Computer Communications 68 (2015) 4–16

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Location privacy without mutual trust: The spatial Bloom filter

Luca Calderoni [a,*], Paolo Palmieri [b], Dario Maio [a]

[a] Department of Computer Science and Engineering, University of Bologna, Cesena, FC, 47521, Italy
[b] Department of Computing and Informatics, Bournemouth University, Poole, Dorset BH12 5BB, UK

ARTICLE INFO

Article history:
Available online 26 June 2015

Keywords:
Location privacy
Bloom filters
Secure multi-party computation

ABSTRACT

Location-aware applications are one of the biggest innovations brought by the smartphone era, and are effectively changing our everyday lives. But we are only starting to grasp the privacy risks associated with constant tracking of our whereabouts. In order to continue using location-based services in the future without compromising our privacy and security, we need new, privacy-friendly applications and protocols. In this paper, we propose a new compact data structure based on Bloom filters, designed to store location information. The spatial Bloom filter (SBF), as we call it, is designed with privacy in mind, and we prove it by presenting two private positioning protocols based on the new primitive. The protocols keep the user's exact position private, but allow the provider of the service to learn when the user is close to specific points of interest, or inside predefined areas. At the same time, the points and areas of interest remain oblivious to the user. The two proposed protocols are aimed at different scenarios: a two-party setting, in which communication happens directly between the user and the service provider, and a three-party setting, in which the service provider outsources to a third party the communication with the user. A detailed evaluation of the efficiency and security of our solution shows that privacy can be achieved with minimal computational and communication overhead. The potential of spatial Bloom filters in terms of generality, security and compactness makes them ready for deployment, and may open the way for privacy preserving location-aware applications.

1. Introduction

Positioning systems are becoming more precise and more portable, and can now be easily embedded into smartphones and other personal devices. The combination of different positioning sources, such as signal strength for cellular phones, the visibility of wireless networks and more traditional satellite-based sources means that an address level precision can now be achieved even in low-cost and low-power devices. Satellite navigation and positioning systems are also seeing renewed interest, after years of stagnation: the deployment of a new global systems, Galileo, is currently being sponsored by the European Union, while regional systems such as the Chinese BeiDou (covering most of Asia), or the projected IRNSS in India are promising an even bigger increase in precision and capabilities.

The ability to know one's position with a certain degree of precision opened the way to the so-called location-aware applications, where users request personalized services based on their geographic position. Location-aware applications and services are now ubiquitous: from cell phone apps to intelligent car navigation systems, they are an integral part of our everyday life. In order to perform their task, location-aware applications usually require the user to disclose her exact position, in order to receive content and information relevant to the user's location. Examples of such location-aware services are local advertising, traffic or weather information, or suggestions about points of interest (PoI) in the user's surroundings [6]. Even existing services are now improved by the addition of location-based data: notable examples are social networks [18] or retail distribution [11].

The ability to track a user's position raises however deep privacy concerns, due to the sensitive nature of location information. In fact, a number of potentially sensitive professional and personal information about an individual can be inferred knowing only her presence at specific places and times [1,4]. Sensitive information such as religious beliefs, sexual preferences or health conditions can be inferred by looking at the mobility trace of an individual, when he attends service at a church or a mosque, visits specific establishments or the practice of a specialized doctor. Even anonymized position data sets (not containing name, phone number or other obvious references to the person) do not prevent precise identification of the user: in fact, just four mobility traces may be enough to identify her. The more users disclose their data, the more providers are able to profile them in an accurate way. This is for instance the case discussed by Wicker in [45], where a marketing company database model is used in conjunction with anonymous mobile phone location traces. While we have

* Corresponding author.
E-mail address: luca.calderoni@unibo.it (L. Calderoni).

http://dx.doi.org/10.1016/j.comcom.2015.06.011
0140-3664/

become so used to smartphones and location-aware services that it would be very hard for a lot of us to give up on them, it is also reasonable to predict that in the coming years users will demand better privacy safeguards for their information with respect to the service provider [40], and more specifically for location information [24,46]. The real challenge is therefore how to protect the user's privacy without losing the ability to deliver services based on her location [35].

A common application scenario of location-based services requires the service provider to learn when the user is close to some sensitive or interesting locations. This is the case, for instance, of "around-me" applications or security and military systems [6]. In this case, the location of the user should be kept private for as long as she is far from one of the areas of interest, and get disclosed to the service provider only when she enters one such area. A similar problem, known as *private proximity testing* has been studied in privacy research literature: Alice can test if she is close to Bob without either party revealing any other information about their location [30]. Narayanan et al. proposed a solution based on location tags (features of the physical environment) and relying on Facebook for the exchange of public keys [30]. Their protocol was later improved in efficiency by Saldamli et al. [36]. Location tags and proximity tests are also used in [19], as a way of providing local authentication, while [47] presents a secure handshake for communication between the two actors in proximity. The security of the basic proximity testing protocol has been further improved in [31]. In [43], Tonicelli et al. propose a solution for proximity testing based on pre-distributed data, secure in the Universal Composability framework. Finally, the problem of checking the proximity in a specific time is addressed in [42].

In this paper we do not focus on proximity testing, but on a broader and more general problem: testing in a private manner whether a user is within one of a set of areas of arbitrary size and shape. By solving this problem and applying an intelligent conformation of areas, we can also solve the proximity testing problem (for one or multiple points simultaneously), and we are actually able to identify with some precision the distance of the user from the point of interest. Given the conceptual similarity of our problem with membership testing in sets, we base our solution on a novel modification of Bloom filters (BF). Bloom filters are a compact data structure that allows to compute whether an element is a member of the set the filter has been built upon, without knowledge of the set itself [2]. Bloom filters have already been used in privacy-preservation protocols, and they are particularly suited to be used in conjunction with the homomorphic properties of certain public key encryption schemes [21].

1.1. Contribution

In this paper we propose a modification of Bloom filters aimed at managing location information, and we present two private positioning protocols for privacy-preserving location-aware applications. Although a preliminary version of the data structure was presented by Palmieri et al. in [34] (which forms the basis of the current work), in this paper we analyze the security and efficiency properties of the structure, and we provide a much greater insight on the usefulness of the construction to actual application scenarios, also by means of practical, real-world examples based on a test implementation.

The novel variant of Bloom filters we introduce, which we call *spatial Bloom filter* (SBF), is specifically designed to deal with location information. In particular, SBF combines multiple superimposed Bloom filters, in conjunction with an ad-hoc spatial representation, to provide a compact data structure for geographical information. Similarly to the classic Bloom filters, SBFs are also well suited to be used in privacy preserving applications, and we show this by presenting two protocols for private positioning. The protocols allow secure computation of location-aware information, while keeping the position of the user private: the only information disclosed to the provider is the user's vicinity to specific points of interest or his presence within predefined areas. At the same time, the areas of interest are not disclosed to the user. Therefore, in both settings we do not assume any trust between the parties. The first protocol is based on a two-party setting, where communication happens directly between the user of a location-based service and the service provider. A more complex scenario is defined in the second protocol, that involves a three-party setting in which the service provider outsources to a third party the communication with the user. Both protocols achieve secure multiparty computation, where all parties have an interest in communicating, but want to keep their information private. In some cases, the privacy of the service provider can in fact be as important as that of the user: military and government applications are just the most immediate examples.

Following the definition of the spatial bloom filter and of the private-positioning protocols, we discuss the security and the computational cost of the proposed schemes, as well the probabilistic and storage properties of the SBF. In order to prove the readiness of the solution for actual deployment, we present the results of a prototype implementation of the filter creation and query routines. We base our tests on the geographic data of two real geographical regions: the metropolitan area of the city of Brussels, and Belgium. For both cases, we estimate optimal sizes for the filter and values for other important parameters.

Location privacy is a fundamental problem of the current age, where ubiquitous computing and unified communications are prevalent. The proposed solution is a solid step in the direction of more privacy-friendly services, and may enable privacy in both existing and future applications.

1.2. Related works

With the recent introduction, and subsequent widespread diffusion of *location-based services* (LBS), the problem of preserving the privacy of the user with regard to his position arose. An early solution addressing this problem was presented in [16] by Gruteser and Grunwald, and consists in the application of k-anonymity to LBS: the location trace for each person should not be distinguished from at least $k-1$ other individuals, thanks to spatial and temporal cloaking of location and timing information. This is just one of the adopted metrics used to quantify privacy of a LBS. A comprehensive discussion of those metrics, including k-anonymity, is provided in [38], where the authors propose to preserve privacy by applying a distortion to the location information. Systems designed to protect location privacy are often referred to as *location-privacy protection mechanisms* (LPPM). Possible attacks to LPPM systems, and a proposal for a general framework able to evaluate the effectiveness of such systems is presented in [39].

While LBS's vary widely in terms of goals, a good number of them follow the model commonly known as *around-me* service. Here the user wants to find points of interests in his surroundings, based on his current position [6]. In order to achieve such goal, a location query is usually performed onto a remote server. In [20] the authors discuss privacy preservation with regard to the location queries used in this kind of service, and a taxonomy of location queries performed on the provider's server is also presented. A different privacy-preserving framework for location-based queries is proposed in [27]: the proposed solution relies on a trusted third party connecting the client with server. The problem of k nearest neighbor (k-NN) in location-based queries is addressed in a privacy-preserving manner in [26], where the authors propose using homomorphic encryption.

While around-me applications are usually designed for end-users, location privacy is also especially important in military and other government settings [12]. In [12] the protocol PRISM is presented. PRISM is designed to achieve privacy-friendly routing in MANETs, mainly for military purposes, using *group signatures*. Possible attacks against routing protocols, aiming at understanding the source