

LocationSafe: Granular Location Privacy for IoT Devices

Joshua Joy
jjoy@cs.ucla.edu

Minh Le
im1nh27@ucla.edu

Mario Gerla
gerla@cs.ucla.edu

ABSTRACT

Today, mobile data owners lack consent and control over the release and utilization of their location data. Third party applications continuously process and access location data without data owners granular control and without knowledge of how location data is being used. The proliferation of GPS enabled IoT devices will lead to larger scale abuses of trust.

In this paper we present the first design and implementation of a privacy module built into the GPSD daemon. The GPSD daemon is a low-level GPS interface that runs on GPS enabled devices. The integration of the privacy module ensures that data owners have granular control over the release of their GPS location. We describe the design of our privacy module integration into the GPSD daemon.

CCS Concepts

•Security and privacy → Mobile and wireless security; •Networks → Location based services;

Keywords

Mobile IoT; Location Privacy

1. INTRODUCTION

Today data owners' personal mobile devices are constantly being tracked and monitored by third party applications without data owners granular consent and control. Data owners' trust is being continuously violated [3].

Data owners have a desire to occasionally share their location data, though desire granular control and approved consent. Third party analysts seek to track data owners continuously. Unfortunately today this tension has resulted in disproportionate control being in favor of the third party analysts.

Recent research has tried to improve user behavior in recognizing permission issues [5], user-defined runtime constraints [9], or tools to help developers identify least-privilege

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

S3'16, October 03-07, 2016, New York City, NY, USA

© 2016 ACM. ISBN 978-1-4503-4255-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2987354.2987365>

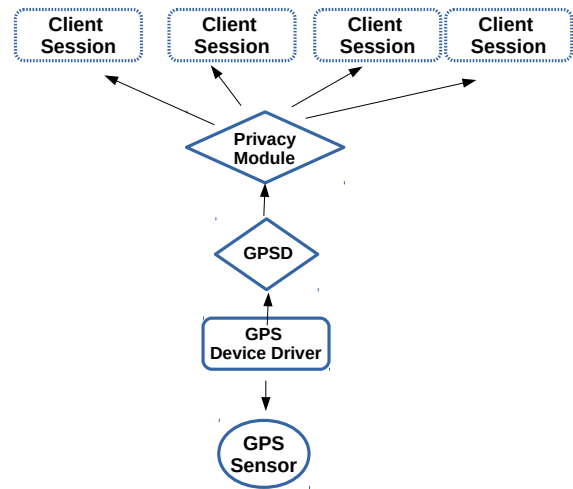


Figure 1: Privatization occurs before data is released to the client application.

[11].

Additionally, permission managers (e.g., Android and iOS) lack fine grained granularity and offer “binary” permissions to disable or enable location services (while the app is running in the background or foreground). While this allows data owners to disable location services for applications that do not require location (e.g., Flashlight application) [6], fine grained granularity over release frequency and accuracy is still missing. An Android modification called CynagonMod has a module called XPrivacy [12]. XPrivacy enables data owners to configure random or a static location, empty cell ID, blocks geofences from being set, prevents sending NMEA data to application, prevents cell tower updates from being sent to an application, prevents aGPS, returns empty Wi-Fi scans, and disables activity recognition. Ultimately, this provides the data owner control at the application layer.

User applications requesting data of users is a binary permission, either I share my data or I don't. However, sensitive data such as location needs finer control on how accurate and how often the location information is released. Users should be able to control the granularity of their personal data that is released. Users require freedom and control over their own personal data.

However, these approaches discards several important facts: 1) these privacy mechanisms protect at the application layer only and system services (e.g., Google Play Services) still

has unfettered access to GPSD 2) granular privacy permission solutions (e.g., XPrivacy) are only for rooted Android phones 3) there is no compromise between third party analyzers and data owners 4) there is no cross-platform approach that can be signed and verified. The expected proliferation of GPS enabled IoT devices will further exacerbate these privacy issues.

In this paper, we present the first (to our knowledge) implementation of a privacy module to GPSD. Figure 1 shows an overview of the flow of queries and responses and demonstrates that the privatization occurs before releasing the data back to the application. The privacy module ensures that all GPS data is released according to the data owner’s consent and choice. We demonstrate that appropriate methodologies can be placed which provides strong location privacy guarantees, yet enable analyzers access to privatized location data.

1. A privacy module that integrates into the GPSD software (runs on every GPS enabled device)
2. A granular privacy interface and control to manage location privacy settings (e.g., location coarseness and release frequency)
3. A performant privacy module with minimal overhead

We first describe the architecture and flow of GPSD, we then describe our privatization algorithms and then we describe our integration with GPSD.

2. RELATED WORK

GPSD is a daemon that network enables the GPS sensor on the majority of mobile embedded systems including Android, iOS, Windows Mobile, UAVs, and driverless cars [7]. On smartphones the network access is limited to localhost applications only (as opposed to remote applications). GPSD enables unfettered access to location data and does not enable or provide any privacy guarantees. LOCATIONSAFE provides a privacy module that provides uniform private access across all platforms.

Mobile device permission systems has received attention in the past. Human interaction studies which seek to enhance reader comprehension have been proposed and evaluated [5, 4]. Such systems lack strong and enforcable privacy guarantees. Static analysis tools have been proposed [13]. Though such systems serve only to notify the data owner of privacy breaches and are unable to enforce any privacy runtime guarantees. However, these solutions modify the underlying OS thus making them specific to a single OS or device [14, 12]. Furthermore, these solutions are unable to balance the privacy and utility tradeoff, ultimately resulting a binary approach to privacy.

To guarantee data owner privacy upon the release of data, various mechanisms have been proposed [10, 8, 1, 2]. Differential privacy has emerged as the strongest of these privacy mechanisms [1, 2]. The core idea of differential privacy is to provide strong bounds and guarantees on the privacy leakage when multiple aggregate analytics are run despite the presence or absence of a single data owner from the dataset. This privacy mechanism is provided by adding differentially private noise to the aggregate answer. As opposed to the originally proposed differentially private mechanism which first collects data in a centralize database and then privatizes the release of the data, LOCATIONSAFE immediately privatizes the data at the data source (sensor) in real-time.

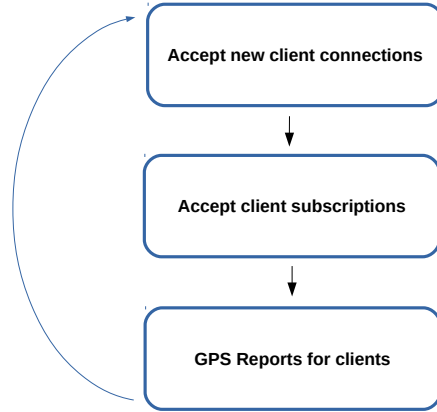


Figure 2: GPSD event loop. Privatization occurs when reporting GPS data to the client.




Figure 3: In the grid privatization a single location may randomize to one or many locations. In the example above two locations are returned. However, in the aggregate the analyst is able to estimate the underlying population value without violating individual privacy.

3. GOALS AND PROBLEM STATEMENT

We now describe the system goals, performance goals, threat model, and privacy goals of LOCATIONSAFE.

3.1 System Goals

There should be well defined and enforced constraints regarding third party application’s (apps) access to location data. The data owner should be able to specify the constraints such as how accurate location information should be disclosed and how frequent the location data should be disclosed.

Apps only have access to the privatized data and are unable to directly access GPSD daemon and data. All location data released must be approved by the data owner.

The system should support applications that need real-time access to location data. The privacy policy defines how frequently the application is allowed to receive updates (express in epochs), how accurate the location data may be, and geographical regions as to where the application is allowed to receive location data from.

We use a social network messaging application as an ex-

ample. The application may want to know which city an individual is in, though pinpoint location information within meter accuracy is not required. The data owner is allowed to define both the radius (e.g., city) that is allowed to be returned as well as the frequency (e.g., say at most every hour).

Ultimately the data owner has final say over how location data and the tradeoff between privacy and utility. The utility has benefits for third party analysts interesting in learning aggregate behavior.

3.2 Performance Goals

The system should scale gracefully as the number of applications connecting to the GPSD daemon increases. Location data will be released within the defined epochs.

3.3 Threat Model

Mobile devices (e.g., smartphones, tablets, wearables) are under the data owner's control. Kernel and underlying OS is vetted and verified (signatures and trusted sources). Focus is not on low level system threats. We assume that the operating system itself is not malicious and provides a mechanism to provide a privacy policy settings manager accessible to the data owner. Secure micro kernels such as seL4 address these issues and are out of scope for this paper. Applications do not have a system exploit (e.g., rootkit) to circumvent the system.

Applications may try to request data more frequently than the defined epoch. LOCATIONSAFE will deny such aggressive requests and ensure that data is only released within the defined epoch.

Applications may act as sybils and send false application IDs in order to confuse the GPSD daemon. LOCATIONSAFE will treat sybil applications accordingly using data owner defined defaults. Thus, sybil applications may either receive location data using default privacy configurations or not at all.

3.4 Privacy Goals

GPS sensor data is only accessible via GPSD. Data owners should be able to limit how frequently an application access location data. Data owners should also be able to define fine-grained access to location data. Applications for which the data owner feels the application does not meter level accuracy, the data owner should be allowed to define a radius from which the location value can be returned from. Additionally, for scenarios where fine-grained location is required, the data owner can define a grid system from which potential locations can be returned from.

4. CONCLUSION

In this paper we present to our knowledge the first software privacy module for GPSD which is a GPS daemon running on the majority of mobile embedded systems today. Data owners are able to express privacy consent and control by enforcing privacy at the lower level of the OS with minimal runtime overhead.

For future work we plan integration with Android and iOS. This will allow us to evaluate the impact and design on location based services.

5. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant Number CNS-1247357.

6. REFERENCES

- [1] C. Dwork. Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006.
- [2] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.
- [3] Facebook changes story, now says phone location not used to recommend friends. <http://fusion.net/story/319712/facebook-now-says-phone-location-not-used-to-recommend-friends/>.
- [4] A. P. Felt, S. Egelman, and D. Wagner. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In T. Yu, W. Enck, and X. Jiang, editors, *SPSM'12, Proceedings of the Workshop on Security and Privacy in Smartphones and Mobile Devices, Co-located with CCS 2012, October 19, 2012, Raleigh, NC, USA*, pages 33–44. ACM, 2012.
- [5] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: user attention, comprehension, and behavior. In L. F. Cranor, editor, *Symposium On Usable Privacy and Security, SOUPS '12, Washington, DC, USA - July 11 - 13, 2012*, page 3. ACM, 2012.
- [6] Android flashlight app tracks users via gps, ftc says hold on. <http://www.techrepublic.com/blog/it-security/why-does-an-android-flashlight-app-need-gps-permission/>.
- [7] gpsd - a gps service daemon. <http://www.catb.org/gpsd/>.
- [8] N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *ICDE*, 2007.
- [9] M. Nauman, S. Khan, and X. Zhang. Apex: extending android permission model and enforcement with user-defined runtime constraints. In D. Feng, D. A. Basin, and P. Liu, editors, *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security, ASIACCS 2010, Beijing, China, April 13-16, 2010*, pages 328–332. ACM, 2010.
- [10] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.
- [11] T. Vidas, N. Christin, and L. Cranor. Curbing Android permission creep. In *Proceedings of the Web 2.0 Security and Privacy 2011 workshop (W2SP 2011)*, Oakland, CA, May 2011.
- [12] Xprivacy. <https://github.com/M66B/XPrivacy>.
- [13] Z. Yang, M. Yang, Y. Zhang, G. Gu, P. Ning, and X. S. Wang. Appintent: analyzing sensitive data transmission in android for privacy leakage detection. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 1043–1054. ACM, 2013.
- [14] Y. Zhou, X. Zhang, X. Jiang, and V. W. Freeh. Taming information-stealing smartphone applications (on android). In J. M. McCune, B. Balacheff, A. Perrig, A. Sadeghi, M. A. Sasse, and Y. Beres, editors, *Trust and Trustworthy Computing - 4th International Conference, TRUST 2011, Pittsburgh, PA, USA, June 22-24, 2011. Proceedings*, volume 6740 of *Lecture Notes in Computer Science*, pages 93–107. Springer, 2011.