Personal Privacy Assistants for the Internet of Things

Anupam Das, Martin Degeling, Daniel Smullen and Norman Sadeh
School of Computer Science
Carnegie Mellon University

***Abstract*—As we interact with an increasingly diverse set of sensing technologies, it becomes more and more difficult to keep up with the different ways data about one's self is collected and used. Study after study has shown that while people care about their privacy, they feel powerless to control what data is collected about them and how it is used. This article summarizes ongoing research to develop and field privacy assistants designed to empower people to regain control over their privacy. Privacy assistants use machine learning to build and refine models of their users' privacy expectations and preferences, selectively inform them about the data practices they care about, and help them configure privacy settings that are available to them. This technology was first demonstrated in the form of assistants that help their users configure their mobile app permission settings, and is now being extended to the Internet of Things (IoT). Herein, we focus on the new infrastructure we have developed and fielded to support IoT privacy assistants. The infrastructure enables the assistants to discover IoT resources (e.g. sensors, apps, services and devices) in the vicinity of their users, and selectively inform users about resources' data practices they would want to know about. The infrastructure also supports the discovery and selection of configurable settings for IoT resources (e.g. opt in, opt out, erase my data), enabling privacy assistants to help users tailor their IoT experience in accordance with their privacy expectations.**

I. Introduction

Information privacy is about giving people meaningful choices when it comes to the collection and use of their data ane about giving them sufficient details about these choices to make informed decisions. In practice, even when browsing the web from desktop or laptop computers, few people find the time to read privacy policies, or exercise choice options available to them. It has been estimated that if users were to read the policies of every website they interact with over the course of a year, they would spend around 244 hours reading them [1]. Research by the authors, as well as others, further shows that users only care to be informed about a small fraction of the statements found in privacy policies [2], [3], [4]. Over the past decade, the challenge of informing users about relevant data collection and use practices has been further exacerbated by the proliferation of smartphones. Reading privacy policies and exercising choices are further hampered by the small form factor of these devices, as well as the added distractions associated with many mobile usage scenarios. But at least smartphone users are generally aware of the majority of mobile apps installed on their devices. Current mobile operating systems also have centralized functionality that provides users with some control over the permissions requested by their mobile apps. With the Internet of Things (IoT), this is no longer the case. Today, users interact with an ever-growing and increasingly diverse collection of IoT technologies, many of which they are unaware of and have no ability to control (e.g. cameras coupled with face recognition and scene recognition functionality, or WiFi location tracking systems that sense the unique device ID's of passerbys). It is no surprise that a November 2014 Pew Internet survey reported 91 percent of adults "agree" or "strongly agree" that consumers have lost control over how personal information is collected and used by companies [5].

What is needed is a new, scalable paradigm that empowers users to regain appropriate control over their data. As part of their work in this area, the authors have been working on the development and evaluation of Personal Privacy Assistants (or PPAs). PPAs are intended to learn models of the preferences and expectations of their users, to selectively inform them about data collection and use practices they would most likely want to be notified about. PPAs also help users configure available privacy settings, wherever possible. An early version of this technology was demonstrated in the form of Mobile PPAs that help their users configure permissions required by the mobile apps on their Android smartphones [6]. These PPAs have been successfully piloted by actual Android users on their personal devices as part of their regular activities [6] In this article, the authors discuss how they are working on similar functionality for the IoT. In particular, we focus on the development and deployment of an IoT privacy infrastructure that enables Personalized Privacy Assistants to discover nearby resources (IoT-connected sensors, systems, services, etc.) collecting information about their users. We discuss the IoT Resource Registries we have developed as part of this infrastructure, as well as functionality designed to

help IoT resource owners to populate entries in these registries with minimal effort. IoT Resource Registries (IRRs) advertise the data collection and use practices of registered resources, enabling PPAs to selectively inform their users about those practices and choice options they are likely to care about. The article further discusses deployment and management options associated with this infrastructure. Specifically, we report on the deployment of this infrastructure on two university campuses in the United States.

II. Related Work

The technical feasibility of PPAs was first evaluated in ubiquitous computing settings during the early 2000s. Langheinrich [7] used beacons and service discovery protocols to advertise the privacy practices of data collection services. In combination with privacy proxies and privacy-preserving databases, this complex infrastructure was intended to tightly control the flows of personal information. Similarly, Sadeh et al. used semantic web technologies to capture and enforce rich collections of privacy preferences in mobile and IoT contexts in their MyCampus project [8]. Sadeh and colleagues also reported on early work to learn people's privacy preferences to automatically or semi-automatically configure privacy settings such as location sharing settings [9].

Individual privacy preferences and expectations have been identified as factors that influence whether one will approve of sharing their personal information. Other factors include transmission principles and social norms [10]. Multiple studies have been conducted to identify individual factors, which include not only what data is shared, but more importantly with whom it is shared [11], [12], [13]. Other factors include the purpose of data sharing, how long the data will be accessible, and how it will be processed. Still, the availability of this information about important factors does not solve a fundamental problem; the amount of privacy decisions that need to be made increases with the diversity of new sensors, services, and apps that collect data. Therefore, a new paradigm in privacy research looks at how machine learning can be used to simplify privacy decision making through recommendations. Liu et al. [6] have shown that recommendations based on clusters of like-minded users and predictive models of people's privacy preferences work to the users satisfaction in the context of mobile app privacy.

In a recent crowd-sourced vignette study [4], we asked participants to assess their comfort and interest in receiving notifications with respect to different hypothetical IoT-related scenarios. These scenarios described up to eight different factors about what data is collected, where, for what purpose, and the data retention period. We found some abstract norms about privacy and physical location still hold true — differentiation between private (at home) and public (in a library, or department store) contexts lead to very distinct privacy decisions for a large majority. However, we also found that in other contexts, such as data being collected in the workplace for purposes like saving energy, individual preferences and values have a higher impact on whether or not someone would want to permit the data collection.

In addition to modeling privacy preferences, we are also examining how to use the same technology to limit and contextualize the notifications a user receives. In this paper, we focus on the development of IoT privacy infrastructure that enables privacy assistants to take advantage of prior findings.

III. Overall Architecture

In the smartphone world, users control the apps they install on their devices and have access to unified privacy management functionality, where they can review and control the permissions granted to apps. The situation in IoT is quite different. Here, users interact with technologies they often did not deploy and are seldom even aware of. This lack of awareness, as well as a dearth of settings available for users who do not own or manage these IoT resources, makes 'Notice and Choice' a significantly more difficult proposition. IoT users generally do not know what devices are around them, what data they collect, and what happens to that data. To remedy this situation, we need an infrastructure that supports the discovery of nearby IoT resources and their data practices. By "nearby" IoT resources, we mean IoT resources that collect data in our physical vicinity. IoT resources may include IoT devices (e.g. smart home assistants, smart doorbells), IoT services (e.g. indoor location tracking systems, video analytics services connected to smart cameras) or IoT apps (e.g. smart TV remote apps) that collect and use data about us. Along with the discovery of these resources, the infrastructure also has to support the discovery of information about the data these resources collect, and how this data is used. Equally important are settings that these resources may expose to

users, such as opt-out settings, opt-in settings, and more. Below, we introduce such an architecture, which we have implemented and deployed on parts of two university campuses in the United States.

We highlight the three main components of our IoT privacy infrastructure: a) Internet Resource Registries (IRR); b) Personal Privacy Assistants for IoT (PPA); and c) Policy Enforcement Points (PEP). We first describe the functionality of each of these components. We then illustrate how these components interact with each other, to notify users of the existence of nearby sensors and privacy settings, and support the configuration of these settings.

Internet of Things Resource Registry (IRR): We have developed IoT privacy infrastructure that is intended to be open and scalable. Any number of actors may be involved in the deployment of IoT resources, and this is captured by our design. Resource owners and deployers include corporations deploying smart locks, HVAC systems, room presence systems, audio/video equipment, scheduling systems, and location tracking. Resources may be deployed in office buildings. Cities may deploy public resources such as traffic monitoring services, computer vision based crime reporting systems, and public health monitoring systems. Malls, stores, and restaurants may deploy IoT systems for security purposes, as well as marketing. Camera- and Bluetooth-based systems can monitor, track, and profile customer behavior. Today, in many homes we see smart door locks, surveillance cameras, thermostats, and voice-enabled home assistants. Scenarios incorporating the Internet of Things involve the deployment of an increasingly diverse array of connected, smart devices designed to capture sensitive data. Thus, there is a need for infrastructure that can, at the very least, inform users what, how, and why they are being sensed by smart devices nearby.

IRRs allow IoT resource owners to publish and disseminate descriptions of their IoT resources. These descriptions include the data practices of these resources. An IoT resource can be an app, a service, a single sensor, a virtual sensor aggregating multiple sensors, as well as any infrastructure element that might collect and/or use user data. The IRR acts as a location-aware lookup service that supports the discovery of nearby IoT resources. Device owners and IRR administrators access IRRs through a secure web-based portal. The portal guides them through a process where they can enter or modify descriptions of IoT resources, including the data they collect, how the data is used, for how long it is retained, and more. The data is stored in a machine-readable format, capable of capturing a rich set of data collection and use practices. Typical resource entries include information about the party that collects data, the purpose of the data collection, retention period, granularity of data collection, and third-party data sharing (if any). Resource owners can also advertise control options that enable users to restrict how their data is used, such as the ability to opt in, opt out, erase data, restrict the retention period, define who the data can be shared with, restrict how it can be used, define whether it needs to be anonymized or aggregated, and more. These settings, where made available, are paired with specifications of APIs and control endpoints that users can access through privacy assistants to configure them. Fig. 1 shows a screenshot of the different policy-related information captured through the IRR user interface. In particular, the top of the screen shows how the resource registration "wizard" guides the user through a succession of steps (or workflow) to define the data practices associated with an IoT resource. For the sake of accommodating a wide range of users and regulatory requirements, the wizard makes minimal assumptions about the particular fields a user needs to fill to specify a valid resource. Most of the available fields are optional. Many of the fields come with predefined options, designed to expose commonly accepted taxonomies used to characterize details of many data practices. For example, predefined options for data retention range from "ephemeral," to "limited," to a specific time period, all the way to "unspecified". This interface is designed to broadly facilitate the registration of resources in IRRs, but is primarily targeted towards professional users, such as system administrators, building managers, and the like. For casual or home users interested in deploying and advertising the presence of commercial off-the-shelf IoT resources in their personal spaces, our infrastructure supports the creation and consumption of vendor-generated resource templates that predefine the specifics of commercial products. Using these templates, vendors predefine the practices and capabilities of their products, reducing the burden on end users. When using templates, end users need only to enter deployment specifics, such as the place in their home where the resource is located. At the time of writing, we have created templates for a dozen popular IoT resources, including Amazon Echo (with Alexa),

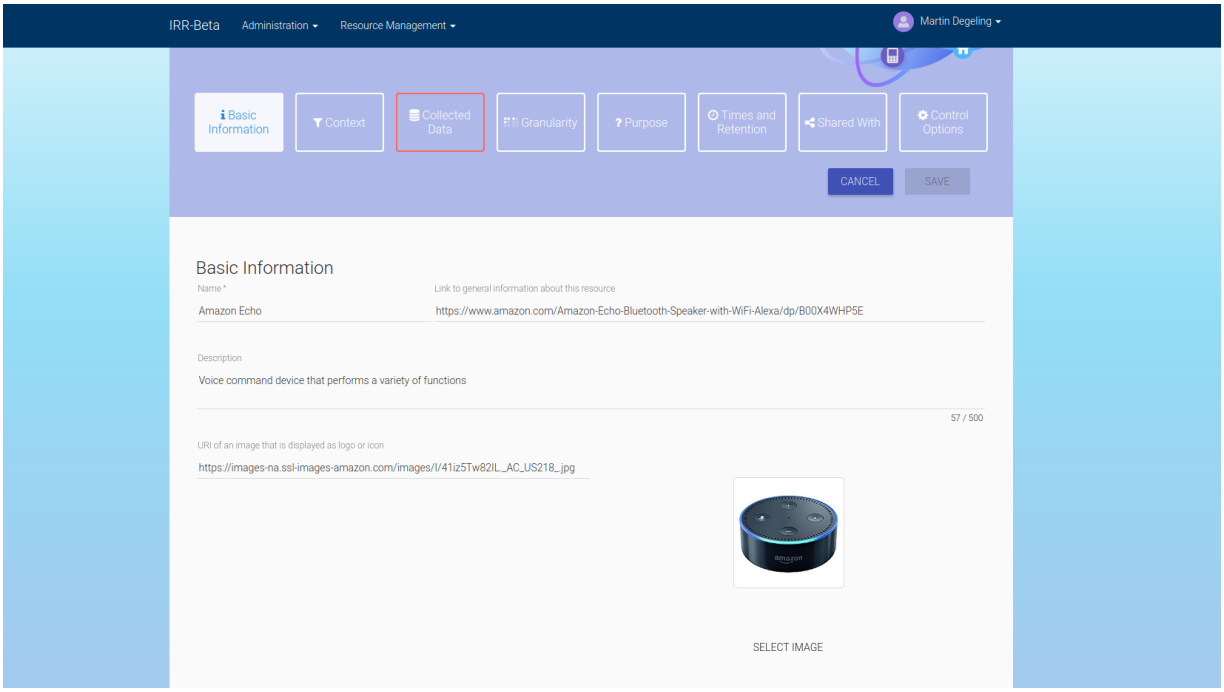

Fig. 1: Screenshot of the IRR portal. It provides a wizard for specifying privacy practices. It allows resource owners to specify information about the IoT resource, including where it is located, what data is collected, in what granularity, and for what purpose. Users can be informed about who the data is shared with. Where available, it also lists individual privacy settings that may be configurable.

Google Smart Home, and Nest Cameras. Our hope is that over time, vendors will develop product templates of their own. This vision is further discussed below.

Our infrastructure is designed to support any number of IRRs concurrently. Different IRRs will be managed according to different policies, by different groups. Some IRRs will be designed to advertise the presence of resources in corporate buildings, and may be very tightly managed. Others, in malls or cities, may be more permissive. Others, managed by communities, may have minimal management and vetting of resource registration. We envision different users will use Personal Privacy Assistants to filter out different IRRs and resources, perhaps based on the entities managing them or other criteria such as their area of coverage. IRRs may have overlapping coverage areas, though some IRRs may be viewed as more authoritative over an area than others. For example, an official IRR for a university campus might be considered more authoritative than an IRR run by a hobby or student organization. Conceivably, some IRRs may charge users for advertising their resources as a way of generating revenue and a possible approach to reducing spam.

Once an IRR is set up, the availability of the IRR can be locally advertised with Bluetooth beacons, discovered through centralized directories of registries covering different geographic areas, or preconfigured by privacy assistants. Users can be directed to a local IRR based on their present location, just as resources can be discovered in IRRs the same way.

As discussed earlier, the IRR infrastructure itself can be managed on different levels. The central directories of IRRs can be curated by different parties to determine which IRRs become publicly available. This process is comparable to ICANN and authorities regulating domain names on the web. This design also allows multiple levels of directories, if the proliferation of IRRs for a given area merits further reorganization.

Individual IRRs may have one more administrators responsible for vetting new resource registrations. Others may have resource owners who publish their IoT resources in their own private IRRs. IRRs may also be completely unmoderated, allowing anyone to advertise IoT resources to potential users. Depending on the

nature and governance of an IRR, whether it is strictly controlled versus whether it is open, some form of administration is required at the individual IRR level to determine which resources get published and prevent abuse.

PPA for IoT: The Personal Privacy Assistant (PPA) for IoT is an application running on a users smartphone that is aimed to assist users to discover IoT resources and services available in their vicinity. It retrieves resource listings from IRRs relevant to a user's current location, and uses their content to produce a privacy short notice. In other words, an organization and a succinct privacy policy, notifying users about the collection of their personal data by IoT resources. The PPA lists resources registered (see Fig. 2) in the IRR and informs users about each resource's functionality, ownership, what data it collects, and what it does with that data. It informs users if data is shared, how long data is retained, if data is aggregated, and so on. The PPA also distinguishes connected services from physical resources, giving an overview of the different apps or websites offering functionality based on data collected by IoT-connected resources.

We envision the PPA will deliver improved functionality by learning users' preferences over time. We plan to model both notification and privacy setting preferences. With the right notification model, the PPA will be able to selectively decide when, how often, and what to show to the user about nearby IoT resources. Modeling privacy preferences will enable the PPA to detect mismatches between the user's privacy expectations and the policies of the resources they engage with. We envision this to be similar to approaches shown to be successful with websites [2]. Previous research in our group has shown that machine learning and clustering techniques can be leveraged to simplify the way profiles are learned and preference mismatches are detected [14], [4]. We believe that in the near future, the availability of privacy settings will become more prevalent in different scenarios, in part because of emerging regulations (such as GDPR) and other new requirements to obtain opt-in consent from users. PPAs could then also be used to semi-automatically configure privacy settings on the user's behalf, where such settings are made available by registered resources.

Policy Enforcement Point (PEP): New regulations such as GDPR, COPPA, GLBA, and CalOPPA (at least under some interpretations) require IoT resource owners or data collectors to expose different privacy settings to their users. In such contexts, there is a need for a PEP which is responsible for both storing users' preferred privacy settings and enforcing those settings accordingly. For example, in the context of deploying cameras equipped with facial recognition, one possible user-configurable privacy setting would allow individual users to opt out of facial recognition during specific times of the day or at a specific location.

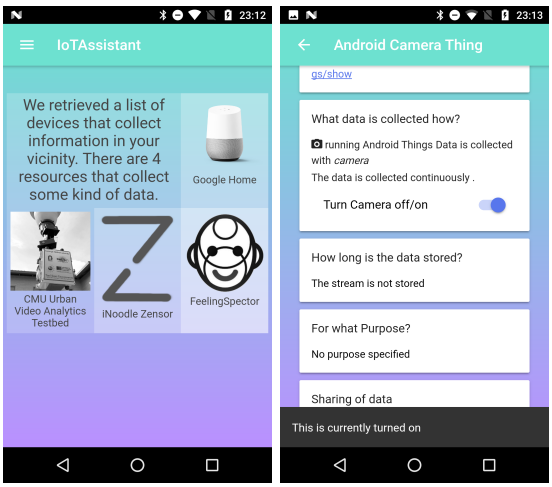

(a) IoT Resources **(b)** Information and options

Fig. 2: PPA for IoT. It lists the resources available in the user's vicinity (left). Details about a data collection and available options for the user (right).

The PPA allows users to configure privacy settings where supported by resources. Changes made through users' PPAs are sent to a PEP for that resource which ensures that their privacy settings are enforced there. The set of privacy choices available depends on the availability of resource-specific services. Our PEP design offers simple RESTful APIs to enforce privacy settings. The URL and availability of these RESTful APIs can be entered, configured, and updated in the IRR by resource owners.

The interaction among the different infrastructural components is shown in Fig. 3. As shown in the figure, IRR resource owners first register their IoT resources with a given IRR (the IRR directory, in this example, lists public IRRs). Access to the portal and administrator privileges are controlled through an authentication system. An IRR resource owner can use predefined templates

to describe their IoT resources. Once IoT resources are registered with an IRR, users can rely on their PPA to discover the resources in their vicinity. PPAs can also help users configure any available privacy settings by brokering access to APIs that interface with the PEP enforcing settings for a resource. All of these parameters are advertised in the IRR entry for that resource. For example, the PPA can expose a facial recognition opt-out API, advertised in the IRR entry of a smart camera system. Perhaps this resource is in a mall, and used for marketing. When a user in the mall opts out, the smart camera resource's PEP ensures that each user's privacy settings are properly applied to the data streams coming out of the camera system, preventing their face from being recognized.

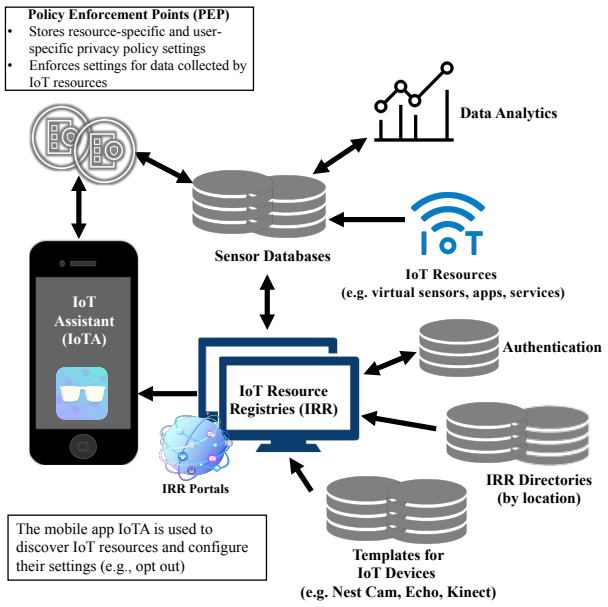

Fig. 3: Interaction among components of our proposed system. The privacy assistant discovers IoT resources through IoT Resource Registries (IRR), and preferences are enforced through Policy Enforcement Points (PEP).

IV. Deployment Scenarios

We envision our infrastructure to be ubiquitous and easily deployable to a variety of public and private settings. In this section, we describe the deployment process that different IoT resource owners would typically use to register their resources with an IRR.

Business Setting: Suppose Jessica, a small business owner, runs a high-fashion clothing shop. In order to track her customers for marketing purposes, she has installed a smart camera system in her store. The system uses facial recognition and behavioral tracking to determine what items her customers linger around, indicating their potential interest. The system contains database of known faces and associated contact information. When the system recognizes a customer interested in a particular item, it sends them a promotional email. Jessica's shop is located in a mall which already has an IRR covering the entire location, so she registers her smart cameras in the mall's IRR. She opens a browser on a laptop, logs into the mall's IRR website, and registers a new IoT resource. She expedites the process by using a template provided by the smart camera system vendor. She then adds details specific to her deployment by updating information about the location of the cameras, and their purpose. Because this camera system is reconfigurable, she exposes privacy settings to her customers. There is a facial recognition opt-out setting, and an opt-in setting that enables users to register their face and contact information to receive promotional emails.

Home Setting: Suppose Alice has purchased a 'Smart Fever Monitor', a new form of IoT system featuring an infrared camera that sends an alert to its owner when it detects unusually high body temperatures. The device comes with a setup guide in which the vendor recommends using an IRR to inform visitors about the infrared video data collection. Alice opens the PPA on her phone and it informs her that she is not the owner of an IRR overseeing her present location (her home). The app forwards her to an IRR directory portal to register a new IRR there. Alice fills out a form requesting a new IRR to oversee her address. After specifying the necessary details signifying her ownership, she encounters a verification step where she is asked to verify her authority and ownership of the location she specified. Verification can be done in different ways. For example, she could be asked to submit a copy of a utility bill for the address, or she could be asked to enter her credit card information for the same billing address. Once the process is completed, Alice can register IoT resources on her IRR. She can rapidly complete the registration process by scanning a QR code provided by the device vendor, which contains a reference to their prefilled template for her new device.

Corporate Setting: Jim is an IT administrator for an enterprise that employs several hundred employees. He

is situated in a shared office building housing other companies on other floors. Jim is responsible for overseeing security for this office branch. The enterprise decided to upgrade its security by installing new security devices around the office. Jim deploys facial recognition cameras, magnetic door locks with smart card and biometric authentication, and alarm buttons with two-way audio recorder intercoms that connect with security guards. Company policy mandates informing employees about the presence of devices that may collect personal information. The company uses strictly curated IRRs. Jim opens a browser and logs into the company's website that lists their deployed IRRs. He requests a new IRR, overseeing the floors where the firm is situated. Jim fills out a form which collects his credentials and corporate email address, verifying his authority over this company-owned space. He receives an email which notifies him that the IRR has been deployed, and links him to its configuration portal. Jim opens the link and enters the details for the new IoT resources that have been deployed around the office.

V. Campus Deployments

So far, we have developed three mobile applications that make use of IoT resources. Two are available on Carnegie Mellon University campus (friend finder and automated class attendance), and one on the University of California Irvine campus (indoor navigator). Both campuses are equipped with indoor location tracking services, using WiFi access points and Bluetooth beacons. WiFi access points offer a coarse grained location (e.g., imprecise location, distinguished by building, wing, or hallway). Fine grained location is based on Bluetooth beacons. Depending on the number and density of beacons that are deployed in a given area, Bluetooth beacons can be used for location detection precise enough to distinguish individual rooms. Pre-registered users of the location service can be located via WiFi access points using mobile phones. Bluetooth tracking requires a location service on a smartphone to scan for nearby Bluetooth beacons. In our deployment, the IoT Assistant notifies users about the availability of apps which use these location services. For example, the location sharing app enables users to share their location with friends, providing settings for controlling location granularity. Additional apps may make use of the location tracking service and share the same infrastructure. To simplify user interaction with the tracking system, allowing them to configure location granularity or arbitrarily disable location tracking at any given time, the PPA exposes simple control options. When users configure these options, their settings are automatically sent to a policy enforcement server that was previously specified as part of the location service's resource registration in the IRR – the privacy policy associated with the resource on the IRR specifies what and how users may configure the resource.

A second application we have implemented uses facial recognition technology to automatically detect and record attendance for university lectures (described in detail in [15]). Participants register their face with the application using their phone. Once registered, as they walk past a camera when entering the lecture room, their attendance is recorded. Lecturers and students may use these records to keep track of who attended the class. Similar to applications that use the location tracking service, users can use the PPA to change their privacy settings for the attendance tracking system. This allows users to opt-in or out of the tracking, during the course of the semester. The application uses the same policy enforcement server as the location tracking service, which controls the facial detection processing service that the attendance tracking relies on. Each of these services may be part of shared infrastructure used to support other applications where facial recognition is required.

VI. Challenges and Discussion

In this articule, we have introduced a novel infrastructure for Privacy Assistants for the Internet of Things. While early deployments of this infrastructure suggests that it can work well, a number of challenges remain. The single most significant challenge is without a doubt to get a critical mass of technology providers (e.g. device manufacturers, app developers, virtual sensor providers) to agree on a common standard to describe the data collection and use practices of their technologies and to adopt protocols such as the ones we have developed to support the advertising and discovery of IoT Resource Registries and IoT resoures.

In addition, determining the right questions and factors to model users' privacy and notification preferences is still an open problem, especially broadening to a large number of IoT scenarios. While researchers have been successful at modeling such preferences in the context of smartphones [6], IoT presents a new set of challenges. Understanding the relevant context is very important to

make the right decision – automatically turning off Alexa when kids visit our house requires systems to determine that kids are present. Moreover, the ever changing IoT landscape also poses a challenge for modeling user's preferences. The underlying factors of the user model may change frequently, if the right level of abstraction is not used. It is therefore necessary to find the right balance between static models and the incorporation of dynamic, context-sensitive factors.

REFERENCES

[1] A. M. McDonald and L. F. Cranor, "The Cost of Reading Privacy Policies," *I/S: A Journal of Law and Policy for the Information Society*, vol. 4, pp. 543–568, 2008.

[2] A. Rao, F. Schaub, N. Sadeh, A. Acquisti, and R. Kang, "Expecting the Unexpected: Understanding Mismatched Privacy Expectations Online," 2016, pp. 77–96. [Online]. Available: https://www.usenix.org/conference/soups2016/technical-sessions/presentation/rao

[3] J. Gluck, F. Schaub, A. Friedman, H. Habib, N. Sadeh, L. F. Cranor, and Y. Agarwal, "How Short Is Too Short? Implications of Length and Framing on the Effectiveness of Privacy Notices," 2016, pp. 321–340. [Online]. Available: https://www.usenix.org/conference/soups2016/technical-sessions/presentation/gluck

[4] P. E. Naeini, S. Bhagavatula, H. Habib, M. Degeling, L. Bauer, L. Cranor, and N. Sadeh, "Privacy Expectations and Preferences in an IoT World," in *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. Santa Clara, CA: Usenix Association, 2017. [Online]. Available: https://www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini

[5] M. Madden, "Public Perceptions of Privacy and Security in the Post-Snowden Era," Pew Research Center, Tech. Rep., 2014. [Online]. Available: http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/

[6] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, "Follow my recommendations: A personalized privacy assistant for mobile app permissions," in *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 2016, pp. 27–41.

[7] M. Langheinrich, "A privacy awareness system for ubiquitous computing environments," in *Proceedings of the 4th International Conference on Ubiquitous Computing*, ser. UbiComp '02, 2002, pp. 237–245.

[8] N. M. Sadeh, E. Chan, and L. Van, "MyCampus: an agent-based environment for context-aware mobile services," *Proc. UBIAGENTS*, pp. 34–39, 2002.

[9] N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao, "Understanding and capturing people's privacy policies in a mobile social networking application," *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 401–412, Aug 2009.

[10] H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.

[11] S. Lederer, J. Mankoff, and A. K. Dey, "Who wants to know what when? privacy preference determinants in ubiquitous computing," in *CHI '03 Extended Abstracts on Human Factors in Computing Systems*, ser. CHI EA '03, 2003, pp. 724–725.

[12] M. Benisch, P. G. Kelley, N. Sadeh, and L. F. Cranor, "Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs," *Personal and Ubiquitous Computing*, vol. 15, no. 7, pp. 679–694, Oct. 2011. [Online]. Available: http://link.springer.com/article/10.1007/s00779-010-0346-0

[13] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, "Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy Through Crowdsourcing," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, ser. UbiComp '12. New York, NY, USA: ACM, 2012, pp. 501–510.

[14] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, "Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions," in *Proceedings of the 12th Symposium on Usable Privacy and Security (SOUPS)*, 2016, pp. 27–41.

[15] A. Das, M. Degeling, J. Wang, X. Wang, M. Satyanarayanan, and N. Sadeh, "A Privacy-aware Infrastructure for using Facial Recognition," in *Workshop The Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security (CV-COPS 2017)*, 2017, http://www.cs.cmu.edu/~anupamd/paper/CV-COPS-2017.pdf.