked areas [4, 9] can be used so that the LBS server only learns the area the user is in and then responds with the information that pertains to the whole area. However, a determined attacker could learn which users are located in this area. If, in an arguably rare situation, the user is the only one included in this area, the attacker can easily link the query with the only user “seen” in this area. Location k -anonymity [23, 37] is used to form an obfuscated area involving at least k users, so that the user would not stand out from the obfuscated area. However, merely using k -anonymity could not guarantee an area constructed based on k users large enough to hide symbolic address (e.g., church, shopping mall and gas station). Therefore, in densely populated areas, the perturbed/obfuscated region should be large enough so that all nodes within the area are not closely gathered. At the same time, the quality of the LBS response would be affected if the location in the query is not accurate enough. Thus, it is hard to define how large should k be to ensure an appropriate level of privacy protection. Geo-indistinguishability [1, 3, 13] uses perturbed locations to represent user locations, adding Laplacian noise to user locations based on the required degree of privacy. This trades off accuracy of POI data for higher privacy. Both approaches could fail to protect node privacy when multiple queries can be correlated [1, 56].

Moreover, schemes based on k -anonymity can be vulnerable to the so-called center of Anonymized Spatial Region (ASR) attack: the probability of the requesting node being the node closest to the center of ASR is larger than $1/k$ [19, 33]. Such an attack can be prevented by applying a k -anonymity algorithm to one of the neighbors. These neighbors are discovered by applying the k -anonymity algorithm to itself. This generates a cloaked region that is not necessarily centered around oneself.

Centralized and decentralized approaches: Both centralized and decentralized variants of the above-mentioned privacy enhancing techniques were proposed. Centralized approaches introduce a Trusted Third Party (TTP) that acts as a so-called anonymizer for the queries sent to the LBS server. The anonymizer could find $k - 1$ nearest neighbors of the querying user and send the queries from all k users to the LBS server [4]. The anonymizer can also remove user identities from the queries before sending them to the LBS server, while keeping information that matches user identities and queries locally, so that it can appropriately respond to each user the needed information. However, there can be an inherent controversy here: we assume the LBS server is honest-but-curious, but the anonymizer could be honest-but-curious too. The (centralized) anonymizer could collect user queries and use the information the same way the LBS server would.

Decentralized approaches thwart honest-but-curious infrastructure by relying on users. To achieve k -anonymity with decentralized schemes, users can search for $k - 1$ neighbors in an ad-hoc network and use the area including these k users as the cloaked area [10, 17–19, 27, 29].

However, this incurs high communication overhead among peers: a node needs to initiate the process whenever it wants to query the LBS server. To reduce overhead and complexity of peers, anonymizer peers can be used, that act as a temporary proxy for their neighbors [53]. Therefore, users do not need to search for neighbors everytime. AMOEBA [53] protects user privacy by forming groups and delegating LBS queries to anonymizer peers, termed group leaders. The group members send the queries to the group leaders that act as the TTP in centralized schemes: the group leaders query the LBS server with collected queries and distribute the responses to each group member. Such group formation and provision of k -anonymity (both centralized and decentralized) are orthogonal to our work here. In fact, they could possibly co-exist with and even be facilitated by our scheme, which explicitly addresses trust assumptions and provides a security architecture.

Caching and sharing: Caching and sharing POI data [39, 42, 44, 48, 55] are orthogonal to above-mentioned protection approaches. Cached POI data can serve future queries and reduce direct queries to the LBS server. Road-side infrastructure can proactively cache POI data [42, 48] and serve user queries; however, users would be vulnerable to honest-but-curious road-side infrastructure entities. Nodes can cache information received from the LBS server and pass to neighbors when requested [39, 44, 55]. This is the approach we extend in this paper. Gathered partial POI data from neighbors can form more complete POI data [39, 44], but it is hard to examine whether the gathered POI data covers all that could be obtained directly from the LBS server. Moreover, such approaches are vulnerable to data pollution by malicious or semi-honest nodes, providing masqueraded POI data or partial data (from available own data) to mislead other nodes. Responses signed by the LBS server [55] can be self-verifiable allowing detection of manipulation by a node. However, a misbehaving node passing on tampered responses would remain “invisible” and continue attacking the system, wasting computational power [31]. Moreover, decentralizing functionality could provide an easy way for passive adversaries to learn information from neighbors. Thus, it is important for collaborative approaches to consider internal and external passive or active adversaries. This is exactly where this work comes in: providing a secure and privacy-enhancing LBS while protecting benign nodes against both honest-but-curious and malicious nodes/peers.

Location privacy in other areas: Location privacy has also been widely studied in relevant areas. SPPEAR [20] and SHIELD [21] propose a state of the art architecture and protocols for secure and privacy preserving participatory sensing. Users obtain short-term credentials and contribute sensing data pseudonymously to a central aggregator [20]. The aggregator makes use of the large volume of received sensed data to filter out outliers [21]. The aggregated data is made available through queries submitted to the aggregator; this can be done securely and privately with our proposal here.

In Vehicular Ad-hoc Networks (VANETs), privacy is important, because high rate vehicle-transmitted safety beacons with location information for awareness among neighboring vehicles, expose vehicle activities. More importantly, vehicle messages need to be authenticated, but using traditional certificates would trivially allow linking messages sent from a vehicle. This motivates the use of pseudonymous authentication [50] for message unlinkability while providing message and entity authentication. The most recent proposals protect users from honest-but-curious credential providers and investigate deployment and scalability aspects [2, 22, 34, 35]. In our scheme, pseudonymous authentication is an important component for providing secure and privacy-preserving LBS. We explain how this is customized and integrated in our design in Sec. 4.2.

Query privacy: Apart from location privacy, query privacy is an important issue for LBS: the content of queries could disclose user interests. Location privacy protection schemes based on k -anonymity could potentially protect query privacy: with at least k queries sent together to the LBS server, the LBS server could not figure out which node each query belongs to. However, if a specific POI type has higher popularity than other types, then multiple (among the k) nodes could




Fig. 1. System Architecture (Icons by Freepik, freepik.com)

be interested in the same POI type; this reduces the query diversity of these k queries [12, 43, 51]. Query l -diversity [12] proposes to form a cloaked region involving user queries that have at least l POI types. However, this approach fails to provide location and query privacy protection if a node continuously queries for the same POI type during its trip: intersection of multiple (spatially and temporally) consecutive k -anonymized and l -diversified query sets could result in a pair of distinct node and POI type (i.e., the querying node and the queried POI type). An anonymizer can take past queries into consideration and form the consecutive cloaking areas involving the same set of POI types [12, 51]. Private retrieval of POIs [30, 61] searches an encrypted POI database with anonymized queries that do not disclose any query content to the LBS server. We note that such query privacy protection approaches are orthogonal to sharing-based approaches (including our scheme), which decrease the ratio of queries disclosed to the LBS server, thus can complement each other.

3 SYSTEM MODEL AND REQUIREMENTS

System Model: Fig. 1 illustrates the considered system architecture. Mobile devices (termed *nodes* or *peers* in the rest of the paper), e.g., smartphones and vehicular On-board Units (OBUs), are equipped with multiple communication interfaces, e.g., Wi-Fi and cellular data. They can access LBSs, submitting queries regarding their current locations/regions. They also communicate in a P2P manner over a wireless ad hoc (e.g., IEEE 802.11p) or cellular (e.g., LTE direct) network. Nodes can share POI data and choose to query the LBS server only when no satisfactory response is received from their peers. Nodes are registered with an *identity and credential management* system (i.e., a Public-Key Infrastructure (PKI)) comprising Certification Authorities (CAs) (see Sec. 4 for more details). The PKI issues credentials to the registered nodes and the Service Providers (SPs) (i.e., the LBS servers here), so that SPs and nodes can interact securely. We assume nodes have connectivity to the LBS and the PKI through Internet (e.g, via cellular data, 3G or 4G, or Wi-Fi networks) throughout their trips.

Adversary Model: We assume LBS servers are honest-but-curious: they follow the protocols, responding faithfully to queries, but they may trace the nodes (linking their queries) or even de-anonymize them and infer sensitive data (e.g., home and work sites). This is because queries sent to the LBS servers expose user locations and interests, and can be used to infer additional user information. We maintain and extend the honest-but-curious assumption to cover any TTP, including the ones we introduce in our scheme, notably the PKI (Sec. 4).

Nodes can be also honest-but-curious or deviate from the protocols. Honest-but-curious nodes could stay in the network and collect peer queries and responses. A single honest-but-curious node could have negligible effect on the privacy of legitimate nodes. However, if such honest-but-curious nodes collaborate and merge their transcripts of all queries they overheard from peers in the system, it is very likely they are powerful enough to threaten user privacy. An adversary could

recruit otherwise benign participants (e.g., with some incentives) and then fetch and merge their query and response transcripts; or user equipment (OBUs, smartphones) could be compromised by malware that propagates across the network and enable the adversary access such transcripts. Furthermore, nodes can deviate from the collaborative protocol functionality and policies and attack the system, notably their peers. They can forge or tamper with responses and masquerade other nodes. This could, in turn, affect the quality of service, i.e., the timeliness and authenticity of POI data obtained from peers, and thus force benign nodes, which may get no useful POI data, to query directly (i.e., expose themselves to) the LBS server(s).

Requirements: We require that peer-provided information be verifiable and nodes be accountable for their actions (messages). Nodes should be able to efficiently obtain POI data from their peers with the same quality as that obtained directly from the LBS server. While nodes benefit from P2P POI data sharing, node exposure to neighboring, both assisting and non-assisting, peers and the LBS server should be minimized. Towards these objectives, the following security and privacy requirements need to be met:

Authentication and integrity - Node messages should allow their receivers to corroborate the legitimacy of their senders (Note: not their identity; please see anonymity/pseudonymity below.) and verify they were not modified or replayed.

Accountability - Message senders should not be able to deny having sent a message (non-repudiation). Any node can be tied to its actions, and, if need arises, be held accountable and possibly have its long-term identity revealed and have itself evicted from the system.

Anonymity/Pseudonymity and unlinkability - Node actual (long-term) identities should not be linked to their P2P and node-to-LBS messages. Anonymity should be conditional, allowing the system to identify a misbehaving node and evict it. Ideally, it should be impossible for any observer to link any two or more messages (e.g., queries) by the same node. However, for efficiency reasons, messages can be linkable at most over a protocol selectable period, τ .

Confidentiality and reduced exposure - POI data and sensitive user information (e.g., node queries) should be accessible only by authorized entities; the amount of information revealed to peers and the LBS server should be minimal.

Resilience - Nodes should be resilient to (compromised) malicious nodes that actively disrupt their operations and deviate from protocol definitions. In particular, they should be able to validate correctness of information to reject bogus POI data generated by malicious nodes.

Sybil-resistance - A registered node should be able to participate only with a single identity (pseudonym) at any point in time. It should not be able to be present in the system with multiple legitimate identities (thus, as multiple participants) and inappropriately affect the outcome of protocols.

4 OUR SCHEME

In this section, we introduce our decentralized privacy protection scheme in detail. Our design is driven by privacy, resilience and efficiency considerations. Our approach significantly extends P2P data sharing LBS privacy schemes [31, 55], addressing a broad(er) set of requirements (Sec. 3) and contributing the following main ideas: (1) Each node is equipped with short-term anonymous credentials, to authenticate all node-to-node and node-to-LBS interactions. (2) Peer-provided POI data can be drawn from a large volume of POI data, proactively distributed by the LBS server to a small fraction of randomly chosen nodes, termed *serving nodes*. (3) Nodes submit queries to serving nodes, which periodically announce their presence and available POI data (i.e., POI data for their regions). (4) Encrypted peer query and response messages, to minimize revealed information to neighboring nodes. (5) Cross-checking of peer responses and proactive checking with the LBS

server, to validate peer responses and detect/evict malicious nodes. Table 1 summarizes the used notation.

Table 1. Notation

U	<i>A node(/user)</i>
$LTCA$	<i>Long-Term Certification Authority</i>
PCA	<i>Pseudonymous Certification Authority</i>
RA	<i>Resolution Authority</i>
Lk/LK	<i>Long-term Private/Public Key</i>
LTC	<i>Long-Term Certificate</i>
Sk/SK	<i>Short-term Private/Public Key</i>
PC	<i>Short-term/Pseudonymous Certificate</i>
K_s	<i>A symmetric session key</i>
SN	<i>Serial Number</i>
$(n/q/r)id$	<i>(Node/Query/Region) identifier</i>
L	<i>Square region side length</i>
G	<i>Number of POI type groups</i>
N	<i>Maximum peer requests per node query</i>
Γ/τ	<i>Pseudonym request interval/Pseudonym lifetime</i>
Pr_{serve}	<i>Probability of serving node assignment</i>
T_{serve}	<i>Serving period of each serving node assignment</i>
Pr_{check}	<i>Probability of checking with the LBS server</i>
T_{beacon}	<i>Beacon interval</i>
T_{POI}	<i>POI update interval</i>
T_{wait}	<i>Beacon waiting time before requesting the LBS server</i>
Q_{rid}	<i>A query for POI data in a region, r id</i>
POI_{rid}	<i>POI data covering an entire region, r id</i>
t/t_{now}	<i>Timestamp/Fresh timestamp indicating current time</i>
$\{msg\}_{\sigma_{LTC/PC}}$	<i>Signed message with signature and LTC/PC attached</i>
$E_K(msg)$	<i>Encryption of message with K</i>
$Sign(msg, LTC/PC)$	<i>Signature on a message with corresponding Lk/Sk under LTC/PC</i>
$H()$	<i>Hash function</i>
$Send(msg, E)$	<i>Send a message to an entity E</i>
$Receive(msg, E)$	<i>Receive a message from an entity E</i>

4.1 Overview

Nodes are interested in diverse POI data throughout their trips. The LBS server maintains a database of real-time POI data, subject to change over time. Without loss of generality, we assume: POI data is refreshed/updated every T_{POI} ; the area (e.g., of a city) is divided into a number of equally sized regions; nodes are interested in POI data for regions they are currently in.

The LBS server and the nodes in the system are registered with an *identity and credential management facility*, i.e., a PKI. A *Long-Term Certification Authority (LTCA)* issues a *Long-Term Certificate (LTC)* for each registered node, used as a long-term identity for the node. With the LTC, a node can obtain an LTCA-issued ticket, presented to a *Pseudonymous Certification Authority (PCA)* for obtaining *Pseudonymous Certificates (PCs)/pseudonyms* (if the ticket is validated). The ticket is authenticated by the LTCA but it is *anonymized*: it does not reveal the long-term (real) node identity (referred as *node identity* in the rest of the paper) to the PCA [34, 35]. Therefore, a single LTCA or a single PCA cannot link the node identity to the issued pseudonyms. Thus, the messages signed under PCs are also unlinkable to the node identity. To ensure unlinkability after a change of

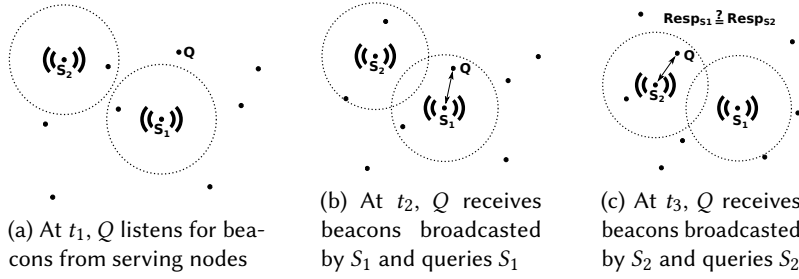


Fig. 2. Illustration of peer query and response collection.

pseudonym, the node can randomly reset its IP and MAC address. This duty separation is based on work done in the context of Vehicular Communication (VC) systems [22, 34, 35]. The ticket and pseudonym acquisition protocols are presented in Sec. 4.2.

Pseudonyms are used by nodes to authenticate themselves to the system entities, including the LBS server and other nodes (peers), and to establish secure communication channel between system entities. Messages are signed under (i.e., with the private keys corresponding to) pseudonyms in order to hide node identities while providing message authentication and integrity.

We assign to the PCA the responsibility for probabilistically choosing the serving nodes. A *serving node* is responsible for requesting extensive POI data, covering an entire region (termed *regional POI data*), consisting of *POI entries* for that region, from the LBS server; storing the POI data in its local storage; and serving other nodes whose requests it receives. The PCA explicitly assigns/binds such a role to PCs (i.e., by setting a specific field in each PC) it issues; thus nodes cannot masquerade serving nodes. Each serving node broadcasts beacons periodically, messages indicating the region it is in. The beacon is signed and the pseudonym is attached, so that any receiver can establish a secure communication channel with the serving node. This prevents exposure of all the communication to nearby nodes, thus disclosure of nodes' interests and presence in the region. A node in need of POI data queries multiple serving nodes; with multiple redundant responses, it can cross-check and validate the obtained data and, at the same time, detect peer node misbehavior (e.g., provision of false, forged, or outdated POI data). Fig. 2 shows an example of the operation concerning peer node queries. At t_1 , Q becomes interested in some POI data and it starts listening for serving node beacons (Fig. 2a). At t_2 , Q moves within the communication range of S_1 ; it queries S_1 upon a receipt of one of its broadcasted beacons (Fig. 2b); at t_3 , Q encounters another serving node, S_2 , it queries S_2 , and then cross-checks the two responses (Fig. 2c).

The LBS server is also issued an LTC, used to authenticate itself to the nodes and infrastructure entities. The LBS server and the nodes are involved in (i.e., registered with) the same PKI architecture, so that the LBS server is also able to authenticate the nodes with their PCs. The LBS server authenticates itself with the LTC and provides message integrity and authentication by signing the messages.

Furthermore, to prevent abuse of the node anonymity, our scheme provides *conditional anonymity* and allows *revocation of anonymity and eviction of nodes*. The node interactions with the facility entities are explained throughout this section below (Sec. 4.2). While we assume the certificates of system entities (i.e., authorities and the LBS server) are pre-installed in each node, all the messages signed under PCs should be attached with the corresponding PCs. We mandate that all signatures must be verified before the messages can be processed. However, for simplicity, we skip the steps for signature verifications in the algorithms and protocols: progress in protocol execution indicates signatures are already validated.

4.2 Pseudonymous Authentication




Fig. 3. Example of ticket and pseudonym set acquisitions with $\Gamma = 2\tau$ for various node trips

Ticket and Pseudonym Acquisition: We refer the reader to [35], for a detailed protocol description of LTC issuance. Here, we focus on our customized pseudonym issuance protocol. We follow the privacy-enhancing pseudonym issuance policy proposed in [34, 35]: independently of the time of contacting the LTCA/PCA, all pseudonym lifetimes are aligned. This implies that nodes change their pseudonyms at the same time, and as a result, the successive pseudonyms of a specific node would not be linked based on pseudonym lifetimes. In our scheme, we do not restrict the ticket lifetimes to be started at same time, since the anonymity set (all nodes registered with an LTCA) for ticket acquisition is large enough considering an expected large number of nodes in the system. Fig. 3 shows an example of pseudonym issuance for nodes with various trips, for $\Gamma = 2\tau$. For each pseudonym request, a node always requests Γ/τ pseudonyms. In general, a node is more likely to start its trip in the middle of a τ period. As shown in Fig. 3, the first pseudonym of a node is not valid from the beginning of a τ period, but from the beginning of the trip until the end of that τ period. For the rest of the trip, the node always requests Γ/τ pseudonyms from the PCA. Although a node (e.g., U_1) could finish its trip shortly after its last pseudonym acquisition, the node always requests pseudonyms with lifetimes covering the entire Γ period: there is a practical uncertainty for the trip duration and this hides the trip ending time from both the LTCA and the PCA.

$$U \rightarrow LTCA : ticket_req\{t_s\}_{\sigma_{LTCU}} \quad (1)$$

$$LTCA : check(nid_U, t_s) \quad (2)$$

$$LTCA \rightarrow C : ticket = \{SN_{ticket}, t_s\}_{\sigma_{LTC_{LTCA}}} \quad (3)$$

$$U : n = \lceil \frac{\Gamma}{\tau} \rceil \quad (4)$$

$$U : \{SK_i, SK_i\}, i \in [0, n) \quad (5)$$

$$U \rightarrow PCA : pseudonym_req\{ticket, \{SK_0\}_{\sigma_{SK_1}} \dots \{SK_{n-1}\}_{\sigma_{SK_{n-1}}}\} \quad (6)$$

$$PCA : t_{s_0} = t_s, t_{e_0} = t_s - (t_s \bmod \tau) + \tau \quad (7)$$

$$PCA : t_{s_i} = t_{e_{i-1}}, t_{e_i} = t_{s_i} + \tau, i \in [1, n) \quad (8)$$

$$PCA : s = 1 \text{ with probability } Pr_{serve}, s = 0 \text{ otherwise.} \quad (9)$$

$$PCA \rightarrow U : PC_i = \{SN_{PC_i}, SK_i, s, t_{s_i}, t_{e_i}\}_{\sigma_{LTC_{PCA}}}, i \in [0, n) \quad (10)$$

When a node needs to request pseudonyms, it sends a request to the LTCA with a pseudonym validity starting time, t_s (1). For a node starting its new trip, t_s is the trip beginning (e.g., the beginning of the first dashed circle for U_1 in Fig. 3); while for a node in the middle of its trip, t_s is the starting point of next Γ period (e.g., the beginning of the second dashed circle). The pseudonyms are issued for the remaining time of each Γ , as we described earlier. The LTCA checks if a ticket

was issued with an overlapping lifetime, based on the locally stored ticket issuance history for each node (2); if not, it issues a ticket valid from t_s (3). With the ticket in hand, the node obtains a set of pseudonyms from the PCA: the node generates the needed public/private key pairs and sends self-signed public keys with the ticket to the PCA (4-6). The anonymized ticket does not reveal the node identity to the PCA. The PCA calculates the lifetimes of pseudonyms according to the pseudonym lifetime policy (7-8). The assignment of serving nodes is done during the pseudonym issuance process. The PCA assigns the requesting node as a serving node with probability Pr_{serve} . This is achieved by setting a field, s , in each PC (9-10). It indicates that the node possessing the pseudonym is responsible for requesting POI data from the LBS server and serving other nodes. We align the serving period, T_{serve} , of a serving node with Γ , thus all pseudonyms for the Γ period will be assigned the same s . Both ticket and pseudonym acquisitions are done across secure channels. We explain our resilient and privacy-enhancing LBS functionality in more detail in Sec. 4.3.

Reporting Misbehavior: When misbehavior (of any type) is detected by a node, the node sends to the Resolution Authority (RA) the messages related to the misbehavior, with the corresponding pseudonyms attached. In case the messages are proved to be related to misbehavior, the LTC of the misbehaving node is exposed (and the node possibly evicted from the system) by the RA with the help of the LTCA and PCA. The revoked LTCs/PCs are published through a Certificate Revocation List (CRL). We describe the response validation process in Sec. 4.3 in more detail.

4.3 Privacy-enhancing LBS

LBS Query Model: We divide a large area (e.g., a city) into equally-sized regions, each with a unique identifier, rid . A node is interested in POI data, of various types, around its current location. This can be any subset of the regional POI data, which comprise various POI data types. For example, a node query could be *What are the menus of nearby restaurants?* and *How many parking spaces are available nearby?* Without loss of generality, we assume that POI data in the LBS server is updated every T_{POI} , with the expiry time explicitly indicated in returned POI data; which have varying lifetimes, depending on their types. Basic information of some POI entries (e.g., name and address of a restaurant) could remain unchanged for several years, while others, e.g., weather and traffic conditions or parking availability could be updated on a frequent basis (e.g., every 10 minutes). We capture a generalized POI update frequency with a constant T_{POI} and we are mainly concerned with frequently updated POI data. POI data with longer lifetimes could actually be preloaded to user devices. For example, Google Maps¹ provide offline caching of POI data (e.g., names and addresses), so that they can be queried locally, without an Internet connection. However, frequently updated POI data (e.g., traffic or road conditions, or available parking slots) cannot be preloaded and should be obtained on-demand. Each node can cache POI data it received and use the information as long as it is not expired: a node does not initiate a query for the same piece of data it cached before the data expires; new queries seek POI data for a different location or POI type than those of non-expired POI data in its cache.

Regional POI Data: Once cached by a serving node, they can be used to serve other nodes. For region-based queries [48, 55], nodes are assumed to be interested in POI data for their current regions. Thus, regional POI data includes POI data within the corresponding region. A limitation of such a query model is that, for some POI types, no response is available in the querier's current region, although the nearest out-of-region POI data can be still valuable. Typically, in this situation, a direct query to the LBS server could be serviced with the nearest POI entries in other regions. Our scheme can be readily adapted to address this issue, by generating and distributing to the serving nodes regional POI data consisting of additional nearest POI entries (when some POI types are

¹<https://www.google.com/maps>

unavailable in a region or some out-of-region POI entries are closer to the queriers at the border of the region), instead of strictly including only POI entries for the current region of the serving node(s). A k -nearest-neighbor query model [61] can be natural here; returning the k nearest POI data, including possibly out-of-region POI data.

Protocol 1 Beacon, repeat every T_{beacon}

- 1: Node is in region rid , possessing a valid serving pseudonym, PC with $s = 1$
 - 2: $Broadcast(\{rid, t_{exp}\}_{\sigma_{PC}})$
 - 3: **return**
-

Protocol 2 POI data update

- 1: Serving node, U , entering a new region rid ,
 - 2: or reaching a POI update point (i.e., t_{exp} for POI_{rid}) for its current region, rid .
 - 3: **if** $Cache \neq \{\}$ **then**
 - 4: **for each** $\{POI_{rid_i}, t_{exp_i}\} \in Cache$ **do**
 - 5: **if** $rid_i == rid$ and $t_{exp_i} > t_{now}$ **then**
 - 6: **return**
 - 7: **end if**
 - 8: **end for**
 - 9: **end if**
 - 10: $Send(\{qid, rid\}_{\sigma_{PC}}, LBS) *$
 - 11: $Receive(\{qid, \{POI_{rid}, t_{exp}\}_{\sigma_{LTC_{LBS}}}\}, LBS) *$
 - 12: Update $Cache$ with $\{POI_{rid}, t_{exp}\}$
 - 13: **return**
-

Serving Node: As mentioned earlier, the PCA assigns the role of serving node with probability, Pr_{serve} , when a node requests pseudonyms from the PCA. Each pseudonym, PC , of a serving node has the field s set to 1 (Sec. 4.2). After a node becomes a serving one, it broadcasts a beacon every T_{beacon} , as per Protocol 1. A beacon, signed under PC , includes an rid , indicating the region the serving node is in, thus the POI data it can provide, and the expiry time of the POI data it has (i.e., obtained from LBS and stored locally) for that region.

Serving nodes retrieve regional POI data from the LBS server whenever they enter new regions (including the case they become serving nodes in their current regions). As shown in Protocol 2, when a serving node enters a new region, it checks whether the latest POI data, covering that entire region, has been cached (during a previous visit to the same region). If not (or the cached POI data is expired), it requests from the LBS server the latest POI data of that region. The serving node is also responsible for updating POI data from the LBS server to have the latest POI data cached locally and serve other nodes. The communication between nodes and the LBS server is carried out over a secure channel, e.g., a TLS channel; we mark steps over a secure channel with asterisk (*).²

²In general, the format of PCs does not need to comply with the format of a standard X.509 certificate, because PCs can only include minimal necessary information to keep communication overhead low and to reduce node identifiability [6, 34, 35], while LTCs can be standard X.509 certificates. As a result, PCs are not compatible with SSL/TLS protocols, but nodes are required to provide message integrity and authentication by signing the messages with their private keys corresponding to the PCs. When a node is required to present its LTC (e.g., for a ticket acquisition), an SSL/TLS channel can be established with the LTC.

Protocol 3 A node query

```
1: Node,  $U$ , interested in POI data in region  $rid$ ; interest expressed by query  $Q_{rid}$ 
2:  $Result = Search(Cache, Q_{rid})$ 
3: if  $Result \neq \phi$  then
4:   return
5: end if
6:  $j = 0, t = t_{now} + T_{wait}$ 
7: while  $j < N$  do
8:   Listen for beacons; wait until  $t$ ; if  $t$  passed, break
9:   Received a beacon  $\{rid', t'_{exp}\}_{\sigma_{PC_{U_j}}}$  from  $U_j$ .
10:  if  $s == 1$  in  $PC_{U_j}$ ,  $rid' == rid$  and  $t'_{exp} > t_{now}$  then
11:     $Send(\{E_{K_{s_j}}(qid_j, rid, Q_{rid}, PC_U), E_{SK_{U_j}}(K_{s_j}), t_{now}\}_{\sigma_{PC_U}}, U_j)$ 
12:     $Receive(\{E_{K_{s_j}}(qid_j, Resp_j), Auth_j, t_j\}_{\sigma_{PC_{U_j}}}, U_j)$ 
13:     $j = j + 1$ 
14:  end if
15: end while
16: if  $j == 0$  then
17:    $Send(\{qid, Q_{rid}\}_{\sigma_{PC_U}}, LBS) *$ 
18:    $Receive(\{qid, Resp\}_{\sigma_{LTC_{LBS}}}, LBS) *$ 
19: else if  $j == 1$  then
20:    $b = 1$  with probability  $Pr_{check}$  or  $b = 0$  with probability  $1 - Pr_{check}$ 
21:   if  $b == 1$  then
22:      $Send(\{qid_0, Q_{rid}, H(Resp_0), Auth_0, t_0, PC_{U_0}\}, LBS) *$ 
23:      $Receive(\{qid_0, Auth_0, result\}_{\sigma_{LTC_{LBS}}}, LBS) *$ 
24:     if  $result$  is negative then
25:        $Report = \{K_{s_0}, \{E_{K_{s_0}}(qid_0, Resp_0), Auth_0, t_0\}_{\sigma_{PC_{U_1}}}\}$ 
26:        $Send(\{\{Q_{rid}, Report\}, t_{now}\}_{\sigma_{PC_U}}, RA) *$ 
27:     end if
28:   end if
29: else if Conflict exists among  $\{Resp_0 \dots Resp_{j-1}\}$  then
30:    $Report = \{\{K_{s_0}, \{E_{K_{s_0}}(qid_0, Resp_0), Auth_0, t_0\}_{\sigma_{PC_{U_0}}}\} \dots$ 
31:      $\{K_{s_{j-1}}, \{E_{K_{s_{j-1}}}(qid_{j-1}, Resp_{j-1}), Auth_{j-1}, t_{j-1}\}_{\sigma_{PC_{U_{j-1}}}}\}\}$ 
32:    $Send(\{\{Q_{rid}, Report\}, t_{now}\}_{\sigma_{PC_U}}, RA) *$ 
33: end if
34: return
```

LBS Query: Protocol 3 and Protocol 4 illustrate the querying and the serving process respectively. When a node is interested in POI data (for its current region rid), it first searches in the local cache whether the data is cached (while it is/was a serving node) and is still valid (i.e., not expired). If yes, then the node interest (i.e., the node query) is fulfilled and the process is finished. Otherwise, the node starts listening to beacons in the network for the period of T_{wait} at most. If a beacon including rid is received from a serving node (i.e., $s = 1$ in PC), the regular node sends the query to

Protocol 4 Serving of a node query

```
1: A serving node,  $U$ , waits for incoming peer query in region  $rid$ 
2:  $Receive(\{E_{K_s}(qid, rid', Q_{rid'}, PC_{U'}), E_{SK_U}(K_s), t_{U'}\}_{\sigma_{PC_{U'}}}, U')$ 
3: if  $rid == rid'$  then
4:    $Resp = Search(POI_{rid}, Q_{rid'})$ 
5:   if  $Resp \neq \phi$  then
6:      $Auth = Sign(\{qid, Q_{rid}, H(Resp), t_{now}\}, PC_U)$ 
7:      $Send(\{E_{K_s}(qid, Resp), Auth, t_{now}\}_{\sigma_{PC_U}}, U')$ 
8:   end if
9: end if
10: return
```

and receives the response from the beacon sender. If a beacon from a non-serving node is received, this is considered as misbehavior and reported to the RA.

The query-response process is encrypted with a session key, K_s , generated by the querying node for each peer query, encrypted with the public key in the serving node's PC. Thus, the communication is kept confidential between the two nodes. The node listens and waits for at most N beacons with the region identifier (rid) of interest, and requests information from the beacon senders. The sought number of responses, N , is a protocol selectable parameter, used to provide redundancy for response validation. We explain it in more detail (in *response validation*) below.

Protocol 4 shows the action of a serving node when a peer query is received. It first decrypts the session key, K_s , with its private key, and then decrypts the query and the pseudonym of the querying node with K_s . Once the query is verified, it searches in the local cache according to the query. If the search yields a result, then the serving node encrypts the response with K_s , and generates an *authenticator*, $Auth$, which can be used to check the correctness of the response with the LBS server.

Response Validation: Steps 16-33 in Protocol 3 show the response validation approach. To protect the querying node from false information, forged by malicious serving nodes, and detect/reveal such misbehavior, the querying node queries $N > 1$ (discovered) serving nodes, with the same query within a T_{wait} period. If the waiting timer expires before N responses are received, the node concludes this process. In an apparently benign system setting (e.g., no malicious node detected recently), the node could use the information received from the first request and use the rest for cross-checking. Each (legitimate) serving node by default cached the same (LBS-obtained) POI data for that region, thus their responses to the same query should be the same (i.e., a given search on the same data should return the same result, as long as the responders are legitimate). In order to make sure that any two honest responses from the serving nodes are identical, we assume a query won't span across a POI data update point: if a node receives two responses from two serving nodes before and after the POI data originating from the LBS server were updated, these two honest responses could result in conflicting data. Without loss of generality, we assume a node waits (i.e., does not initiate a query in order to obtain fresher POI data) until the POI is updated (based on the t_{exp} value learned from beacons) if the remaining time is less than $T_{wait}/2$; in contrast, it initiates the query after t_{exp} is passed. If the remaining time is within $[T_{wait}/2, T_{wait}]$, the node initiates the query immediately but it only waits until the end of the current T_{POI} period (i.e., t_{exp}), because any peer response after t_{exp} could potentially conflict with previous peer responses.

At the very least, a node query succeeds if at least one serving node is discovered and queried successfully (i.e., it produces an authenticated response). Nonetheless, this does not provide any

Protocol 5 Processing of a misbehavior report

```
1: Receive( $\{\{Q_{rid}, Report\}, t\}_{\sigma_{PC_U}}, U\)$  *)
2: Send( $\{qid, Q_{rid}\}_{\sigma_{LTC_{RA}}}, LBS\)$  *)
3: Receive( $\{qid, Resp\}_{\sigma_{LTC_{LBS}}}, LBS\)$  *)
4: for Each  $\{K_{s_i}, \{E_{K_{s_i}}(qid_i, Resp_i), Auth_i, t_i\}_{\sigma_{PC_{U_i}}}\}$  in Report do
5:   | if  $Resp_i \neq Resp$  then
6:   |   | Initiates a pseudonym resolution for  $PC_{U_i}$ 
7:   | end if
8: end for
9: return
```

protection in the event this responder is malicious. POI data from additional serving nodes can be used for cross-checking: any conflicting responses will be reported (with the originally attached signatures and pseudonyms) to the RA. The higher N is, the faster the detection of malicious nodes can be expected, especially in an environment with a high malicious node ratio. However, there is no guarantee that a querying node will discover N serving nodes (before T_{wait} expires). In a less densely populated area, the number of serving nodes would be also low. Thus, a querying node might be able to discover only one (in some cases, possibly) serving node. In this case, it will check the correctness of the single peer response with the LBS server with probability Pr_{check} .

For each peer query, the serving node generates an authenticator and send to the querying node with the peer response. An authenticator is a signature on the concatenation of query id (qid), the LBS query (Q_{rid}), the hash value of the response ($H(Resp)$) and the timestamp on peer response (t). Once the LBS server receives the request, it first verifies the signature (i.e., the $Auth$ itself) with the attached serving node's pseudonym. Once the signature is verified, it fetches the response based on Q_{rid} from its own database and check the correctness of $H(Resp)$. We use $H(Resp)$ instead of $Resp$ in order to decrease the communication overhead (typically the size of $Resp$ could be much larger than a hash value). Again, such validation is possible because the serving node has requested the POI data covering the entire region; thus, the peer response and the LBS response should be the same. If they are different, the querying node receives a negative result and reports the peer response to the RA.

Once the RA receives any report on misbehavior, the RA continues with processing the report as shown in Protocol 5. The RA checks with the LBS server, submitting the same query as that by node reporting the misbehavior. If any dishonest peer response is discovered, the RA reveals misbehaving node(s) through pseudonym resolution (Sec. 4.2).³

POI Type Division: The overhead for obtaining and caching complete regional POI data could vary, depending on POI density and number of available POI types. A parameter, G , can be used to adjust the regional POI data overhead for each serving node. G indicates the number of groups available POI types are divided into. By assigning randomly one of the groups to each serving node at the time of PC issuance (i.e., role assignment), each serving node is only responsible for that subset of POI data. To achieve this, PCs of serving nodes should be augmented with the assigned group(s) (i.e., with an index of each group). The division of POI types can be pre-configured and be known to all participating nodes (that can map a specific POI type to a specific group). Therefore,

³A malicious serving node could increase the beacon rate attempting to increase the probability to be chosen by a querying node. Therefore, during the collection of responses from N serving nodes, if the querying node detects any abnormally high beacon rate from a specific serving node (under the same PC), such misbehavior (i.e., the authenticated beacons at a high rate) should be reported to the RA.

when a querying node receives a beacon signed under a PC, it knows, by looking at the group index in the PC, whether the beacon sender can respond to its query.

5 SECURITY AND PRIVACY EVALUATION

In this section, we analyze the achieved security and privacy properties with our scheme. We first provide a qualitative analysis as per the requirements in Sec. 3. Then, we provide an extensive quantitative analysis on protecting node privacy against the LBS server and honest-but-curious nodes and on resilience to malicious nodes.

5.1 Qualitative Analysis

Authentication and integrity: Entity and message authentication and message integrity are achieved thanks to message digital signature verifications. Message (e.g., beacon) timestamps prevent replays of old messages. A peer query is encrypted and bound to a specific serving node, thus it cannot be meaningfully replayed and any other serving node would fast reject it. More important, the query and response identifiers can trivially allow the (given) same serving (querying) node to reject and not serve (accept) them, at the expense of modest local memory.

Accountability: In spite of the pseudo-/ano-nymity, upon detection of misbehavior reported to the RA, the RA can reveal the actual, long-term identity of the reported misbehaving node, through pseudonym resolution and possibly evict the node from the system. The RA can verify the truthfulness of the reported misbehavior with the help of the LBS server.

Confidentiality: Communication among nodes and infrastructure entities (LTCA, PCA and LBS) is kept confidential by using public key cryptography and (symmetric) session keys.

Sybil-resilience: The LTCA and the PCA issue *tickets* and pseudonyms with non-overlapping lifetimes, ensuring a node is equipped with only one valid pseudonym at any point in time.

Response validation: Our scheme prevents malicious nodes from providing false information, using probabilistic serving node assignment, cross-checking and proactive LBS checking. First, each node is chosen as serving node probabilistically by the PCA and such role is explicitly bound to the provided pseudonyms. Therefore, a malicious node has no control on whether it becomes a serving node, which is the only possibility to provide false information to its peers. For a non-serving node, the waiting time to be assigned as serving node follows a geometric distribution with parameter Pr_{serve} . Thus, the expected waiting time is $\frac{1-Pr_{serve}}{Pr_{serve}} \cdot \Gamma$. For example, when $Pr_{serve} = 0.05$ and $T_{serve} = \Gamma$, an adversary has to wait for $19 * \Gamma$ on average before being selected as a serving node. Second, a malicious serving node has to be chosen by the querying node; again, this selection is not under the control of any malicious node. There may exist multiple serving nodes around a querying node that takes the initiative to choose one or multiple (N) serving node(s). The use of multiple (redundant) serving nodes can reveal a malicious node by cross-checking their responses, given all (benign) serving nodes have the same POI data for the region. Moreover, with the help of the proactive LBS checking mechanism, benign nodes can protect themselves against malicious nodes when only one peer response is received. Such proactive check does not expose long-term or short-term identities of the querying node, because the authenticator in a peer response is a signature generated by the serving node. We provide quantitative evaluation of resilience of our scheme in Sec. 5.2.

Exposure reduction: Our scheme reduces the exposure to the LBS server by sharing information among the peers, as its most closely related predecessors [31, 55] did. In addition, through the controlled selection of serving nodes and encrypted P2P communication, only the selected serving nodes learn queries. Queries by one node, U_q , to a certain serving node, U_s , can be synthetically linked only while U_q uses the same pseudonym. Moreover, U_q could choose among several U_{s_i}

nodes even within a given region for successive queries. Mobility of all nodes (serving or not), short-lived pseudonyms, encryption of query-response process, and rotating assignment of serving nodes minimize exposure to any curious serving node acting alone.

Table 2. Linked queries for different collusion cases

<i>Case</i>	<i>Linked queries</i>	<i>Colluding entities</i>
C_1	Same Id_{PC}	No collusion with CA
C_2	Same Id_{ticket}	Collusion with PCA
C_3	Same Id_{LTC}	Collusion with PCA and LTCA

However, colluding serving nodes could merge the queries they received, attempting to link them the same way the curious LBS server would do. Moreover, collusion with the CAs could allow further linking of queries by linking pseudonyms of the same node. Given the honest-but-curious assumption for the security infrastructure, the RA is not able to initiate a pseudonym resolution without any reported and confirmed misbehavior (through a check with the LBS server). Table 2 shows the queries that can be linked for different collusion cases. We provide quantitative evaluation on node exposure for different collusion cases in Sec. 5.2. We refer to [34, 35] for the information disclosed to honest-but-curious PKI entities.

The PCA, when colluding with honest-but-curious nodes, could assign the role of serving nodes to the colluding nodes or increase Pr_{serve} within a reasonable range (e.g., from 0.06 to 0.08), attempting to collect more node queries through the colluding nodes. However, this deviates from the honest-but-curious assumption for the security infrastructure, because the honest PCA should randomly assign serving nodes with a given Pr_{serve} (defined by system policy).

Jamming: An attacker could jam beacons from benign serving nodes to hinder the provision of correct responses. This is possible because, by default, beacons from a specific serving node are predictable, given a specific T_{beacon} (e.g., observed based on previously overheard beacons or specified by system policy). Jamming, i.e., “erasing” benign serving node beacons makes it harder for malicious serving nodes to be detected through response cross-checks. This can be mitigated with a random (e.g., uniformly distributed) beaconing interval, so that upcoming beacons from a specific serving node are not predictable. We show that a random beacon interval has negligible effect on the performance of our scheme in Sec. 5.2.

5.2 Quantitative Analysis

We further evaluate our scheme through simulations, to show its effectiveness in a real-world scenario. *Exposure* to the LBS server and honest-but-curious nodes is quantitatively evaluated through two metrics: peer hit ratio and exposure degree. The peer hit ratio shows the ratio of node queries that are hidden from the LBS server, while the exposure degree indicates the accuracy of reconstructed user trajectories based on exposed user locations. The resiliency of our scheme, in the presence of malicious nodes, is also evaluated, based on the ratio of affected node queries by (forged) false responses. We also compare, thorough the two quantitative privacy measurements, with MobiCrowd [55].

We find that with 6% of the peer nodes acting as serving nodes (i.e., $Pr_{serve} = 0.06$), around 50% of the queries can be hidden from the LBS server. At the same time, the exposure degree to the LBS server is significantly decreased, from around 0.6 to 0.16. This is achieved thanks to the aforementioned low ratio of query exposure to the LBS server and the use of pseudonymous authentication with relatively short-lived pseudonyms. Even in the presence of a high ratio (e.g., 20%) of colluding honest-but-curious nodes, the exposure degree is kept very low thanks to the

encrypted peer query-response process. From a different viewpoint, a high ratio (e.g., 20%) of colluding malicious nodes (i.e., malicious nodes acting as serving nodes providing the same false response to a node query, in an attempt to defeat the response validation) can affect only less than 1.5% of the node queries (submitted to nearby serving nodes); which can be further mitigated or even eliminated by proactive LBS checking mechanism with different Pr_{check} . Moreover, when achieving a similar peer hit ratio to that of MobiCrowd, we show our scheme achieve a much lower exposure degree to honest-but-curious nodes and impose a much lower communication overhead.

Table 3. Simulation Parameters (**Bold** for Default Settings)

L	1, 2, 3 <i>km</i>	T_{query}	3 <i>min</i>
T_{serve}	10 <i>min</i>	N	2, 3
Γ	10 <i>min</i>	Pr_{serve}	0.02, 0.04, 0.06 , 0.08, 0.1, 0.12, 0.18
τ	1, 2.5, 5 , 10 <i>min</i>	$Ratio_{adv}$	0.05, 0.1, 0.2 , 0.3, 0.4, 0.5
T_{poi}	20 , 30, 40, 50, 60 <i>min</i>	Pr_{check}	0 , 0.1, 0.2, 0.3, 0.5, 0.7, 1
T_{wait}	5, 10, 20, 30, 60 , 90, 120 <i>s</i>	G	1 , 2, 3
T_{beacon}	5, 10 , 15, 20, uniform(5, 15) <i>s</i>	$Ratio_{coop}$	0.5 , 1 (MobiCrowd Only)

Simulation setup: We simulate our scheme with OmNET++⁴ and the TraCI mobility interface in Veins [57], connected to the SUMO [38] traffic simulator. We use the Luxembourg SUMO Traffic (LuST) internal mobility scenario (i.e., with source, destination, or both points internal to the city) [11] and the TAPAS Cologne mobility scenario [60]. We assume a penetration ratio of 40%, i.e., 40% of the mobile nodes use LBSs and participate in our collaborative scheme. For each mobility scenario, we use the traces for the 12:30 pm – 2:00 pm period and use the 1:00 pm – 2:00 pm part for the evaluation. We assume ideal wireless connection for ad-hoc node-to-node communication, with a range of 200 *m*. The results are averaged over 5 seeded simulation runs. In the simulation, we divide the whole area into equally sized regions, with a region size $L \times L$. Table 3 shows the parameters of our simulation, bold values indicating default simulation settings, where T_{query} is the query interval of a node. For example, if Pr_{serve} is the parameter we evaluate, then we set the rest of the parameters as: $L = 2$ *km*, $T_{serve} = \Gamma = 10$ *min*, $\tau = 5$ *min*, $T_{poi} = 20$ *min*, $T_{wait} = 60$ *s*, $T_{beacon} = 10$ *s*, $T_{query} = 3$ *min*, $N = 3$, and $G = 1$. For the evaluation of exposure to honest-but-curious nodes and resilience to malicious nodes, we set $Ratio_{adv} = 20\%$ and $Pr_{check} = 0$ by default.

For MobiCrowd, we assume a querying node broadcasts its query every 10 *s* and consider collaboration/cooperation ratio ($Ratio_{coop}$) values of 0.5 and 1. We set default $Ratio_{coop}$ to 0.5 for MobiCrowd: a node would respond with a probability of 0.5 to peer queries. In our scheme, we do not consider $Ratio_{coop}$, because Pr_{serve} can be considered as a combination of collaboration ratio and actual serving node assignment probability. For example, an actual serving node assignment probability of 0.12 and a collaboration ratio of 0.5 could result in the equivalent $Pr_{serve} = 0.06$.

Peer Hit Ratio: The peer hit ratio reflects the ratio of hidden node queries from the honest-but-curious LBS server. It is defined as the ratio of the node queries for which a response is obtained using local or peer caches (of serving nodes), while the remaining is responded by the LBS server itself. Fig. 4 shows the peer hit ratio as a function of Pr_{serve} for the LuST scenario. Figs. 4a to 4c show peer hit ratio with $L = 1, 2, 3$ *km* respectively and the rest of the parameters having the default values in Table 3. Peer hit ratios improve with increasing Pr_{serve} . For example, we see a significant increase in peer hit ratio when Pr_{serve} increases from 0.02 to 0.04, while such improvement becomes moderate for high Pr_{serve} values (e.g., from 0.1 to 0.12). Thus, modest Pr_{serve} (e.g., 0.06 or 0.08)

⁴omnetpp.org




Fig. 4. LuST: Peer hit ratio as a function of P_{r_serve} with $L =$ (a) 1 km, (b) 2 km, (c) 3 km, and (d) $T_{beacon} \sim \text{uniform}(5, 15)$ s.




Fig. 5. TAPASCologne: Peer hit ratio as a function of P_{r_serve} with $L =$ (a) 1 km, (b) 2 km, and (c) 3 km.

is enough to hide a significant amount (e.g., more than 50%) of queries from the LBS server. We see a slight increase in peer hit ratio with larger region sizes, because more serving nodes exist in each region. Moreover, with a small region size, more queries are initiated at the borders of regions, thus querying nodes are more likely to cross to other regions while waiting for a response. Fig. 4d shows that a uniformly distributed T_{beacon} does not affect the peer hit ratio; while it protects benign serving nodes from being jammed (Sec. 5.1).

Fig. 5 shows the peer hit ratio for the TAPASCologne scenario: it follows the same trend as in Fig. 4, but values are slightly higher for the TAPASCologne scenario than those for the LuST scenario. This is because the node density in the central part of TAPASCologne is higher than those in the central part of LuST: higher density results in higher numbers of serving nodes given a P_{r_serve} . This is confirmed through Fig. 6 and Fig. 7. Fig. 6a and Fig. 7a show node density maps (nodes per $1 \text{ km} \times 1 \text{ km}$ region) at 1 pm for the LuST and TAPASCologne scenarios, and Figs. 6b to 6d and Figs. 7b to 7d show maps of the peer hit ratio with different L . We see a clear correlation between node density and peer hit ratio in each region; as expected, the peer hit ratio is roughly proportional to the node density. For the LuST scenario and the default settings, in the central area, the peer hit ratio exceeds 0.5; it approaches 0.8 in the densest region: the higher the node density, the higher the exposure reduction to the LBS server. For TAPASCologne, the peer hit ratio even approaches 95% in the densest areas. For a low-density region, a local relative increase of the system parameter, P_{r_serve} , could improve the peer hit ratio with a modest increase in overhead.

Figs. 8a to 8d show peer hit ratios as a function of T_{wait} , T_{beacon} , τ and T_{POI} , respectively. Higher T_{wait} allows a querying node more time to discover serving nodes, and lower T_{beacon} for a serving node results in a higher probability to be discovered by nearby querying nodes; thus both result in a higher peer hit ratio. However, a change in τ or T_{POI} does not affect the number of serving nodes in the system or the probability of discovering serving nodes, thus the peer hit ratio remains roughly the same as τ or T_{POI} changes.




Fig. 6. LuST: (a) Node density for each region at 1 pm with $L = 1$ km. Peer hit ratio for each region with (b) $L = 1$ km, (c) $L = 2$ km, and (d) $L = 3$ km.




Fig. 7. TAPASCologne: (a) Node density for each region at 1 pm with $L = 1$ km. Peer hit ratio for each region with (b) $L = 1$ km, (c) $L = 2$ km, and (d) $L = 3$ km.




Fig. 8. LuST: Peer hit ratio as a function of (a) T_{wait} , (b) T_{beacon} , (c) τ , and (d) T_{POI} . (Default: $T_{wait} = 60$ s, $T_{beacon} = 10$ s, $\tau = 5$ min and $T_{POI} = 20$ min.)




Fig. 9. LuST: Peer hit ratio as a function of Pr_{serve} with (a) $G = 2$, and (b) $G = 3$. Peer hit ratio for each region with (c) $G = 2$ and $Pr_{serve} = 0.12$, and (d) $G = 3$ and $Pr_{serve} = 0.18$.

We evaluate further our scheme considering multiple groups of POI types (i.e., $G > 1$). Fig. 9 shows the peer hit ratio as a function of Pr_{serve} , with $G = 2, 3$. Given a Pr_{serve} , the peer hit ratio




Fig. 10. LuST with MobiCrowd: Peer hit ratio as a function of (a) T_{POI} , and T_{wait} with (b) $Ratio_{coop} = 1$ and (c) $Ratio_{coop} = 0.5$. (Default: $T_{POI} = 20$ min, $T_{wait} = 60$ s and $Ratio_{coop} = 0.5$.)

decreases as G grows (Figs. 9a and 9b), because less serving nodes exist for a queried POI type. In order to achieve the same level of cache hit ratio, Pr_{serve} should increase accordingly with G . When a serving node is interested in a POI type that it is not responsible for (thus, not cached), it queries the LBS server directly, because it had already exposed itself to the LBS server while obtaining regional POI data. Therefore, the cache hit ratio with $G = 3$ and $Pr_{serve} = 0.18$ is lower than that with $G = 2$ and $Pr_{serve} = 0.12$; the reason is that more serving node queries are sent directly to the LBS server with $G = 3$ and $Pr_{serve} = 0.18$.

Fig. 10 shows the peer hit ratio achieved by MobiCrowd for the LuST scenario. From Fig. 10a, we see that the peer hit ratio slightly increases as T_{POI} increases, because cached POI data is valid for a longer period, thus serving more peer (and own) queries. Figs. 10b and 10c show the peer hit ratio as a function of T_{wait} when the collaboration ratios are 1 and 0.5 respectively. A higher collaboration ratio results in a higher peer hit ratio. However, a longer T_{wait} period does not necessarily improve the peer hit ratio. For example, it slightly increases as T_{wait} increases from 10 s to 60 s, but it decreases for $T_{wait} = 90$ s. Long T_{wait} increases the probability to encounter nodes that cached the required POI data, but it decreases the chance to serve other nodes with its own cached POI data (that could have been obtained earlier from the LBS server with a shorter T_{wait}). The peer hit ratio is around 0.7 with MobiCrowd when $Ratio_{coop} = 0.5$ and $T_{wait} = 10$ s. Our scheme provides roughly the same peer hit ratio when $Pr_{serve} = 0.12$ (Fig. 4b) with the default simulation settings. However, our scheme significantly improves over MobiCrowd in terms of exposure and overhead. MobiCrowd trades off higher node exposure to neighboring honest-but-curious nodes and higher communication overhead for increased peer hit ratio.

Node Exposure: After gauging node privacy through the peer hit ratio, we refine the measurement of node exposure to curious LBS servers and curious peer nodes with the help of the *exposure degree* (*ExpoDeg*), defined for any node as:

$$ExpoDeg(Id_{LTC}, C) = \sum_{Id_i \in ID(Id_{LTC}, C)} \frac{T(Id_i)}{T(Id_{LTC})} * \frac{R_H(Id_i)}{R(Id_{LTC})}.$$

Id_{LTC} is the node long-term identity, corresponding to a whole series of node actions in the system. $ID(Id_{LTC}, C)$ is a set of identities, corresponding to Id_{LTC} , exposed to the honest-but-curious (possibly colluding) entities for collusion case C (Table 2). $ID(Id_{LTC}, C)$ differs for different collusion cases. $T(Id_i)$ is the corresponding trip duration of a node under identity $Id_i \in ID(Id_{LTC}, C)$. $R(Id_i)$ is the number of regions the node visits during its trip under identity Id_i and $R_H(Id_i)$ is the number of visited regions exposed to honest-but-curious entities under the same identity Id_i . $\frac{R_H(Id_i)}{R(Id_{LTC})}$ indicates the exposure degree under a single identity Id_i . To derive the exposure degree

of a node, the exposure degrees under each Id_i are weighted by a time parameter $\frac{T(Id_i)}{T(Id_{LTC})}$: the ratio of the (partial) trip duration under identity Id_i over the total trip time. The exposure degree indicates the accuracy of reconstructed node trajectories based on recorded node queries, taking into consideration the effect of pseudonymous authentication on location privacy protection. We measure exposure to colluding curious nodes through the aggregation of the recorded queries.

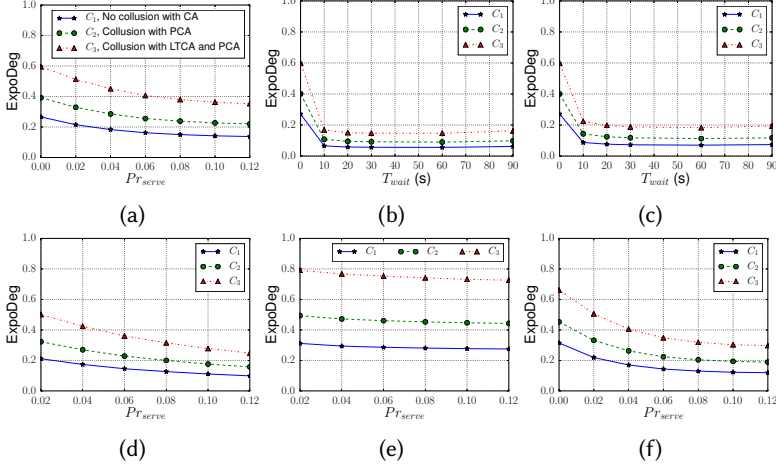


Fig. 11. LuST: *ExpoDeg* to the LBS server as a function of (a) Pr_{serve} . *ExpoDeg* to the LBS server as a function of T_{wait} with MobiCrowd when (b) $Ratio_{coop} = 1$ and (c) $Ratio_{coop} = 0.5$. *ExpoDeg* of (d) non-serving and (e) serving nodes to the LBS server as a function of Pr_{serve} . TAPASCologne: (f) *ExpoDeg* to the LBS server as a function of Pr_{serve} . (Default: $T_{wait} = 60$ s and $Pr_{serve} = 0.06$.)

Figs. 11a and 11f show the average *ExpoDeg* (i.e., the average over all nodes) for exposure to the LBS server as a function of Pr_{serve} , for different collusion cases, for the LuST and the TAPASCologne scenarios respectively. The collusion case C_3 is equivalent to the case that messages are authenticated with node LTCs (i.e., no pseudonymous authentication). $Pr_{serve} = 0$ is equivalent to the case that all queries are sent to the LBS server. For the LuST scenario, *ExpoDeg* is around 0.6 without any protection in place: even if all the queries are sent to the LBS server, *ExpoDeg* is not 1 because a node enters and exits one or more regions between two successive queries. With pseudonymous authentication only, *ExpoDeg* drops below 0.3 and bounces back to around 0.4 for the collusion case C_2 . With the decentralized information sharing scheme in use, the exposure decreases further. For example, when $Pr_{serve} = 0.06$, *ExpoDeg* is around 0.16 for C_1 , but rises to around 0.4 for C_3 .

Fig. 11b and fig. 11c show *ExpoDeg* to the LBS server with MobiCrowd when $Ratio_{coop} = 1$ and $Ratio_{coop} = 0.5$ respectively. We see a lower $Ratio_{coop}$ could result in higher *ExpoDeg*. However, *ExpoDeg* with $T_{wait} = 10$ s and $Ratio_{coop} = 0.5$ is lower than *ExpoDeg* with $Pr_{serve} = 0.12$ in our scheme, although these two settings result in a similar peer hit ratio (Fig. 4b and Fig. 10c). This is because serving nodes in our scheme need to continuously expose their locations to obtain regional POI data. Figs. 11d and 11e show the *ExpoDeg* for non-serving nodes and serving nodes respectively. The *ExpoDeg* for non-serving nodes in our scheme is roughly the same as the *ExpoDeg* for MobiCrowd in the above two settings. Although serving nodes have much higher *ExpoDeg*, a probabilistic and periodical assignment of serving nodes balances high exposure among nodes throughout their trips. However, the *ExpoDeg* to honest-but-curious nodes for MobiCrowd is much higher than the *ExpoDeg* to honest-but-curious nodes for our scheme.

Fig. 12 and Fig. 13 show *ExpoDeg* to honest-but-curious nodes. We assume that the honest-but-curious nodes could collude, merge recorded queries and link the queries (as the LBS could do). Fig. 12 shows *ExpoDeg* as a function of $Ratio_{adv}$. Fig. 12a and Fig. 12d show *ExpoDeg* when $N = 2$, with and without P2P encryption respectively. Fig. 12a shows that a modest realistic $Ratio_{adv}$ (e.g., 0.05 and 0.1) results in relatively low *ExpoDeg*. For example, when $Ratio_{adv} = 0.05$, *ExpoDeg* is lower than 0.05 for C_1 and C_2 , and it is slightly higher for C_3 . However, without P2P encryption (Fig. 12d), *ExpoDeg* significantly increases, because all queries within the communication range of an honest-but-curious (serving or non-serving) node can be recorded. For example, *ExpoDeg* for $Ratio_{adv} = 0.05$ and $Ratio_{adv} = 0.1$ without encryption are almost the same as those for $Ratio_{adv} = 0.3$ and $Ratio_{adv} = 0.5$ with encryption, respectively. This shows the importance of query encryption in terms of reducing node exposure. We see the same effect (mentioned above for $N = 2$) in Fig. 12b and Fig. 12e, when $N = 3$. Fig. 12d and Fig. 12e show that *ExpoDeg* steadily increases as $Ratio_{adv}$ increases, and rather “flattens” out as $Ratio_{adv}$ grows from 0.2 to 0.5. This shows that non-encrypted P2P communication significantly increases *ExpoDeg*, as it strengthens the capability of honest-but-curious nodes even in settings with relatively low $Ratio_{adv}$. *ExpoDeg* only slightly increases from $N = 2$ (i.e., Fig. 12a and Fig. 12d) to $N = 3$ (i.e., Fig. 12b and Fig. 12e), because a querying node cannot always discover exactly N serving nodes within T_{wait} . An increase in N could help evicting malicious nodes from the system more efficiently with a slightly higher overhead (see the evaluation for *resilience* below).

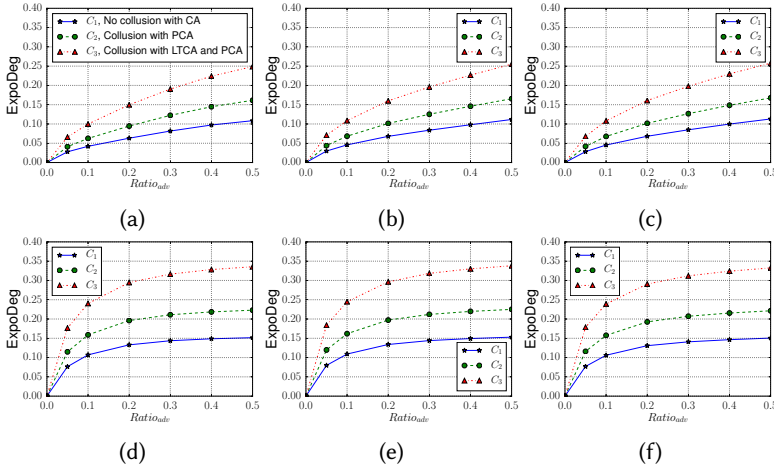


Fig. 12. LuST: *ExpoDeg* to colluding honest-but-curious nodes as a function of $Ratio_{adv}$ with (first row) and without (second row) P2P encryption when $N = 2$ (first column), $N = 3$ (second column) and $N = 3$ with $T_{beacon} \sim \text{uniform}(5, 15)s$ (third column).

Fig. 13 shows *ExpoDeg* to colluding honest-but-curious nodes as a function of Pr_{serve} . Again, the encrypted P2P communication has a significant effect on reducing *ExpoDeg*. *ExpoDeg* for both encrypted and non-encrypted P2P communication increases with increased Pr_{serve} , because higher Pr_{serve} results in a higher number of honest-but-curious serving nodes in the system given a $Ratio_{adv}$. Therefore, more peer queries would be disclosed to honest-but-curious serving nodes. The third columns of Fig. 12 and Fig. 13 show *ExpoDeg* with uniformly distributed T_{beacon} for the default settings; again, this has negligible effect on *ExpoDeg*. Fig. 14 shows the same trend for *ExpoDeg* for the TAPASCologne scenario with $N = 3$, with a similar *ExpoDeg* increase from the LuST scenario to the TAPASCologne scenario, which can be observed from Fig. 11a and Fig. 11f.




Fig. 13. LuST: *ExpoDeg* to colluding honest-but-curious nodes as a function of Pr_{serve} with (first row) and without (second row) P2P encryption when $N = 2$ (first column), $N = 3$ (second column), and $N = 3$ with $T_{beacon} \sim \text{uniform}(5, 15)s$ (third column).




Fig. 14. TAPAS Cologne: *ExpoDeg* to colluding honest-but-curious nodes with and without P2P encryption as a function of (a, b) $Ratio_{adv}$ and (c, d) Pr_{serve} . (Default: $Ratio_{adv} = 0.2$ and $Pr_{serve} = 0.06$.)




Fig. 15. LuST: *ExpoDeg* to colluding honest-but-curious nodes as a function of Pr_{serve} with and without P2P encryption for (a, b) $L = 1 \text{ km}$ and (c, d) $L = 3 \text{ km}$.

Fig. 15 show *ExpoDeg* for different L . A smaller region decreases *ExpoDeg*, because the simulated area is divided into more regions for smaller L (thus more regions a node visits during its trip), while the number of exposed regions could remain roughly the same.

We continue the evaluation with $G > 1$. Fig. 16 shows that *ExpoDeg* decreases as G increases for the same Pr_{serve} , because less queries are responded by the serving nodes. *ExpoDeg* increases with higher G under the comparable settings (i.e., $G = 1$ and $Pr_{serve} = 0.06$, $G = 2$ and $Pr_{serve} = 0.12$,




Fig. 16. LuST: *ExpoDeg* to colluding honest-but-curious nodes as a function of Pr_{serve} with and without P2P encryption for (a, b) $G = 2$ and (c, d) $G = 3$.

and $G = 3$ and $Pr_{serve} = 0.18$), because more serving nodes expose themselves continuously with their beacons. For example, for these settings, *ExpoDeg* for C_1 is around 0.07, 0.09 and 0.11 respectively in Fig. 13b, Fig. 16a, Fig. 16c.




Fig. 17. LuST: *ExpoDeg* to colluding honest-but-curious nodes with P2P encryption as a function of (a) T_{wait} , (b) T_{beacon} , (c) τ and (d) T_{POI} . (Default: $T_{wait} = 60$ s, $T_{beacon} = 10$ s, $\tau = 5$ min and $T_{POI} = 20$ min.)




Fig. 18. LuST with MobiCrowd: *ExpoDeg* to colluding honest-but-curious nodes with (a, b) $Ratio_{coop} = 1$ and (c, d) $Ratio_{coop} = 0.5$ as a function of T_{wait} (a, c) and $Ratio_{adv}$ (b, d). (Default: $T_{wait} = 60$ s and $Ratio_{adv} = 0.2$.)

Figs. 17a and 17b show *ExpoDeg* to colluding honest-but-curious nodes with P2P encryption as a function of T_{wait} and T_{beacon} respectively. *ExpoDeg* slightly increases with longer T_{wait} and shorter T_{beacon} ; in accordance with the peer hit ratio shown in Figs. 8a and 8b. Fig. 17c show *ExpoDeg* as a function of τ with a fixed $\Gamma = 10$ min. *ExpoDeg* significantly increases with increased τ for C_1 , because more queries can be linked together under the same pseudonym when its lifetime is longer. For C_2 , the trend is that *ExpoDeg* slightly decreases with increased τ , given the same Γ , because lower τ results in less pseudonym requests per trip (see Sec. 4.2). For example, consider a node trip duration of 15 min. With a fixed $\Gamma = 10$ min, the nodes needs two pseudonym requests




Fig. 19. LuST: Ratio of malicious serving nodes with (a) $N = 2$, and $N = 3$ when $L =$ (b) 1 km, (c) 2 km, and (d) 3 km.

when $\tau = 1 \text{ min}$. For example, the first pseudonym request covers a duration of 9.5 min and the second pseudonym request covers the rest of 5.5 min. However, when $\tau = 10 \text{ min}$, if there is, e.g., 2 min, remaining in the current τ , three pseudonym requests would be needed for the same example. The node would use the pseudonyms obtained through three requests for 2 min, 10 min and 3 min respectively. Thus, the latter case results in lower *ExpoDeg*, because the node exposure is “split” under three Id_{ticket} rather than two for the former case. *ExpoDeg* remains the same under different T_{POI} (Fig. 17d), because this only affects the communication overhead of serving nodes for updating POI data and does not change the probability of encountering serving nodes by querying nodes.

Fig. 18 shows *ExpoDeg* for MobiCrowd. *ExpoDeg* increases with longer T_{wait} , because more queries are broadcasted before the LBS server is queried. Higher $Ratio_{adv}$ also increases *ExpoDeg* as expected. Fig. 4b and Fig. 10c showed that MobiCrowd with $Ratio_{coop} = 0.5$ and $T_{wait} = 10 \text{ s}$ provides a similar peer hit ratio as our scheme with $Pr_{serve} = 0.12$ and $T_{wait} = 60 \text{ s}$ for the LuST scenario. However, from Figs. 18c and 18d, we see that for MobiCrowd, *ExpoDeg* is higher than for our scheme. For example, with the above settings and $Ratio_{adv} = 0.2$, *ExpoDeg* for MobiCrowd is around 0.16, 0.25 and 0.38 for C_1 , C_2 and C_3 respectively: that is, higher than the 0.12, 0.17 and 0.26 values achieved by our scheme.

Resilience: We evaluate the resilience of our scheme against malicious nodes. We consider the same $Ratio_{adv}$ values as those for the evaluation of the exposure. Although a ratio e.g., 20 %, of adversarial nodes (thus, node owners) may be unrealistic, we consider such a rather harsh setting to capture situations with extensive node infection by malware [8, 52] while node owners can be benign and unsuspecting. We assume that, node owners can be notified and malicious nodes be reinstated as benign nodes through, e.g., diagnostics and updates, once misbehavior is reported and malicious nodes are identified. A recovered benign node would be issued a new LTC and obtain new PCs. Thus, old PCs from detected malicious nodes will be revoked and published through the latest CRL. In the simulation, we assume the malicious serving nodes collude to provide identical false responses (based on falsified data stored locally) to a specific node query, so that the false responses appear consistent and thus be accepted by the querying node. Otherwise, even the responses from independently acting malicious serving nodes could be conflicting and be easily detected.

Fig. 19 shows the ratio of the adversarial serving nodes over time, when $Pr_{check} = 0$. Even though 20 % of the nodes are compromised, the ratio of non-detected malicious serving nodes (over all serving nodes) is less than 10 %, for the majority of time for both $N = 2$ (Fig. 19a) and $N = 3$, with different region sizes (Figs. 19b to 19d). The periodical ratio peaks are due to the simultaneous reassignment of roles to the participating nodes. Without misbehavior detection (i.e., cross-checking), the ratio would have roughly remained the same as $Ratio_{adv}$ (i.e., 20 %). However, the controlled selection of serving nodes effectively limits active participation of malicious nodes, and the cross-checking mechanism further helps detecting and evicting them from the system.




Fig. 20. LuST: Affected node query ratio as a function of $Ratio_{adv}$ with (a) $N = 2$, and $N = 3$ when (b) $L = 1$ km, (c) $L = 2$ km, and (d) $L = 3$ km.




Fig. 21. TAPASCologne: Ratio of malicious serving nodes with (a) $N = 2$, and (b) $N = 3$. Affected node query ratio as a function of $Ratio_{adv}$ with (c) $N = 2$, and (d) $N = 3$.

Moreover, the ratio of “active” malicious nodes (i.e., ones that could actually provide false responses) is always less than 1 % for all cases: a significant decrease from the original $Ratio_{adv}$. This shows that even if 20 % of the nodes are compromised (by a “master” adversary), the ratio of actual usable compromised nodes is considerably lower (i.e., less than 1 %). Such low ratio of “active” malicious nodes is reflected on the ratio of the affected node queries in Fig. 20: for example, when $Ratio_{adv} = 0.2$, around only 1.5 % of the node queries get false peer responses that are not detected. This happens when all peer responses to a node query are given by malicious serving nodes, thus, there is no conflict. If conflicting responses are received, the querying node checks the correctness with the RA and reports the misbehavior accordingly (Sec. 4). The adversarial serving node ratios and affected node query ratios are overall lower for the TAPASCologne scenario (Fig. 21), because higher node (thus benign serving node) density facilitates the detection of false peer responses (thus the eviction of malicious serving nodes).

For the evaluation above, we did not presume the availability of prior information on the security status of malicious peers. This can be available through a periodically publicized CRL by the PKI, and in turn, can further improve resilience to faulty peers. We consider next the effect of a CRL on malicious node detection: for simplicity, assume that the CRL is readily available to nodes, eliminating the effect of delay for obtaining the CRL (which can be practically downloaded or distributed [36]). Consider, a query from U_1 was responded by two malicious serving nodes U_{ms_1} and U_{ms_2} , while U_{ms_2} was detected later as misbehaving node through a query from U_2 responded by U_{ms_2} and a benign serving node U_{bs_1} . Once U_1 obtains the latest CRL, it can do a post-checking for its queries responded by U_{ms_2} . This could further reveal U_{ms_1} as a misbehavior, which, in turn, could disclose more node queries served by single false peer responses from U_{ms_1} . Figs. 22a and 22d show the ratio of non-detected attacked (referred as *attacked* in the rest of the paper) node queries when $N = 3$ with the above mentioned post-checking. For example, when $Ratio_{adv} = 0.2$ for the

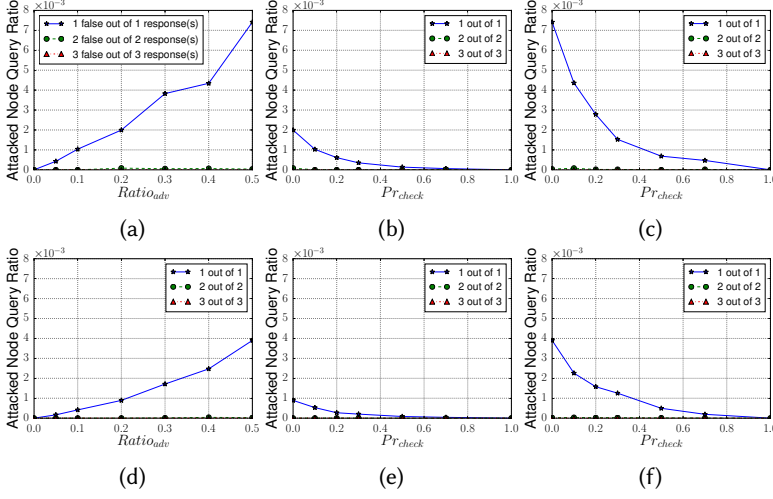


Fig. 22. LuST: (first row) and TAPASCologne (second row): Non-detected attacked node query ratio as a function of (first column) $Ratio_{adv}$, and Pr_{check} when (second column) $Ratio_{adv} = 0.2$ and (third column) $Ratio_{adv} = 0.5$. (Default: $Ratio_{adv} = 0.2$ and $Pr_{check} = 0$.)

LuST scenario, the ratio of *attacked* queries is around 0.2%: a significant decrease from 1.5%, as shown in Fig. 22a, thanks to the access to the CRL.

Figs. 22b and 22e show the ratio of attacked node queries as a function of Pr_{check} with $Ratio_{adv} = 0.2$. We see the attacked node query ratio almost reaches 0 when Pr_{check} is higher than 0.5. When $Ratio_{adv} = 0.5$ (Figs. 22c and 22f), Pr_{check} should be set to higher than 0.7 (or even 1), in order to achieve the same performance. With such a high ratio of serving nodes proven to be malicious, the system should instruct the benign nodes to proactively defend themselves (and the system) against the malicious nodes. Moreover, as described earlier, proactive LBS checks do not expose benign querying nodes (i.e., do not harm their privacy), because the authenticators are the signatures generated by the serving nodes (Sec. 4.3).

Fig. 23 shows the resultant resilience with $G > 1$. From Fig. 19c, Fig. 23a and Fig. 23d, we see the trend that the overall adversarial serving node ratio is higher as G grows, due to the higher Pr_{serve} being set. As a result, affected and attacked node query ratios also increase with higher G . For example, the attacked node query ratio is around 0.2%, 0.3% and 0.5%, with $Ratio_{adv} = 0.2$, in Fig. 22a, Fig. 23c and Fig. 23f respectively.

Communication Overhead: We close this section with a comparative evaluation of the communication overhead based on simulation results for our scheme and MobiCrowd. In our scheme, only serving nodes need to request regional POI data, and querying nodes only obtain responses with POI entries that match their queries. However, in MobiCrowd, nodes always obtain regional POI data (that is, not the subset of precise interest) either from their peers or from the LBS server, in order to keep the POI data verifiable with the LBS-provided signatures.

Fig. 24 shows the communication overhead in terms of the average number of regional POI data obtained from the LBS server (and peers, for MobiCrowd only). Figs. 24a to 24c show the number of regions with $L = 1, 2, 3$ km respectively. Around 6 regional POI data are obtained by each serving node when $L = 1$ km, while around 3.5 regional POI data are obtained when $L = 2$ km. This reduces to around 0.6 and 0.4 for $L = 1$ km and $L = 2$ km if averaged over all nodes (when $Pr_{serve} = 0.06$): a previously non-serving node could be chosen as the serving node in a

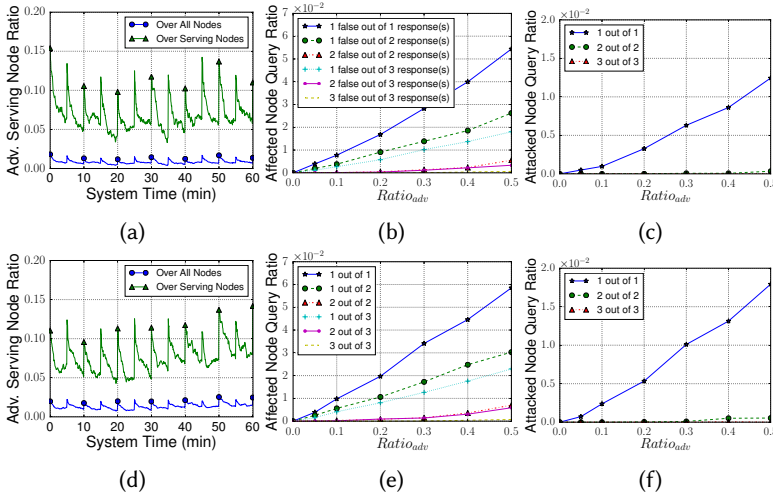


Fig. 23. LuST: (First column) Ratio of malicious serving nodes, (second column) affected node query ratio as a function of $Ratio_{adv}$, and (third column) non-detected attacked node query ratio as a function of $Ratio_{adv}$, with (first row) $G = 2$ and $Pr_{serve} = 0.12$, and (second row) $G = 3$ and $Pr_{serve} = 0.18$.

future trip, thus the load is balanced among all nodes. A region size with $L = 2 \text{ km}$ ($L = 3 \text{ km}$) is 4 (9) times as large as region size with $L = 1 \text{ km}$. Therefore, larger L results in higher communication overhead for obtaining regional POI data, although the number of obtained regions is decreased.

The actual communication overhead depends on actual POI entry size and POI density in each region. Consider a typical value from the literature [42, 49], e.g., 500 bytes for each POI entry. If there are 10 000 POI entries in each regional POI data set, then the communication overhead for obtaining regional POI data is around 5 *Mbytes*, while a region with 100 000 POI entries results in 50 *Mbytes*. Consider regional POI data of 5 *Mbytes*. For the default settings, a serving node needs to proactively obtain around 17.5 *Mbytes* of regional POI data (for around 3.5 regions, as shown in Fig. 24b) during their trips to effectively serve other nodes. This reduces to around 2 *Mbytes* (for around 0.4 region) if averaged over all nodes. For peer queries and direct LBS queries, only the required POI entries need to be obtained. For example, if 10 POI entries are returned for each query, then around 5 *Kbytes* of POI data needs to be obtained: significantly lower than the communication overhead for obtaining the entire regional POI data. Figs. 24d and 24e show a comparison of communication overhead between our scheme and MobiCrowd. With MobiCrowd, each node need to obtain around 3 regional POI data from the LBS server and their peers: significantly higher than 0.4 in our scheme. With higher G (Figs. 24f and 24g), the average number of obtained regional POI data remains roughly the same as in Fig. 24b, because the average node mobility remains the same with even if G changes. However, the overhead for obtaining regional POI data decreases by a factor of $1/G$ accordingly.

6 PERFORMANCE EVALUATION

We further evaluate experimentally the performance of our scheme with vehicular OBUs. An OBU has a 1.66 GHz dual-core CPU and it is equipped with an IEEE 802.11p interface for P2P communication. We implement our scheme with C++. Elliptic Curve (EC) public/private key pairs are used for asymmetric cryptography and Advanced Encryption Standard (AES) with 256-bit

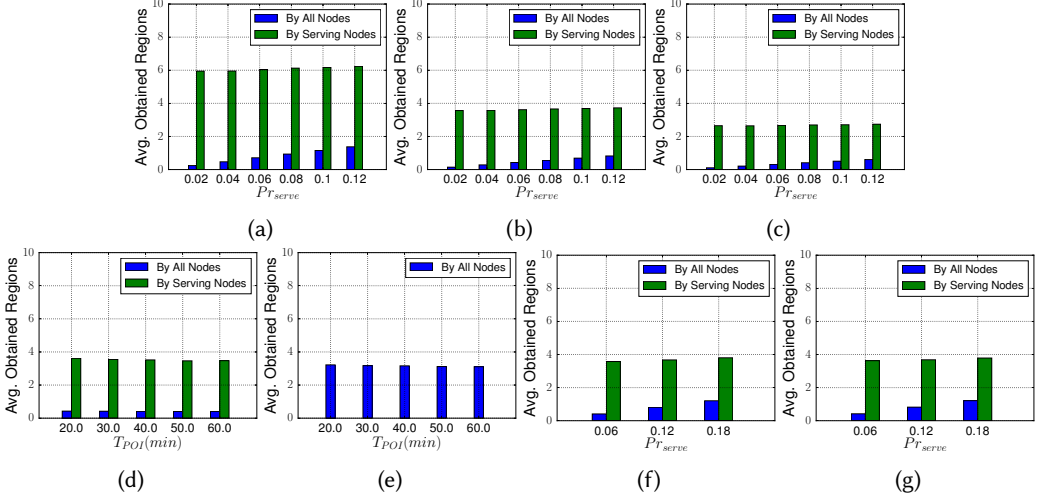


Fig. 24. LuST: Average number of obtained regional POI data as a function of (a, b, c) Pr_{serve} when $L = 1, 2, 3$ km, as a function of T_{POI} with (d) our scheme and (e) MobiCrowd, and as a function of Pr_{serve} with (f) $G = 2$ and (g) $G = 3$.

session key is used for symmetric key encryption. We use Crypto++⁵ for the Elliptic Curve Integrated Encryption Scheme (ECIES) and OpenSSL⁶ for the rest of cryptographic algorithms (e.g., Elliptic Curve Digital Signature Algorithm (ECDSA) and AES). Table 4 shows the parameters we use in the experiment.

Table 4. Experimental Parameters

	Serving	Querying	Query Generator	Beacon Generator
Node num.	5	1	2	2
T_{beacon}	$Exp(\frac{1}{5s})$	-	-	-
T_{query}	-	5 s	$Exp(\frac{1}{0.5s}), Exp(\frac{1}{6s})$	-
γ	1, 2.5, 5 Hz	-	-	40, 100, 200 Hz
EC size	192, 224, 256 bit			256 bit
N	-	1, 2, 3		-

We run experiments in two different settings. In the first setting, we use five OBUs as serving nodes, each broadcasting beacons at a rate of $\frac{1}{T_{beacon}}$. We use one OBU as a querying node, with a node query rate of $\frac{1}{T_{query}}$. Two additional query generators are introduced to emulate a large number of querying nodes in the network, each with a much higher node query rate ($T_{query} = Exp(\frac{1}{0.5s})$) than the aforementioned querying node. The beacon interval and the query generator intervals are exponentially distributed, to emulate beacons and queries from mobile neighboring serving and querying nodes, from the perspective of the querying node we evaluate. We randomize the two aforementioned intervals so that network load changes over time and across experiments.

We also run experiments in a VC specific setting with each serving node and querying node running a transportation safety beacon application: broadcasting signed safety beacons (termed

⁵www.cryptopp.com

⁶www.openssl.org

Cooperative Awareness Messages (CAMs)) and verifying received safety beacons. We consider three CAM rate values, 1, 2.5 and 5 Hz. Although a maximum $\gamma = 10$ Hz is possible according to the standard [14], CAM rate adaptation schemes [47, 54, 58] show that an adaptive CAM rates (generally lower than 10 Hz), considering channel load and vehicle mobility, leads to higher CAM reception ratio. Therefore, we consider modest CAM rates rather than the high-load approach (10 Hz). We add in our experimental setup two CAM generators, each broadcasting CAMs at a rate equivalent to an aggregate of 40 nodes (thus 80 nodes in total). T_{query} is also adjusted to roughly match the aggregate query rate of 80 nodes. With $T_{query} = Exp(\frac{1}{6s})$ for each query generator, and $T_{query} = 5$ s for the evaluated querying node, the average T_{query} is roughly 2.5 min for each node (among the 81 emulated nodes). In this setting, we only use EC-256 keys (which is the level in the corresponding standard [14]).

Moreover, we evaluate processing delays with an implementation of MobiCrowd. With MobiCrowd, once a querying node obtains and caches the regional POI data, it no longer needs to query other nodes. With the limited number of OBUs in the testbed, all nodes would have the regional POI data cached locally within a short period. However, to benchmark the processing delays with MobiCrowd in our experiment, the querying nodes continuously broadcast queries (without turning themselves into cached nodes, i.e., nodes that cached the local POI data).

We evaluate a setting with 3 querying nodes and 3 serving nodes (i.e., nodes that cached the regional POI data), with the same settings for CAM rates for the six evaluated nodes and the two CAM generators. Essentially, we emulate a network of 86 nodes, among which three nodes cached the regional POI data and three querying nodes are interested in the POI data. We downloaded the OpenStreetMap POI database for Luxembourg ⁷, and found that the highest number of POIs in a region with $L = 2$ km is around 5000. Consider a size of 500 bytes for each POI, we assume the regional POI data in our experiment is 2.5 Mbytes.

Table 5. Computation operations for each message type

	Our scheme		MobiCrowd	
	Sender	Receiver	Sender	Receiver
Beacon	ECDSA signing	ECDSA verification	-	-
Peer query	ECIES encryption, AES encryption, ECDSA signing	ECIES decryption, AES decryption, ECDSA verification	ECDSA signing	ECDSA verification
Peer response	AES encryption, ECDSA signing	AES decryption, ECDSA verification	None	ECDSA verification

Table 6. EC cryptographic benchmarks on OBU and P2P query/response sizes

Security	Sign (ms)	Verify (ms)	Encryption (ms)	Decryption (ms)	Query (byte)	Response (byte)
Plaintext	-	-	-	-	500	5000
EC-192	1.75	1.95	6.52	4.73	1263	6877
EC-224	2.65	2.94	8.13	5.54	1303	6897
EC-256	3.03	3.42	9.43	6.22	1343	6921

Table 5 shows computational overhead for each message type generation and validation. Table 6 shows the benchmark for EC cryptographic operations on an OBU, and query/response sizes in

⁷<https://osm.kewl.lu/luxembourg.osm/>

their plaintext, as well as signed and encrypted forms. We use UDP broadcasting for beacons from serving nodes and TCP for P2P query-response process. In the experiment, the beacon size is 340 bytes (with a signature and a PC attached) and we consider pure query and response sizes of 500 bytes and 5000 bytes (in accordance with the communication overhead in Sec. 5) respectively. The extra communication overhead for secured formats is due to the attached (encrypted) symmetric key, the digital signature and the signer’s PC. We encode all binary data (including keys, signatures and cipher texts) into Base64 format, which further increases the size by 4/3.

Figs. 25a and 25e show the average processing delay for the querying node and the serving node, for peer query and response with different combinations of N and EC key size values. $T_{query} = \text{Exp}(\frac{1}{0.5s})$ is used for this experiment. The query sending delay includes the TCP handshake delay; the query reception delay on serving nodes does not include the TCP handshake delay, because it is not possible to record the starting time point of the TCP handshake (i.e., the starting time point of accepting a connection request by the *listening socket* in C++) on the serving nodes.

For the querying node, we see the communication delays with different key sizes, given the same N , are roughly the same, because the difference in communication overhead (with different key sizes) is negligible compared to the total secured query/response size (see Table 6). For example, when $N = 3$, the average communication delay (i.e., the sum of query sending and response receiving delays) of the querying node with EC-192, EC-224 and EC-256 keys are around 68 ms, 66 ms and 67 ms respectively. At the same time, the average total processing delay slightly increases with increased key size, due to higher cryptographic processing latencies.

However, given a key size, the average communication delay experienced by the querying node significantly increases as N increases, because more peer responses are required for the same needed node queries. For example, with EC-192, the average processing delay at the querying node, with $N = 1$, $N = 2$, and $N = 3$, is around 55 ms, 68 ms and 88 ms respectively.

For the serving nodes, the average processing delay does not change significantly for different N , because a serving node needs to handle high rates of incoming query even when $N = 1$, thus may be saturated by high communication overload. From the experiments, we find that the TCP handshake delay is much higher than the networking delay for sending query/response messages. In our evaluation, the query delay includes the TCP handshake delay; while a response delay does not include a TCP handshake, because a TCP connection was already established upon receipt of the query. Figs. 25a and 25e show that the former is significantly higher than the latter one for any combination of N and EC key size values.

The communication delays in Figs. 25b and 25f are higher than those in Figs. 25a and 25e even with lower query rate values for the query generators, because the network is loaded with CAMs. Moreover, the processing delay slightly increases because a safety beacon application is running in parallel with the LBS application: CAM generation and reception need frequent signature generations and verifications. Overall, we see only a modest increase in networking and processing delays even in such demanding network setting (note: we rely on software-only cryptographic operations, refraining from using a hardware accelerator that could validate several hundreds of signatures per second).

Figs. 25c, 25d, 25g and 25h show processing delays for MobiCrowd. From Figs. 25c and 25g, we see the computational delays for MobiCrowd are lower than those in our scheme, because only signature verifications are required, without any encryption. However, we found the transmission of 2.5 Mbytes regional POI data over a TCP connection requires more than 10 s. In a realistic network with highly mobile nodes, transmission delays could be possibly higher, and essentially not workable with the current OBUs and protocol stack. Although such transmissions can be potentially done through a different type of connection, or over a Roadside Unit (RSU) as the




Fig. 25. Average processing delay on (first row) querying node and (second row) serving node. (First column) Without CAM application. (Second column) EC-256 keys with CAM application. (Third column) Computational and (fourth column) communication delays for MobiCrowd with EC-256 keys and CAM application.

proxy, our scheme or any scheme that requires exchange of much smaller data sizes is definitely advantageous, especially in a network with highly mobile nodes.

Discussion: Our scheme leverages P2P interactions, requiring node-to-node communication within the network. In this experiment, our scheme is evaluated with vehicular OBUs, each equipped with an IEEE 802.11p interface, as user devices. IEEE 802.11p is designed mainly for safety applications in VC domain. Apart from 802.11p, many network technologies support P2P (or device-to-device (D2D)) communication. WiFi-Direct [7] and Bluetooth [25] can be used for P2P communication, supported by most off-the-shelf smartphones nowadays. However, they do not support fully decentralized ad-hoc networking, i.e., a master node is needed and other nodes can join the network by connecting to the master. Therefore, they only support communication with a star topology, which is not fully decentralized. Moreover, the number of slaves a master can maintain is limited. LTE-Direct [41] and upcoming 5G networks [59] support D2D communication with the help of network infrastructure, while research on fully ad-hoc (direct) D2D communication has been ongoing [41]. While 802.11p is designed for VC, D2D communication in cellular networks could be supported by smartphones or any 4G/5G enabled device, which could promote the adoption of P2P-based technologies, as is the proposed scheme here. An alternative approach, before the full deployment and support for direct D2D communication, is using infrastructure mode WiFi networking, readily supported by off-the-shelf smartphones nowadays. Smartphone users can connect to an access point, and broadcast/multicast messages (e.g., beacons) and carry out P2P interactions (e.g., query and response) via the access point. This requires users to connect to the same access point. In pedestrian settings, e.g., in a campus, this could be an option, but this is less so in a vehicular setting, considering the limited coverage of an access point and the mobility of vehicles: frequent switching from one access point to another would be necessary and users could be connected to different access points even if they are close to each other.

7 CONCLUSION

In this paper, we proposed a secure decentralized privacy protection scheme for LBS. Our approach extends the recent P2P LBS privacy protection approach, addressing a number of practical open issues. More important, it ensures resilience to malicious nodes, and low exposure to honest-but-curious nodes and LBS servers even if they collude with an honest-but-curious identity management facility. With a simulation-based quantitative analysis, we showed that the exposure to honest-but-curious nodes is low even if a high ratio (e.g., 20 %) of nodes are honest-but-curious and collude; while the same ratio of malicious nodes could only affect less than 1.5 % of the peer-responded node queries. We showed the deployability of our scheme with an implementation on an automotive testbed and discussed the alternative approaches for deployment.

ACKNOWLEDGMENTS

This work has been supported by the Swedish Foundation for Strategic Research (SSF) SURPRISE project and the KAW Academy Fellow Trustworthy IoT project.

REFERENCES

- [1] R. Al-Dhubbani and J. Cazalas. 2017. Correlation analysis for geo-indistinguishability based continuous LBS queries. In *International Conference on Anti-Cyber Crimes (ICACC)*. Abha, Saudi Arabia, 203–208.
- [2] N. Alexiou, M. Laganà, S. Gisdakis, M. Khodaei, and P. Papadimitratos. 2013. VeSPA: Vehicular Security and Privacy-preserving Architecture. In *ACM HotWiSec*. Budapest, Hungary, 19–24.
- [3] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. 2013. Geo-indistinguishability: Differential privacy for location-based systems. In *ACM CCS*. Berlin, Germany, 901–914.
- [4] C. Ardagna, M. Cremonini, E. Damiani, S. De Capitani di Vimercati, and P. Samarati. 2007. Location privacy protection through obfuscation-based techniques. *Data and Applications Security XXI*, 47–60.
- [5] L. Barkhuus and A. K. Dey. 2003. Location-Based Services for Mobile Telephony: a Study of Users' Privacy Concerns.. In *INTERACT*. Zurich, Switzerland, 702–712.
- [6] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Liou. 2011. On the performance of secure vehicular communication systems. *IEEE Transactions on Dependable and Secure Computing* 8, 6 (2011), 898–912.
- [7] D. Camps-Mur, A. Garcia-Saavedra, and P. Serrano. 2013. Device-to-device communications with Wi-Fi Direct: overview and experimentation. *IEEE Wireless Communications* 20, 3 (2013), 96–104.
- [8] S.-M. Cheng, W. C. Ao, P.-Y. Chen, and K.-C. Chen. 2011. On modeling malware propagation in generalized social networks. *IEEE Communications Letters* 15, 1 (2011), 25–27.
- [9] C.-Y. Chow, M. F. Mokbel, and X. Liu. 2006. A peer-to-peer spatial cloaking algorithm for anonymous location-based service. In *ACM international symposium on Advances in geographic information systems*. Arlington, VA, USA, 171–178.
- [10] C.-Y. Chow, M. F. Mokbel, and X. Liu. 2011. Spatial cloaking for anonymous location-based services in mobile peer-to-peer environments. *GeoInformatica* 15, 2 (2011), 351–380.
- [11] L. Codeca, R. Frank, and T. Engel. 2015. Luxembourg SUMO traffic (LuST) scenario: 24 hours of mobility for vehicular networking research. In *IEEE Vehicular Networking Conference*. Paderborn, Germany, 1–8.
- [12] R. Dewri, I. Ray, I. Ray, and D. Whitley. 2010. Query m-invariance: Preventing query disclosures in continuous location-based services. In *Mobile Data Management (MDM) conference*. Kansas City, MO, USA, 95–104.
- [13] W. Eltarjman, R. Dewri, and R. Thurimella. 2017. Private Retrieval of POI Details in Top-K Queries. *IEEE Transactions on Mobile Computing* 16, 9 (2017), 2611–2624.
- [14] ETSI EN 302 637-2. 2014. Intelligent Transport Systems; Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service.
- [15] J. Fogel and E. Nehmad. 2009. Internet social network communities: Risk taking, trust, and privacy concerns. *Computers in human behavior* 25, 1 (2009), 153–160.
- [16] B. Gedik and L. Liu. 2008. Protecting Location Privacy with Personalized k-Anonymity: Architecture and Algorithms. *IEEE Transaction on Mobile Computing* 7, 1 (2008), 1–18.
- [17] M. Ghaffari, N. Ghadiri, M. H. Manshaei, and M. S. Lahijani. 2017. P⁴QS: A Peer-to-Peer Privacy Preserving Query Service for Location-Based Mobile Applications. *IEEE Transactions on Vehicular Technology* 66, 10 (2017), 9458–9469.
- [18] G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K.-L. Tan. 2008. Private queries in location based services: anonymizers are not necessary. In *ACM SIGMOD*. Vancouver, Canada, 121–132.
- [19] G. Ghinita, P. Kalnis, and S. Skiadopoulos. 2007. MOBIHIDE: a mobile peer-to-peer system for anonymous location-based queries. In *Advances in Spatial and Temporal Databases*. Boston, MA, USA, 221–238.

- [20] S. Gisdakis, T. Giannetsos, and P. Papadimitratos. 2014. SPPEAR: security & privacy-preserving architecture for participatory-sensing applications. In *ACM WiSec*. Oxford, UK, 39–50.
- [21] S. Gisdakis, T. Giannetsos, and P. Papadimitratos. 2015. SHIELD: A Data Verification Framework for Participatory Sensing Systems. In *ACM WiSec*. New York, NY, USA.
- [22] S. Gisdakis, M. Laganà, T. Giannetsos, and P. Papadimitratos. 2013. SEROSA: SERVICE Oriented Security Architecture for Vehicular Communications. In *IEEE Vehicular Networking Conference*. Boston, MA, USA, 111–118.
- [23] A. Gkoulalas-Divanis, P. Kalnis, and V. S. Verykios. 2010. Providing k-anonymity in location based services. *ACM SIGKDD explorations newsletter* 12, 1 (2010), 3–10.
- [24] P. Golle and K. Partridge. 2009. On the anonymity of home/work location pairs. In *PerCom*. Nara, Japan, 390–397.
- [25] C. Gomez, J. Oller, and J. Paradells. 2012. Overview and evaluation of bluetooth low energy: An emerging low-power wireless technology. *Sensors* 12, 9 (2012), 11734–11753.
- [26] M. Gruteser and D. Grunwald. 2003. Anonymous usage of location-based services through spatial and temporal cloaking. In *ACM MobiSys*. San Francisco, CA, USA, 31–42.
- [27] T. Hashem and L. Kulik. 2007. Safeguarding location privacy in wireless ad-hoc networks. In *International Conference on Ubiquitous Computing*. Innsbruck, Austria, 372–390.
- [28] B. Hoh and M. Gruteser. 2005. Protecting location privacy through path confusion. In *International Conference on Security and Privacy for Emerging Areas in Communications Networks*. Athens, Greece, 194–205.
- [29] H. Hu and J. Xu. 2009. Non-exposure location anonymity. In *ICDE*. Shanghai, China, 1120–1131.
- [30] Q. Hu, S. Wang, C. Hu, J. Huang, W. Li, and X. Cheng. 2018. Messages in a Concealed Bottle: Achieving Query Content Privacy with Accurate Location-Based Services. *IEEE Transactions on Vehicular Technology* 67, 8 (2018), 7698–7711.
- [31] H. Jin and P. Papadimitratos. 2015. Resilient Collaborative Privacy for Location-Based Services. In *Nordic Conference on Secure IT Systems*. Stockholm, Sweden, 47–63.
- [32] M. E. Johnson, D. McGuire, and N. D. Willey. 2008. The Evolution of the Peer-to-Peer File Sharing Industry and the Security Risks for Users. In *Hawaii International Conference on System Sciences*. Waikoloa, HI, USA, 1–10.
- [33] P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias. 2007. Preventing location-based identity inference in anonymous spatial queries. *IEEE Transactions on Knowledge and Data Engineering* 19, 12 (2007), 1719–1733.
- [34] M. Khodaei, H. Jin, and P. Papadimitratos. 2014. Towards deploying a scalable & robust vehicular identity and credential management infrastructure. In *IEEE Vehicular Networking Conference*. Paderborn, Germany, 33–40.
- [35] M. Khodaei, H. Jin, and P. Papadimitratos. 2018. SECMAE: Scalable and Robust Identity and Credential Management Infrastructure in Vehicular Communication Systems. *IEEE Transactions on ITS* 19, 5 (May 2018), 1430–1444.
- [36] M. Khodaei and P. Papadimitratos. 2018. Efficient, Scalable, and Resilient Vehicle-Centric Certificate Revocation List Distribution in VANETs. In *ACM WiSec*. Stockholm, Sweden, 172–183.
- [37] H. Kido, Y. Yanagisawa, and T. Satoh. 2005. An anonymous communication technique using dummies for location-based services. In *International Conference on Pervasive Services*. Santorini, Greece, 88–97.
- [38] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker. 2012. Recent Development and Applications of SUMO - Simulation of Urban Mobility. *International Journal On Advances in Systems and Measurements* 5, 3&4 (2012), 128–138.
- [39] W.-S. Ku and R. Zimmermann. 2008. Nearest neighbor queries with peer-to-peer data sharing in mobile environments. *Pervasive and Mobile Computing* 4, 5 (2008), 775–788.
- [40] S. H. Kwok, K. R. Lang, and K. Y. Tam. 2002. Peer-to-peer technology business and service models: risks and opportunities. *Electronic Markets* (2002).
- [41] X. Lin, J. Andrews, A. Ghosh, and R. Ratasuk. 2014. An overview of 3GPP device-to-device proximity services. *IEEE Communications Magazine* 52, 4 (2014), 40–48.
- [42] B. Liu, W. Zhou, T. Zhu, L. Gao, T. H. Luan, and H. Zhou. 2016. Silence is golden: Enhancing privacy of location-based services by content broadcasting and active caching in wireless vehicular networks. *IEEE transactions on vehicular technology* 65, 12 (2016), 9942–9953.
- [43] F. Liu, K. A. Hua, and Y. Cai. 2009. Query l-diversity in location-based services. In *International Conference on Mobile Data Management: Systems, Services and Middleware*. Taipei, Taiwan, 436–442.
- [44] M. T. Mahin, T. Hashem, and S. Kabir. 2017. A crowd enabled approach for processing nearest neighbor and range queries in incomplete databases with accuracy guarantee. *Pervasive and Mobile Computing* 39 (2017), 249–266.
- [45] S. Mascetti, C. Bettini, D. Freni, and X. S. Wang. 2007. Spatial generalisation algorithms for LBS privacy preservation. *Journal of Location Based Services* 1, 3 (2007), 179–207.
- [46] M. F. Mokbel, C.-Y. Chow, and W. G. Aref. 2006. The new Casper: query processing for location services without compromising privacy. In *Very Large Data Bases Conference*. Seoul, South Korea, 763–774.
- [47] H. H. Nguyen and H. Y. Jeong. 2018. Mobility-Adaptive Beacon Broadcast for Vehicular Cooperative Safety-Critical Applications. *IEEE Transactions on Intelligent Transportation Systems* 19, 6 (June 2018), 1996–2010.
- [48] B. Niu, Q. Li, X. Zhu, G. Cao, and H. Li. 2015. Enhancing privacy through caching in location-based services. In *2015 IEEE INFOCOM*. Kowloon, Hong Kong, 1017–1025.

- [49] F. Olumofin, P. K. Tysowski, I. Goldberg, and U. Hengartner. 2010. Achieving efficient query privacy for location based services. In *International Symposium on Privacy Enhancing Technologies*. Berlin, Germany, 93–110.
- [50] P. Papadimitratos, L. Buttyan, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, and J.-P. Hubaux. 2008. Secure vehicular communication systems: design and architecture. *IEEE Comm. Mag.* 46, 11 (2008), 100–109.
- [51] A. Pingley, N. Zhang, X. Fu, H.-A. Choi, S. Subramaniam, and W. Zhao. 2011. Protection of query privacy for continuous location based services. In *IEEE INFOCOM*. Shanghai, China, 1710–1718.
- [52] K. Ramachandran and B. Sikdar. 2006. Modeling malware propagation in gnutella type peer-to-peer networks. In *IEEE International Parallel and Distributed Processing Symposium*. Vancouver, Canada.
- [53] K. Sampigethaya, M. Li, L. Huang, and R. Poovendran. 2007. AMOEBA: Robust location privacy scheme for VANET. *IEEE Journal on Selected Areas in Communications* 25, 8 (2007).
- [54] R. K. Schmidt, T. Leinmuller, E. Schoch, F. Kargl, and G. Schafer. 2010. Exploration of adaptive beaconing for efficient intervehicle safety communication. *IEEE Network* 24, 1 (Jan 2010), 14–19.
- [55] R. Shokri, G. Theodorakopoulos, P. Papadimitratos, E. Kazemi, and J.-P. Hubaux. 2014. Hiding in the mobile crowd: Locationprivacy through collaboration. *IEEE Transactions on Dependable and Secure Computing* 11, 3 (2014), 266–279.
- [56] R. Shokri, C. Troncoso, C. Diaz, J. Freudiger, and J.-P. Hubaux. 2010. Unraveling an old cloak: k-anonymity for location privacy. In *ACM WPES*. Chicago, IL, USA, 115–118.
- [57] C. Sommer, R. German, and F. Dressler. 2011. Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis. *IEEE Transaction on Mobile Computing* 10, 1 (2011), 3–15.
- [58] C. Sommer, O. K. Tonguz, and F. Dressler. 2011. Traffic information systems: efficient message dissemination via adaptive beaconing. *IEEE Communications Magazine* 49, 5 (May 2011), 173–179.
- [59] M. N. Tehrani, M. Uysal, and H. Yanikomeroglu. 2014. Device-to-device communication in 5G cellular networks: challenges, solutions, and future directions. *IEEE Communications Magazine* 52, 5 (2014), 86–92.
- [60] S. Uppoor, O. Trullols-Cruces, M. Fiore, and J. M. Barcelo-Ordinas. 2014. Generation and analysis of a large-scale urban vehicular mobility dataset. *IEEE Transactions on Mobile Computing* 13, 5 (2014), 1061–1075.
- [61] X. Yi, R. Paulet, E. Bertino, V. Varadharajan, et al. 2016. Practical Approximate k Nearest Neighbor Queries with Location and Query Privacy. *IEEE Transactions on Knowledge and Data Engineering* 28, 6 (2016), 1546–1559.
- [62] H. Yun, D. Han, and C. C. Lee. 2013. Understanding the use of location-based service applications: do privacy concerns matter? *Journal of Electronic Commerce Research* 14, 3 (2013), 215.
- [63] L. Zhou, L. Zhang, F. McSherry, N. Immorlica, M. Costa, and S. Chien. 2005. A first look at peer-to-peer worms: threats and defenses. In *International Workshop on Peer-to-Peer Systems*. Ithaca, NY, USA, 24–35.