

Privately Querying Location-based Services with SybilQuery

Pravin Shankar, Vinod Ganapathy, and Liviu Iftode

Department of Computer Science, Rutgers University

{spravin,vinodg,iftode}@cs.rutgers.edu

ABSTRACT

To usefully query a location-based service, a mobile device must typically present its own location in its query to the server. This may not be acceptable to clients that wish to protect the privacy of their location. This paper presents the design and implementation of SybilQuery, a fully decentralized and autonomous k -anonymity-based scheme to privately query location-based services. SybilQuery is a client-side tool that generates $k - 1$ Sybil queries for each query by the client. The location-based server is presented with a set of k queries and is unable to distinguish between the client's query and the Sybil queries, thereby achieving k -anonymity. We tested our implementation of SybilQuery on real mobility traces of approximately 500 cabs in the San Francisco Bay area. Our experiments show that SybilQuery can efficiently generate Sybil queries and that these queries are indistinguishable from real queries.

Author Keywords

Location-based Services, Vehicular Computing, Privacy, Anonymity.

ACM Classification Keywords

K.6.5 Management of Computing and Information Systems: Security and Protection

General Terms

Design, Experimentation.

INTRODUCTION

Modern mobile devices contain global positioning systems (GPS) that let these devices precisely determine their location. Coupled with Internet connectivity via 3G and WiFi, GPS helps transform these mobile devices to truly ubiquitous devices that can make *location-based queries*. A querying client device sends its location to a server, which responds with answers specific to the client's location. Examples of such queries include real-time traffic queries, *e.g.*, "how congested is the route from my home to my office?" and points-of-interest queries, *e.g.*, "what are the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Ubicomp 2009, Sep 30 – Oct 3, 2009, Orlando, Florida, USA.

Copyright 2009 ACM 978-1-60558-431-7/09/09...\$10.00

restaurants close to my current location?" Several modern devices, such as the iPhone, are pre-installed with query interfaces to location-based services, such as Google Maps.

However, clients that wish to protect their privacy may not want to reveal their current location to a location-based server (LBS). Client location can be misused by LBSs in ways that range from serious breaches of privacy, such as spying on the client's daily commuting patterns, to simple annoyances, such as flooding the client with location-based spam and advertising in response to a query. Recent evidence suggests that privacy concerns can deter the widespread use of LBSs. For example, EZPass RFID tags have seen a low adoption rate due to privacy concerns among drivers [10, 11]. Similarly, the New York City cab drivers' group called a strike upon the introduction of a plan to fit all cabs with GPS tracking devices [36]. Such concerns motivate the need for techniques that protect the privacy of client location from LBSs.

Several prior techniques [1, 5, 8, 13, 14, 16, 17, 19, 23, 25, 34, 45] to protect the privacy of client location fall under the broad umbrella of k -anonymity [39, 42]. In a k -anonymity-based scheme, an LBS is unable to distinguish a querying client from a group of at least k clients, where k is a security parameter. For example, Spatial cloaking [19] is one such scheme that sends cloaked regions (*e.g.*, rectangular blocks) containing at least k clients to the LBS. Unfortunately, most previously proposed k -anonymity-based schemes require a trusted third party, called an *anonymizer* that accepts and forwards queries from a client to an LBS and ensures the k -anonymity property. Because anonymizers must compute cloaked regions for each client and process the LBS' query results for these regions, they can cause performance and scalability bottlenecks. In addition, most of these schemes do not work for mobile clients. Recently proposed alternatives to eliminate anonymizers include peer-to-peer techniques [5, 16, 17] and protocols based on private information retrieval (PIR) [15, 22]. Peer-to-peer techniques rely on the participation of k peers to ensure k -anonymity, in a manner akin to Crowds [37]. However, reliance on peers restricts the autonomy of such systems. Techniques based on PIR provide strong cryptographic guarantees, but are currently computationally expensive.

This paper presents the design and implementation of *SybilQuery*, a fully decentralized and autonomous k -anonymity-based system for location-based queries. SybilQuery is a client-side tool that operates by synthetically generating $k - 1$ *Sybil queries* for each location-based query by a client.

Sybil queries contain locations that resemble the client’s actual location. For example, Sybil queries for a real query from a busy downtown area will also be from areas with similar traffic conditions. SybilQuery sends these k queries to the LBS, which is unable to identify the original query from the synthetic queries, thereby ensuring k -anonymity of the client’s location. SybilQuery’s design offers several advantages:

(a) Performance. SybilQuery eliminates the need for anonymizing servers, and the scalability and performance bottlenecks associated with centralized designs. For example, choosing higher values of k for better privacy leads to greater computational overheads at the anonymizer [1, 34]. SybilQuery’s decentralized design offloads the computations needed to achieve k -anonymity to the clients themselves. Our experiments show that the costs of generating Sybil queries at the client are negligible.

(b) Autonomy. In contrast to prior work where queries generated by an anonymizer (or peer-to-peer techniques) depend on *other* clients of the LBS, SybilQuery generates Sybil queries autonomously. For example, cloaked regions generated by spatial cloaking techniques [19] must have at least k clients; this in turn depends on the geographic distribution of clients [26].

(c) Ease of deployment. A decentralized and autonomous design lends itself to easy deployment. Clients do not have to rely on peer participation or on service providers to deploy anonymizers in order to achieve k -anonymity. Unlike PIR-based techniques, SybilQuery does not require any changes to the LBS and only requires minor modifications to the querying client. We have integrated SybilQuery with LBSs such as Google Maps, Live Maps and Yahoo! Maps.

We have implemented a prototype of SybilQuery for vehicular networking applications. Vehicles today are increasingly being equipped with computing device having GPS capability and internet connectivity, which makes vehicles an interesting ubiquitous computing environment. Our SybilQuery system takes as input a path to be followed by a vehicle along which the vehicle may issue several queries to LBSs. SybilQuery outputs $k - 1$ Sybil paths that statistically resemble the input path. The vehicle sends k queries when it accesses an LBS—its actual location, as well as $k - 1$ waypoints derived from the Sybil paths. We evaluated this prototype using real mobility traces of approximately 500 cabs in San Francisco [2]. We also conducted a user study with 15 volunteers which confirmed that Sybil paths produced by SybilQuery closely resemble real paths.

PROBLEM DEFINITION AND ASSUMPTIONS

An LBS is a database that stores a set of tuples $\langle \ell, v \rangle$, where ℓ ’s are geographic locations, *e.g.*, represented using latitudes/longitudes, and each v denotes value(s) associated with location ℓ , *e.g.*, restaurants or current traffic conditions at ℓ . We consider a model in which mobile clients periodically query the LBS as they move from a source to a destination. Each query contains the current location of the client, and the LBS returns the set of values associated with that location.

The problem of privately querying an LBS is for a client to issue location-based queries without revealing its current location to an adversary. The SybilQuery system aims to achieve this goal by sending $k - 1$ *Sybil queries* along with each real query by the client. These queries contain locations along $k - 1$ *Sybil paths*, which are synthetically generated by SybilQuery. To prevent an adversary from identifying the client’s actual location, SybilQuery must generate Sybil paths (and thereby, Sybil queries) that closely resemble real paths.

We make the following assumptions about the adversary:

(a) Adversary cannot access client’s identifiers. A fully-constructed client query to the LBS logically consists of three entities: (i) *a network identifier*, such as the IP address of the client’s device that issues the location-based query; (ii) *an explicit user identifier*, such as the user’s registered account name at the LBS or an HTTP cookie; and (iii) *the client’s location*, as may be obtained by the client using GPS. We assume that the adversary does not have access to the client’s network identifiers. To hide network identifiers from the adversary, we assume that the client’s service provider (*e.g.*, the cellular provider) and local access points, which know the client’s current location, do not collude with the adversary (*e.g.*, the LBS). We also assume that the LBS does not require the user to send explicit user identifiers in order to use the service. This is indeed a practical assumption (and is also a standard assumption in prior work); several popular LBSs, such as Google Maps, Microsoft Live Maps and Yahoo! Maps, do not require a registered user or HTTP cookies for a client to usefully query them. We further assume that the adversary does not use the query content to infer the client’s location. Thus, the adversary can only observe the location field in each client query.

(b) Adversary could be active or passive. A passive adversary can observe the contents of all location-based queries as well as the responses received from the LBS. On the other hand, an active adversary can modify responses from the LBS as a way to compromise client privacy. For example, consider an adversarial LBS that reports real-time traffic information. Upon receiving a request for traffic conditions at a particular location ℓ , the LBS could return false information, *e.g.*, by reporting heavy traffic at ℓ . This response from the LBS may cause the client to behave differently, *e.g.*, take a detour to avoid heavy traffic. The SybilQuery system must therefore respond appropriately, *e.g.*, by appropriately modifying the Sybil queries to statistically match the real query.

(c) Statistical background knowledge. The adversary could have access to *global* statistical knowledge, such as the average traffic densities at different locations at different times during the day. This information can potentially be used to compromise client privacy. The SybilQuery prototype can currently defend against several such statistical background knowledge attacks, particularly those that use knowledge about regional traffic. However, because we cannot anticipate all such attacks, we have designed SybilQuery to support an extensible architecture. In particular, it can incorporate several statistical features in its algorithm to generate Sybil paths, thereby making




Figure 1: Basic design of SybilQuery.

the system robust to future attacks. SybilQuery however *cannot* defend against adversarial LBSs that have *targetted* background information about a specific client. If the LBS knows a particular client's daily commuting patterns or knows specific locations that the client visits, it can differentiate between real queries and Sybil queries sent by that client with high probability. For example, if the LBS knows a client's residential address (R) and his work address (W), it can identify the real path followed by the client by searching for a path that starts at R and ends at W .

HIDING CLIENT LOCATION USING SYBIL QUERIES

The SybilQuery system presents an interface akin to existing navigation systems. A user enters the address of the source and destination of a trip, and the system outputs a path P for the user to follow. A destination prediction scheme, such as the one proposed by Krumm and Horvitz [28], can be optionally used to alleviate the need for the user to enter the destination. Additionally, SybilQuery requires as input a security parameter k , that it uses to generate $k - 1$ Sybil paths. Each path is represented by the system as a sequence of waypoints enroute from the source to the destination. As depicted in Figure 1, SybilQuery consists of three modules: an endpoint generator, a path generator, and a query generator. We describe each of these modules in detail below.

The endpoint generator produces $k - 1$ synthetic start and end points that statistically resemble the source and destination input by the user so that the real trip does not stand out from the synthetic trips. To do so, the endpoint generator requires a database of regional traffic statistics. This database reflects historic traffic trends at various geographic locations in the area (it is *not* a database of real-time traffic conditions). The high level idea is that the endpoint generator processes this database to identify clusters of locations that share similar features, such as traffic density (details of specific features used in our implementation are deferred to the next section). This process is a one-time activity and suffices to produce synthetic start and end points for multiple trips, *e.g.*, until the database is refreshed with newer statistics. The endpoint generator chooses synthetic start and end points from the clusters to which the actual source and destination belong. In doing so, it also ensures that the Euclidean distance between the source and destination of the actual path are within a threshold of the Euclidean distance between the synthetic start and end points.

The path generator uses the k start and end points, including the actual source and destination, to produce k

paths, each represented as a sequence of waypoints. In generating paths, it consults a database of regional maps (as do existing on-board navigation systems). In fact, the path generator may be implemented using existing navigation systems, such as Google Maps, Yahoo! Maps, and on-board GPS devices, because their interfaces are similar. Both the endpoint generator and path generator need to be invoked only once per trip.

The query generator is triggered when the user queries the LBS with his current location ℓ . Intuitively, it simulates the motion of users along the $k - 1$ Sybil paths and generates $k - 1$ Sybil locations. In its simplest form, the query generator simply computes the offset of ℓ from the source of the user's path P , and applies similar offsets to the sources of the Sybil paths to produce $k - 1$ Sybil locations. However, it can also use current traffic conditions to more accurately simulate user movement along Sybil paths (*e.g.*, simulate slower movement if traffic is congested at one of the Sybil locations).

The modular design of SybilQuery offers several advantages. It lends itself to easy deployment, even with legacy devices. The endpoint generator is a standalone component that can possibly be installed as a user-level application either on the user's desktop or on his mobile device. The path generator can be implemented using off-the-shelf navigation systems. Using a navigation system enables robust path generation based on regional maps, which allows SybilQuery to generate Sybil paths that respect features such as one-ways and road closures. It also allows SybilQuery to robustly handle detours from the path P suggested by the navigation system. Upon a detour, SybilQuery can simply trigger the path generator to produce a new path to the destination. Because paths are generated independent of each other, SybilQuery can also simulate detours in Sybil paths. SybilQuery's design only requires the query generator to be modified to send k queries to the LBS instead of one. Based upon the deployment scenario, even this modification should be relatively easy. For example, if an end user employs a Web browser on his mobile phone to send location-based queries to an LBS, the query generator could be implemented as a browser extension.

Several extensions to the basic design above can improve the robustness of the system to potential attacks.

(a) Randomizing path selection. Path generators implemented using navigation systems typically return the shortest route to the destination. However, real users may not always follow the shortest path to the destination because of factors such as detours and road closures. If SybilQuery always produces shortest Sybil paths and the user chooses a longer path to the destination, an adversary will be able to differentiate the real path from Sybil paths with high probability.

This problem is addressed with path generators that can compute multiple paths to the destination (each with varying lengths). Instead of choosing the shortest path, the path generator instead uses a probability distribution (of the frequency with which users choose paths other than the

shortest path) to make an appropriate choice from one of the available paths to the destination.

(b) Handling active adversaries. An actively adversarial LBS may return doctored query responses as an attempt to differentiate Sybil paths from a client’s real path. For example, suppose that an adversarial LBS falsely reports traffic congestion at a particular location in response to a user query. In response, a real user will likely take a detour. If the SybilQuery system does not respond similarly to doctored responses (*e.g.*, by checking for congestion and triggering detours in Sybil paths) an adversary could likely distinguish the real path from Sybil paths.

SybilQuery handles active adversaries using N -variant queries. In this technique, SybilQuery queries multiple LBSs and compares their responses in a manner akin to N -variant systems [6]. Unless all the LBSs queried collude with each other, adversarial LBSs are likely to be detected.

(c) Endpoint caching. Suppose that a real path P frequented by the user (*e.g.*, commuter paths, such as home to office) is associated with multiple sets of Sybil paths. An LBS that observes paths over a period of time can statistically identify P as the real path. Similarly, when a user travels to a new location shortly after arriving at the first destination, if the second set of Sybil trips are randomly generated, then the real paths can be easily distinguished from the Sybil paths since the real paths share an endpoint. SybilQuery handles the above attacks by performing three types of caching: (1) for the most common trips by the user, the Sybil endpoints are cached; (2) if the user makes multiple trips from one common endpoint (*e.g.*, his home or office), the corresponding Sybil endpoints are cached; and (3) when the user embarks on a multi-destination trip, the start points of the Sybil trips are cached so that the endpoint of a trip is the same as the startpoint of the following trip.

(d) Providing path continuity. Although the paths generated by SybilQuery statistically resemble each other, it is possible for the trip durations to differ (for both real and Sybil trips). If a real trip ends before some of the Sybil trips end, and the system stops sending queries, the LBS can differentiate the real path from Sybil paths. SybilQuery guards against this by being an “always on” tool that continues to simulate movement along Sybil paths even when the user’s real trip is complete. This feature ensures security even when the lengths/durations of Sybil trips do not match those of real trips.

(e) Adding GPS sensor noise. The location returned by GPS devices is typically inaccurate, with the errors following a Gaussian distribution [7]. If the Sybil locations always fall on a road segment, an adversary can differentiate them from locations returned by real GPS devices. SybilQuery can prevent this by adding a random noise to each Sybil location.

THE SybilQuery PROTOTYPE

In this section, we describe the implementation of our SybilQuery prototype. The bulk of this section is devoted to the description of the endpoint generator, which is implemented as a Python client that uses a PostgreSQL database backend with PostGIS spatial extensions to store regional traffic information. The path generator is implemented as a client that queries an off-the-shelf service [35] to find the

set of waypoints corresponding to the shortest path, given a startpoint and an endpoint. Finally, the query generator simulates user movement along all the k paths.

The Endpoint Generator

As discussed in the previous section, the endpoint generator uses the source and destination addresses input by the user to generate $k - 1$ synthetic endpoints. To ensure that these synthetic endpoints statistically resemble the actual endpoints of the user’s trip, the endpoint generator uses a regional traffic database. The endpoint generator first preprocesses the database to identify key features of each geographic location.

Regional Traffic Database. For our prototype, we used a month-long (August 25, 2008–September 24, 2008) set of GPS traces from the Cabspotting project [2] as regional traffic database. The Cabspotting project tracks the mobility of cabs in the San Francisco Bay area.¹ Each cab that participated in this project was outfit with a GPS device that updated its location with a server each minute. Each of these updates was a quadruple:

\langle timestamp, cab ID, current location, flag \rangle

Here, flag indicates whether the cab was metered (*i.e.*, in a trip) or empty. We used flag to convert the GPS traces into trips (*i.e.*, to demarcate sources and destinations during which the cab was metered). Overall, our database contains a total of 529,533 trips by 530 unique cabs; we observed at least 447 cabs on any given day. Although we tailor our discussion in the rest of this section to this database, we emphasize that the techniques employed by SybilQuery are applicable to any traffic database.

Preprocessing the Traffic Database. An ideal implementation of the endpoint generator would use an annotated database of the local region to identify synthetic endpoints that resemble the endpoints input by the user. The annotated database would contain descriptive tags for each geographic location, such as “parking lot,” “downtown office building,” or “freeway.” The endpoint generator would output synthetic sources and destinations whose tags match the corresponding tags of the source and destination supplied by the user. However, such annotated databases are laborious to create and even so are unlikely to be comprehensive or contain tags suitable for LBSs in different application domains.

The preprocessing step addresses this problem by automatically computing a feature set that describes each geographic location using information from the traffic database. These automatically extracted features then serve as tags that the endpoint generator can use to find synthetic sources and destinations that statistically resemble the user’s input.

Our implementation uses *traffic density* as the feature that characterizes each geographic location. We define the traffic density τ_ℓ of a geographic location ℓ as the number of trips that start, end, or traverse through ℓ in a fixed interval of time

¹Consequently, our experiments also focus on trips in this geographic region. However, SybilQuery can automatically generate Sybil queries for any geographic region if provided with a similar traffic database for that region.

(we describe in detail below how our system represents geographic locations; for simplicity, it suffices to think of them as fixed-size rectangular regions). We calculated τ_ℓ for each location ℓ using the PostGIS spatial extension of PostgreSQL. Because τ_ℓ can acquire a large number of discrete values, we used a simple clustering algorithm to group “similar” values of τ_ℓ into a single cluster. Although any clustering technique may be used, we found that even simple clustering algorithms, such as separating values of τ_ℓ into buckets, serve our purpose well. We empirically verified that these clusters of the San Francisco Bay area share similar semantic properties. For example, different shopping centers were grouped together into the same cluster as were residential areas.

In addition to computing the traffic density of each geographic location ℓ , the preprocessing step also computes a probability distribution function (PDF) π_ℓ of the length of a trip that originates at ℓ . In particular, it computes the following quantity for each location ℓ for all values of len .

$$\pi_\ell(len) = \frac{\# \text{ trips of length } len \text{ that start in } \ell}{\# \text{ trips that start in } \ell}$$

This PDF is used by the endpoint generation algorithm to choose appropriate geographic locations as sources for Sybil paths.

Temporal traffic patterns. Traffic densities at a geographic location ℓ vary depending on the hour of the day and the day of the week, which results in temporal patterns in the values of τ_ℓ . To reflect these temporal patterns in the features of each geographic location, we define six kinds of temporal states based on the patterns in our dataset, namely, “peak interval/weekday” (7pm-11pm), “normal interval/weekday” (6am-7pm and 11pm-3am), “off-peak interval/weekday” (3am-6am), “peak interval/weekend” (7am-11pm), “normal interval/weekend” (6am-7pm and 11pm-3am) and “off-peak interval/weekend” (3am-6am). The precomputation step computes six values of τ_ℓ for each location, corresponding to each of the temporal states above. When the user supplies a source and a destination, the endpoint generator chooses the value of τ_ℓ to compute synthetic endpoints based upon the timestamp in the user’s input.

Representing geographic locations. To compute τ_ℓ , the preprocessing algorithm needs an appropriate data structure to represent the location ℓ . This data structure must simultaneously balance precision and scalability, *i.e.*, it must store precise values of τ_ℓ for each location ℓ and must readily scale to large geographic regions, such as the San Francisco Bay area. To better illustrate this tradeoff, suppose that geographic locations are represented using fixed-size rectangular blocks, as was the case in an early implementation of our prototype. We found that a block size of 400m×400m resulted in a manageable number of blocks for the entire Bay area (about 100,000 blocks), but did not accurately represent traffic densities in crowded areas, such as the downtown region and the airport. On the other hand, a block size of 25m×25m accurately represented traffic densities, but resulted in a large number of blocks (about 6.4 million blocks).

SybilQuery therefore uses an adaptive data structure, the Quadtree [12], to represent traffic densities. This data structure represents geographic locations using blocks of varying sizes (smallest size is 25m×25m) depending on the traffic density of that location. In particular, geographic regions with non-uniform distribution of traffic are represented using small block sizes, while regions with a uniform distribution of traffic are represented using larger block sizes. Using the Quadtree data structure, we were able to represent traffic densities for the entire San Francisco Bay area using just 16,000 blocks.




Figure 2: Quadtree representations of geographic locations.

Output of preprocessing.

Figure 2 pictorially represents the output of the preprocessing step. Figure 2(a) shows the quadtree representation of the entire San Francisco Bay area. The adaptive nature of the Quadtree data structure results in different block sizes for different geographic locations. Although each of these blocks has uniform traffic density, in our experience, areas with high traffic density tend to be represented with smaller block sizes. This effect is most pronounced in areas with dense traffic, such as the airport, depicted in Figure 2(b). Note that the regions with highest traffic density, such as the freeway and the pickup area in front of the airport are represented using small blocks, while areas inside the airport that have no traffic are represented using larger block sizes. Not all blocks have the same traffic density; in Figure 2(b), blocks with the highest traffic densities are filled in black.

Generating Sybil Endpoints. When provided with a source and a destination input by the user, the endpoint generator computes a Sybil source / destination pair as shown in Algorithm 1 (our prototype adapts the same algorithm to generate $k - 1$ Sybil source / destination pairs; for brevity, we only illustrate the algorithm for one pair). It first finds the geographic locations (in the Quadtree representation) ℓ_{src} and ℓ_{dst} that contain the source and destination addresses input by the user. Next, it computes the set of all source locations ℓ that satisfy two conditions: (1) the traffic density of ℓ is approximately that of ℓ_{src} ; and (2) the probability of a trip of length $dist$ originating from ℓ matches that of ℓ_{src} . The key intuition behind this step is to find the set of source locations that closely resemble the source location input by the user. It randomly chooses a source location $\ell_{src'}$ from this set of locations, and finds a destination location $\ell_{dst'}$ within a radius $dist$ of $\ell_{src'}$ whose traffic density matches that of ℓ_{dst} . The algorithm generates end points that are of similar lengths but not exactly the same length as the real

```

Algorithm : Generate_Sybil_Endpoints(src, dst)
Input : (a) src: Street address of user's source;
          (b) dst: Street address of user's destination.
Output : (src', dst'): A Sybil source/destination pair.
1 dist = Euclidean distance between src and dst;
2  $\ell_{src}$ ,  $\ell_{dst}$  = geographic locations that contain src, dst;
3 Sources = Set of all locations  $\ell$  with (i)  $\tau_\ell \approx \tau_{\ell_{src}}$  and (ii)  $\pi_\ell(\text{dist}) \approx \pi_{\ell_{src}}(\text{dist})$ ;
4  $\ell_{src'}$  = random location  $\in$  Sources;
5 Destinations = Set of all locations  $\ell$  with (i)  $\tau_\ell \approx \tau_{\ell_{dst}}$  and (ii) Euclidean distance
between  $\ell_{src'}$  and  $\ell \approx \text{dist}$ ;
6  $\ell_{dst'}$  = random location  $\in$  Destinations;
7 src', dst' = Reverse geocode random points in  $\ell_{src'}$ ,  $\ell_{dst'}$ ;
8 return (src', dst');

```

Algorithm 1: Generating Sybil endpoints.

client path. This is to allow for the client path or the Sybil paths to make occasional detours in the path generation step.

The last step is to identify actual addresses within the Sybil source/destination locations just identified. To do so, it randomly chooses a point within the geographic location. However, this point may reside in non-driveable terrain, as shown in Figure 3(a). If this point were returned as a source/destination, an adversary can easily identify this point as a Sybil address. To avoid such cases, we reverse geocode this point (using an API provided by Google Maps [18]) to find the street address closest to the point identified. Doing so greatly improves the quality of the endpoints generated by our algorithm. Figure 3(b) shows the closest street address of the point identified in Figure 3(a).

The Path Generator

The SybilQuery prototype implements the path generator using the Microsoft Multimap API [35]. When supplied with a source and a destination, this API produces a sequence of waypoints representing the path to the destination. As noted earlier, SybilQuery admits the use of any off-the-shelf path generator, such as those used by on-board navigation systems. This feature helps SybilQuery produce paths that automatically account for environmental factors, such as road closures and one ways.

The Multimap API currently only produces the shortest path to the destination, thus making our prototype vulnerable to attack if the user follows a longer path to the destination. However, as noted earlier, this problem can easily be fixed with the use of an API that generates a choice of paths to the destination (or an API that allows the user to specify a waypoint that must be included enroute the destination). The





Figure 3: Finding a point in a block using reverse geocoding.



(a) Real trip. The values of τ_ℓ for the source and destination were 22 and 247, for the source and destination were 22 and 255, respectively.

Figure 4: A real user trip and a Sybil trip generated by SybilQuery.

path generation algorithm could then choose paths (both real and Sybil paths) based upon a probability distribution of the length of path that users typically follow to their destination.

The Query Generator

The basic version of SybilQuery's query generator simulates user movement along each path. It simulates the movement of users along Sybil paths at approximately the projected speed along that path. (Information about the average speed along a path is obtained by most off-the-shelf path generators). When the user sends a query to the LBS, the query generator obtains the current location of the simulated users and sends these as Sybil queries to the LBS. We have also interfaced the query generator with the Yahoo! local API [44] to more accurately simulate movement along Sybil paths under the constraints of current traffic. Thus, if there is traffic congestion at a particular location, SybilQuery can either simulate slower movement in that location or simulate a detour (and request the path generator to generate a fresh path to the destination). Note that Sybil paths get congested and decongested independently from real paths. When querying an LBS for online information, the query is made for all the Sybil locations as well as the real client location, so that the uncertainty at the LBS about the client location remains the same. Although querying an LBS that reports current traffic conditions renders the prototype vulnerable to malicious LBSs that report false traffic data, the query generator can use the N -variant queries approach described earlier to counter this threat.

Example

Figure 4 presents SybilQuery in action. It shows a real user trip (Figure 4(a)), obtained from the Cabspotter traces for another month and one of the Sybil trips generated by SybilQuery (Figure 4(b)); k was set to 4 for this example. The real trip (at approximately 6pm) originates at a home in Daly City and ends at a shopping area in downtown San Francisco. All the Sybil trips generated by SybilQuery also started in residential areas with similar traffic densities as the source of the real trip, and ended in a crowded region of San Francisco downtown. Figure 4 also reports the traffic densities for the sources and destinations for the real and Sybil trip generated by SybilQuery. A key point to note from this example is that traffic densities accurately capture

k	# questions	# correct	Probability
4	75	20	0.26
6	75	14	0.19

Figure 5: Results of a user study with 15 volunteers.

semantic properties (e.g., “residential area,” “downtown,” “shopping area”) of geographic regions.

EVALUATION

In this section, we present the results of our evaluation of the SybilQuery prototype. We evaluated both the privacy and the performance offered by SybilQuery.

Privacy

To evaluate the quality of the Sybil paths generated by our system, we conducted a user study. The high-level idea behind the user study was to give the working system to adversarial users, who would try to break the system by trying to find the real user path hidden between Sybil paths. Our methodology was as follows—we picked real paths at random from the Cabspotter traces (different from the month-long traces that were used to seed our prototype’s endpoint generator) and used SybilQuery to generate Sybil paths with different values of k . We also performed a quantitative evaluation of privacy by devising a new metric. The intuition behind this metric is that an adversary who has access to a database \mathcal{R} of real paths and a database \mathcal{S} of Sybil paths corresponding to paths in \mathcal{R} should be unable to differentiate between the two databases. We present the results from this evaluation in our detailed technical report [40].

User Study. We conducted the user study with 15 volunteers, each of whom had to answer ten questions. Each question presented k paths (one real path and $k - 1$ Sybil paths) to the volunteer, who had to identify the real path from the Sybil paths. Five questions had $k = 4$ while the remaining five had $k = 6$. Each volunteer was presented with a different set of ten questions, *i.e.*, our study obtained responses for 150 questions in all. Volunteers could view paths using the Google Maps API, and were instructed to use any tools at their disposal, such as zooming into the map, street views, and local information about the San Francisco Bay area, in their attempt to distinguish real paths from Sybil paths. All our volunteers were computer science graduates familiar with the basic geography of the San Francisco Bay area.

Figure 5 presents the results of this study and shows the number of questions for which users were able to correctly identify the real path from the Sybil paths. As this figure indicates, volunteers were able to correctly guess the real path with a probability of 0.26 for $k = 4$ and 0.19 for $k = 6$. These probabilities are close to their expected values (0.25 and 0.17, respectively). We further calculated the per-user standard deviation of the probability of correct guesses and found these values to be quite low (0.19 and 0.15, respectively). These results lead us to conclude that the Sybil paths qualitatively resemble real paths.




Figure 6: Query/response Increase of cloaked region latency of SybilQuery.

Figure 7: Spatial cloaking: Increase of cloaked region radius with trip length.

Sample responses from the user study gave us an interesting perspective into the mind of an adversary trying to break the SybilQuery system. Several users selected a false trip as the “*odd man out*” since all the other trips (including the real trip) were similar. Different users used contrasting rationale to guess real user trips. For example, some users wrongly selected shortest paths since the real path “*looked too circuitous to be true*”. Yet, others wrongly selected circuitous paths as they felt that “*outside human knowledge could be directing this route*”. Users typically found it hard to guess when the trip started late (*i.e.*, if the first query was made after travelling some distance), since “*the starting point seemed a bit odd*”. Sybil paths with a prominent endpoint such as a University or a Metro station were often selected by users with responses such as “*destination is University (different from others) and more likely to be the real user’s query*”.

To summarize the user study, the reasonings used by the adversaries were extremely diverse, which seems to explain why the overall probability of guessing correctly was close to that of random guessing.

Performance

We report the performance of several aspects of SybilQuery, including the time needed to preprocess the traffic database and the real-time performance of querying an LBS. We also compare the performance of SybilQuery against an implementation of spatial cloaking. All the results presented below were obtained on a 1.73GHz Pentium M laptop with 512 MB RAM. Each experiment was repeated 50 times, and the value of the privacy parameter k was fixed to be 4, unless otherwise indicated.

One-Time and Once-Per-Trip Costs. The one-time cost (offline step) for SybilQuery involves preprocessing of the traffic database. This step processed 529,533 trips and took approximately 2 hours and 16 minutes.

The once-per-trip costs for SybilQuery involve end-point generation and path generation. Generating a single pair of endpoints at the beginning of a user trip took an average of 5.47 seconds with a standard deviation of 1.02 seconds. To measure the cost of generating paths, we randomly chose trips with lengths varying between 5 and 30 minutes. The mean cost of computing a path (including the network latency to query the Microsoft MultiMap API) was 360 milliseconds, with a standard deviation of 150 milliseconds.




Figure 8: Response size returned by SybilQuery and spatial cloaking, for medium POI density.




Figure 9: Response size returned by spatial cloaking, for different POI densities.




Figure 10: Comparing the performance of SybilQuery with spatial cloaking for range queries.

Query/Response Performance. We measured the response latency of SybilQuery by integrating it with the Yahoo! Maps local search API [43]. Figure 6 shows the latency of sending k queries and receiving responses. As expected, the cost increases linearly with k , since k requests are sent to the LBS each time the client makes a query.

Comparison with Spatial Cloaking. Spatial cloaking [19] is a state of the art technique that achieves k -anonymity by using an anonymizer to send the LBS cloaked regions containing at least k users. The LBS returns all points of interest (POIs) within the cloaked region that match the query to the anonymizer. Although spatial cloaking was originally proposed for static users, we considered a simple variant for mobile users in which cloaked regions grow as users move.²

We conducted experiments to compare spatial cloaking with SybilQuery. Because spatial cloaking uses an anonymizer to send cloaked regions containing k clients to the LBS, we can expect the cloaked regions to increase in size as clients move (because the cloaked region must contain the *same* set of clients to preserve k -anonymity). In addition, because an LBS returns query results for the cloaked region, we can expect that increased POI density will lead to increased query/response processing time at the anonymizer. Our experiments, reported below, confirmed these expectations.

Figure 7 presents the results of the experiment that shows that the size of cloaked regions increases as clients move. We randomly chose trips varying in duration from 1 minute to 30 minutes from the Cabspotter database. We then fixed $k = 4$, and selected the smallest cloaked region containing at least $k - 1$ other clients. As Figure 7 shows, the size of the cloaked region increases with the duration of the trip.

To understand how this increase translates to query/response processing overhead, we studied the size of the response (in kilobytes) received from the LBS for a fixed nearest neighbor query: “return the Chinese restaurant closest to my current location.” In spatial cloaking, the anonymizer sends the entire cloaked region to the LBS, and must process responses from the LBS to identify the closest Chinese

²Although we have not carefully analyzed the privacy offered by this scheme, it suffices to compare the query/response performance of spatial cloaking with that of SybilQuery.

restaurant for the client issuing the query. However, because SybilQuery sends exact street addresses, the LBS sends the closest Chinese restaurant. Figure 8 compares the sizes of the query response for spatial cloaking and SybilQuery. As this figure shows, the query response size for SybilQuery is nearly a constant (at approximately 2KB); it only varies only with the value of k . However, the response size for spatial cloaking increases linearly. This is because the size of the cloaked region increases with trip length, which in turn directly translates to an increased number of POIs that match the user’s query.

We also conducted an experiment to study the effect of increased POI densities on query/response performance. We chose three representative nearest neighbor queries that required the LBS to return the closest “restaurant,” “Chinese restaurant,” and “Hunan Chinese restaurant” to the client’s location. For spatial cloaking, these queries represent, respectively, high, medium, and low POI densities. As Figure 9 shows, with spatial cloaking the size of the query results depends both on POI density *and* on the duration of the trip, and rises to about 5MB within 30 minutes for high POI densities. With SybilQuery, the query response size remains fixed at 2KB (because only the closest restaurant is returned).

The experiments above used nearest-neighbour queries, for which the difference in query responses for spatial cloaking and SybilQuery are most pronounced. In contrast, range queries, such as “return the list of all Chinese restaurants in an x -mile radius,” require the LBS to process the query over an entire geographic region. If a client issuing range queries uses SybilQuery to protect his location, the LBS must process k geographic regions, each of x -mile radius. On the other hand, if the client uses spatial cloaking, the LBS must only process one geographic region whose radius is x miles larger than the cloaked region. We can therefore expect spatial cloaking to outperform SybilQuery as x increases.

Figure 10 shows that this is indeed the case. This figure plots the query response size against increasing values of x for both SybilQuery and spatial cloaking. For this experiment, we assumed a circular cloaked region with a fixed radius 5 kilometers, and used SybilQuery with $k = 4$. For smaller values of x , SybilQuery outperforms spatial cloaking because the combined area of k regions of radius x

is smaller than the cloaked region. However, as x increases, spatial cloaking outperforms SybilQuery.

Nevertheless, we note that the query/response performance of SybilQuery can never be worse than k times the performance of spatial cloaking even for range queries.

RELATED WORK

- **Synthetic locations.** The idea of augmenting client location with synthetically generated locations to achieve location privacy was also recently proposed by Krumm [27] in a concurrently developed system. Although our work is similar in spirit, there are a number of differences between the two systems. Krumm uses a probabilistic model of driving behavior which depends on having data about the road network and ground cover. SybilQuery on the other hand uses statistical clustering techniques for end point generation and an off the shelf path generation tool, which makes it easily scale to large regions. Using false trips for location privacy was earlier proposed by Kino [21, 20]. Kino’s work focused on reducing the computation cost, so they use random walk based models for false path generation, which can be easily reidentified.

- **Cloaking schemes.** Originally introduced by Gruteser and Grunwald [19], cloaking aims to achieve k -anonymity by hiding client location from the LBS both in space and time. These systems achieve spatial anonymity by sending a cloaked region containing at least k users to the LBS. They achieve temporal cloaking by delaying a query until at least k other users have also issued queries. This basic model has been extended in several ways. For example, work by Duckham and Kulik [8], Gedik and Liu [13, 14], Bamba *et al.* [1] and Mokbel *et al.* [34] allow users to personalize their privacy requirements, *e.g.*, by letting them decide thresholds for spatial and temporal resolution of cloaked regions. Similarly, work by Bamba *et al.* [1] extends the basic model to allow for ℓ -diverse [32] queries. Other related approaches include achieving k -anonymity using path confusion [23, 25], where the anonymizing system attempts to confuse the adversary by crossing paths where at least two users meet.

A common theme in all the above systems is the need for a trusted third-party anonymizer that ensures k -anonymity (or ℓ -diversity). Using a third-party anonymizer has two disadvantages. First, the anonymizer is the central point of failure and presents scalability and performance bottlenecks. For example, an adversary could cripple the system with a denial-of-service attack on the anonymizer. Moreover, because the anonymizer has access to sensitive information from clients, a compromise of the anonymizer would result in a privacy breach. By offering a decentralized design, SybilQuery avoids these shortcomings. Second, most previous approaches only provide security guarantees for static snapshots and do not consider history of user movement (with the exception of work by Chow and Mokbel [4], which also uses an anonymizer). Therefore, if a user asks the same query from multiple cloaked regions as she moves, the LBS could compromise her privacy by correlating queries from these regions. These attacks are called correlation attacks [15]. SybilQuery offers improved

protection than cloaking schemes because it does not attempt to hide the actual path of the user from the LBS. Instead, it ensures that the adversary is unable to differentiate user paths from Sybil paths.

- **Peer-to-peer schemes.** Recent research has developed techniques based on peer-to-peer techniques to eliminate the need for an anonymizer [17, 16, 5]. Also related, although proposed as a mechanism for anonymous Web access, is Crowds [37], in which a query originates from a “crowd” of k users and the adversary is unable to identify the source of the query. These techniques have the advantage of being decentralized in nature. However, they still rely on the presence of at least k participating peers; in contrast, SybilQuery operates autonomously, independent of other querying peers.

- **Private Information Retrieval (PIR).** Cryptographic techniques based on PIR [29] offer another alternative to eliminate anonymizers [15, 22]. Using PIR to protect client location offers strong cryptographic guarantees on privacy that SybilQuery unfortunately cannot offer. However, the PIR-based scheme suffers from two drawbacks that limit its applicability. First, in spite of impressive recent advances, PIR remains computationally expensive [41]. For example, Ghinita *et al.* [15] employ a PIR protocol that imposes a cost of $O(n)$ at the server, in addition to client/server communication costs of $O(\sqrt{n})$. Practically, this translated to about 30 to 60 seconds of server processing time for each location query and communication of a few megabytes of data to process a single location-based query in their implementation. In contrast, SybilQuery is computationally cheap and generates queries with sub-second latency. Second, PIR-based schemes require code modifications at both the client and the server to implement the PIR protocol. Thus, in contrast to SybilQuery, PIR-based schemes cannot easily be applied to legacy systems.

- **Other related research.** Zhong *et al.* [45] recently proposed a distributed k -anonymity scheme for location privacy that makes use of an additive homomorphic cryptosystem. The main limitation of such a system is that it requires location brokers (cellular and/or WiFi providers) to be modified in order to implement their scheme. SpaceTwist [31] is a recently proposed location privacy system that uses a cloaking-like scheme to obfuscate the location of the client from the LBS and uses an incremental algorithm to process nearest neighbor queries. Like SybilQuery, SpaceTwist is also decentralized and autonomous. However, unlike SybilQuery, it does not ensure k -anonymity, does not handle mobile users, and only supports nearest neighbor queries.

The idea of introducing synthetic information to achieve privacy has previously been explored for anonymous Web access [24, 38]. In these systems, a user query to a Web server (*e.g.*, a search request) is hidden amongst synthetically generated queries. Similarly, Chuffnes *et al.* recently developed SwarmScreen [3], a system which protects user privacy in P2P systems by adding additional random connections that are statistically indistinguishable from natural ones.

Recently, Machanavajjhala *et al.* [33] developed rigorous techniques based upon a variant of differential privacy [9] to add synthetic information to released databases, and applied it to a database of traffic commuting patterns. Adapting these techniques to enable real-time generation of synthetic queries is an interesting direction for future work.

SUMMARY AND FUTURE WORK

SybilQuery is an efficient, decentralized technique to hide user location from LBSs. Its modular design allows SybilQuery to be deployed with off-the-shelf client devices and without any changes to the LBS. We implemented a prototype of SybilQuery and integrated it with LBSs such as Google Maps, Yahoo! Maps and Microsoft Live Maps. Our experiments—both a qualitative user study as well as a performance evaluation—show that Sybil queries can be efficiently generated and they resemble real user queries. In future work, we plan to enhance the SybilQuery prototype to generate Sybil queries that also achieve stronger privacy guarantees, such as ℓ -diversity [32], t -closeness [30] and differential privacy [9].

ACKNOWLEDGEMENTS.

We thank our shepherd Alex Varshavsky, Arati Baliga, Mohan Dhawan, Stephen Smaldone, Rebecca Wright and all the anonymous reviewers for valuable feedback on this paper. We also thank all the participants of our user study. The work has been supported in part by the National Science Foundation under awards CNS-0520123 and CNS-0831268.

REFERENCES

- B. Bamba, L. Liu, P. Pesti, and T. Wang. Supporting anonymous location queries in mobile environments with PrivacyGrid. In *Proc. WWW*, 2008.
- Cabspotting project. <http://cabspotting.org/>.
- D. Choffnes, J. Duch, D. Malmgren, R. Guierma, F. Bustamante, and L. Amaral. Swarmscreen: Privacy through plausible deniability in p2p systems. In *Northwestern EECS Technical Report*, 2009.
- C.-Y. Chow and M. F. Mokbel. Enabling private continuous queries for revealed user locations. In *Proc. SSTD'07: Advances in Spatial and Temporal Databases*, 2007.
- C.-Y. Chow, M. F. Mokbel, and X. Liu. A peer-to-peer spatial cloaking algorithm for anonymous location-based services. In *Proc. GIS'06: ACM International Symposium on Advances in Geographic Information Systems*, 2006.
- B. Cox, D. Evans, A. Filipi, J. Rowanhill, W. Hu, J. Davidson, J. Knight, A. Nguyen-Tuong, and J. Hiser. N-Variant systems: A secretless framework for security through diversity. In *Proc. USENIX Security Symposium*, 2006.
- F. Diggelen. Gnss accuracy: Lies, damn lies, and statistics. *GPS World*, 2007.
- M. Duckham and L. Kulik. A formal model of obfuscation and negotiation for location privacy. In *Proc. Pervasive*, pp. 152-170, 2005.
- C. Dwork. Differential privacy. In *Proc. ICALP'06: Intl. Colloquium on Automata, Languages and Programming*, 2006.
- EZPass: A pass on privacy? <http://www.nytimes.com/2005/07/17/magazine/17WWLN.html>.
- EZPass records out cheaters in divorce court. <http://www.msnbc.msn.com/id/20216302>.
- R. Finkel and J. Bentley. Quad trees: A data structure for retrieval on composite keys. *Acta Informatica*, 4(1):1-9, 1974.
- B. Gedik and L. Liu. Location privacy in mobile systems: A personalized anonymization model. In *Proc. ICDCS*, 2005.
- B. Gedik and L. Liu. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Transactions on Mobile Computing*, 2007.
- G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K.-L. Tan. Private queries in location-based services: Anonymizers are not necessary. In *Proc. ACM SIGMOD*, 2008.
- G. Ghinita, P. Kalnis, and S. Skiadopoulos. Mobihide: A mobile peer-to-peer system for anonymous location-based queries. In *Proc. SSTD'07: International Symposium on Spatial and Temporal Databases*, 2007.
- G. Ghinita, P. Kalnis, and S. Skiadopoulos. PRIVE: Anonymous location-based queries in distributed mobile systems. In *Proc. WWW*, 2007.
- Google maps geo API. <http://code.google.com/apis/maps/>.
- M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proc. MobiSys*, 2003.
- K. H. Y. Yanagisawa, and T. Satoh. An anonymous communication technique using dummies for location-based services. In *Proc. IEEE International Conference on Pervasive Services*, 2005.
- K. H. Y. Yanagisawa, and T. Satoh. Protection of location privacy using dummies for location-based services. In *Proc. ICDE Workshops*, 2005.
- U. Hengartner. Hiding location information from location-based services. In *Proc. International Workshop on Privacy-Aware Location-based Mobile Services (PALMS)*, 2007.
- B. Hoh and M. Gruteser. Location privacy through path confusion. In *Proc. SecureComm*, 2005.
- D. Howe and H. Nissenbaum. TrackMeNot: Resisting surveillance in Web search. *On the Identity Trail: Privacy, Anonymity and Identity in a Networked Society (Oxford University Press)*, 2008.
- R. R. C. J. Meyerowitz. Realtime location privacy via mobility prediction: Creating confusion at crossroads. In *Proc. HotMobile*, 2009.
- P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias. Preventing location-based identity inference in anonymous spatial queries. *IEEE Transactions on Knowledge and Data Engineering*, 19(12):1719-1733, 2007.
- J. Krumm. Realistic driving trips for location privacy. In *Proc. Pervasive*, 2009.
- J. Krumm and E. Horvitz. Predestination: Where do you want to go today? In *IEEE Computer Magazine*, vol. 40, no. 4, pp. 105-107, 2007.
- E. Kushilevitz and R. Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *Proc. FOCS'97: IEEE Symposium on Foundations of Computer Science*, 1997.
- N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *Proc. ICDE*, 2007.
- M. L. Liu, C. S. Jensen, X. Huang, and H. Lu. Spacetwist: Managing the tradeoffs among location privacy, query performance, and query accuracy in mobile services. In *Proc. ICDE*, 2008.
- A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramanian. ℓ -Diversity: Privacy beyond k-anonymity. In *Proc. ICDE*, 2006.
- A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber. Privacy: Theory meets practice on the map. In *Proc. ICDE*, 2008.
- M. F. Mokbel, C.-Y. Chow, and W. G. Aref. The new Casper: Query processing for location services without compromising privacy. In *Proc. VLDB*, 2006.
- Microsoft multimap API. <http://www.multimap.com/>.
- NYC cabs strike over GPS system plans. <http://www.engadget.com/2007/03/09/nyc-cab-drivers-say-no-thanks-to-gps-installation/>.
- M. K. Reiter and A. D. Rubin. Crowds: Anonymity for Web transactions. *ACM Transactions on Information and System Security*, 1(1):66-92, 1998.
- F. Saint-Jean, A. Johnson, D. Boneh, and J. Feigenbaum. Private Web search. In *Proc. WPES'07: ACM Workshop on Privacy in the Electronic Society*, 2007.
- P. Samarati. Protecting respondents identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, 13(6):1010-1027, 2001.
- P. Shankar, V. Ganapathy, and L. Iftode. Privately querying location-based services with sybilquery. In *Rutgers University Technical Report DCS-TR-652*, 2009.
- R. Sion and B. Carbunar. On the computational practicality of private information retrieval. In *Proc. NDSS*, 2007.
- L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557-570, 2002.
- Yahoo maps local search API. <http://developer.yahoo.com/maps/>.
- Yahoo! Local. <http://traffic.yahoo.com/traffic/>.
- G. Zhong and U. Hengartner. A distributed k-anonymity protocol for location privacy. In *Proc. IEEE PerCom*, 2009.