ACM HotMobile 2013 Poster: An energy efficient semantic context model for managing privacy on smartphones

Prajit Kumar Das *prajit1@umbc.edu*
Dibyajyoti Ghosh *dg9@umbc.edu*
Anupam Joshi *joshi@umbc.edu*
Tim Finin *finin@umbc.edu*

University of Maryland, Baltimore County
Baltimore MD, USA

We describe a method to carry out energy efficient privacy preservation on a mobile smartphone. Our work is based on a study of an Android smartphone's component-wise energy consumption pattern and is based on a three-fold approach to ensure efficient execution of privacy policies, based on user and app context modeled using semantic web technologies.

I. Introduction

Modern smartphones are capable of gathering massive amounts of data about a user and her context. While this data is primarily used to provide services adapted to the user, user data and context leakage from smart phones can have disastrous results. This is especially true as many enterprises are accommodating or even adopting a Bring-Your-Own-Device (BYOD) model for mobile devices such as smartphones. We recognize this change as a potential threat to data privacy, both of the user and their organizations whose data is also on the devices.

II. Context modeling

Recent advances in context modeling, tracking and collaborative localization has led to the emergence of a new class of smartphone applications that can access and share embedded sensor and context data. Unfortunately, existing security mechanisms in Android and other mobile operating systems are not geared toward protecting such dynamic data. In ongoing work, we have shown application and user context-dependent information sharing policies that dynamically control data flow among applications at a fine-grained level. Our approach differs from those in the literature [7, 8, 10] on context based privacy and security. We create semantically rich policies, and reason over them and the user and application context to either release or obfuscate the sensor/context data being shared with the application [2, 3, 9].

Our context representation includes location and surroundings, the presence of people and devices, inferred activities and the roles people fill in them. These facts are inferred by a model that is created by using a machine learning system trained on data for each user [6]. Our context is realized as a dynamic knowledge base of RDF triples grounded in an ontology expressed in the semantic web language OWL.

Policy rules are encoded in SWRL [4] and use conjunctions of facts in the context knowledge base in their conditions. The rules monitor and control application access to sensitive information and sensor data. The policies filter data flowing from sensor resources to applications to reduce disclosure by generalizing or obfuscating data. Our ontology includes the ability to represent application provenance and other metadata that can be used by the policies. The resulting system provides fine-grained, context-dependent control to sensitive user data [2]. For instance, we can have a policy that says that a location based information service should only get block level data, and not get data at all when the user is attending a confidential meeting.

III. Energy issues and a possible solution

Unfortunately, the process of gathering context and applying policies has a significant impact on energy consumption, since the system needs fresh sensor data to keep the user's context updated at all times. Current work on the energy consumption focuses on exact battery utilization of specific applications and also refers to tail energy issues [1], but has not dealt with creating an energy efficient context inference system that can be used for security. Our approach addresses this problem using three methods.

First, during each time period, we only enable the sensors required to satisfy the antecedents of relevant policy rules. And, if a set of the policy rules being enforced require the same sensor data (e.g., location) then we take one reading and use it for all of the rules.

Second, if certain information can be gathered from multiple sensors, we use the sensor with the lowest energy footprint or one that is already being used. We do this even if there is some loss of accuracy. For example, we prefer to use a location estimate from a Wi-Fi sensor over a GPS if the Wi-Fi system is on and the GPS is not.

Support for this work was provided by NSF grant 0910838, MURI award FA9550-08-1-0265 from the Air Force Office of Scientific Research.

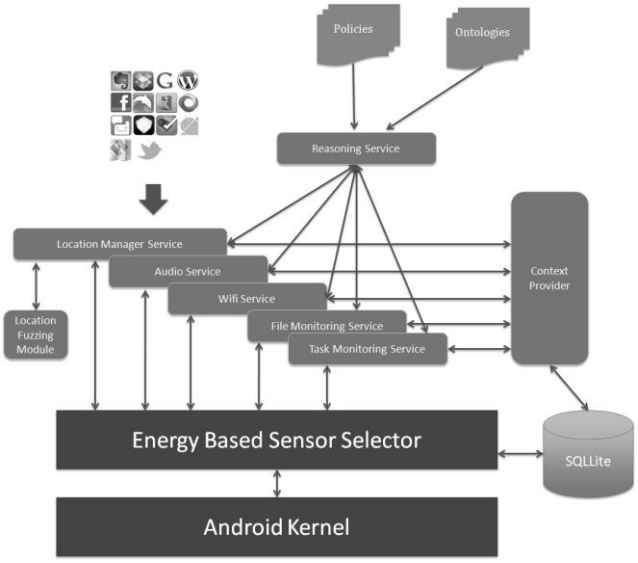

Fig-1: Our architecture for the modified Android framework enforces policies that use context information to control the information flow between applications.

Third, reorder the conditions in the rule's antecedents in order of their energy usage. We are currently gathering data on energy usage of the different sensors on a typical Android smartphone.

Our experiments were done by toggling on/off various sensors of the device and letting the battery drain. The time required to drain the battery is noted for specific sensors. These timings are then compared with the baseline operation when the phone is in "airplane mode". The time difference for the battery drain of each component is used as the parameter for deciding which sensor to use in the second method.

In course of our research, we observed that there is a significant impact on the battery of an Android phone if the frequency of location context update was high. We also observed that at the same update frequency, GPS consumed battery faster than Wi-Fi. However, justifiably the GPS unit took a significant amount of time to get an initial position fix [5] but gave more precise location updates once it was obtained. Wi-Fi, on the other hand, had less precise location updates but an initial position fix was much faster. We plan to take advantage of this trade-off of location precision versus energy measurements to optimize our energy efficiency algorithm. We have created the policy-based security mechanism in the Android framework [2, 3] and are currently adding the energy optimization modules.

IV. Conclusion and Future work

The security mechanism we have developed is extremely fine grained and suitable for smart phones. Nevertheless, it creates the overhead of high energy consumption, which is a cause for concern for the phone's battery life. In our ongoing work we are implementing an Android framework for energy efficient security (Fig-1). As a future work we plan to experimentally evaluate the proposed, energy efficient security mechanism.

References

[1] Abhinav Pathak, Y. Charlie Hu, Ming Zhang, Paramvir Bahl and Yi-Min Wang. "Fine-grained power modeling for smartphones using system call tracing." In *Proc. of the Sixth European conference on Computer systems, 2011.*

[2] Dibyajyoti Ghosh, "Context based privacy and security in smartphones." *MS thesis, UMBC, 2012.*

[3] Dibyajyoti Ghosh, Anupam Joshi, Tim Finin and Pramod Jagtap. "Privacy control in smart phones using semantically rich reasoning and context modeling." In *Proc. IEEE Workshop on Semantic Computing and Security,* pp. 82-85. IEEE, 2012.

[4] Ian Horrocks, Peter F. Patel-Schneider, Harold Boley, Said Tabet, Benjamin Grosof and Mike Dean. "SWRL: A semantic web rule language combining OWL and RuleML." *W3C submission 21, 2004.*

[5] Jie Liu, Bodhi Priyantha, Ted Hart, Heitor S. Ramos, Antonio A. F. Loureiro, and Qiang Wang. 2012. Energy efficient GPS sensing with cloud offloading. In *Proc. of the 10th ACM Conference on Embedded Network Sensor Systems* (SenSys '12). ACM, New York, NY, USA, 85-98.

[6]Laura Zavala, Radhika Dharurkar, Pramod Jagtap, Tim Finin and Anupam Joshi. "Mobile, Collaborative, Context-Aware Systems." In *Proc. AAAI Workshop on Activity Context Representation: Techniques and Languages,* AAAI Press. 2011.

[7] Mauro Conti, Vu Thien Nga Nguyen and Bruno Crispo. "CRePe: Context-related policy enforcement for Android," In *Lecture Notes in Computer Science*, M. Burmester, G. Tsudik, S. S. Magliveras, and I. Ilic, Eds., v. 6531. Springer, 2010, pp. 331–345.

[8] Norman Sadeh, "A semantic web environment for context-aware mobile services," In *Proc. Wireless World Research Forum, 2001.*

[9] Palanivel Kodeswaran, Vikrant Nandakumar, Shalini Kapoor, Pavan Kamaraju, Anupam Joshi and Sougata Mukherjea. "Securing Enterprise Data on Smartphones using Run Time Information Flow Control" In *Proc. 13th International Conference on Mobile Data Management, 2012.*

[10] William Enck, Peter Gilbert, Byung-Gon Chun, Landon Cox, Jaeyeon Jung, Patrick McDaniel and Anmol N. Sheth. "TaintDroid: an information-flow tracking system for real-time privacy monitoring on smartphones," In *Proc. 9th USENIX conference on Operating systems design and implementation.* Berkeley: USENIX Association, 2010, pp. 1–6.