Verifiably accountable surveillance

Mark D. Ryan and Jia Liu
University of Birmingham

March 13, 2014

Abstract

Individual privacy is a core human need, but society sometimes has the requirement to do targetted, proportionate investigations in order to provide security. To reconcile individual privacy and societal security, we explore whether we can have surveillance in a form that is verifiably accountable to citizens. This means that citizens get verifiable proofs of how much surveillance actually takes place.

1 Introduction

Surveillance and security. Mobile phones, wireless transport ticket systems, bank cards, email and web are systems which each moment of every day log the minutiae of our daily lives. As a result of the global surveillance disclosures in 2013, we know that the vast amount of data produced in this way is collected *en masse* by intelligence organisations the world over. This is done in several ways. The large Internet companies contribute their users' data to the NSA, via a programme called PRISM. But not all of the data can be collected so simply. The NSA and GCHQ have also covertly weakened the encryption implementations in commercial software products, for example by weakening the randomness of generated keys, in order to gain access to more data. All this is done for the purpose of identifying and combatting threats to national security. None of it, however, has anything like a democratic mandate since it was (until 2013) unknown by most people and politicians [5].

The need for privacy. This gathering of data about people's private lives has been perceived as a great threat to individual privacy. Much outrage has been expressed by academics [1, 2], politicians [3], and, somewhat hypocritically, by the very companies that enabled it to happen [4]. The maintainance of individual privacy appears to be a core human need. People have a need to keep secrets, in order to maintain purposeful relationships with others. For example, to maintain her credibility as a professional, a dentist prefers to keep secret from her patients and colleagues the details of illnesses she may have, the nature and frequency of her sexual fantasies, her financial profile, her opinions about religion, and the conversations she has with her partner. If these were disclosed to her colleagues and patients, it would change their attitude toward

her in a way she would not like, and could put her in danger. People might jump to incorrect conclusions, if the data they have is erroneous or incomplete. They may try to blackmail her, or, more mundanely, merely spam her, based on their impression of her vulnerabilities. In summary, it seems that humans want to have privacy in order to avoid:

- The consequences of incorrect conclusions that result from deliberate or accidental errors in the data, or misinterpretations, or prejudice;
- Blackmail or extortion, or other abuse of power, by people with access to data;
- Commercial and other kinds of pestering (e.g., spam)

For these reasons, the right to privacy is enshrined in the European Convention on Human Rights; its Article 8 accords the citizen the "right to respect for his private and family life, his home and his correspondence". Similar foundational legislation exists in the USA. The fourth amendment to the US Constitution provides the "right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures". In the UK, privacy from data collection and processing by companies is supported by the Data Protection Act 1998.

Reconciling the two. It seems necessary, therefore, to find ways to reconcile the requirement of societal security and that of individual privacy, and when that's impossible, to balance them appropriately. Society needs to agree a set of principles that govern when and how data about communications, finance, and internet usage should be used for preventing and detecting crime; these principles must express the sort of balance between the conflicting requirements that society wants to have. Rogaway's statement [1] distinguishes mass surveillance, which he condemns, and targeted surveillance, which he accepts. Other principles include the idea that surveillance should be proportionate.

A major challenge for society is, therefore, finding ways to reconcile and balance the requirements of individual privacy and societal security. Social scientists, policy makers and legislators need to consider the ways in which this balance should be achieved. Computer scientists need to propose ways in which the balance can be realised in practice, so that the debates by policy makers are properly informed.

2 Surveillance with verifiable accountability

It appears to be impossible to codify precisely the circumstances in which law enforcement agencies should be considered entitled to access the data stored in the logs created by usage of mobile phones, the internet, payment systems and transport systems. One could instead try to frame legislation in terms of quantities of information; the law enforcement agencies may access a specified proportion of information about an individual, an event, or during a period. But stipulating such proportions is unlikely to be satisfactory, either from the

citizen's point of view or that of the law enforcers. A third possibility is to rely on judges and other trusted parties to consider requests, one by one. This is roughly what happened up until the aftermath of the September 2001 attacks; since then, as the Snowden revelations have shown, law enforcers have obtained much broader permissions to undertake surveillance. The approach using judges to consider requests one by one doesn't scale up well; more and more information is being created, and therefore there is an ever-increasing set of opportunities for it to be used, or abused.

The approach we suggest is to supplement the procedural checks-and-balances approach with verifiable and quantitative accountability that allows citizens to understand how much surveillance is being carried out. Under the proposal, there would still be legislation and procedures for determining whether access is allowed, on a case-by-case basis; but it would be supported by quantitive information about actual access that took place, against which citizens can hold politicians accountable. We provide a means for individuals and society as a whole to obtain verifiable evidence about what the degree and nature of the surveillance that has taken place, and to vote for governments and officers that demonstrate proportionality in the way they use the data.

"Verifiable evidence" means that citizens have a means to check the veracity of the levels of surveillance that are reported. This is achieved using cryptographic protocols that produce data which can be subjected to tests by citizens. In principle, any citizen can verify the data, although it might be technically difficult and/or expensive to do so. It is sufficient if some trustworthy organisations (such as universities, charities, or journalists) do so on behalf of everyone else.

3 Example: wireless tickets

Wireless ticket systems (such as the London Oyster card or the Paris Navigo card) allow passengers to travel on city-wide transport by presenting a contactless smartcard at the time of taking a journey. The traditional paper-ticket system that preceded it allowed perfectly anonymous travel, but the wireless card has made transport ticketing into a classical privacy-invasive technology. With a wireless ticket, a passenger's journeys are logged and stored in perpetuity. To combat terrorism, and to avoid the need to obtain court orders each time, the UK intelligence agencies MI5 and MI6 have sought full automated access to Transport for London's Oyster smart card database. The data could potentially be used not just for law enforcement but potentially for advertising purposes, or even criminal stalking and harassment.

We describe a possible mechanism for making accountable the accesses to the wireless ticketing database. Assume pk is a public key belonging to a *decrypting party* (DP). The entity DP can be a TPM-attested service, or can be distributed as several independent parties where we assume that at least one of them is honest. We consider these ways to securely implement DP in more detail later; for now, we assume DP is a trusted party.

The passenger's data may be produced directly by the wireless ticket, and

transmitted encrypted to the reader installed at the transport station. The reader informs the wireless ticket about the name of the station and the current time. The wireless card appends the information about the user identity, and encrypts the whole with pk before sending the encrypted packet back to the reader. Thus, the information is held in a database, encrypted by pk.

The core idea is that DP will decrypt the data requested by the authorities, but will create a log of all the decryptions it performs. The passenger may inspect the log, perhaps after a time interval, and perhaps mediated by an access control system, to verify how much of the information about her has been decrypted. For efficiency, the log may be stored as a binary ordered search tree. A node of the tree contains data that was decrypted by DP in response to an access request. The log tree is organised as a Merkle tree. The hash of the log tree is the hash value stored at its root. This arrangement allows the operations, like insertion of new data and proof of the set of decryptions for a passenger's journeys to be efficiently.

On the assumption that DP acts correctly, the logs provide complete information about all the data accesses made for surveillance purposes. However, as mentioned, it might not be desirable to allow citizens to have immediate and direct access to the logs; for example, this could alert criminals to the fact of an investigation concerning them. We envisage there being rules, decided democratically, for what sort of access to the logs is allowed. It is unlikely that it would be appropriate for individuals to obtain information about surveillance accesses concerning themselves in real-time, since such information may alert criminals to the fact that they are being investigated. Information concerning data accesses about oneself may therefore be made available only after a certain period of time, perhaps two years. But coarser-grained information about what proportion of each day's records are accessed across a city or country might be made available immediately. In this way, citizens can hold the authorities and the government accountable for surveillance accesses.

The job of the decrypting party DP is very simple, and fully automatable. DP is not required to make any judgments about whether access should be allowed; it blindly decrypts every authorised request. However, DP is required to insert every decryption into the log, and therefore the system relies on the trustworthiness of DP. There are several ways in which this trustworthiness can be assured.

The most promising method is to distribute DP across several parties, in such a way that the system is secure provided that at least one of the parties is honest. Each party making up DP holds part of the decryption key, and performs part of the decryption. Each party must also insert information about the decryption into the log. If any party doesn't do so, it is exposed as possibly dishonest.

Another method is to use trusted computing hardware, such as the Trusted Platform Module (TPM, [6]), ARM's TrustZone [7], or Intel's SGX [8]. The idea is that the key pk under which the data is held encrypted can be verifiably bound to particular code base which encodes the behaviour of DP. Note that the trustworthy hardware must provide facilities that prevent roll-back attacks, in which DP is presented with an out-of-date version of the log.

These two methods can be combined together. If DP is a set of parties, one or more of them can choose to perform its role using trustworthy hardware.

Cryptographically more sophisticated implementations using fully-homorphic encryption [9] or functional encryption [10] may also be possible.

4 Discussion

This paper sets out the idea that surveillance could be made accountable to citizens, in such a way that people could decide through the democratic process how much and what kind of surveillance they want to allow. Moreover, the quantity and nature of the surveillance is *verifiable* by citizens; rather than merely having to believe statements about it, they obtain proof that the statements are correct.

We only scratch the surface. There is still a vast amount to do, to refine the ideas and propose mechanisms for realising them. The problem that these ideas address can only get much worse over the coming decades, as the Internet of Things generates vastly more data about our lives than before, and more and more ways emerge to use the data in privacy invasive ways.

References

[1] Phil Rogaway. A cryptographer's Statement of condemnation of U.S. mass-surveillance programs, and a reminder of our ethical responsibilities as computer scientists. `http://www.cs.ucdavis.edu/~rogaway/`.

[2] Academics Against Mass Surveillance. `http://www.academicsagainstsurveillance.net/`.

[3] BBC 5 Live. George Galloway MP says Edward Snowden deserves a medal. `https://twitter.com/bbc5live/status/345300092300976128`.

[4] Reform Government Surveillance. `http://reformgovernmentsurveillance.com/`.

[5] Cabinet was told nothing about GCHQ spying programmes, says Chris Huhne. `http://www.theguardian.com/uk-news/2013/oct/06/cabinet-gchq-surveillance-spying-huhne`.

[6] TPM main specification, Part 1: Design principles, Part 2: TPM structures, Part 3: Commands. Version 1.2:116. 2011.

[7] ARM Limited. *ARM TrustZone API Specification, Version 3.0*, 2009. ARM PRD29-USGC-000089 3.1.

[8] Innovative Instructions and Software Model for Isolated Execution. In *HASP*, 2013.

[9] Craig Gentry. 2009. A Fully Homomorphic Encryption Scheme. Ph.D. Dissertation. Stanford University, Stanford, CA, USA. Advisor(s) Dan Boneh. AAI3382729.

[10] D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In Theory of Cryptography (pp. 253-273), 2011. Springer Berlin Heidelberg.

[11] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. EUROCRYPT, volume 3027 of Lecture Notes in Computer Science, page 506-522. Springer, (2004)