3

Location privacy and location-aware computing

Matt Duckham & Lars Kulik
University of Melbourne, Australia

CONTENTS

3.1 Introduction

Combined technological advances in location-sensing, mobile computing, and wireless communication are opening up new and exciting opportunities in the domain of location-aware computing. Many of these opportunities are explored elsewhere in this book (e.g., chapters 2, 11–13); others are already being developed into practical applications that will provide benefit to a wide cross-section of society, such as elder care [64], emergency response and E911 systems [70], and navigation systems for the visually impaired [32].

Despite the undoubted future potential of location-aware computing, location-awareness also presents inherent future threats, perhaps the most important of which is location privacy. Most people would not feel comfortable if regularly updated information about their current location were made public, any more than we would feel comfortable if information about our home address, telephone number, age, or medical history were public. Our precise location uniquely identifies us, more so than our names or even our genetic profile.

This chapter examines the foundations of location privacy: the factors that affect location privacy and the strategies for managing location privacy. The development of location-aware computing technology and mobile GIS is changing forever the way we interact with information, our physical environment, and one another. How we deal with location privacy issues will be a determining factor in the ultimate direction of those changes.

This chapter begins by exploring the different concepts of privacy and their relevance to location-aware computing and mobile GIS (section 3.2). Section 3.3 reviews the important privacy characteristics of one of the key enabling technologies for location-aware computing: positioning systems. The four classes of privacy protection strategy, which form the basis of any location privacy protection system, are introduced and described in section 3.4. Section 3.5 concludes the chapter with an examination of some future challenges for location privacy research.

3.2 Background and definitions

The term “privacy” covers a wide range of concepts, and many different definitions of privacy have been proposed. An initial distinction is often made between *bodily privacy* (concerned with protection from physically invasive procedures, such as genetic testing), *communication privacy* (concerned with security of communications, like mail and email), *territorial privacy* (concerned with intrusions into physical space, like homes and workplaces), and *information privacy* (concerned with the collection and handling of personal data) [55]. Under the heading of “information privacy,” one of the the most influential and commonly quoted definitions was developed by the privacy pioneer Alan Westin:

> Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others. [71, p7]

Correspondingly, *location privacy* can be defined as a special type of information privacy which concerns the claim of individuals to determine for themselves when, how, and to what extent location information about them is communicated to others. In short, *control* of location information is the central issue in location privacy.

Location privacy is especially important (to this book, specifically, and at this time, generally) as a result of the development of location-aware computing. *Location-awareness* concerns the use of information about an individual’s current location to provide more relevant information and services to that individual [73]. Location-awareness is a special type of *context-awareness*. The term “context” is used to encompass the entire characteristics of an individual’s physical, social, physiological, or emotional circumstances [59]. Location information is one of the most important aspects of an individual’s (physical) context (see, for example, Ljungstrand’s discussion of context awareness and mobile phones [45]). Thus, location-aware computing environments offer the capability for automatic, regular, and real-time sensing of a person’s location with a high degree of spatial and temporal precision and accuracy. Together with technological advances in mobile computing and wireless communication, which enable rapid processing and communication of location information, these developments allow the location of mobile individuals to be tracked in a way never before possible.

3.2.1 The right to location privacy

Privacy is regarded as a fundamental human right, internationally recognized in Article 12 of the UN Universal Declaration of Human Rights [24]. The history and development of privacy rights have been examined from many different perspectives in the literature (e.g., see [42] for a concise overview of the history of privacy from the perspective of ubiquitous and location-aware computing).

Not all authors agree that privacy should be regarded as an inalienable right. Some authors (e.g., [10, 21]) have argued for greater transparency in place of privacy. Proponents of greater transparency cite the practical difficulties of protecting privacy in the face of changing technological capabilities (encapsulated in the now infamous remark by Sun CEO Scott McNealy: "You have zero privacy anyway, get over it!" [63]) and the public benefits that may be accrued through the relaxation of some privacy protections (for example, saving infant lives through the disclosure of positive HIV test results of pregnant mothers [22]).

Studies of users' attitudes to location privacy issues often provide some support for these views. Evidence presented in [6, 39] indicates a lack of awareness or even moderate indifference to location privacy issues amongst the general public. Other studies have painted a more complex picture. For example, Barkuus and Dey [5] found that concern about location privacy can be dependent on the type of application, with applications that track users' movements over a period of time causing more concern than simple positioning applications.

Attitudes to privacy have changed in the past and will continue to change over time[*]. Although the need for a right to privacy will continue to be debated, in the shorter term at least there would seem to be a pressing need for privacy protection measures able to cope with a rapidly changing technological landscape. Concerns about protecting the individual's right to privacy have previously appeared in connection with numerous other new technologies, including GIS [49], the Internet [1], and collaborative user interfaces [35]. The need for location privacy is recognized in some of the earliest literature on information privacy (e.g., [71]) and location-aware computing (e.g., [30, 31, 58]). Looking at more recent literature, it is possible to identify at least three key negative effects associated with failures to protect location privacy within a location-aware computing environment (e.g., [28, 57, 39]):

1. *Location-based "spam"*: Location could be used by unscrupulous businesses to bombard an individual with unsolicited marketing for products or services

[*] As an example of how attitudes have changed in the past, J.B. Rule quotes the 1753 bill to establish a census in Britain [56]: the bill was defeated as being "totally subversive of the last remains of English liberty." In the same 1973 book, Rule himself discards as "unhelpfully rash speculations" Westin's vision of a future credit system, in which all transactions are digital and individuals can be tracked through their spending habits. By today's standards, this "future" credit system seems rather conventional and unremarkable.

related to that individual's location.

2. *Personal wellbeing and safety*: Location is inextricably linked to personal safety. Unrestricted access to information about an individual's location could potentially lead to harmful encounters, for example stalking or physical attacks.

3. *Intrusive inferences*: Location constrains our access to spatiotemporal resources, like meetings, medical facilities, our homes, or even crime scenes. Therefore, location can be used to infer other personal information about an individual, such as that individual's political views, state of health, or personal preferences.

High profile media coverage of accusations of location privacy infringements are indicative of increasing public awareness of location privacy issues. For example, rental companies who use GPS to track their cars and then charge renters for infringements of their rental agreement have resulted in a flush of media articles and legal cases (e.g., James Turner versus Acme car rental [12]). Similarly, Samsung in Korea attracted media attention when it alledgedly used a "Friend finder" service to track its own employees with the aim of blocking the establishment of a labor union [44]. In the future, greater familiarity with cheaper, more reliable location-aware technology is likely to amplify location-privacy concerns. These issues have already created a perception that inadequate privacy protection is retarding the uptake of location-based services, and has led location privacy to be elevated to one of the key research challenges in pervasive computing [47]. In short, there is strong evidence that location privacy will be a key issue for the future of location-aware computing systems, including dynamic and mobile GIS.

3.3 Positioning systems and location privacy

In addition to the social constraints on location privacy, discussed in the previous section, location-aware computing environments place certain technical constraints on location privacy. The primary technical constraints arise from the positioning systems themselves. Hightower and Boriello provide a survey of the wide variety of positioning systems currently in use [33]. In addition to the familiar GPS, positioning systems in the literature and in common usage include triangulation of RF wireless LAN signals (e.g., [4]), proximity to infrared beacons (e.g., [69]), scene analysis and computer vision (e.g., [40]), and inertial tracking (e.g., [61]). New positioning systems, such as audio-based positioning [8, 60] and radio signal profiles [41], are continually being developed.

Positioning systems vary widely in their accuracy and precision characteristics. Accuracy and precision of location have implications for location privacy. For example, a positioning system that locates an individual to a precision of 200m is generating less information about location (and so can potentially be less invasive of location privacy) than a positioning system that locates an individual to a precision of 2m. Other characteristics of the positioning system may also present constraints to location privacy, such as the extent of the coverage of the positioning system (e.g., global or local) or the accuracy and precision of the positioning system relative to the density of geographic features (e.g., a location precision of 100m in a dense downtown area of a city may be considered more private than a location precision of 100m in a desert).

There exist several classifications of positioning systems. For example, a top-level distinction is often made between *active* positioning systems, which rely on the establishment of beacons to operate (such as WiFi signal triangulation, GPS, infrared proximity sensors), and *passive* positioning systems, which require no beacons (such as inertial navigation, scene analysis, and audio-based positioning, see [73] for more information). However, from a privacy perspective, positioning systems are more usefully classified into *client-based*, *network-based*, and *network-assisted* systems [57].

- In client-based positioning systems, mobile clients autonomously compute their own location (for example GPS and inertial navigation). It is technically possible in a client-based positioning system for a client to compute its location, without ever revealing that location to any other entity.

- In network-based positioning systems, the network infrastructure is responsible for computing a mobile client's location. Cell phone phone positioning using CGI (cell global identity) is an example of network-based positioning. In network-based positioning systems, the network infrastructure administrator must hold information about the location of mobile clients.

- In network-assisted positioning systems, a combination of client-based and network-based computation is required to derive a client's location. For example, A-GPS (assisted GPS) combines network-based CGI positioning to increase the speed of GPS positioning. In network-assisted positioning systems, some information about a mobile client's location must reside in the network infrastructure, although this information may be less precise than the information held by the mobile client itself.

Client-based positioning systems inherently allow for greater location privacy than network-assisted or network-based positioning systems. In a client-based positioning system it is technically possible for the client to have complete control over information about its location, possibly to the extent that the client becomes the only entity with information about its own position.

One potential solution to location privacy issues, therefore, is to use only client-based positioning, perform all processing of location information locally on the mobile device, and *never* share any personal location information with other entities, whether centralized servers of peer-to-peer clients (cf., [46]). However, adopting this completely client-oriented, centralized model of mobile computing presents several drawbacks:

- Mobile devices typically possess limited processing and storage capacity, making it inefficient to perform complex calculations on voluminous spatial data directly on the mobile device.
- Spatial data sets remain expensive to collect and collate, despite continuing advances in positioning systems. The companies who collect this data would usually be reluctant to make their valuable data sets available in their entirety to mobile users.
- Downloading spatial data sets from a remote service provider will be subject to wireless network bandwidth limitations and may provide an indication of the user's location (either by inferring location from knowledge of the data sets of interest to the user or by positioning using a client's mobile IP address, as in [15]). Alternatively, storing all potentially useful spatial data in a user's mobile device leads to the data integrity and currency issues that are inevitably associated with maintaining copies of the same data sets across multiple clients.

In summary, the different types of positioning system place some inherent constraints on the privacy characteristics of location-aware computing environments. Irrespective of these constraints, as mobile computing environments move toward increasingly distributed models of computation, the need to share personal information about location with a variety of remote location-based service providers increases correspondingly.

3.4 Location privacy protection strategies

Having identified location privacy as a key issue for location-aware computing and outlined some of the technical aspects of location privacy, the next step is to ask what mechanisms exist for location privacy protection. The different strategies that exist for protecting a mobile individual's location privacy can be classified into four categories: *regulatory*, *privacy policies*, *anonymity*, and *obfuscation* strategies. In this section each type of strategy is reviewed in turn.

3.4.1 Regulatory strategies

Regulatory approaches to privacy involve the development of rules to govern fair use of personal information. Most privacy regulation can be summarized by the five principles of *fair information practices* (originally developed as the basis of the US privacy legislation [68, 67]):

1. *Notice and transparency*: Individuals must be aware of who is collecting personal information about them and for what purpose.

2. *Consent and use limitation*: Individuals must consent to personal information being collected for particular purposes, and the use of personal information is limited to those purposes.

3. *Access and participation*: Individuals must be able to access stored personal data which refers to them, and may require that any errors be corrected.

4. *Integrity and security*: Collectors must ensure personal data is accurate and up-to-date and protect against unauthorized access, disclosure, or use.

5. *Enforcement and accountability*: Collectors must be accountable for any failures to comply with the other principles.

Although these principles of fair information practice are at the core of most privacy regulation (e.g., [50, 66]), there are a variety of ways in which these rules have been implemented. In general, regulatory frameworks aim to adequately guarantee privacy protection for individuals without stifling enterprise and technology. The concept of *co-regulation*, which aims to encourage flexible self-regulation on top of legal enforcement of minimum privacy standards, is one example of a mechanism for achieving such a balance [13].

The concept of fair information practices is usually applied to "personal information" in general, not specifically to location information. Personal information can be defined as "information ... about an individual who's identity is apparent, or can reasonably be ascertained, from the information ..." [3]. In this respect, location information is usually treated as one type of personal information, like age, gender, or address. A small number of privacy regulations have been developed to address location privacy issues explicitly (for example, proposed location tracking legislation in Korea [51] and the discontinued AT&T "Find Friends" location-based service [65]).

Although regulation lies at the foundations of any privacy protection system, there are at least four reasons for believing that, on their own, regulations do not represent a complete solution to location-privacy concerns. First, regulation itself does not prevent invasions of privacy, it simply ensures that there exist mechanisms for "enforcement and accountability" when unfair information practices are detected. Second, the development of regulation may lag behind innovation and new technology. Third, regulation applies "across the board," making a satisfactory balance between guaranteed levels of privacy protection and freedom to innovate and develop new technology difficult to achieve, even using models such as co-regulation. As a consequence, other privacy protection mechanisms are needed in addition to regulation. Finally, abiding by fair information practice principles can give rise to practical problems with respect to location-awareness. For example, Ackerman et al. [2] examine the difficulties created by the requirements for notice and consent for user interfaces and HCI in context-aware computing environments (e.g., overwhelming users with frequent, disruptive, and complex consent forms or notice information).

3.4.2 Privacy policies

Privacy policies are trust-based mechanisms for proscribing certain uses of location information. Whereas regulation aims to provide global or group-based guarantees of privacy, privacy policies aim to provide privacy protection that is flexible enough to be adapted to the requirements of individual users and even individual situations and transactions. Overviews of a range of different privacy policy systems can be found in [26, 48]. In this section we summarize three of the major privacy policy initiatives currently underway that illustrate the range of approaches that privacy policies can take.

3.4.2.1 IETF GeoPriv

The Internet Engineering Task Force (IETF) is an international consortium concerned with future Internet architectures. The IETF's GeoPriv working group is adapting PIDF (presence information data format) as a privacy policy system for location privacy. PIDF is an IETF XML dialect for instant messaging, which includes a mechanism for exchanging information about the presence of a person (or place or thing) [52]. The GeoPriv specification additionally includes information about the location of that person, effectively annotating location data with metadata about the fair uses of that location data. In order to protect location privacy, the GeoPriv specification defines a *location object* which encapsulates both an individual's location and their privacy policy. At the center of the privacy policy are *usage rules* which describe acceptable usage of the information, such as whether retransmission of the data is allowed or at what date the information expires, and must be discarded. Further, location objects can be digitally signed, making the privacy policy resistant to separation from the location information [48].

3.4.2.2 W3C P3P

The World Wide Web Consortium (W3C) has developed the platform for privacy preferences project (P3P) as a simple mechanism for communicating information about Web-based privacy policies [74]. In contrast to the IETF approach, where users attach privacy policies to their data, the focus of P3P is to enable service providers to publish their data practices. The data practices may include for what uses personal data is collected, for how long it is held, and with what other organizations and entities it may be shared. Users of a particular service can then decide whether these data practices fit with their own requirements [14]. Typically, this process is achieved automatically using software agents with access to users' profiles. P3P does not provide any mechanisms for encrypting privacy protection within location data (like those found in IETF GeoPriv specification) and does not explicitly address location issues. However, because P3P is XML-based it can be easily extended for location-aware computing environments. For example, in [43] Langheinrich describes an architecture (the privacy awareness system, pawS) that uses P3P to enable location-aware system users to keep track of the storage and usage of their personal location information. IBM's enterprise privacy authorization language (EPAL) is a different

XML-based dialect with similar goals to P3P [37].

3.4.2.3 PDRM

Digital rights management (DRM) concerns the technical efforts by some intellectual property vendors and other organizations to enforce intellectual property protection (for example, protection from piracy). PDRM (personal DRM) adopts a similar approach for personal data. When applied to location privacy, the PDRM approach is closer to the "user-oriented" IETF GeoPriv model than the P3P "provider-oriented" model. For location-aware systems, location data is treated as the property of the person to whom that data refers. PDRM then aims to enable that person to "license" the personal data for use by a location-based service provider [29]. So, for example, an entity wishing to use an individual's location data may first need to demonstrate their willingness to agree to the licensing, which may set limits on that entities ability to share or process the data.

Policy-based initiatives for privacy protection, like PDRM, P3P, and GeoPriv, are continuing to develop. However, there are again reasons for believing that policy-based initiatives provide only a partial answer to the question of location privacy protection. First, privacy policies are often highly complex and their practicality for use in location-aware environments with frequently updated highly dynamic information remains, as yet, unproven. Second, privacy policies systems generally cannot enforce privacy, instead relying on economic, social, and regulatory pressures to ensure privacy policies are adhered to. Consequently, privacy policies are ultimately vulnerable to inadvertent or malicious disclosure of personal information [28, 75].

3.4.3 Anonymity

Anonymity concerns the dissociation of information about an individual, such as location, from that individual's actual identity. A special type of anonymity is *pseudonymity*, where an individual is anonymous, but maintains a persistent identity (a pseudonym) [53]. For example, [20] describe a location-aware system for allowing users to leave and read digital notes at specific locations ("geonotes"). One of the ways users can protect their privacy is to associate an alias (pseudonym) with a note in place of their real name.

An explicitly spatial approach to providing anonymity in location-aware computing environments is presented in [27]. Gruteser and Grunwald used a quadtree-based data structure to examine the effects of adapting the spatial precision of information about about an individual's location according to the number of other individuals within the same quadrant, termed "spatial cloaking." Individuals are defined as *k-anonymous* if their location information is sufficiently imprecise in order to make them indistinguishable from at least $k-1$ other individuals. The authors also explore the orthogonal process of reducing the frequency of temporal information, termed "temporal cloaking."

There are several disadvantages to using anonymity-based approaches. First, anonymity based approaches often rely on the use of a trusted anonymity "broker," which retains information about the true identity of a mobile individual, but does not reveal that identity to third party service providers (e.g., [28]). Second, anonymity often presents a barrier to authentication and personalization, which are required for a range of applications [42, 34]. Pseudonymity does allow some personalization and is therefore sometimes preferred to general anonymity in order to combat this problem. For example, Rodden et al. [54] use a randomly generated pseudonym which is held by a trusted information broker and persists only for the duration of the provision of a particular service (like a location-aware taxi collection system). A promising new research direction that may help overcome these limitations is *zero-knowledge interactive proof systems* (see [25], described in more detail below).

3.4.3.1 Zero knowledge proofs

The idea of a zero-knowledge proof is to prove the knowledge of a certain fact without actually revealing this fact. Zero-knowledge proofs (ZKPs) involve a *prover*, who attempts to prove a fact, and a *verifier*, who validates the prover's proof. The verifier may determine the correctness of the proof, but not does learn *how* to prove the fact or anything about the fact itself. Fiat and Shamir developed the first practical zero-knowledge proof system in 1987 [23].

ZKPs often appear somewhat counter-intuitive at first, so consider the following simple example. Person *A* claims to know the secret combination to a safe. Person *B* deposits a valuable item in the safe, locks the safe, and leaves the room without the safe. Person *B* does not know the combination to the safe. If person *A* is able to present the item locked in the safe to *B*, then *A* has proven to *B* that *A* knows the combination to the safe without revealing the actual combination. In ZKP terminology, the proof is interactive because the verifier (person *B*) *challenged* the prover (person *A*) and the prover must *respond* to the verifier.

In a ZKP, a prover *may* provide the correct response to a challenge purely by chance. To combat this possibility, there are usually several rounds of challenges and responses in a ZKP. As the number of rounds increases, the probability that the prover will give the correct answer in every round decreases. Typical ZKPs will verify a proof with a probability of $1-1/2^n$, where n is proportional to the number of rounds used.

There are two distinct application scenarios for ZKPs:

- *Authentication*: Prover P is able to prove to verifier V that P is authorized to access information without requiring any knowledge about P's identity.

- *Identification*: Prover P can prove to verifier V that P is P, but no party Q is able to prove to V that Q is P.

The first application scenario that uses ZKPs without revealing an individual's identity is *anonymous digital cash* [9]. To date, ZKPs have not been widely researched within the domain of location-aware computing. However, clearly ZKP-based authentication and identification might also be used with location-based services, and initial work in this area is beginning to appear (e.g., [11]).

There is one further, explicitly spatial problem facing any anonymity-based system for location privacy: a person's identity can often be inferred from his or her location. Consequently, anonymity strategies (even those employing pseudonymity or ZKPs) are vulnerable to data mining [19]. Beresford and Stajano [7] have used simulated historical data about anonymized individual's movements to investigate ways of subverting anonymity-based privacy protection. Their results show how simple heuristics can be used to de-anonymize pseudonyms, providing users with much lower levels of location privacy than might naively expected. Thus, anonymity alone cannot hope to provide total location privacy protection.

3.4.4 Obfuscation

Obfuscation is the process of degrading the quality of information about a person's location, with the aim of protecting that person's location privacy. The term "obfuscation" is introduced in [16, 17], but several closely related concepts have been proposed in previous work. The "need-to-know principle" aims to ensure that individuals release only enough information that a service provider needs to know in order to provide the required service [36]. The idea of a need-to-know principle is closely related both to obfuscation and the fundamental fair information practice principle of consent and use limitation (section 3.4.1). Snekennes investigates a privacy policy-based approach to enforcing the need-to-know principle in location-aware computing by adjusting precision of location information [62]. In the domain of anonymity-based approaches, the work of Gruteser and Grunwald (discussed in section 3.4.3) aims to enforce the "principle of minimal collection" [27], again akin to obfuscation. On a slightly different theme, Jiang et al. discuss the "principle of minimal asymmetry," which aims to ensure that the flow of personal information away from an individual is more closely matched by the information flow back to that individual about who is using that information for what purposes [38].

It is possible to identify three distinct mechanisms (types of imperfection) in the literature for degrading the quality of location information: *inaccuracy*, *imprecision*, and *vagueness* (see [72, 18, 73]). Inaccuracy concerns a lack of correspondence between information and reality; imprecision concerns a lack of specificity in information; vagueness concerns the existence of boundary cases in information. Any combination of inaccuracy, imprecision, and vagueness may be used as the basis for an obfuscation system. An inaccurate description of an agent's location means that the agent's actual location differs from the conveyed location: the agent is "lying" about its current location. An imprecise description of location might be a region including the actual location (instead of the location itself). A vague description would involve linguistic terms, for example that the agent is "far" from a certain location. Most research to date has looked at the use of imprecision to degrade the quality of location information (e.g., [62, 27, 34, 16]). However, the use of inaccuracy has also been investigated and compared with imprecision in [17]

The work in [16] develops and tests an algorithmic approach to obfuscating *proximity queries* (e.g., “where is the closest ... ?”) based on imprecision. A simplified version of the algorithm introduced in [16] is summarized in figure 3.1. The algorithm assumes a graph-based representation of a geographic environment (for example, a road network). An individual protects his or her location privacy by only reporting a set O of locations (an *obfuscation set*), one of which is that individual’s actual location (figure 3.1a). For an obfuscation set O, the location-based service provider must compute the relation δ (figure 3.1b), where $o\delta p$ means $o, p \in O$ are most proximal to the same point of interest (POI). The algorithm then proceeds according to three possibilities. First, all the locations in the obfuscation set may be most proximal to a single POI ($O \in O/\delta$), in which case that POI can be returned to the user (figure 3.1c). Second, the individual may agree to reveal a more precise representation of his or her location, in which case the algorithm can reiterate (figure 3.1d). Otherwise, the best estimate of the most proximal POI us returned (figure 3.1e). The analysis in [16] shows that efficient mechanisms for computing the relation δ can ensure that the entire algorithm has the same computational (time) complexity as a conventional algorithms for proximity queries, and that the algorithm must terminate in a finite number of iterations.

Obfuscation has several important advantages that complement the other privacy protection strategies. Obfuscation and anonymity are similar, in that both strategies attempt to hide data in order to protect privacy. The crucial difference between obfuscation and anonymity is that while anonymity aims to hide a person’s identity, obfuscation is an explicitly spatial approach to location privacy that aims to allow a person’s identity to be revealed. Potentially, this combats one of the key limitations of anonymity approaches: the need to authenticate users. At the same time, degrading the quality of location information makes inferring identity from location more difficult. Obfuscation is flexible enough to be tailored to specific user requirements and contexts, unlike regulatory strategies; does not require high levels of complex infrastructure and is less vulnerable to inadvertant disclosure of personal information, unlike privacy policies; and is lightweight enough to be used without the need for trusted privacy brokers, unlike many anonymity approaches.

Obfuscation aims to achieve a balance between the level of privacy of personal information and the quality of service of a location-based service. Current research has indicated that there exist many situations where it is possible to expect high quality location-based services based on low quality positional information (see [17]). Consequently, in situations where the user requires a higher quality of service than can be achieved at a user’s minimum acceptable level of privacy, then other privacy protection strategies must be relied upon instead. Further, obfuscation assumes that the individual is able to choose what information about his or her location to reveal to a service provider. While this may be realistic when using client-based or network-assisted positioning systems and when sharing location information with a third party location-based service provider, dealing with the entities that administer network-based positioning systems still requires privacy protection based on regula-

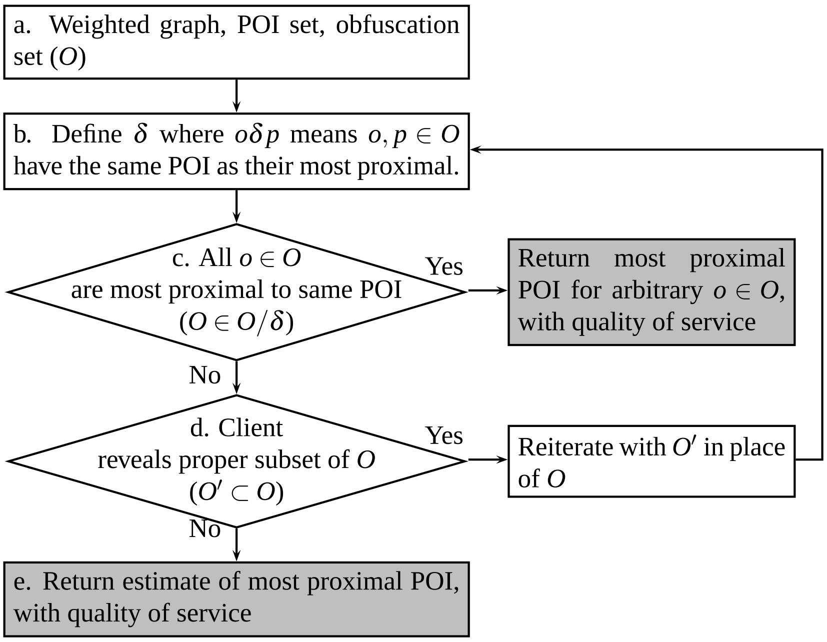
a. Weighted graph, POI set, obfuscation set (O)
b. Define δ where $o\delta p$ means $o, p \in O$ have the same POI as their most proximal.
c. All $o \in O$ are most proximal to same POI ($O \in O/\delta$)
Yes
Return most proximal POI for arbitrary $o \in O$, with quality of service
No
d. Client reveals proper subset of O ($O' \subset O$)
Yes
Reiterate with O' in place of O
No
e. Return estimate of most proximal POI, with quality of service

FIGURE 3.1

Summary of simplified obfuscation algorithm, after[16]

tory or privacy policy approaches.

3.5 Conclusions

Location privacy lies at the intersection of society and technology. This chapter has reviewed the reasons why location privacy is becoming such an important topic in society, and the technological constraints to location privacy. When considering the strategies that can be used to protect an individual's location privacy, it becomes clear that no single strategy currently available is capable of providing a complete solution to location privacy protection. Each approach has distinct advantages and disadvantages. Therefore, it seems likely that the future of location privacy protection involves combinations of the approaches: regulation, privacy policies, anonymity, and obfuscation.

There remain many challenges for privacy researchers. For example, for information to be worth protecting, it must also be worth attacking. Current research tends to be biased toward privacy protection. By contrast, it is also important to understand the techniques a hostile agent might employ in order to invade a person's privacy (circumventing location privacy protection and attempting to discover an individual's exact location). In this respect, privacy research is analogous to cryptology, which comprises both cryptography (code making) and cryptanalysis (code breaking).

As this chapter has shown, location information differs from many other types of personal information. Consequently, future research aimed specifically at *location* privacy will need to focus on specialized privacy protection techniques for several reasons. First, unlike many other types of personal information, identity may be inferred from location. Such inferences are especially likely where a history of locations can be derived (for example, my patterns of movement over the course of a week). These types of inferences make anonymity and pseudonymity much harder to maintain than in other privacy applications, such as Internet use.

Second, information about personal location is highly dynamic. By contrast, current research approaches to location privacy are usually fundamentally *static* in nature, modeling the movement of an individual as a sequence of static snapshot locations. Many aspects of location privacy demand models that provide a more faithful representation of the temporal aspects of LBS. For example, counter-strategies for invading an individual's privacy can be devised by making assumptions about an individual's maximum or minimum speeds of movement. Understanding such counter-strategies requires requires the development of truly spatiotemporal models of location privacy. Further, the potential uses and privacy implications of dynamic location information change over time. Current privacy protection strategies, such as regulation and privacy policies, tend to make no distinction between static information (such as an individual's date of birth) and dynamic information (such as an individual's location). Thus, these approaches may ignore the dynamic aspects of location information, making it difficult to definite privacy policies that have a temporal component, for example, where acceptable uses change over time.

Finally, the potential uses of spatial information are highly varied. Correspondingly, the potential benefits of invading an individual's location privacy may be higher than for some other types of information. Without proper protection, the location information generated by location-aware systems could conceivably be abused or unfairly used in almost any domain of human, social, or economic activity, including marketing, insurance, surveillance, harassment, social security, politics, law enforcement, health, or employment. Indeed, it is this very feature of location information that makes location information so vital to our future information systems.

Acknowledgments

Dr Duckham is partially supported by funding from the Australian Academy of Science, an Edward Clarence Dyason Fellowship, and an International Collaborative Research Grant from the University of Melbourne. Dr Duckham and Dr Kulik are further supported by Early Career Researcher Grants from the University of Melbourne and by an ARC Linkage Grant, entitled “CEWAY: Cognitively ergonomic wayfinding directions for location-based services.”

References

[1] M. S. Ackerman, L. F. Cranor, and J. Reagle. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proc. 1st ACM conference on Electronic Commerce*, pages 1–8. ACM Press, 1999.

[2] M. S. Ackerman, T. Darrell, and D. J. Weitzner. Privacy in context. *Human-Computer Interaction*, 16(2, 3, & 4):167–176, 2001.

[3] Australian Government. Privacy act. http://www.privacy.gov.au/act/, 1988. Accessed 25 July 2005.

[4] P. Bahl and V.N. Padmanabhan. Radar: An in-building RF-based user lcoation and tracking system. In *Proceedings IEEE INFOCOM 2000*, volume 2, pages 775–784, 2000.

[5] L. Barkuus and A. Dey. Location-based services for mobile telephony: A study of users’ privacy concerns. In *Proc. INTERACT 2003, 9th IFIP TC13 International Conference on Human-Computer Interaction*, 2003.

[6] R. Beckwith. Designing for ubiquity: The perception of privacy. *IEEE Pervasive Computing*, 2(2):40–46, 2003.

[7] A.R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, 2003.

[8] X. Bian, G. D. Abowd, and J. M. Rehg. Using sound source localization in a home environment. In H. W. Gellersen, R. Want, and A. Schmidt, editors, *Pervasive 2005*, volume 3468, pages 19–36. Springer, Berlin, 2005.

[9] S. Brands. Untraceable off-line cash in wallet with observers. In *CRYPTO ’93: Proc. 13th annual international cryptology conference on Advances in cryptology*, pages 302–318, Berlin, 1994. Springer.

[10] D. Brin. *The transparent society*. Perseus, 1999.

[11] J. Canny. Some techniques for privacy in ubicomp and context-aware applications. In *Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing*, Göteborg, Sweden, September 2002.

[12] Chicago Tribune. Rental firm uses GPS in speeding fine. July 2nd, p9. Associated Press: Chicago, IL, 2001.

[13] R. Clarke. Internet privacy concerns confirm the case for intervention. *Communications of the ACM*, 42(2):60–67, 1999.

[14] L. F. Cranor. P3P: The platform for privacy preferences project. In S. Garfinkel and G. Spafford, editors, *Web Security, Privacy, and Commerce*, pages 699–707. O'Reilly, Sebastopol, CA, 2nd edition, 2001.

[15] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation Onion router. In *Proc. 13th USENIX Security Symposium*, 2004.

[16] M. Duckham and L. Kulik. A formal model of obfuscation and negotiation for location privacy. In H. W. Gellersen, R. Want, and A. Schmidt, editors, *Pervasive 2005*, volume 3468 of *Lecture Notes in Computer Science*, pages 152–170. Springer, Berlin, 2005.

[17] M. Duckham and L. Kulik. Simulation of obfuscation and negotiation for location privacy. In D.M. Mark and A.G. Cohn, editors, *COSIT 2005*, volume 3693 of *Lecture Notes in Computer Science*, pages 31–48. Springer, Berlin, 2005.

[18] M. Duckham, K. Mason, J. Stell, and M. Worboys. A formal approach to imperfection in geographic information. *Computers, Environment and Urban Systems*, 25:89–103, 2001.

[19] S. Duri, M. Gruteser, X. Liu, P. Moskowitz, R. Perez, M. Singh, and J-M. Tang. Framework for security and privacy in automotive telematics. In *Proc. 2nd International Workshop on Mobile Commerce*, pages 25–32. ACM Press, 2002.

[20] F. Espinoza, P. Persson, A. Sandin, H. Nyström, E. Cacciatore, and M. Bylund. GeoNotes: Social and navigational aspects of location-based information systems. In G. D. Abowd, B. Brumitt, and S. Shafer, editors, *Ubicomp 2001: Ubiquitous Computing*, volume 2201 of *Lecture Notes in Computer Science*, pages 2–17. Springer, 2001.

[21] A. Etzioni. A contemporary conception of privacy. *Telecommunications and Space Journal*, 6:81–114, 1999.

[22] A. Etzioni. Less privacy is good for us (and you). *Privacy Journal*, April:3–5, 1999.

[23] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Proc. on Advances in cryptology—CRYPTO '86*, pages 186–194, Berlin, 1987. Springer.

[24] General Assembly of the United Nations. Universal declaration of human rights. United Nations Resolution 217 A (III), December 1948.

[25] S Goldwasser, S Micali, and C Rackoff. The knowledge complexity of interactive proof-systems. In *STOC '85: Proceedings of the seventeenth annual ACM*

symposium on Theory of computing, pages 291–304, New York, NY, 1985. ACM Press.

[26] W. W. Görlach, A. Terpstra and A. Heinemann. Survey on location privacy in pervasive computing. In *Proc. First Workshop on Security and Privacy at the Conference on Pervasive Computing (SPPC)*, 2004.

[27] M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proc. MobiSys '03*, pages 31–42, 2003.

[28] M. Gruteser and D. Grunwald. A methodological assessment of location privacy risks in wireless hotspot networks. In D. Hutter, G. Müller, and W. Stephan, editors, *Security in Pervasive Computing*, volume 2802 of *Lecture Notes in Computer Science*, pages 10–24. Springer, 2004.

[29] C. A. Gunter, M. J. May, and S. G. Stubblebine. A formal privacy systems and its application to location-based services. In *Proc. Workshop on Privacy Enhancing Technologies*, Toronto, Canada, 2004.

[30] R. H. R. Harper. Looking at ourselves: An examination of the social organisation of two research laboratories. In *Proc. 1992 ACM conference on Computer-Supported Cooperative Work*, New York, 1992. ACM Press.

[31] R. H. R. Harper, M.G. Lamming, and W. M. Newman. Locating systems at work: Implications for the development of active badge applications. *Interacting with Computers*, 4(3):343–363, 1992.

[32] A. Helal, S. Moore, and B. Ramachandran. Drishti: An integrated navigation system for visually impaired and disabled. In *Proceedings Fifth International Symposium on Wearable Computer*, Zurich, Switzerland, 2001.

[33] J. Hightower and G. Boriello. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, 2001.

[34] J. I. Hong and J. A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *Proc. 2nd International Conference on Mobile Systems, Applications, and Services*, pages 177–189. ACM Press, 2004.

[35] S. E. Hudson and I. Smith. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proc. ACM conference on Computer Supported Cooperative Work*, pages 248–257. ACM Press, 1996.

[36] D. Hutter, W. Stephan, and M. Ullmann. Security and privacy in pervasive computing: State of the art and future directions. In D. Hutter, G. Müller, and W. Stephan, editors, *Security in Pervasive Computing*, volume 2802 of *Lecture Notes in Computer Science*, pages 284–289. Springer, 2004.

[37] IBM. The enterprise privacy authorization language (epal 1.1). http://www.zurich.ibm.com/security/enterprise-privacy/epal/, 2005. Accessed 2 August 2005.

[38] X. Jiang, J. I. Hong, and J. A. Landay. Approximate information flows:socially-based modeling of privacy in ubiquitous computing. In G. Borriello and L. E. Holmquist, editors, *Proc. 4th international conference on Ubiquitous Computing*, volume 2498 of *Lecture Notes in Computer Science*, pages 176–193. Springer, Berlin, 2002.

[39] E. Kaasinen. User needs for location-aware mobile services. *Personal and Ubiquitous Computing*, 7(1):70–79, 2003.

[40] J. Krumm, S. Harris, B. Meyers, B. Brumitt, M. Hale, and S. Shafer. Multi-camera multi-person tracking for EasyLiving. In *Proceedings Third IEEE Workshop on Visual Surveillance VS2000*, pages 3–10, 2000.

[41] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, H. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. N. Schilit. Place lab: Device positioning using radio beacons in the wild. In H. W. Gellersen, R. Want, and A. Schmidt, editors, *Pervasive 2005*, volume 3468, pages 116–133. Springer, Berlin, 2005.

[42] M. Langheinrich. Privacy by design—principles of privacy-aware ubiquitous systems. In G. D. Abowd, B. Brumitt, and S. Shafer, editors, *Ubicomp 2001: Ubiquitous Computing*, volume 2201 of *Lecture Notes in Computer Science*, pages 273–291. Springer, 2001.

[43] M. Langheinrich. A privacy awareness system for ubiquitous computing environments. In G. Borriello and L. E. Holmquist, editors, *UbiComp 2002: Ubiquitous Computing*, volume 2498 of *Lecture Notes in Computer Science*, pages 237–245. Springer, 2002.

[44] J-W Lee. Location-tracing sparks privacy concerns. Korea Times, http://times.hankooki.com, 16 November 2004. Accessed 26 July 2005.

[45] P. Ljungstrand. Context awareness and mobile phones. *Personal and Ubiquitous Computing*, 5(1):58–61, 2001.

[46] N. Marmasse and C. Schmandt. Location-aware information delivery with comMotion. In *Proceedings 2nd International Symposium on Handheld and Ubiquitous Computing (HUC)*, pages 157–171, Bristol, UK, 2000.

[47] R.R. Muntz, T. Barclay, J. Dozier, C. Faloutsos, A.M. Maceachren, J.L. Martin, C.M. Pancake, and M Satyanarayanan. *IT Roadmap to a Geospatial Future*. The National Academies Press, Washington, DC, 2003.

[48] G. Myles, A. Friday, and N. Davies. Preserving privacy in environments with location-based applications. *Pervasive Computing*, 2(1):56–64, 2003.

[49] H. J. Onsrud, J. Johnson, and X. Lopez. Protecting personal privacy in using geographic information systems. *Photogrammetric Engineering and Remote Sensing*, 60(9):1083–1095, 1994.

[50] Organisation for Economic Co-operation and Development (OECD). Guidelines on the protection of privacy and transborder flows of personal data. http://www.oecd.org, 1980. Accessed 25 July 2005.

[51] C Park. Location-based information service due next year. Korea Times, http://times.hankooki.com, 2 July 2004. Accessed 26 July 2005.

[52] J. Peterson. A presence-based GEOPRIV location object format. http://www.ietf.org/internet-drafts/draft-ietf-geopriv-pidf-lo-03.txt, September 2004.

[53] A. Pfitzmann and M. Köhntopp. Anonymity, unobservability, and pseudonymity—a proposal for terminology. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2001.

[54] T. Rodden, A. Friday, H. Muller, and A. Dix. A lightweight approach to managing privacy in location-based services. Technical Report Equator-02-058, University of Nottingham, Lancaster University, University of Bristol, 2002.

[55] M. Rotenberg and C. Laurant. Privacy and human rights 2004: An international survey ofprivacy laws and developments. http://www.privacyinternational.org/survey/, Privacy International, 2004. Accessed 26 July 2005.

[56] J. B. Rule. *Private lives and public surveillance*. Allen Lane, London, 1973.

[57] B. N. Schilit, J.I. Hong, and M. Gruteser. Wireless location privacy protection. *IEEE Computer*, 36(12):135–137, 2003.

[58] B.N. Schilit and M. M. Theimer. Disseminating active map informationto mobile hosts. *IEEE Network*, 8(5):22–32, 1994.

[59] A. Schmidt, M. Beigl, and H-W. Gellerson. There is more to context than location. *Computer and Graphics Journal*, 23(6):893–902, 1999.

[60] J. Scott and B. Dragovic. Audio location: Accurate low-cost location sensing. In H. W. Gellersen, R. Want, and A. Schmidt, editors, *Pervasive 2005*, volume 3468, pages 1–18. Springer, Berlin, 2005.

[61] S. Scott-Young and A. Kealy. An intelligent navigation solution for land mobile location based services. *Journal of Navigation*, 55:225–240, 2002.

[62] E. Snekkenes. Concepts for personal location privacy policies. In *Proc. 3rd ACM conference on Electronic Commerce*, pages 48–57. ACM Press, 2001.

[63] P. Sprenger. Sun on privacy: "Get over it". Wired, Janurary 26 1999.

[64] V. Stanford. Using pervasive computing to deliver elder care. *IEEE Pervasive Computing*, 1(1):10–13, 2002.

[65] M. Strassman and C. Collier. Case study: Development of the *Find Friend* application. In J. Schiller and A. Voisard, editors, *Location-based services*, chapter 2, pages 27–39. Morgan Kaufmann, 2004.

[66] U.K. Government. *Data Protection Act*. HMSO, London, 1998.

[67] U.S. Department of Justice, Office of Information and Privacy. Overview of the Privacy Act of 1974, May 2004.

[68] U.S. Deptartment of Health, Education and Welfare, Secretary's Advisory Committee on Automated Personal Data Systems. *Records, Computers, and the Rights of Citizens*. MIT Press, Cambridge, MA, 1973.

[69] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The Active Badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, 1992.

[70] K. Werbach. Location-based computing: Wherever you go, there you are. *Release 1.0*, 18(6):1–26, 2000.

[71] A. F. Westin. *Privacy and freedom*. Atheneum, New York, 1967.

[72] M. F. Worboys and E. Clementini. Integration of imperfect spatial information. *Journal of Visual Languages and Computing*, 12:61–80, 2001.

[73] M.F. Worboys and M. Duckham. *GIS: A Computing Perspective*. CRC Press, Boca Raton, FL, 2nd edition, 2004.

[74] World Wide Web Consortium (W3C). Platform for privacy preferences project (p3p). http://www.w3.org/P3P/, 2005. Accessed 2 August 2005.

[75] M. Wu and A. Friday. Integrating privacy enhancing services in ubiquitous computing environments. In *Proc. Workshop on Security in Ubiquitous Computing,4th Intl. UbiComp Conference*, 2002.