When Things are Sensors for Cloud AI: Protecting Privacy Through Data Collection Transparency in the Age of Digital Assistants

Paul G. Flikkema and Bertrand Cambou
School of Informatics, Computing and Cyber Systems
Northern Arizona University
Email: {paul.flikkema, bertrand.cambou}@nau.edu

***Abstract*—The rapidly increasing number of intelligent, cloud-connected things that are embedded in our daily lives raises legitimate concerns about the privacy costs paid for the benefits these technologies provide. In this paper, we argue that this is a false choice, and motivate and describe a technology that enables citizens to effectively and conveniently monitor data collected about them. Our overarching goal is to expand awareness of the need to move from the current state of consumer ignorance or apprehensive trust to an era of data collection transparency (DCT), where consumers understand the data that is collected about them and make informed decisions based on that understanding. We show that DCT can be achieved with a suite of technologies built around recent developments in secure elements, as well as a virtual token methodology with public keys using addressable cryptographic tables.**

***Index Terms*—privacy; data collection transparency; personal assistant**

I. Introduction

> *You are recovering from a cold and still coughing. Your home digital assistants skill of recognizing coughing captures this in its database and is later harvested by your health insurance provider. Your health insurance premiums inexplicably go up, and there is no way to trace why.*
>
> *At a group conversation in your home, someone makes a sarcastic joke about a government, either domestic or foreign, that includes certain key words on a watch list. The next evening, law enforcement officials ring your doorbell.*

New technologies based on personal and in-home devices coupled with cloud-based AI can provide a dazzling array of convenient on-demand services, including entertainment, shopping, smart home control, and interactive research by voice. These services are already embodied in networked things such as Amazon Echo, google Home, and internet-connected televisions, and we can expect these assistant capabilities to move rapidly into myriad devices, including appliances and automobiles. These devices may eventually incorporate cameras, which are already omnipresent in computers and tablets. And they are rapidly being joined by wearable devices that can gather information about both general personal health and specific medical conditions. Many of these devices will also provide information leading to commands from the cloud that affect the consumer's residence (e.g., home security settings and thermostats), or that may shape the actuation of therapeutic or prosthetic devices.

Citizens are rightly concerned about surveillance of their activities and communication with others. Ostensibly, one of government's roles is to enact and enforce laws, including those that impact privacy rights. However, this is not a given in many nation-states. And domestic laws may not be enough, since cross-border entities or governments may find personal data collection very useful, especially in times of international conflict. Moreover, private enterprises have a very clear mission: to maximize market share and/or profits, and the profit-making potential of these technologies is completely dependent on the monitoring of individuals. These devices and associated cloud-based AI technologies can richly catalog almost everything that can be sensed about individual human activities, oral and written communication, behavior, and health. Because these devices may be sensing at any time without the individual's knowledge, the potential for surveillance is alarming.

Already, the great majority of consumers feel unable to grasp what information is collected about them, and by whom; in the U.S., a recent survey found that "91% of adults agree or strongly agree that consumers have lost control of how personal information is collected and used by companies" [1]. Consumers are increasingly facing a deep conflict: in order to gain the benefits of smart things, they may be forced to give up the expectation of privacy in their daily activities. The conflict often results in one of two reactions: either a "shrug" as the consumer approves a dauntingly long and opaque End User Licence Agreement (EULA), or a simple refusal to purchase the assistant and its associated services. As we outline below, this is a false trade-off: with the proper deployment of technological solutions, perhaps in tandem with rights-based regulations or laws, individuals should be able to enjoy both the benefits of these technologies and their rights to privacy. A reasonable goal is to maximize the benefits that come from companies collecting, storing, and analyzing our data, while minimizing the harms [2]. This paper describes an approach to the development and deployment of *Data Collection Transparency* (DCT), a first step to protecting privacy in the coming age of ubiquitous personal monitoring.

II. Current State of Affairs

A number of authors, government and quasi-government agencies, industry consortia, and non-governmental organiza-

978-1-5090-5279-6/17/$31.00 ©2017 IEEE

tions have weighed in on the protection of digital privacy. As a result, there is already a large body of guiding concepts and principles in both technological and social/governmental spheres.

In a comprehensive report [3], The Broadband Internet Technical Advisory Group (BITAG) recommends (among others) that software updates of things be automated and secure, use strong authentication by default, and use best cybersecurity practices. The recommendations include functional robustness to failures, i.e., that the devices should remain able to perform their primary functions and services in the event of internet connectivity losses or failures of the cloud-based components of services. However, these recommendations rely on the underlying assumption that these things and their associated services are operated by trustworthy entities. Indeed, one recommendation is that service providers should provide notification at time of purchase that the functionality of the device may be remotely decreased. While this may cause an economic loss for the consumer, the opposite problem—increased functionality—is where threats to privacy may emerge.

Relevant to privacy are the principles for national application of collection limitation, purpose specification, use limitation, and individual participation outlined in the 2013 OECD Privacy Framework [4]. However, they are stubbornly difficult to enforce—if and when they are encoded into statutory or regulatory law.

A Consumer Privacy Bill of Rights proposed in the U.S., described in [5], builds on fair information practice principles (FIPPS) and addresses how private-sector entities should handle personal data. One proposed right is focused collection, or reasonable limits on the personal data that companies collect and retain [5]. Perhaps more important is individual control, wherein consumers have a right to exercise control over what personal data companies collect from them and how they use it [5]. However, we cannot control what we cannot measure or monitor. Regulation of the exploitation of privacy-relevant information can occur at either the point (place and time) of collection or the point of use. There is an advantage to regulation at the point of use: it is conceptually far more efficient. However, it is extremely difficult to track the use of information, perhaps aggregated from heterogeneous sources over months or years and passed from entity to entity. And it is very difficult to assess harm or risk, because it might occur far into the future and due to an untraceable chain of events.

Recent recommendations for privacy protection targeting wearables [6] emphasize that privacy policies should specify what data is collected and how it is collected, stored, used, secured, and disclosed. Notably, this recommendation is strengthened for EU countries with the added conditions that categories of collected data, as well as recipients of the data, be disclosed. A related recommendation explicitly distinguishes data that is only stored locally on the device from data that is transferred to other parties.

Schneier [2] emphasizes the overall need in a call for transparency: "people should be entitled to know what data is being collected about them, what data is being archived about them, and how data about them is being used—and by whom". What is needed is the combination of oversight of what information is collected and accountability for the collection and use of that information [2]. As we will see, we also need accountability of the party trusted with that oversight.

Policing, audits, oversight panels, and fiduciary rules are just a few examples of how we recognize that the best interests of an organization may be in conflict with the best interests of the people they serve. The overarching principle of our approach is that personal assistant technologies should enable monitoring and independent, accountable oversight of what is being collected. Thus requiring that digital assistants provide a separate secure stream to a monitoring device or service is not sufficient, because it is not possible to know if the device is only sending partial information to the monitor.

Providing people with the services provided by cloud-connected things while simultaneously assuring privacy is an unsolved problem: we do not yet have a workable combination of policies and tools to protect privacy and preserve (or restore) trust in this evolving landscape [4]. But we can start by setting a goal of understanding what is collected; to paraphrase Lord Kelvin, we cannot know what we cannot measure. After a brief look at suggested solutions and current practices, this paper outlines an approach to DCT.

III. Other Solutions

One proposed method for protecting electronic privacy is certification. Because the obvious method is a one-time certification at the point of manufacture or software loading prior to shipment, certification is fundamentally problematic: systems of any complexity have multiple software updates during their lifetimes. This requires security verification during the lifetime of the device and clean shutdown at end-of-life. Because the notions of certification and verification center around what data is collected and transferred to the cloud, the approach we outline below can be seen as a method of continuous independent verification, allowing unlimited updates of the assistant device's software.

Operating systems, especially mobile OS's, increasingly have permission-based mechanisms to allow user control of applications' access to data. Work is on-going to improve usability of access control; see, e.g., [7]. However, these mechanisms rely entirely on the user's trust of the application and OS software, since there is no provision for accountable oversight. This is particularly troubling when the OS and the application are from the same provider.

IV. Current Approaches

Personal assistants that transmit sensed information to the cloud provide authentication and protection using the well-known TLS [8]. A Diffie-Helman exchange or public-key asymmetric cryptography is used only for the initialization of the communication because of its high computational complexity of decryption. In the initialization or handshake

protocol, pairs of symmetric keys are shared for the use in the session to lessen the computational burden.

TLS prevents the consumer from monitoring the information being collected, since the symmetric key cannot be revealed to anyone. Indeed, the process of revealing this key to the consumer could lead to security breaches if it is intercepted or falls into wrong hands by any number of means. Thus, while assuring security and privacy against passive third-party attacks, it cannot provide data collection transparency.

In our approach, a trusted device and software, enabled by a new approach to key management and distribution, is dedicated to the mission of monitoring the data sent from a person or family's devices to the cloud.

V. Towards Data Collection Transparency

Figure 1 depicts the forces that can drive DCT and its key technological components. We describe here an approach to data collection transparency. To the usual pair of a client C (the personal assistant) and the server S, we add a monitor M that is designed to enable DCT. The critical elements of the proposed system are (1) a key distribution mechanism; (2) secure/verifiable hardware; (3) open-source software; and (4) an (optional) cloud service for DCT.

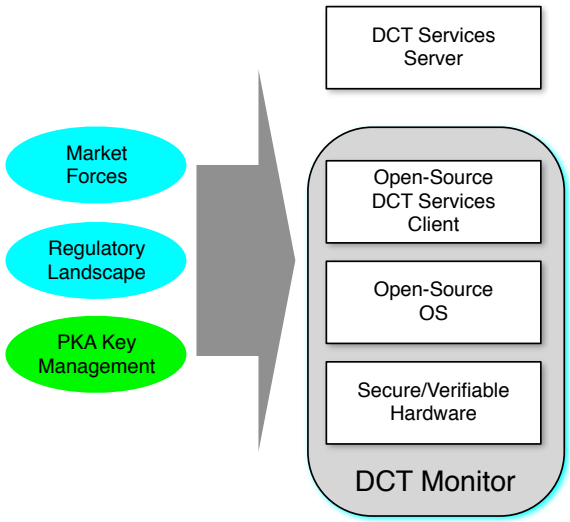

Fig. 1. Multiple forces will drive development of personal assistant monitoring technology, which will build on progress in key management algorithms and hardware, secure/verifiable hardware, open-source software, and cloud-based services.

A. PKA Technology

Public Keys that are Addressable (PKA) is a scheme that complements or replaces other encryption methods such as RSA to simultaneously deliver DCT in addition to the secure, private communication between personal assistants and the cloud that conventional approaches provide. One component of PKA is a pair of identical hardware secure elements: one for the personal assistant C and one for the monitor M; see Figure 2. The provider of the personal assistant and its cloud-based services generates a virtual token, a cryptographic table, and distributes it simultaneously to the secure elements of C and M. This distribution must be performed once, during an initial setup, also called personalization, via a highly secure means. For example, the cryptographic tables can be downloaded to secure USB tokens using non-volatile memories that are mailed to the purchasing consumer. They can be delivered in a separate package via a shipping method that is independent of the method used for shipping the device.

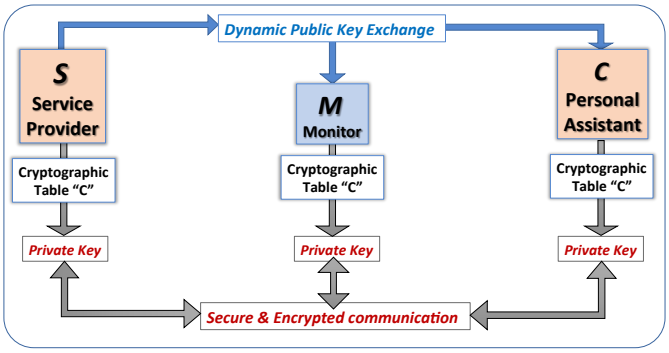

Fig. 2. PKA-based security protocol.

In their simplest form, the cryptographic tables can contain a truly random collection of bit values [9] generated by the service provider. As an example, each table can have 256 rows and 256 columns for a total memory of 8 Kb. The cryptographic table can also be generated by an array of physically unclonable functions (PUF) that is part of one of the two secure elements. This array generates digital finger prints of the 8 Kb cells, called PUF challenges, which replace the random collection of bits. The value of a PUF-based cryptographic table is its relative strength against side channel analysis.

The dynamic public key of the PKA encryption method is randomly generated by the server of the service provider, and is used to generate a private key from the cryptographic table (Figure 2). After sharing the public key through an insecure channel, S, C, and M can independently generate the private key, a fixed-length data stream of typically 128 to 256 bits. Symmetrical encryption schemes such as AES can then be used with this private key for C to encrypt messages and both M and S to decrypt them, and perform authentication cycles. Crucially, unlike M, a fourth party having access to the public key cannot decrypt information exchanged between C and S without the cryptographic table. PKA can thereby establish a tri-party secure communication channel between S, C, and M.

To increase the level of security, hash functions, such as the standard hash algorithm (SHA), can be used to protect the public keys with an additional password or pin code. Such an architecture hides the public key from third parties. The scheme can also be improved by adding a level of randomness to the private key. A second random number, added to the public key, can modify the data stream generated by the cryptographic table to generate a more complex private key.

In addition to offering a tri-party encrypted communication protocol between S, C, and M, the PKA can strengthen access control and authentication. During an authentication cycle, the generation of a new public key by the server can be followed by the encryption of authentication patterns by C and M using the corresponding private key. Candidates for these authentication patterns include passwords, pin codes, biometric prints, and PUF responses of the secure elements. For additional protection, S can also send an encrypted authentication pattern to C and M. The scheme can be considered as a multi-factor authentication method protected by the PKA and independent authentication patterns.

As a result, the monitor M can securely collect all data transmitted to and from the personal assistant C, enabling DCT as part of the monitor M and in concert with cloud-based services, as described below.

B. Secure/Verifiable Hardware

While a number of open-source cores exist, backdoors can be inserted, frustrating assurance that a consumer's DCT monitor is trustworthy. Potential backdoors exist at multiple levels, e.g., in proprietary FPGA synthesis toolchains and at fabrication [10].

Potential solutions exist. Multi-chip designs would assemble the needed hardware from multiple sources, partitioning the functionality so that the monitoring could not be performed without all chips; simpler chips are also easier to verify. For example, one chip might be dedicated to packet analysis that allows forwarding only to addresses on a hardware-based whitelist. Independent power measurement techniques in tandem with benchmarking via power use modeling open-source software might enable detection of hardware trojans activated long after device activation. However, further work is needed as more challenging threats emerge [11].

C. Open-Source Software and Cloud-Based DCT Services

The final component of a DCT monitor is its resident software, supported by an operating system that provides networking and web services. This open-source software can interact with cloud-based DCT services to monitor collected data, interpret how it can be used, and issue appropriate alerts. For example, a software agent running on the monitor platform can advise a consumer on how to negotiate the trade of data for valuable services. As the economic value of personal data security increases (approaching the value of physical security of the home), we expect that provision of these services will lead to new business opportunities.

VI. Initial Results

Two PKA prototypes were developed to demonstrate the PKA-based security protocol (see Section V-A): a server-to-terminal prototype, and a server-to-secure Java card prototype. In both cases, cryptographic tables of 8 Kb were generated from the server and downloaded to the client devices. Public keys were generated with random number generators, and private keys were generated concurrently by the server and the various client devices. The server-to-terminal prototype used AES 256 encryption to securely transmit files between the server and the terminals. The server to Java card prototype used AES 128 encryption to securely authenticate several Java cards. We benchmarked this approach against different versions of RSA and Elliptic Curve Cryptography (ECC) for the computational efficiency of key generation; the results are shown in Table I.

TABLE I
Computational efficiency of key generation
(smaller is better)

Technique	Number of Bits	CPU Clock Cycles
RSA	1024	23,443
	2048	59,845
	3072	417,166
	4096	1,492,687
ECC $GF(P)$	128	313
	256	1,204
	521	4,696
ECC $GF(2^N)$	113	981
	233	4,393
	409	12,294
PKA	256	100

VII. Conclusion

We cannot understand what we cannot monitor. To achieve the ultimate goal of privacy while using cloud-connected devices, the first step is capturing the data leaving and entering our homes and devices. We have presented a technological roadmap on which work can be started almost immediately, possibly as part of of an Industry Cybersecurity Program as suggested in [3].

References

[1] L. Rainie. (2016, September) The state of privacy in post-Snowden America. [Online]. Available: http://www.pewresearch.org/fact-tank/2016/09/21/the-state-of-privacy-in-america/

[2] B. Schneier, *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World*. W. W. Norton & Company, 2016.

[3] Broadband Internet Technical Advisory Group, "Internet of things (IoT) security and privacy recommendations," Tech. Rep., November 2016.

[4] Organization for Economic Co-operation and Development, "The OECD privacy framework," 2013.

[5] White House, United States of America, "Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy," February 2012.

[6] Future of Privacy Forum, "Best practices for consumer wearables & wellness apps & devices," August 2016.

[7] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. Zhang, N. Sadeh, A. Acquisti, and Y. Agarwal, "Follow my recommendations: A personalized privacy assistant for mobile app permissions," in *2016 Symposium on Usable Privacy and Security (SOUPS) 2016*. USENIX Association, June 2016.

[8] T. Dierks, "The transport layer security (TLS) protocol version 1.2," IETF, RFC 5246, August 2008.

[9] B. Cambou, "A XOR data compiler combined with physical unclonable function for true random number generation," in *SAI/IEEE Computing Conference*, July 2017, (accepted).

[10] M. Rostami, F. Koushanfar, and R. Karri, "A primer on hardware security: Models, methods, and metrics," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1283–1295, Aug 2014.

[11] K. Yang, M. Hicks, Q. Dong, T. Austin, and D. Sylvester, "A2: Analog malicious hardware," in *2016 IEEE Symposium on Security and Privacy (SP)*, May 2016, pp. 18–37.