

Privacy-Conscious Location-Based Queries in Mobile Environments

Jianliang Xu, *Senior Member, IEEE*, Xueyan Tang, *Senior Member, IEEE*, Haibo Hu, and Jing Du

Abstract—In location-based services, users with location-aware mobile devices are able to make queries about their surroundings anywhere and at any time. While this ubiquitous computing paradigm brings great convenience for information access, it also raises concerns over potential intrusion into user location privacy. To protect location privacy, one typical approach is to cloak user locations into spatial regions based on user-specified privacy requirements, and to transform location-based queries into region-based queries. In this paper, we identify and address three new issues concerning this location cloaking approach. First, we study the representation of cloaking regions and show that a circular region generally leads to a small result size for region-based queries. Second, we develop a mobility-aware location cloaking technique to resist trace analysis attacks. Two cloaking algorithms, namely *MaxAccu_Cloak* and *MinComm_Cloak*, are designed based on different performance objectives. Finally, we develop an efficient polynomial algorithm for evaluating circular-region-based k NN queries. Two query processing modes, namely *bulk* and *progressive*, are presented to return query results either all at once or in an incremental manner. Experimental results show that our proposed mobility-aware cloaking algorithms significantly improve the quality of location cloaking in terms of an entropy measure without compromising much on query latency or communication cost. Moreover, the progressive query processing mode achieves a shorter response time than the bulk mode by parallelizing the query evaluation and result transmission.

Index Terms—Location-based services, location privacy, query processing, mobile computing.

1 INTRODUCTION

LOCATION-BASED services (LBS) are emerging as a major application of mobile geospatial technologies [1], [10], [16], [21]. In LBS, users with location-aware mobile devices are able to make queries about their surroundings anywhere and at any time. Spatial range queries and k -nearest-neighbor (k NN) queries are two types of the most commonly used queries in LBS. For example, a user can make a range query to find out all shopping centers within a certain distance of her current location, or make a k NN query to find out the k -nearest gas stations. In these queries, the user has to provide the LBS server with her current location. The disclosure of location information to the server, however, raises privacy concerns, which have hampered the widespread use of LBS [25]. Thus, how to provision location-based services while protecting user location privacy has recently become a hot research topic [10], [12], [18], [19], [20].

Location cloaking is one typical approach to protecting user location privacy in LBS. Upon receiving a location-based spatial query (e.g., a range query or a k NN query) from the user, the system cloaks the user's current location into a *cloaking region* based on the user's privacy requirement. The location-based spatial query is, thus, transformed into a *region-based spatial query* before being sent to the LBS server. The LBS server then evaluates the region-based query and returns a *result superset*, which contains the query results for

all possible location points in the cloaking region. Finally, the system refines the result superset to generate the exact results for the query location. Fig. 1a shows a sample NN query. Instead of providing the exact location l , the system submits a cloaking region R to the LBS server, which then returns the set of objects $\{b, c, d\}$ that are the nearest neighbors of at least one point in R . Finally, among $\{b, c, d\}$, the system finds out the true nearest neighbor b of query location l . Throughout this query processing procedure, the LBS server knows only the region R in which the user is located, not the exact location l . In the literature, a variety of cloaking algorithms based on snapshot user locations have been developed for different privacy metrics (e.g., [6], [10], [11], [18], [20]).

In this paper, we identify and address three new issues concerning the location cloaking approach. We first show that the representation of a cloaking region has an impact on the result superset size of the region-based query. In general, a small result superset is preferred for saving the cost of data transmission and reducing the workload of the result refinement process (especially if this process is implemented on the mobile client). Our findings indicate that, given a privacy requirement, representing the cloaking region with a circle generally leads to a smaller result superset than using other shapes.

Second, we consider the location cloaking problem for continuous LBS queries. In such scenarios, trace analysis attacks are possible by linking historical cloaking regions with user mobility patterns. Assume that in our previous example, the user issues a second query at location l' with a cloaking region R' (see Fig. 1b). If the LBS server somehow learns the user's maximum possible moving speed v_m , the server can draw a region R^c (the shaded area in Fig. 1b) expanded from the last cloaking region R based on v_m and the interval t between the two queries. The server is then able to infer that the user must be located in the intersection area of R^c and R' , which degrades the quality of location cloaking and may fail to meet the expected privacy

• J. Xu, H. Hu, and J. Du are with the Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, KLN, Hong Kong.
E-mail: {xujl, haibo, jdu}@comp.hkbu.edu.hk.

• X. Tang is with the School of Computer Engineering, Nanyang Technological University, Nanyang Avenue, Singapore 639798.
E-mail: asxytang@ntu.edu.sg.

Manuscript received 9 Mar. 2007; revised 17 Dec. 2008; accepted 31 Mar. 2009; published online 16 Apr. 2009.

Recommended for acceptance by S. Olariu.

For information on obtaining reprints of this article, please send e-mail to: tpds@computer.org, and reference IEEECS Log Number TPDS-2007-03-0073. Digital Object Identifier no. 10.1109/TPDS.2009.65.




Fig. 1. Dynamic location cloaking. (a) Location cloaking. (b) Isolated cloaking.

requirement. The cloaking quality will further deteriorate with the analysis of more successive queries and cloaking regions. To address this issue, we develop a mobility-aware location cloaking technique that resists trace analysis attacks. Given that the server observes a cloaking region together with any series of historical cloaking regions, our proposed technique makes equal the derivable probability that the user will be located at any one point within the cloaking region. To achieve this, we leverage the probability theory to control the generation of cloaking regions and design two cloaking algorithms, namely *MaxAccu_Cloak* and *MinComm_Cloak*, based on different performance objectives. *MaxAccu_Cloak* is designed to maximize the accuracy of query results, while *MinComm_Cloak* attempts to reduce the network communication cost.

Finally, we investigate how to evaluate efficiently circular-region-based spatial queries on the LBS server. While the evaluation of circular-region-based range queries is straightforward, we develop an efficient $\mathcal{O}(kM^3)$ algorithm for evaluating circular-region-based k NN queries, where M is the cardinality of the spatial object set. In addition, we present two query processing modes, namely *bulk* and *progressive*, which return query results either all at once or in an incremental manner.

We conduct simulation experiments to evaluate the performance of the proposed location cloaking and query processing algorithms. The results show that the proposed mobility-aware cloaking algorithms outperform an isolated cloaking algorithm in terms of an entropy measure of cloaking quality, without compromising much on query latency or communication cost (and sometimes performing even better). Regarding the end-to-end system performance, *MaxAccu_Cloak* results in a very high query accuracy, while *MinComm_Cloak* achieves a good balance between communication cost and query accuracy. When the result superset size is small, the bulk and progressive modes of query progressing perform similarly. For large result supersets that require a long time to evaluate and transmit, the progressive mode achieves a shorter user-perceived response time than the bulk mode by parallelizing the query evaluation and result transmission.

The rest of this paper is organized as follows: Section 2 surveys the related work on location privacy protection and spatial query processing. Section 3 gives an overview of our system model and location privacy metrics. Section 4 studies the representation of cloaking regions, followed by Section 5, which presents the mobility-aware location cloaking algorithms. The processing of circular-region-based queries is discussed in Section 6. Section 7 experimentally evaluates the proposed location cloaking and query processing algorithms. Finally, Section 8 concludes this paper.

2 RELATED WORK

Location privacy protection. There are two main approaches to protecting location privacy in LBS. The first approach relies on a trusted LBS server to restrict access to location data based on rule-based policies [8], [27]. The second approach runs a trustworthy agent between the client and the LBS server. Every time the user makes a location-based query, the agent anonymizes the user identity and/or location before forwarding the query to the LBS server [4], [10], [20]. Our study is under the framework of the second approach.

Early studies on location privacy protection considered object tracking applications, where a proxy server is employed to collect exact locations from moving clients and to anonymize location data through depersonalization before release. In [4], once a client enters a predefined zone, its identity is mixed with all other clients in the same zone. It appears that this idea can be extended to deal with trace analysis attacks by associating each LBS request with a different pseudonym. Unfortunately, this approach may not be effective because historical user locations are highly correlative and, hence, they could be relinked using trajectory tracking methods (e.g., multitarget tracking [22]) even without knowing any identity [26].

Gruteser and Grunwald [10] proposed to achieve identity anonymity in LBS by spatiotemporal cloaking based on a k -anonymity model, that is, the cloaked location is made indistinguishable from the location information of at least $k - 1$ other users. To perform the spatial cloaking, they used a Quad-tree-like algorithm. Gedik and Liu [11], [12] extended this to a personalized k -anonymity model, in which users can specify the parameter k at a per-message level. They also developed a new cloaking algorithm called CliqueCloak. Ghinita et al. [9] proposed a location cloaking algorithm called hilbASR, in which all user locations are sorted and grouped by Hilbert space-filling curve ordering. Bettini et al. [2] presented a framework to model various background attacks in LBS and discussed defense techniques to guarantee user anonymity. Bamba et al. [1] developed location cloaking algorithms for both k -anonymity and l -diversity models in a PrivacyGrid framework. While the above cloaking algorithms require to know the accurate locations of all users, in a more recent study [16], we developed a novel cloaking algorithm based on proximity information among mobile users, without the need to know their accurate locations. But like most existing cloaking algorithms, we considered only snapshot user locations in [16]. In this paper, we investigate the location cloaking problem for continuous LBS queries. In particular, we focus on trace analysis attacks and propose a new mobility-aware cloaking technique to resist them.

Xu and Cai [26] recently developed a trajectory cloaking algorithm that aims to reduce the cloaking area and the frequency of location updates. The idea is to use historical user locations as footprints in performing k -anonymity cloaking. In contrast, we aim to prevent an adversary from utilizing historical cloaking regions to degrade the quality of current location cloaking.

Spatial query processing. A large body of research has investigated spatial query processing, in particular k NN queries. Most k NN query algorithms have focused on disk access methods based on R-tree-like index structures [13]. The branch-and-bound approach is often employed in query evaluation to traverse the index and prune search space. Various query evaluation algorithms differ in terms of the

visiting order of index nodes and the metric used to prune search space [14], [23]. Whereas the previous studies investigated the k NN problem for a location point or a line segment only, our recent work has developed an evaluation strategy for rectangular-region-based k NN queries that retrieve the k -nearest neighbors of all possible location points in a rectangular region [15]. We remark that the strategy developed in [15] is based on the fact that the perimeter of a rectangle can be decomposed into a set of straight-line segments. But because such decomposition is infeasible for a circle, the strategy of [15] cannot be extended to evaluate circular-region-based k NN queries. In another related work [5], Cheng et al. developed algorithms for evaluating probabilistic queries over imprecise object locations. In contrast, we are interested in using imprecise locations to retrieve result supersets through region-based spatial queries.

Parallel to our work, Mokbel et al. [20] and Kalnis et al. [18] have investigated both the location cloaking and query processing problems. But our work differs from theirs in several respects. First, like other previous studies [10], [11], [12], the location cloaking algorithms in [18] and [20] account for snapshot user locations only. Neither of them considers continuous queries and trace analysis attacks. In contrast, we focus on location cloaking for protecting against trace analysis attacks for continuous queries. Second, Kalnis et al. [18] and Mokbel et al. [20] did not study the issue of how to represent a cloaking region. In this paper, we show that a circular cloaking region generally leads to a small result superset size and, thus, we focus on query processing algorithms for circular regions. Finally, Mokbel et al. [20] investigated *bulk* query processing for rectangular regions only. Though Kalnis et al. [18] developed a *bulk* processing algorithm for circular-region-based k NN queries, the algorithm has an exponential time complexity of $\mathcal{O}(M^k)$, where M is the cardinality of the spatial object set. In this paper, we propose a polynomial $\mathcal{O}(kM^3)$ algorithm for circular-region-based k NN queries. Furthermore, we develop a novel *progressive* query processing algorithm, which is favorable to slow mobile networks.

3 SYSTEM MODEL AND PRIVACY METRICS

3.1 System Model

This section describes the system model under our study. We consider mobile clients that are equipped with wireless interfaces to communicate with the Internet. We assume that mobile clients are location aware, that is, they are able to position their locations at any time. The users of mobile clients are interested in querying public spatial objects (e.g., hotels, restaurants, gas stations, etc.) related to their current locations. We consider two types of location-based spatial queries. A *range query*, specified with the user's current location l and a distance d_r , retrieves all the objects lying in the circle centered at l with radius d_r . A *k NN query*, specified with the user's current location l and a parameter k , retrieves the k -nearest objects to l .

Fig. 2 illustrates the procedure for processing a location-based query. After the user issues such a query, the mobile client sends the query $Q = \{l, q\}$, where l is the current location and q includes other query parameter(s), to a location cloaking agent. The cloaking agent then cloaks the location l into a region R ($l \in R$) based on the user's privacy requirement, and forwards the modified query $Q' = \{R, q\}$ to the LBS server. The LBS server evaluates Q' and returns the




Fig. 2. System architecture.

results of Q' to the cloaking agent. Since the results of Q' are a superset of the results of Q , the cloaking agent refines the results of Q' to obtain the exact results of Q and finally returns them to the mobile client. In this procedure, we focus on two performance objectives: 1) to optimize the quality of location cloaking with respect to trace analysis attacks while satisfying the user-specified privacy requirement, and 2) to make the size of the results of Q' as small as possible for saving the cost of data transmission and the workload of the cloaking agent in downloading and refining them.

We remark that in the system architecture, the location cloaking agent runs between the mobile client and the LBS server. It may be implemented on an Internet-resident proxy or incorporated into the mobile client. These two solutions have different performance trade-offs. The first proxy-based solution greatly alleviates the workload of the mobile client by delegating the tasks of location cloaking and result refinement to the resource-richer proxy. But, implementing the proxy-based solution is not cost free. First, the connection between the mobile client and the proxy has to be secured to prevent disclosure of location data in network transmission (e.g., by applying proper encryption and authentication protocols), which incurs extra processing overhead at the mobile client. These measures are not needed, however, in the second client-based solution. Second, since the proxy owns the private information about mobile users (including their privacy preferences as well as current and historical locations), more security risks would be introduced owing to the presence of the proxy. The proxy can become a new target of attacks and a potential performance bottleneck. A system administrator can determine where to implement the location cloaking agent by taking into consideration the bandwidth budget, client capabilities, and security requirements.

Yet, regardless of which solution the system adopts, the following issues arising from the location cloaking approach deserve our investigation: 1) how to represent cloaking regions in terms of shape such that the result size of the region-based query Q' is minimized (Section 4); 2) how to effectively perform location cloaking on the location cloaking agent so that the cloaking quality is optimized against trace analysis attacks (Section 5); and 3) how to efficiently evaluate region-based spatial queries (on the LBS server) to reduce the query response time (Section 6). It is worth noting that the techniques proposed in this paper are beneficial to both proxy-based and client-based solutions.

3.2 Privacy Metrics

We employ an intuitive privacy metric for location anonymity, that is, the area of the cloaking region (or briefly, *the cloaking area*). A user specifies a minimum acceptable cloaking area for each query. For example, a user can set the minimum acceptable cloaking area to one square mile. To consider resistance to trace analysis attacks, the




Fig. 3. Proof of Theorem 1.

quality of location cloaking is measured by *entropy*, a well-known metric for quantifying the amount of uncertainty in information theory. Suppose it can be derived that the probability density function for the user to be at location l in cloaking region R is $p(l)$, the entropy is then defined by

$$-\int_{l \in R} p(l) \ln p(l) dl. \quad (1)$$

Given a cloaking region, entropy will be zero if it is derived that the user is at some location with 100 percent probability. Entropy will increase if the user location is more uncertain, and will be maximized when the derivable probability for the user to be at any location in the region is equal.

4 REPRESENTATION OF CLOAKING REGIONS


In this section, we study the representation of cloaking regions. Given a cloaking area, we are interested in finding out how to represent the cloaking region in terms of shape such that the result size of the region-based query is minimized. It is worth noting that the representation of a cloaking region is independent of the issue of maximizing entropy in location cloaking. For any cloaking region of a given area, irrespective of its shape, entropy is maximized when the derivable probability for the user to be at any location in the region is uniform across that region.

Consider a region-based k NN query that retrieves the k -nearest neighbors of all the points in the region. The following theorem shows that the result of a region-based k NN query should include all objects in the region as well as the k NNs of the points on the perimeter of the region.

Theorem 1. *An object o is in the k NN results of region R if and only if: i) $o \in R$, or ii) o is in the k NN results of some point on the perimeter of R .*

Proof. Obviously any object inside R is the NN of the same point it occupies. Next, we use a proof-by-contradiction approach to show that if an object outside R is the i th NN ($i \leq k$) of a point inside R , this object must be in the i NN results (and hence the k NN results) of some point on the perimeter of R .

As shown in Fig. 3, suppose that object a is the i th NN of point p inside R . Assume, on the contrary, that a is not in the i NN results of any point on the perimeter of R . Consider the intersecting point p' of the segment \overline{pa} and the perimeter of R . It follows that a is not in the i NN results of p' . Thus, the i NN results of p' and p overlap by at most $i - 1$ objects. As a result, there must exist an object b in the i NN results of p' that is not in the i NN results of p . This implies $|p'b| \leq |p'a|$. Thus, we have $|pb| < |p'b| + |pp'| \leq |p'a| + |pp'| = |pa|$. This means that b is closer to p than a , which contradicts the

Fig. 4. Solution space of a region-based k NN query. (a) Convex region. (b) Circular region.

hypothesis that a is the i th NN of p , and that b is not in the i NN results of p . Hence, the theorem is proven. \square

To simplify our analysis, we follow the previous work of Berchtold et al. [3] and assume that the spatial objects to be queried are uniformly distributed in the search space. Denote by ρ the object density. According to Berchtold et al. [3], the average distance between a query point and its k th NN is given by

$$d_{kNN} = \sqrt{\frac{k}{\pi\rho}}. \quad (2)$$

Following Theorem 1, the solution space for a region-based k NN query can be approximated by the area extended from the query region by a distance of d_{kNN} (see the shaded areas in Fig. 4).¹ Thus, we estimate the size of the k NN results $|\mathcal{R}_{kNN}|$ by the number of objects lying in the approximated solution space. Let A and P , respectively, be the area and the perimeter length of the query region. For a general convex region (Fig. 4a), we obtain

$$\begin{aligned} |\mathcal{R}_{kNN}| &\doteq \left(A + P \cdot d_{kNN} + \sum_i \frac{1}{2} \theta_i d_{kNN}^2 \right) \cdot \rho \\ &= A \cdot \rho + P \cdot \sqrt{\frac{k\rho}{\pi}} + k. \end{aligned} \quad (3)$$

Similarly, for a region-based range query, we can estimate the size of its query results as

$$|\mathcal{R}_{range}| = (A + P \cdot d_r + \pi d_r^2) \cdot \rho, \quad (4)$$

where d_r is the radius of the query range.

Theorem 2. *Comparing different shapes for a cloaking region of area A , a circle gives the smallest value for both $|\mathcal{R}_{kNN}|$ in (3) and $|\mathcal{R}_{range}|$ in (4).*

Proof. Given a cloaking area A , from (3) (or (4)), the relative value of $|\mathcal{R}_{kNN}|$ (or $|\mathcal{R}_{range}|$) is determined by the perimeter length P . It is well known that a circle (see Fig. 4b) has the shortest perimeter under a fixed area. \square

Theorem 2 implies that given a cloaking area, a circular region is expected to give the smallest result set for both range and k NN queries under a uniform distribution of spatial objects. Fig. 5a compares the result sizes obtained by using both the circular and square cloaking regions of area 10^{-5} for

1. Note that this is neither a necessary nor a sufficient condition for an object to be part of the k NN results.




Fig. 5. Size of the k -NN results. (a) Uniform data set. (b) California data set.

k -NN queries on a data set containing 300,000 objects randomly distributed in a unit space. The simulation results in Fig. 5a are the average of 1,000 random queries on the data set²; the analytical results are computed using (3). It can be seen that the analytical results well match the simulation results, and the average result size given by a circular cloaking region is less than that given by a square region of the same area. We also compare circular and square cloaking regions for a real California data set where the objects are not uniformly distributed (see Section 7 for more details about this data set). As shown in Fig. 5b, a circular cloaking region again leads to a smaller result size than a square cloaking region. Thus, in the rest of this paper, we will use circles to represent cloaking regions.

5 MOBILITY-AWARE LOCATION CLOAKING

We now study how to generate circular cloaking regions based on privacy requirements. Under isolated cloaking, for each query with a cloaking area requirement A_{min} , a circle with radius $\sqrt{A_{min}/\pi}$ covering the user location l is randomly generated to serve as the cloaking region. But this scheme is vulnerable to trace analysis attacks. As discussed in Section 1, by correlating the query trace and some knowledge of mobility pattern, the LBS server (adversary) is able to derive possible user locations in the cloaking region. This leads to a significant degradation of the quality of location cloaking. In this section, we develop an optimal mobility-aware cloaking technique that works as follows: The cloaking region of the first query is generated randomly. For each subsequent query, we control the generation of cloaking regions to maximize the cloaking quality in terms of entropy as defined in (1), that is, given that the server observes a cloaking region together with any series of historical cloaking regions, the derivable probability for the user to be located at any point in the cloaking region is equal.

5.1 Problem Formulation

We consider a general user mobility pattern that is known to the mobile client. We assume, in the worst case, that the adversary also knows the user mobility pattern and, thus, has the potential to conduct trace analysis attacks. The user mobility pattern may be built by the adversary based on traces (of nonprivacy-conscious users of the same type) or mobility scenarios (e.g., the random walk model is good to model the mobility pattern of pedestrians in small-scale urban areas) [17].

Denote by O the center of the cloaking region produced for the last query (with a radius of $r = \sqrt{A_{min}/\pi}$). Assume

2. To allow for fair comparison, both the circular and square cloaking regions are formed with the query point at the centroid.




Fig. 6. Last and new cloaking regions.

that at the time of the last query, the probability for the user to be located at any point in the last cloaking region is equal. Suppose that the user moves in all directions with equal probability. Let $u(x)$ be the probability density function of the new user location being distance x away from O at the time of the new query. It follows that

$$\int_0^D u(x)dx = 1, \quad (5)$$

where D is the farthest possible distance that the user can travel since the last query, $D = \min\{y \mid \forall x \geq y, u(x) = 0\}$.


Denote by O' the center of the new cloaking region. Suppose that the generation of the new cloaking region is governed by the probability density function $p(z \mid y)$ for making O' distance z away from O given that the user is distance y away from O at the time of the new query (see Fig. 6). In order for the new cloaking region to cover the user, O' must be within a distance of r from the user's new location, i.e., $\max\{0, y - r\} \leq z \leq \min\{D - r, y + r\}$. Thus, we have

$$\int_{\max\{0, y - r\}}^{\min\{D - r, y + r\}} p(z \mid y)dz = 1. \quad (6)$$

Our goal is to determine the $p(z \mid y)$ function such that the derivable probability for the new user to be located at any point is uniform across the new cloaking region, i.e., to maximize entropy. To mathematically characterize this objective, we define $q(y \mid z)$ as the derivable probability density function of the new user location being distance y away from the center O of the last cloaking region given that the center O' of the new cloaking region is distance z away from O . For the user to be equally possible at any point in the new cloaking region, $q(y \mid z)$ should be proportional to the length of the arc (centered at O and with radius y) overlapping with the new cloaking region (as indicated by the solid arc in Fig. 7b; hereafter, referred to as the *overlapping arc length*). Below we analyze the value of $q(y \mid z)$ under maximum-entropy cloaking.

- Assume $z \geq r$. If $0 \leq y \leq z - r$ (Fig. 7a) or $y \geq z + r$ (Fig. 7c), the overlapping arc length is 0. If $z - r \leq y \leq z + r$ (Fig. 7b), the overlapping arc length is $2\alpha y = 2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y$. Therefore, after normalizing by the integration over the possible range of y values in which $q(y \mid z)$ is not zero, we obtain

$$q(y \mid z) = \begin{cases} \frac{2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y}{\int_{z-r}^{z+r} 2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y dy}, & \text{if } z - r \leq y \leq z + r, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Fig. 7. $z \geq r$.

- Assume $z < r$. If $0 \leq y \leq r - z$ (Fig. 8a), the overlapping arc length is $2\pi y$. If $r - z \leq y \leq z + r$ (Fig. 8b), the overlapping arc length is $2\alpha y = 2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y$. If $y \geq z + r$ (Fig. 8c), the overlapping arc length is 0.

Therefore, after normalization, we obtain

$$q(y | z) = \begin{cases} \frac{2\pi y}{\pi(r-z)^2 + \int_{r-z}^{z+r} 2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y dy}, & \text{if } 0 \leq y \leq r - z, \\ \frac{2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y}{\pi(r-z)^2 + \int_{r-z}^{z+r} 2 \cdot \arccos \frac{y^2 + z^2 - r^2}{2yz} \cdot y dy}, & \text{if } r - z \leq y \leq z + r, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

Having known $q(y | z)$ as expressed in (7) and (8) under maximum-entropy cloaking, our problem becomes to determine $p(z | y)$ given $q(y | z)$. Note that the relation between $p(z | y)$ and $q(y | z)$ can be established by the Bayes' rule, i.e.,


$$q(y | z) = \frac{p(z | y) \cdot u(y)}{\int_{\max\{0, x-r\}}^{\min\{R-r, x+r\}} p(z | x) \cdot u(x) dx}. \quad (9)$$

We will discuss how to solve $p(z | y)$ from (6), (7), (8), and (9) in the next section.

We remark that in our approach, only the last cloaking region is needed to generate a new maximum-entropy cloaking region. The following theorem shows the correctness of this approach.

Theorem 3. *Given that the server observes the new cloaking region and all historical cloaking regions, the derivable probability is equal for the user to be located at any point in the new cloaking region.*

Proof. Denote by (x_n, y_n) the user's location and G_n the cloaking region at the time of the n th query. Define $p(x_n, y_n)$ as the derivable probability density function of the user being at location (x_n, y_n) in region G_n . We prove the claim by induction: given that the server observes $G_1 - G_n$, for any two points $(x_n, y_n), (x'_n, y'_n)$ in G_n , $p(x_n, y_n) = p(x'_n, y'_n)$.

Fig. 8. $z < r$.

First, it is obvious that $p(x_1, y_1) = p(x'_1, y'_1)$ for $n = 1$ since the first cloaking region is randomly generated.

Next, we assume that the claim holds for some n ($n \geq 1$). Then, $p(x_n, y_n)$ is a constant $\frac{1}{\pi r^2}$. We are going to prove that the claim also holds for $n + 1$. Given $G_1 - G_{n+1}$, for any two points $(x_{n+1}, y_{n+1}), (x'_{n+1}, y'_{n+1})$ in G_{n+1} , we have

$$\begin{aligned} p(x_{n+1}, y_{n+1}) &= \iint_{G_n} p((x_{n+1}, y_{n+1}) | (x_n, y_n)) dx_n dy_n \\ &= \iint_{G_n} p((x_{n+1}, y_{n+1}) | (x_n, y_n)) \cdot p(x_n, y_n) dx_n dy_n \\ &= \iint_{G_n} p((x_{n+1}, y_{n+1}) | (x_n, y_n)) \cdot \frac{1}{\pi r^2} dx_n dy_n, \end{aligned}$$

and similarly,

$$p(x'_{n+1}, y'_{n+1}) = \iint_{G_n} p((x'_{n+1}, y'_{n+1}) | (x_n, y_n)) \cdot \frac{1}{\pi r^2} dx_n dy_n.$$

Satisfying (7) and (8), our cloaking approach ensures

$$\begin{aligned} &\iint_{G_n} p((x_{n+1}, y_{n+1}) | (x_n, y_n)) \cdot \frac{1}{\pi r^2} dx_n dy_n \\ &= \iint_{G_n} p((x'_{n+1}, y'_{n+1}) | (x_n, y_n)) \cdot \frac{1}{\pi r^2} dx_n dy_n. \end{aligned}$$

Hence, $p(x_{n+1}, y_{n+1}) = p(x'_{n+1}, y'_{n+1})$, and the theorem follows. \square

5.2 Problem Discretization

Now what we are left is to solve $p(z | y)$ from (6), (7), (8), and (9). Unfortunately, a closed-form solution is difficult to obtain. In this section, we present a discretization-based numerical method. We divide the plane into a set of rings of sufficiently small width Δ . The rings are centered at O . As shown in Fig. 9, ring 1 is enclosed by a circle centered at O with a radius of Δ , i.e., ring 1 contains all points that are within distance Δ from O . For each $i > 1$, ring i is enclosed by two circles centered at O with radii of $(i-1)\Delta$ and $i\Delta$, respectively, i.e., ring i includes all points that are $(i-1)\Delta$ to $i\Delta$ away from O .

Without loss of generality, we assume that the radius of a region $r = K\Delta$, and the farthest possible distance that the user can travel since the last query $D = L\Delta$, where K and L are integers. Based on the assumption of mobility pattern, the probability $U(i)$ of the new user location being in ring i is given by $U(i) = \int_{(i-1)\Delta}^{i\Delta} u(x) dx$, and it follows that

$$U(1) + U(2) + \dots + U(L) = 1.$$

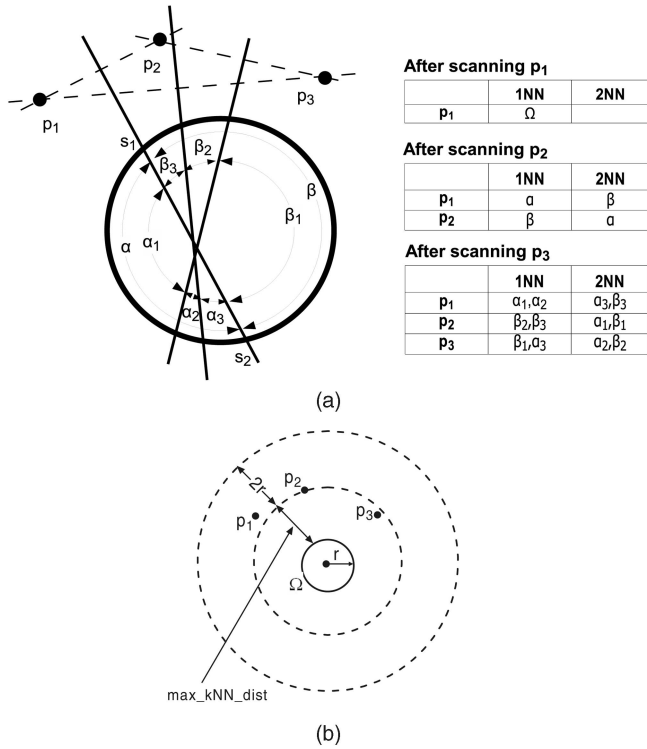


Fig. 12. Finding k CRNN results. (a) Query processing. (b) Stop condition.

and we check it against p_1 and p_2 . The perpendicular bisectors further split α into α_1, α_2 , and α_3 , and split β into β_1, β_2 , and β_3 . Now p_3 takes the place of p_1 to be the first NN for α_3 and the second NN for β_2 ; p_3 takes the place of p_2 to be the second NN for α_2 and the first NN for β_1 . In general, when object p_i is scanned, initially we assume that p_i is farther away from any arc than any candidate k CRNN result. Afterward, we check p_i against each p_j in the candidate k CRNN result set. Given a p_j , the perpendicular bisector of p_j and p_i splits the existing arcs at two points at most. For each of the arcs located on the p_i side of the perpendicular bisector, p_j moves backward in the k NN list (e.g., the second NN becomes the third NN), while p_i advances in the k NN list (e.g., the third NN becomes the second NN). After each scan, those objects which have at least one l th-NN arc ($l \leq k$) constitute the candidate set of k CRNN results. The final k CRNN results of Ω are obtained by scanning the entire set of objects. Recall that the results of a k CRNN query also include all the objects in the circular region. Thus, when scanning the objects, the algorithm also checks whether they are in the circular region and if so, includes them in the final k CRNN results.

Furthermore, in order to speed up the convergence of the k CRNN candidate set, we sort the objects and apply a heuristic (Heuristic 1) to scan the objects closest to Ω first because they are most likely to appear in the final k CRNN results.

Heuristic 1. *The objects are sorted and scanned in the increasing order of their minimum distances to Ω . And those objects whose minimum distances to Ω are $2r$ (r is the radius of Ω) farther than the k th NN of Ω are removed from scanning.*

The second statement of Heuristic 1 sets a stop condition for the scan, because those objects that are $2r$ farther from Ω

than the current k th NN of Ω must be farther away from any point on Ω than the current k NNs of Ω . This can be explained in Fig. 12b for the same 2CRNN query as in Fig. 12a. For the moment, p_1, p_2 , and p_3 are the 2CRNNs of Ω and p_2 is the second NN of Ω . The minimum distance from p_2 to Ω , denoted by max_kNN_dist , is used to prune faraway objects in future search. More specifically, if any object is more than $2r + max_kNN_dist$ away from any point on Ω (i.e., outside the outermost circle in Fig. 12b), the object does not need to be scanned since it must be farther away from any point on Ω than p_2 and p_3 .

The complete query processing algorithm is described in Algorithm 2, where the data structure $\mathcal{F}(p_i, \hat{ab})$ maintains the order of object p_i to arc \hat{ab} . We call it a *bulk* algorithm as all the candidate k CRNN results are finalized at the end of query evaluation.

Algorithm 2. Bulk Query Processing of k CRNN

Input: query circle Ω with radius r , spatial object set \mathcal{S}

Output: the k CRNN results of Ω

Procedure:

- 1: enqueue the objects in \mathcal{S} into a priority queue in increasing order of their (minimum) distances to Ω (denoted by $minDist(p_i, \Omega)$)
- 2: dequeue the first object p_1
- 3: $\mathcal{F}(p_1, \Omega) := 1$ // $\mathcal{F}(p_i, \hat{ab})$ maintains order of p_i to \hat{ab}
- 4: $cand_kCRNN_results := \{p_1\}$
- 5: $max_kNN_dist := \infty$
- 6: dequeue the next object p_i
- 7: **while** $minDist(p_i, \Omega) < 2r + max_kNN_dist$ **do**
- 8: **if** p_i is inside Ω **then**
- 9: $in_circle_results := in_circle_results \cup p_i$
- 10: initialize $\mathcal{F}(p_i, \hat{ab}) := |cand_kCRNN_results| + 1$ for any arc \hat{ab} // initially assuming p_i is farther // than any candidate k CRNN result
- 11: **for** each object p_j in $cand_kCRNN_result$ **do**
- 12: split existing arcs by $\perp p_i p_j$ —the perpendicular bisector of p_i and p_j
- 13: **for** any arc \hat{ab} located on p_i side of $\perp p_i p_j$ **do**
- 14: **if** entry $\mathcal{F}(p_j, \hat{ab})$ exists **then** $\mathcal{F}(p_j, \hat{ab})++$ // move p_j backward in the k NN list
- 15: $\mathcal{F}(p_i, \hat{ab})--$ // move p_i forward in the list
- 16: let S be the set of scanned objects including p_i ; $cand_kCRNN_results := \{p | p \in S, \exists \text{ an arc } \hat{ab}, \mathcal{F}(p, \hat{ab}) \leq k\}$
- 17: remove $\mathcal{F}()$ entries with $\mathcal{F}(p, \hat{ab}) > k$ for $p \in S$
- 18: **if** $i = k$ **then** $max_kNN_dist = minDist(p_i, \Omega)$
- 19: dequeue the next object p_i
- 20: **return** $in_circle_results \cup cand_kNN_results$ as the final results

We now analyze the complexity of Algorithm 2. Given a k CRNN query with M objects, the while loop iterates through at most M scans. Since each scan increases the number of arcs by $2M$ in the worst case, the total number of arcs for all candidate k CRNN results is bounded by $\mathcal{O}(M^2)$. Each arc can appear in the arc sets of at most k objects. Thus, the worst case storage complexity is $\mathcal{O}(kM^2)$. For the time complexity, each $\mathcal{F}(p_j, \hat{ab})$ entry may be updated at most M times, once in each scan. Hence, the worst-case time complexity is $\mathcal{O}(kM^3)$. Nevertheless, in practice the cost would be far less because the candidate set of k CRNN results

is normally not large and the scanning may terminate early with the stop condition proposed in Heuristic 1.

6.2 Progressive Query Processing of k CRNN

The bulk query processing algorithm generates the k CRNN results at the end of query evaluation. Therefore, the server cannot start transmitting the results to the client until the end of query evaluation. We now propose an alternative progressive query processing algorithm to parallelize the query evaluation and result transmission.

The idea is to determine whether an object will be a final k CRNN result earlier. The query progressing procedure remains the same as in Algorithm 2 except that 1) any object in *in_circle_results* is immediately returned to the client when it is scanned, and 2) after the scan of each object, we add a checking procedure (see Algorithm 3). We randomly pick an unchecked split point on Ω as the check point, and go through the list of unscanned objects to compute its full k NN results. If any of the k NN results has not been returned to the client, it is output for immediate transmission. For our running example shown in Fig. 12a, after scanning p_2 , we may select s_1 as the check point. We then compute s_1 's 2NN results as p_1 and p_2 and return them immediately. Compared to the bulk query processing, since the checking procedure here incurs extra overhead, the overall query processing time would be increased. Nevertheless, the worst case time complexity remains $\mathcal{O}(kM^3)$ since the checking procedure adds a complexity of $\mathcal{O}(kM^2)$ only. On the other hand, the progressive algorithm can start returning the k CRNN results earlier and, hence, is likely to result in a shorter user-perceived response time, as will be demonstrated in Section 7.4.

Algorithm 3. Checking Procedure in Progressive Query Processing of k CRNN

Input: $\mathcal{F}(p_i, \hat{ab})$ entries, the queue of unscanned objects
Output: the k NN results of a check point
Procedure: // this procedure is added to between // lines 17 and 18 of Algorithm 2

- 1: randomly select an unchecked split point s as the check point
- 2: retrieve the tentative k NN results of s : $\{p | \mathcal{F}(p, \hat{ab}) \leq k, s \in \hat{ab}\}$
- 3: $current_kNN_distance :=$ distance of the current k th NN
- 4: dequeue the next object p_i
- 5: **while** $minDist(p_i, \Omega) < current_kNN_distance$ **do**
- 6: **if** $Dist(p_i, s) < current_kNN_distance$ **then**
- 7: update the tentative k NN results
- 8: update $current_kNN_distance$
- 9: dequeue the next object p_i
- 10: return the final k NN results if not yet

7 PERFORMANCE EVALUATION

7.1 Experiment Setup

We have developed a testbed [7] to evaluate the performance of the proposed location cloaking and query processing algorithms. The client-side query interface and location cloaking agent were implemented on an O2 Xda Atom Exec PDA with Intel PXA 27 \times 520 MHz Processor and 64 MB RAM. The PDA supports GSM/GPRS/EGDE and WiFi communications. The LBS server was implemented on a

TABLE 1
Default Parameter Settings

Parameter	Setting
k NN Query (k)	5
Query Interval (I)	4 min
Privacy Requirement (r)	0.001
Spatial Object Set Size	2,249,727 objects
Object Record Size	1 kb
Query Size	20 bytes
Data Transfer Rate	114 kbps

Redhat 7.3 Linux server with Intel Xeon 2.80 GHz processor and 2 GB RAM. We assume that the client and the server communicate through a wireless network at a data transfer rate of 114 kbps.

The spatial object set used in the experiments contains 2,249,727 objects representing the centroids of the street segments in California [24]. We normalize the data space to a unit space and index the spatial objects by an R-tree (with a page fan-out of 200 and a page occupancy of 70 percent) [13]. The size of an object record is set at 1 kb. To evaluate a k CRNN query of a circle Ω , we first use our previously developed method [15] as a preprocess to retrieve the k NN results for the minimum bounding rectangle of Ω . By definition, this set of k NN results is a superset of Ω 's k CRNN results. The k CRNN processing algorithms developed in Section 6 are then applied on this superset to get the k CRNN results of Ω .

We simulate a well-known random walk model [17], in which the user moves in steps. In each step, the user selects a speed and travels along an arbitrary direction for a duration of 2 min. We test two speed settings: 1) constant speed: the moving speed is fixed at 0.0003/min; 2) random speed: for each step, a speed is randomly selected from a range of [0.0001/min, 0.0005/min]. By default, the random speed setting is adopted. The user makes privacy-conscious k NN queries from time to time. The query interval I is set at 4 min by default. The user specifies the privacy requirement by a radius r (i.e., the minimum acceptable cloaking area is πr^2), with a default setting of 0.001. The size of a k NN query message is set at 20 bytes. For the numerical method of optimal location cloaking, Δ is set at 0.0001, and the Simplex algorithm is employed to solve the linear program. The default parameter settings are summarized in Table 1. The experimental results reported below are averaged over 1,000 randomly generated queries.

7.2 Effectiveness of Mobility-Aware Cloaking

In this section, we compare the proposed optimal mobility-aware cloaking technique (Algorithm 1) against the isolated cloaking scheme (described at the beginning of Section 5). For both the optimal and isolated cloaking techniques, initially a cloaking region is randomly generated based on the user location. In other words, the user is equally likely to be at any location in the cloaking region. We measure the quality of the cloaking region for a subsequent query in terms of entropy based on 1,500 sample locations and 1,000 random queries. As shown in Figs. 13a and 13b, when the query interval is small (i.e., 1 min), the entropy of isolated cloaking is nearly 20 percent lower than that of optimal cloaking for all queries tested. With increasing query interval, the average entropy of isolated cloaking improves (see Fig. 13c) but is still far lower than that of optimal cloaking. When the query interval is

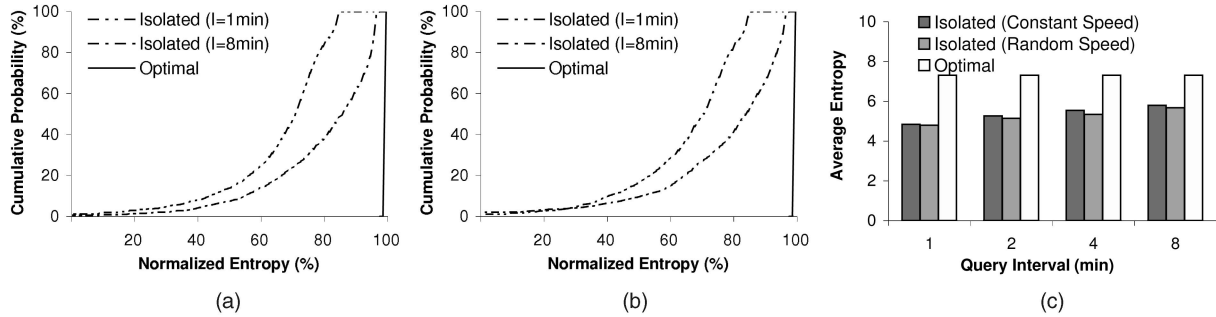


Fig. 13. Performance of cloaking quality (entropy). (a) Percentile entropy (constant speed). (b) Percentile entropy (random speed). (c) Average entropy.

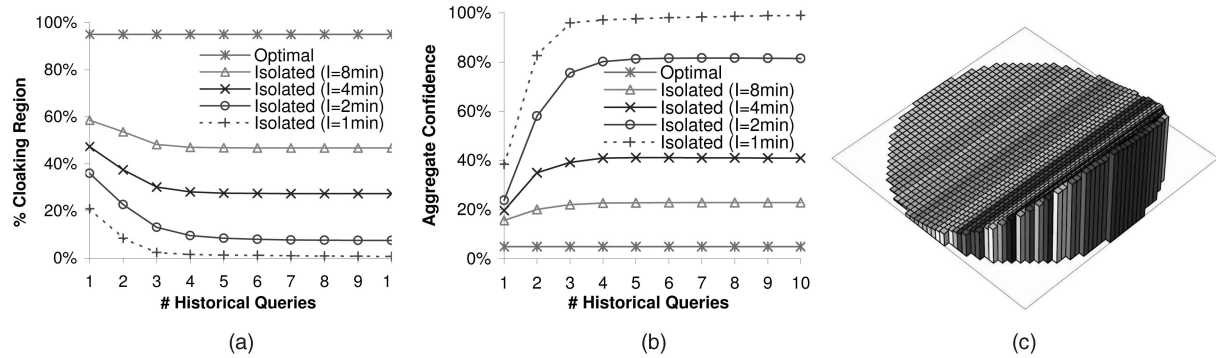


Fig. 14. Trace analysis attacks. (a) First attack (95 percent confidence). (b) Second attack (5 percent region). (c) A probability distribution of isolated cloaking.

8 min, Figs. 13a and 13b show that the entropy of isolated cloaking is 40 percent lower than that of optimal cloaking for over 15 percent of the queries tested and 20 percent lower for over 40 percent of the queries tested. Note that the results shown here are for one successive query only. With more successive queries, the quality of isolated cloaking would further degrade.

To highlight the benefit of achieving higher entropy, we conduct two possible attacks. Recall that through trace analysis attacks, the LBS server can derive the probabilities of the user being at different locations in a cloaking region. The first attack attempts to limit the possible user location to a subregion with 95 percent confidence. The second attack calculates the highest aggregate probability for any subregion with size equal to 5 percent of the cloaking region. Figs. 14a and 14b show the results when the number of historical cloaking regions used in trace analysis attacks is increased from 1 to 10. We can see from the results that our optimal cloaking is robust against the attacks: for example, concerning the first attack (Fig. 14a), the subregion size is as large as 95 percent of the cloaking region since the derivable probability for the user to be at any location is uniform across the region. In contrast, with the same level of confidence, the subregion size under isolated cloaking could be much smaller due to a skewed probability distribution (see Fig. 14c for a sample distribution we observed in the experiment). Similarly, concerning the second attack (Fig. 14b), under isolated cloaking, the server would be able to derive the probability for the user to be in a subregion of 5 percent size of the cloaking region with a confidence of 16-99 percent. The confidence for the same subregion is only 5 percent under optimal cloaking.

7.3 Comparison of Mobility-Aware Cloaking Algorithms

This section compares the two cloaking algorithms developed in Section 5 based on the optimal cloaking technique, namely MaxAccu_Cloak (abbreviated as *MaxAccu*) and MinComm_Cloak (abbreviated as *MinComm*). Recall that MaxAccu aims at a higher query accuracy by minimizing the outer query blocking probability while MinComm attempts to achieve a balance between communication cost and query accuracy by maximizing the inner query blocking probability at the same time.

As shown in Fig. 15a, MaxAccu has an outer query blocking probability of zero. Hence, its query results are 100 percent accurate as shown in Fig. 15b. In contrast, MinComm has an outer query blocking probability of about 15 percent. For those blocked outer queries, approximate results are obtained based on cached result supersets. Fig. 15b shows that the average error (measured by the ratio of the distance of an approximate k NN result to the actual k NN distance) is pretty small. In the worst case, the approximate k NN distance is no more than 2.3 times of the actual distance.

Fig. 15a also shows that MinComm has a much higher inner query blocking probability than MaxAccu. Recall that inner queries are sent to the server for evaluation merely for the purpose of optimal cloaking. They do not affect query accuracy but communication cost. With more queries (including both inner and outer queries) being blocked, the communication cost incurred by MinComm is about half that of MaxAccu (see Fig. 15c).

7.4 Comparison of k CRNN Query Algorithms

In this section, we evaluate the performance of the *bulk* and *progressive* k CRNN query processing algorithms developed in Section 6. Fig. 16 shows the user-perceived response




Fig. 15. Performance comparison of cloaking algorithms. (a) Query blocking probability. (b) Query accuracy. (c) Communication cost.




Fig. 16. Performance comparison of k CRNN query processing algorithms. (a) Varying k ($r = 0.001$). (b) Varying region size ($k = 5$).




Fig. 17. Timeline performance. (a) $k = 5, r = 0.00025$ and (b) $k = 5, r = 0.004$.

time for both algorithms. When k or r is small, the bulk and progressive algorithms perform similarly. However, when k or r is large, the progressive algorithm clearly outperforms the bulk algorithm. To explain, we show in Figs. 17a and 17b the timeline performance of two sample queries with $r = 0.25 \times 10^{-3}$ and $r = 4 \times 10^{-3}$, respectively. For the query with $r = 0.25 \times 10^{-3}$ (Fig. 17a), both the bulk and progressive algorithms took a short time to process. Thus, parallelizing the query evaluation and result transmission does not help a lot in user-perceived response time. On the other hand, when $r = 4 \times 10^{-3}$ (Fig. 17b), the result superset size is large and the query requires a long time (over 1,000 ms) to evaluate; then by returning the k CRNN results incrementally, the progressive algorithm completes the result transmission earlier.

7.5 End-to-End System Performance

This section evaluates the end-to-end system performance. In this set of experiments, we used the progressive k CRNN query processing. In addition to the MaxAccu and MinComm cloaking algorithms, the existing isolated cloaking method is also included for comparison. For all the

cloaking algorithms, the inner queries (i.e., the queries inside the last cloaking region) reuse cached result supersets and compute their answers immediately. The inner queries are not sent to the server by default. However, with MaxAccu and MinComm, some of them might need to be sent to the server to achieve optimal cloaking, depending on the inner query blocking probability. Moreover, as discussed before, the blocked outer queries for MaxAccu and MinComm compute their approximate k NN results based on cached result supersets.

Figs. 18a and 18b show the average end-to-end query latency, which is defined as the period from the time when the user issues a location-based query to the time when the exact query results are obtained. It can be seen that MinComm outperforms MaxAccu in all cases tested. This is explained as follows: As shown in Fig. 18c, the query evaluation and result transmission time dominates the overall query latency. Since MinComm has a higher outer query blocking rate and hence a higher result reuse rate (Fig. 18d), it results in a lower average query latency. For the same reason, MinComm outperforms Isolated when the region size is small (see Fig. 18b). When the

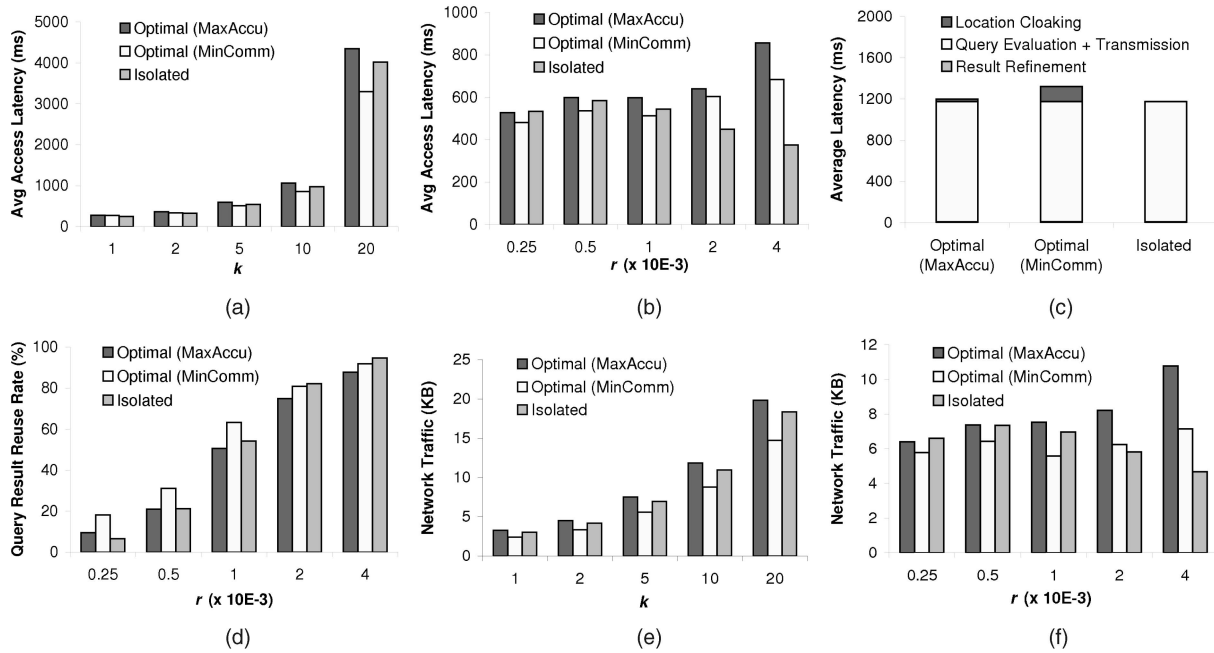


Fig. 18. End-to-end performance. (a) Latency: Varying k ($r = 0.001$). (b) Latency: Varying region size ($k = 5$). (c) Latency breakdown ($k = 5, r = 0.001$). (d) Result reuse rate ($k = 5$). (e) Traffic: Varying k ($r = 0.001$). (f) Traffic: Varying region size ($k = 5$).

region size is large, Isolated performs better than both MaxAccu and MinComm due to a higher result reuse rate. Their relative performance is insensitive to the value of k (see Fig. 18a).

Figs. 18e and 18f show the average amount of network traffic incurred for each query, which is a good indicator of energy consumption on the client. The less is the network traffic, the lower is the energy consumption. As expected, MinComm incurs less network traffic than MaxAccu due to a higher query blocking rate. Both MaxAccu and MinComm are competitive compared to Isolated; in particular, MinComm outperforms Isolated for most cases tested. Summarizing the results of Figs. 13 and 18, it can be concluded that the price to pay for resisting trace analysis attacks is not high. Our proposed MaxAccu and MinComm cloaking algorithms improve the cloaking quality over Isolated without compromising much on query latency or communication cost (and sometimes performing even better).

8 CONCLUSION

This paper has presented a complete study on processing privacy-conscious location-based queries in mobile environments. We have studied the representation of cloaking regions and showed that a circular region generally leads to a small result superset. We have developed a mobility-aware location cloaking technique to resist trace analysis attacks. Two cloaking algorithms, namely *MaxAccu_Cloak* and *MinComm_Cloak*, have been designed to favor different performance objectives. We have also developed two efficient polynomial algorithms, namely *bulk* and *progressive*, for processing circular-region-based k NN queries. In addition, we have conducted simulation experiments to evaluate the proposed algorithms. Experimental results show that the mobility-aware cloaking algorithms are robust against trace analysis attacks without compromising much on query

latency or communication cost. MaxAccu_Cloak gets a 100 percent query accuracy while MinComm_Cloak achieves a good balance between communication cost and query accuracy. It is also shown that the progressive query processing algorithm generally achieves a shorter user-perceived response time than the bulk algorithm.

As for future work, we are going to extend the mobility-aware location cloaking technique to other privacy metrics (e.g., the k -anonymity model and the l -diversity model) and road networks. We are also interested in investigating mobility-aware peer-to-peer cloaking techniques.


ACKNOWLEDGMENTS

The authors would like to thank the editor and anonymous reviewers for their valuable suggestions that significantly improved the quality of this paper. This work was supported by the Research Grants Council, Hong Kong SAR, China, under Projects HKBU211206, HKBU210808, FRG/07-08/II-23, and FRG/08-09/II-48.


REFERENCES

- [1] B. Bamba, L. Liu, P. Pesti, and T. Wang, "Supporting Anonymous Location Queries in Mobile Environments with Privacygrid," *Proc. Int'l World Wide Web Conf. (WWW)*, 2008.
- [2] C. Bettini, S. Mascetti, X.S. Wang, and S. Jajodia, "Anonymity in Location-Based Services: Towards a General Framework," *Proc. Eighth Int'l Conf. Mobile Data Management (MDM)*, May 2007.
- [3] S. Berchtold, C. Böhm, D.A. Keim, F. Krebs, and H.-P. Kriegel, "On Optimizing Nearest Neighbor Queries in High-Dimensional Data Spaces," *Proc. Int'l Conf. Database Theory (ICDT)*, 2001.
- [4] A.R. Beresford and F. Stajano, "Location Privacy in Pervasive Computing," *IEEE Pervasive Computing*, vol. 2, no. 1, pp. 46-55, Jan.-Mar. 2003.
- [5] R. Cheng, Y. Zhang, E. Bertino, and S. Prabhakar, "Preserving User Location Privacy in Mobile Data Management Infrastructures," *Proc. Privacy Enhancing Technology Workshop*, June 2006.

- [6] C.-Y. Chow, M.F. Mokbel, and X. Liu, "A Peer-to-Peer Spatial Cloaking Algorithm for Anonymous Location-Based Services," *Proc. ACM Int'l Symp. Advances in Geographic Information Systems (GIS)*, 2006.
- [7] J. Du, J. Xu, X. Tang, and H. Hu, "iPDA: Supporting Privacy-Preserving Location-Based Mobile Services (Demonstration)," *Proc. Eighth Int'l Conf. Mobile Data Management (MDM)*, May 2007.
- [8] G. Myles, A. Friday, and N. Davies, "Preserving Privacy in Environments with Location-Based Applications," *IEEE Pervasive Computing*, vol. 2, no. 1, pp. 56-64, Jan.-Mar. 2003.
- [9] G. Ghinita, P. Kalnis, and S. Skiadopoulos, "Prive: Anonymous Location-Based Queries in Distributed Mobile Systems," *Proc. Int'l World Wide Web Conf. (WWW '07)*, pp. 371-380, 2007.
- [10] M. Gruteser and D. Grunwald, "Anonymous Usage of Location-Based Services through Spatial and Temporal Cloaking," *Proc. ACM MobiSys*, 2003.
- [11] B. Gedik and L. Liu, "A Customizable k-Anonymity Model for Protecting Location Privacy," *Proc. IEEE Int'l Conf. Distributed Computing Systems (ICDCS)*, 2005.
- [12] B. Gedik and L. Liu, "Protecting Location Privacy with Personalized k-Anonymity: Architecture and Algorithms," *IEEE Trans. Mobile Computing*, vol. 7, no. 1, pp. 1-18, Jan. 2008.
- [13] A. Guttman, "R-Trees: A Dynamic Index Structure for Spatial Searching," *Proc. ACM SIGMOD*, 1984.
- [14] G.R. Hjaltason and H. Samet, "Distance Browsing in Spatial Databases," *ACM Trans. Database Systems*, vol. 24, no. 2, pp. 265-318, 1999.
- [15] H. Hu and D. Lee, "Range Nearest Neighbor Queries," *IEEE Trans. Knowledge and Data Eng.*, vol. 18, no. 1, pp. 78-91, Jan. 2006.
- [16] H. Hu and J. Xu, "Non-Exposure Location Anonymity," *Proc. IEEE Int'l Conf. Data Eng. (ICDE)*, 2009.
- [17] B.D. Hughes, *Random Walks and Random Environments*. Oxford Univ. Press, 1996.
- [18] P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias, "Preserving Anonymity in Location Based Services," technical report, School of Computing, The Nat'l Univ. of Singapore, 2006.
- [19] H. Kido, Y. Yanagisawa, and T. Satoh, "An Anonymous Communication Technique Using Dummies for Location-Based Services," *Proc. Int'l Conf. Pervasive Services (ICPS)*, 2005.
- [20] M.F. Mokbel, C.-Y. Chow, and W.G. Aref, "The New Casper: Query Processing for Location Services without Compromising Privacy," *Proc. Int'l Conf. Very Large Data Bases (VLDB)*, 2006.
- [21] X. Pan, J. Xu, and X. Meng, "Protecting Location Privacy against Location-Dependent Attack in Mobile Services," *Proc. ACM Conf. Information and Knowledge Management (CIKM)*, 2008.
- [22] D. Reid, "An Algorithm for Tracking Multiple Targets," *IEEE Trans. Automatic Control*, vol. 24, no. 6, pp. 843-854, Dec. 1979.
- [23] N. Roussopoulos, S. Kelley, and F. Vincent, "Nearest Neighbor Queries," *Proc. ACM SIGMOD*, 1995.
- [24] *The R-Tree Portal*, <http://www.rtreeportal.org/>, 2009.
- [25] B. Schilit, J. Hong, and M. Gruteser, "Wireless Location Privacy Protection," *Computer*, vol. 36, no. 12, pp. 135-137, Dec. 2003.
- [26] T. Xu and Y. Cai, "Exploring Historical Location Data for Anonymity Preservation in Location-Based Services," *Proc. IEEE INFOCOM*, 2008.
- [27] M. Youssef, V. Atluri, and N.R. Adam, "Preserving Mobile Customer Privacy: An Access Control System for Moving Objects and Custom Profiles," *Proc. Sixth Int'l Conf. Mobile Data Management (MDM)*, 2005.




Jianliang Xu received the BEng degree in computer science and engineering from Zhejiang University, Hangzhou, China, in 1998 and the PhD degree in computer science from Hong Kong University of Science and Technology in 2002. He is an associate professor in the Department of Computer Science, Hong Kong Baptist University. He was a visiting scholar in the Department of Computer Science and Engineering, Pennsylvania State University, University Park. His research interests include data management, mobile/pervasive computing, wireless sensor networks, and distributed systems. He currently serves as a vice chairman of the ACM Hong Kong Chapter. He is a senior member of the IEEE.



Xueyan Tang received the BEng degree in computer science and engineering from Shanghai Jiao Tong University, Shanghai, China, in 1998 and the PhD degree in computer science from Hong Kong University of Science and Technology in 2003. He is currently an associate professor in the School of Computer Engineering, Nanyang Technological University, Singapore. His research interests include distributed systems, mobile and pervasive computing, wireless sensor networks, Web, and Internet. He has published more than 40 technical papers in the above areas, mostly in prestigious journals and conference proceedings. He has served as a program committee member for many international conferences. He is a senior member of the IEEE.



Haibo Hu received the PhD degree in computer science from Hong Kong University of Science and Technology in 2005. He is an assistant professor in the Department of Computer Science, Hong Kong Baptist University. Prior to this, he held several research and teaching posts at the HKUST and the HKBU. His research interests include mobile and wireless data management, location-based services, and privacy-aware computing. He has published 20 research papers in international conferences, journals, and book chapters. He is also the recipient of many awards, including the ACM Best PhD Paper Award and the Microsoft Imagine Cup.



Jing Du is an MPhil student in the Department of Computer Science, Hong Kong Baptist University. His research interests include location-based services and data privacy.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.