Binding Privacy Rules to Location on the Web

Alissa Cooper
Center for Democracy & Technology
1634 Eye St. NW, Suite 1100
Washington, DC 20006
+1 202 637 9800 x110
acooper@cdt.org

John Morris
Center for Democracy & Technology
1634 Eye St. NW, Suite 1100
Washington, DC 20006
+1 202 637 9800 x116
jmorris@cdt.org

ABSTRACT

As a tool for mitigating the potential privacy risks of gathering and transmitting location information on the Web, we suggest in this paper a model for conveying location information together with privacy rules to govern the use of that information. Binding privacy rules to the conveyance of location information is one useful tool to help developers build location-based systems and services that comport with the concept of *fair information practices* (FIPs)—a set of widely accepted principles that create a basis for privacy-protective systems. We offer as a model one fully developed set of standards for binding location data conveyed across IP networks to privacy rules.

Categories and Subject Descriptors

H.3.5 [**Online Information Services**]: Web-based services; Data sharing; K.4 [**Computers and Society**]: Public Policy Issues – *Privacy; Regulation; Ethics*; K.5 [**Legal Aspects of Computing**]: Governmental Issues - *Regulation*;

General Terms

Design, Human Factors, Standardization, Legal Aspects

Keywords

Privacy, location, policy

1. INTRODUCTION

The ubiquity of increasingly high-powered mobile devices has already spawned the Web's first generation of location-based services and applications. As the accuracy of location data improves and the expense of calculating and obtaining it declines, location may well come to pervade the Web experience. While the increasing availability of location information paves the way for exciting new applications and services, it comes with potential privacy concerns.

As a tool for mitigating the potential privacy risks of gathering and transmitting location information on the Web, we suggest in this paper a model for conveying location information together with privacy rules to govern the use of that information. Binding privacy rules to the conveyance of location information is one useful tool to help developers build location-based systems and services that comport with the concept of *fair information practices* (FIPs)—a set of widely accepted principles that create a basis for privacy-protective systems [1]. We offer as a model one fully developed set of standards for binding location data conveyed across IP networks to privacy rules.

Binding rules to location is suggested merely as one component of a broader privacy protection scheme for location-based applications. There are many other possible approaches to location privacy. Duckham and Kulik have defined four categories of strategies [8]: obfuscation (e.g., location “fuzzing”), de-identification of location data, regulatory approaches, and privacy policies (which the transmission of privacy rules may arguably fall under). While the discussion here is limited to rules, this focus is not meant to preclude other approaches; indeed, it is likely that a combination of techniques will provide the most robust privacy assurances.

2. LOCATION PRIVACY CONCERNS

The increasingly easy availability of location information raises several different kinds of privacy concerns. While some of these are common to all forms of personal data, many of them are heightened or uniquely applicable to location information.

Because individuals often carry their mobile devices with them, location data may be used to form a comprehensive record of an individual's movements and activities. While other kinds of data could arguably be considered more sensitive than location information in certain contexts—an individual's medical record or bank statement, for instance—these kinds of data provide mere snapshots of an individual's activities at discrete moments in time, or within discrete aspects of their lives. Location data, on the other hand, may be collected everywhere and at any time, often without user interaction, and it may potentially describe both what a person is doing and where he or she is doing it. The fact that an individual's mobile phone location is triangulated when he is at the bank can reveal the fact that he was at the bank, when he was there, and which branch he uses. Amassing such data points about an individual's every movement allows for the creation of richly detailed profiles of individual behavior.

The availability of location information may also allow an individual's whereabouts to unwittingly become more public than desired, with potentially grave consequences. Location information may reveal the fact that an individual was in a particular medical clinic or government building, for example, implying information about the individual that was not meant to be shared. The ubiquity of location information may also increase the risks of stalking and domestic violence if perpetrators are able to use (or abuse) location-based services to gain access to location information about their victims. Location information also raises

LocWeb 2009, April 4, 2009 Boston, MA, USA.

enormous child safety concerns as more and more children access mobile devices.

Even when location information is not being used nefariously, its mere availability opens the door for applications that could easily spark consumer backlash. Highly targeted location-based advertising campaigns may seem attractive to advertisers, but when implemented without appropriate user input or involvement, they run the risk of seeming creepy or invasive.

Finally, location information is and will continue to be of particular interest to governments and law enforcers around the world. In jurisdictions such as the United States, the standards for government access to location information held by companies are unclear at best, and far too low at worst [2]. The existence of detailed records of individuals' movements should not automatically facilitate the ability for governments to track their citizens, but in many cases, laws dictating what government agents must do to obtain location data have not kept pace with technological evolution.

3. FAIR INFORMATION PRACTICES

Given the unique privacy challenges that the availability of location information presents, the question of how to convey location information in a privacy-sensitive manner across the Web is central. But what does it mean to be privacy-sensitive? Over the past several decades, numerous governments, industry bodies, and privacy experts have coalesced around the notion of *fair information practice (FIP)* principles that describe how data about individuals can be handled in a privacy-protective way. The principles are intended to foster individuals' control over their personal information, limit data collection, and place responsibilities on data collectors. These principles are the basis for current privacy laws and policies.

3.1 FIP Examples

While all of the FIP principles are relevant to location-based services and applications, the following subset is illustrative of how FIPs may broadly apply to location information on the Web:

Individual Participation. The principle of individual participation states that individuals should have a right to view the information that is collected about them. They must also be able to correct or remove data that is not timely, accurate, relevant, or complete.

Individual participation is ultimately about involving individuals in the process of data collection, use and storage. The seamlessness of current location technologies allows location information to be continuously collected without the user's knowledge. This means that those developing location-based systems need to put in extra effort to promote awareness of location collection among users, whether through UI design, educational efforts, or other means.

Collection Limitation. The collection limitation principle states that there should exist restrictions on the collection of personal data. Data should be collected by lawful and fair means and should be collected, where appropriate, with the knowledge or consent of the subject.

In the location context, collection limitation means only collecting information in the amount and level of detail that is necessary to implement a particular service or application. For example, if "fuzzed" location suffices for a particular Web application, or if location can be sampled daily rather than hourly, location-based systems should be designed to reduce potential privacy impact by taking these considerations into account.

Use Limitation. This principle states that there should be limits to the use and disclosure of personal data. Data should be used only for purposes specified at the time of collection. Data should not be otherwise disclosed without consent.

Once the operator of a service or application obtains location information, all kinds of new uses and disclosures may become more attractive. For example, the Web site of a pizza restaurant may collect an individual's location to display a map of the nearest chain locations. Users may expect the restaurant to then discard the data, but the restaurant may have incentives to sell the data to marketers, use the data to market back to individuals, or even disclose it to law enforcement should it continue to exist in storage when a subpoena or other request comes in. Limiting uses to those specified at the outset or permitted by the individuals involved serves as an important check on the potential misuses of the masses of location data being collected.

3.2 FIP Enforceability Through Rules

Enforcing compliance with the FIP principles has historically been accomplished in large part through legal and policy means rather than technical measures. But the emergence of ubiquitous location information provides a new opportunity to supplement these tools with the force of technology. By designing Web applications and services that bind location information together with privacy rules, the operators of these systems will have additional means at their disposal to help them comply with the FIPs and better protect privacy in the process. The most prominent and widely accepted model for technologically binding privacy rules to location conveyance is the IETF's Geopriv work, as explored in the next section.

4. GEOPRIV

The IETF's Geopriv standards comprise a helpful example for understanding how location conveyance and privacy rules may be bound together. This approach provides a useful starting point for thinking about conveying location information in a privacy-sensitive way on the Web.

4.1 Geopriv History

When discussions in the IETF about standardizing location conveyance first began nearly a decade ago, there were a number of efforts that sought to standardize location conveyance without fully addressing the privacy concerns raised by such conveyance. After much debate, the IETF's Internet Engineering Steering Group concluded that privacy and security had to be an integral part of any standard to send or carry location information, which is why the Geopriv work focuses heavily on privacy concerns [4].

As location-based Web applications proliferate, the same pressures and tensions that reared their heads in the early days of the Geopriv effort are likely to emerge once again. For Web applications that "just want coordinates," the prospect of also handling and complying with privacy rules may seem difficult or unnecessary. But the unique sensitivity of location information demands a more rigorous approach to privacy, and the burgeoning array of location applications provide a crucial opportunity to improve upon the status quo for how all personal data is handled.

4.2 Location Object

The Geopriv standards call for the creation of location objects, which contain a location along with a limited set of rules that can point to an external set of more complex rules, if necessary. The location information itself can be expressed in multiple different formats, including latitude/longitude/altitude coordinates via the GML Markup Language [3] or more traditional civic location identifiers (street, city, region, etc.).

The Geopriv standard describes two sets of privacy rules—a limited set that any Geopriv object "must carry," and a more robust set that can be stored or referenced externally. By requiring a limited set of rules to be bound to the location object itself, Geopriv ensures that no recipient can claim ignorance of the basic privacy rules that apply to that information.

Using a minimal set of required rules with an optional extended set of rules has additional benefits in the Web context. The minimal rule set is lightweight, making it easy for Web developers to implement a privacy-protective framework without having to create a complex privacy rule scheme. Conversely, the existence of an extensible framework for more granular rules allows privacy-aware developers ample space to offer exactly the kinds of privacy options they desire to their users.

4.3 Privacy Rules

The Geopriv standards bind the following privacy rules into the location object [7]:

Retention limit date and time ("retention-expires"). The time limit is an indication of how long location information can be retained by its recipient (not how long it remains valid). This rule squarely supports the FIP principle of use limitation by restricting the length of time that location information is held.

Indication of consent (or lack thereof) to retransmit location information ("retransmission-allowed"). For many simple location transactions (such as, "Where is the closest pizza place to where I am right now?"), a denial of retransmission consent coupled with a very short retention time limit effectively conveys that the recipient should respond to the immediate query and then discard location information. This rule also supports the portion of the use limitation principle focused on limiting disclosure.

Pointer to external, fuller set of privacy rules for retransmission of location information ("ruleset-reference"). In the pizza place example, no pointer is needed because no permission is granted for information retransmission. But in many other cases, it may be useful or desirable to have a more granular set of privacy rules associated with location information. For an application that allows for geo-tagging photos, for example, users may want to allow certain other users or services to retransmit their location information (associated with a photo), but not extend this privilege to all recipients of the photos. Geopriv has defined a flexible, extensible standard for expressing such rules.

Free-form text area ("note-well"). This area can convey the privacy policy in a human-readable format.

Although limited, these rules are sufficient to cover many forms of consumer-oriented location services, including those in which information or immediate service is based on location (such that no continuing services are sought or expected).

4.4 Geopriv Entities

The Geopriv standard refers to four primary logical entities [6]:

The *location generator* (LG) determines the location of an individual (known in Geopriv parlance as a "target"). In some scenarios the LG acts as a proxy for the target (such as the mobile phone that the target carries). The LG can determine its own location (using on-board GPS, for example), or from another source (such as a network access provider or even manual human entry of location).

The *rule holder* (RH) stores the privacy rules to be applied to location information. One or more *rule makers* (RMs) create the rules and set the policies governing location information. In the Web context, an RH could be an individual's browser or OS, or a remote server designated by the user to store his or her privacy rules.

The *location server* (LS) receives location objects from location generators and privacy rules from rule holders. The location server responds to requests for location information by applying the appropriate privacy rules and determining whether a particular request is authorized.

The *location recipient* (LR) receives the location object from the LS (or can subscribe to receive regular location updates).

Under this architecture, the LS essentially serves as a clearinghouse for location information and a broker for privacy rules. When an LR wants to receive a target's location, the LS must first obtain the location from an LG, obtain the rules from an RH, and apply the rules to determine whether and to what extent the LR is authorized to receive the target's location. The architecture is pictured in Figure 1.

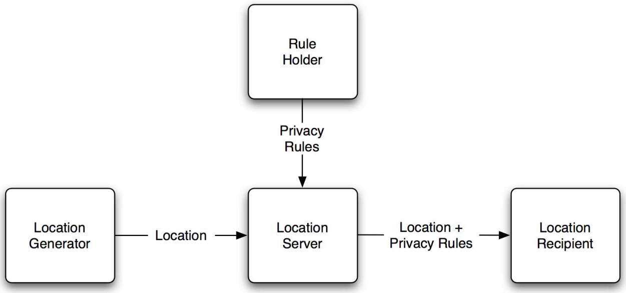

Figure 1. A basic version of the Geopriv architecture.

4.5 Geopriv Example for the Web

Since the Geopriv entities are merely logical, there are many ways in which the Geopriv architecture could map to the Web. In one plausible scenario, the user's browser acts as both her location server and and her rule holder. The browser receives the user's location from a location generator (perhaps the GPS or WiFi interface on the device where the browser is running). The browser also exposes an interface that allows the user to set rules about which Web applications may access her location, and it retains those rules. When a request for the user's location comes in from a Web application, the browser checks the user's privacy rules, and if it determines that the application has authorization, it obtains the location from the location generator and delivers the location, together with the rules, to the location recipient. This architecture is pictured in Figure 2.

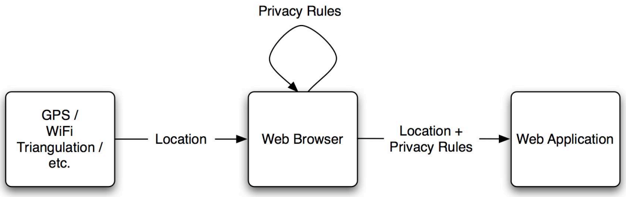

Figure 2. An example Geopriv architecture for the Web.

Focusing solely on the "retention-expires" and "retransmission-allowed" rules for the moment, imagine that the browser implements a set of global privacy rules that apply to all requesting Web applications, with strong defaults: retransmission prohibited and retention set to expire 24 hours after transmission. The browser could also allow users to set these rules differently on a per-domain basis (similar to how many browsers have global cookie-setting rules but allow users to make exceptions on a per-domain basis). With this simple rule structure in place, the user can convey a wide array of preferences.

Under these circumstances, the user might set different preferences for different Web applications based on what the applications do. For a Web application that finds the nearest pizza place, the default rules are likely appropriate: the pizza place has no reason to retransmit the user's location nor retain it for any length of time, because it is merely responding to a one-time request for the nearest restaurant. For an applications that displays local tourist attractions, the user might set the retention period to be longer with the thought that she could use the application whenever she goes on vacation and for the duration of her trip. For an application that allows users to geo-tag photos and post them publicly on the Web, the user would likely allow retransmission and set retention to be indefinite, since the whole point of the application is to create a public log of the photos' locations.

4.6 Interplay with Privacy Law

From a privacy perspective, Geopriv offers the opportunity to convey fairly robust and potentially complex privacy rules along with location information. It cannot, however, provide guarantees that those rules will be honored or followed in any given situation. Yet, Geopriv could be a critical element of a larger policy framework (perhaps created by local law) that provides such guarantees through legal rather than technical means.

This interplay between law and technology could prove beneficial in various ways. A law could decree that no location information be distributed without the express permission of the person being tracked (aside from emergency and authorized law enforcement cases), and Geopriv could provide the means to grant or deny such permission. Alternatively, a law might allow such distribution unless the consumer takes a proactive step to deny permission; again, Geopriv could be the consumer's vehicle. In this way, technical approaches to privacy like Geopriv can serve to supplement existing protections in law and policy.

5. LOCATION ON THE WEB

Binding privacy rules to location information forces the recipient of the location information to confront the privacy rules. Recipients may choose to ignore the rules, as there may be no technical way to ensure that the rules are honored. In the Web context, any approach that squarely confronts Web developers with clear privacy expectations will lead some to focus on privacy, even if some will not. As location-based services and applications proliferate, it is a worthwhile endeavor to spur such confrontation, even while acknowledging that perfect protection is out of reach.

The growth of location-based Web services and applications also provides a unique opportunity to improve over the status quo for privacy protection on the Web. Traditionally, much of the sensitive data that Web users transmit and share online has been governed first and foremost by Web site privacy policies that most users never read and cannot understand. The availability of user tools and controls that provide individuals with the ability to participate in decision-making about their own data has been limited. Designing in user empowerment mechanisms – one component of what has been referred to as "ethics-driven design" [5] – helps to reverse this trend.

Location information is available through more interfaces and devices than ever before. Yet because of its sensitive properties, it deserves a model for conveyance that has privacy protections built in. Binding privacy rules to location information requires recipients of location information to confront privacy head-on. Given current trends towards involving users in decisions about who can see their data and how it can be used, it is high time for Web-based location solutions to consider incorporating strong technical privacy protections.

6. REFERENCES

[1] Center for Democracy & Technology. CDT's Privacy Guide: Privacy Basics: Fair Information Practices. 2000. http://www.cdt.org/privacy/guide/basic/fips.html.

[2] Center for Democracy & Technology. Digital Search & Seizure: Updating Privacy Protections to Keep Pace with Technology. 2006. http://www.cdt.org/publications/digital-search-and-seizure.pdf.

[3] Geography Markup Language. http://www.opengeospatial.org/standards/gml.

[4] Geopriv charter. http://www.ietf.org/html.charters/geopriv-charter.html.

[5] H. Onsurd. 2008. Implementing Geographic Information Technologies Ethically. ArcNews 2008 Fall Issue. http://www.esri.com/news/arcnews/fall08articles/implementing-gi-technologies.html.

[6] J. Cuellar et al. Geopriv Requirements. RFC 3693. Oct. 2003.

[7] J. Peterson. A Presence-based GEOPRIV Location Object Format. RFC 4119. Dec. 2005.

[8] M. Duckham and L. Kulik. 2006. Location privacy and location-aware computing. In Dynamic & Mobile GIS: Investigating Change in Space and Time, J. Drummond et al, Eds. CRC Press, Boca Raton, FL, USA, 34-51.