

A Generic Information and Consent Framework for the IoT

Mathieu Cunche

Univ Lyon, INSA Lyon, Inria, CITI
F-69621 Villeurbanne, France
mathieu.cunche@insa-lyon.fr

Daniel Le Métayer

Univ Lyon, Inria, INSA Lyon, CITI
F-69621 Villeurbanne, France
daniel.le-metayer@inria.fr

Victor Morel

Univ Lyon, Inria, INSA Lyon, CITI
F-69621 Villeurbanne, France
victor.morel@inria.fr

Abstract—The Internet of Things (IoT) raises specific issues in terms of privacy, in particular with respect to information and consent. In this paper, we propose a generic framework for information and consent in the IoT which is protective both for data subjects and for data controllers. We present high level requirements for information and consent and several technical solutions to implement them. We also outline the design space and sketch a prototype implementation.

Index Terms—privacy, IoT, Ubiquitous computing, information, consent, GDPR, regulation

I. INTRODUCTION

The development of the Internet of Things (IoT) raises specific privacy issues especially with respect to information and consent. People are generally unaware of the devices collecting data about them and do not know the organizations operating them. Solutions such as stickers or wall signs are not effective information means in most situations. As far as consent is concerned, individuals do not have simple means to express and communicate it to the entities collecting data. Furthermore, the devices used to collect data in IoT environments have scarce resources; some of them do not have any user interface, are battery-operated or operate passively (they collect data without emitting any signal).

In Europe, the General Data Protection Regulation (GDPR) [9] puts emphasis on the control of data subjects over their personal data. Its application to the IoT is not obvious though. The Working Party 29¹ (WP29) has published guidelines on transparency [16] and consent [15] and an opinion on the development of the IoT [14]. As far as transparency is concerned, the GDPR defines the categories of information to be provided to data subjects: identity of the controller, purpose of the processing, categories of personal data concerned, recipients, etc. It also introduces some requirements on acceptable communication modes including the need to have easily accessible information and to avoid “information fatigue”. In this regard, the WP29 recommends in particular the use of “push” and “pull” notices. The GDPR also defines a number of conditions for the validity of consent: it should be freely given, specific, informed and unambiguous. For example, Recital 42 states that “Consent should not be

regarded as freely given if the data subject has no genuine or free choice or is unable to refuse or withdraw consent without detriment”. In the context of the IoT, this should entail that consent to Wi-Fi based physical tracking [4] is not valid if the only alternative for data subjects is to turn off their Wi-Fi and thereby be deprived of useful services. To ensure the lack of ambiguity, consent should, according to the GDPR, “be given by a clear affirmative act”, which should exclude the collection of identifiers such as MAC addresses for example, without any affirmative action from the user.

As far as the IoT is concerned, the WP29 advocates the design of new consent mechanisms, such as “privacy proxies” [14], on the devices themselves. We agree that this is a fruitful research direction for the effective implementation of consent but this solution, which is discussed further in this paper, also raises challenges in the context of IoT where data subjects generally do not have any prior contact with data controllers.

Starting from these recommendations, we define in Section II general requirements that have to be met to ensure that information and consent are managed in a manner that is satisfactory both for data subjects and for data controllers. In Section III, we introduce several techniques to implement these requirements in different situations, in particular through declaration registers and beacons. Depending on the context and the types of devices involved, not all technical options are always possible. In order to provide guidance to IoT system designers, we outline the main choice factors and the design space in Section IV. These techniques are illustrated with several challenging case studies in Section V. In Section VI, we sketch a prototype implementation of these techniques. We discuss related work in Section VII and conclude with perspectives in Section VIII.

II. GENERAL REQUIREMENTS FOR INFORMATION AND CONSENT

We derive from the recommendations sketched in the introduction a set of technical requirements for implementations of information and consent which are protective both for data subjects (persons from whom data is collected, following the GDPR terminology, hereinafter “DS”) and for data controllers (entities collecting data, following the GDPR terminology, hereinafter “DC”). We assume that DC deploy devices that can

¹One of the roles of the Working Party 29, which is now replaced by the European Data Protection Board, was to make recommendations on data protection and privacy in the European Community.

collect different types of personal data and/or communicate information to DS. For their part, DS may own several devices and at least one of them (typically a smartphone) can be used to consult the information provided by DC and to express their consent. We call this device the Gateway Device. We use the expression “DS privacy policy” to refer to the choices of the data subject regarding his personal data and “DC privacy policy” to refer to the privacy policy declared² by a DC. In order to meet the challenges identified in the introduction, a consent management framework should provide the following facilities for, respectively, information and consent:

Information:

- The declaration by DC of their devices, with all the necessary information, including their position, range, the type of data collected and the associated DC privacy policy.
- The receipt of this information by the device of any DS about whom personal data can be collected (*i.e.* within the range of a DC device).
- The presentation of this information to the DS in forms and at times that should minimize information fatigue and maximize the likelihood that the DS will not miss any useful information.

Consent:

- Means for DS to express their consent in forms and at times that should minimize their fatigue and maximize the likelihood that they make appropriate decisions regarding the protection of their personal data.
- The receipt of this consent by any DC able to collect personal data and the guarantee that they will not collect the data (or will immediately delete it) if this consent is not consistent with their DC policy.
- The possibility for DC to store the consents obtained from DS so they can demonstrate GDPR compliance regarding consent, in particular that it has been provided by the data subject on whom data is held.

The interactions between a DC and a DS can be split into two parts: the interactions of the DS with his Gateway Device (to be informed and to express his consent) and the communications between the DS Gateway Device and the DC devices.

In the rest of this section, we define more precise requirements on the communications between DC devices and DS devices. We first define the operations considered here, which can be triggered either by a DC or by a DS (or their devices):

- $install(\delta, \lambda, \rho, \theta, \pi)$ is the deployment of a collecting device δ at position λ with range ρ , collecting data of type θ with DC privacy policy π . The position and the range define the physical space in which the device can collect data of type θ . The type can be for example MAC address,

sound, or image. We assume without loss of generality that a device is associated with only one type.³

- $declare(\delta, \lambda, \rho, \theta, \pi)$ is the declaration of a collecting device δ at position λ with range ρ , collecting data of type θ with DC privacy policy π .
- $collect(\delta, \sigma, \theta, \pi, \mu)$ is the collection by device δ of value μ of type θ from the DS device σ ; the value is associated with the DS privacy policy π .
- $move(\sigma, \lambda)$ means that the DS device σ moves to position λ .
- $define(\sigma, \theta, \pi, \mu)$ means that the DS privacy policy and the value of data of type θ on device σ are set to π and μ respectively.
- $pair(\sigma_1, \sigma_2)$ is the pairing of the DS device σ_1 to the DS device σ_2 . Pairing is useful in this context to make it possible to define the privacy policy of a device σ_1 with scarce resources on another device σ_2 of the DS (his Gateway Device).

Some of the above operations (*install, move*) are physical while *collect* describes the actual collection of the data. The key operations in terms of information and consent management are *declare* and *define*. Different implementations of these operations are described in the next section. Our goal at this stage is to provide a high-level description of the framework. For example, what is called a “device” here can be any source of personal data, such as a smart phone, a quantified-self device or even the subject himself for data of type voice or image. The identifier δ of a device can be implemented in many different ways (MAC address, plate number, etc.) as discussed in the next section and illustrated in Section V. Some of these devices cannot store their DS privacy policies. Typically, the mobile phone of a DS can play the role of Gateway Device and therefore be used to define and store all his privacy policies (for all his devices).

We do not discuss in this paper the actual content of a privacy policy π . The definition of privacy policies and their semantics are studied in a companion paper [12]. In the present paper, we only assume that DC policies contain all the information required by the GDPR and can be compared (to check whether a DC policy complies with a DS policy).

We have defined the semantics of the above actions, in terms of preconditions and postconditions in the style of Hoare logic. The benefit of this formal semantics is twofold: it ensures that there is no ambiguity about the meaning of the actions and it makes it possible to prove general properties such as the facts that (1) no collection of data about a DS can take place if the DS has not previously received the required information from the DC and (2) a data is never collected from a DS and stored by a DC if the DS privacy policy for this data is not satisfied by the DC privacy policy. The interested reader can find the details about this semantics in a long version of this paper published as a research report [1].

²A declaration can be seen as a commitment of the DC to implement his DC privacy policy but the actual enforcement of this policy is outside the scope of this paper.

³Multi-type devices can be considered as several devices at the same location.

III. TECHNICAL OPTIONS

The requirements introduced in the previous section are very high-level and can be implemented in different ways. In this section, we present technical options to implement them depending on the context and the capacities of the DC and DS devices. The overall design space is presented in Section IV and the techniques are illustrated through several case studies in Section V. We first describe the implementation of the communications between DC devices and DS devices, which corresponds to the *declare* and *collect* operations. Section III-A considers direct communications between devices while Section III-B focuses on indirect communications (through registries). For each solution, we discuss its compliance with the above requirements and its feasibility in terms of cost and deployment effort. Then, we suggest ways to allow DS to interact with their devices through a privacy assistant called “Personal Data Custodian” in Section III-C. These interactions concern in particular the implementation of the *define* operation (expression of consent). They also contribute to the information of the DS (operation *declare*).

A. Direct communications

A first option to implement information and consent is through direct communications between DC devices and the DS Gateway Device. In this option (hereinafter “direct communication”), DC devices use a direct communication channel to advertise their presence and communicate all the parameters of the *declare* operation (position, range, type of data collected and DC privacy policy) within their area of operation. The same communication channel can be used by the DS to transmit his potential consent to the DC.

Direct communication can typically be implemented using medium and short range wireless communication technologies such as Bluetooth or Wi-Fi which are now common place and are embedded in many devices (*e.g.* smartphones). In addition, their range (typically several meters to tenths of meters) matches the scale of the area of operation of IoT systems and their protocol can be leveraged to carry the information required for declaration and consent. In this section, we focus on the BLE technology, but other wireless technologies could be used in a similar way.

Bluetooth Low Energy (BLE) features a discovery mechanism that allows the detection and identification of devices as well as the transmission of small amounts of data. This mechanism can be used to implement direct communications between the DC and the DS Gateway Device. In the following we refer to the element implementing this mechanism on the DC side as the *BLE Privacy beacon*.

In BLE, device discovery is implemented using *Advertisement Packets* [13, Part C, sec. 11] that are broadcast at regular intervals and can be received by any BLE device in range. Those packets can be configured to carry data necessary for the declaration of DC devices (parameters of the *declare* operation). A DS Gateway Device in the range of the DC *BLE Privacy beacon* will thus be able to passively retrieve the declaration data by collecting the advertising packets.

Another feature offered by BLE is the Attribute Protocol (ATT) [13, Part A, sec. 6.4] that allows the exposure of services and the transmission of small amounts of information through a lightweight connection. This feature can be leveraged to implement the communication of consent: the DS Gateway Device connects to the *BLE Privacy Beacon* (this is a lightweight process) and sends the consent data (parameters of the *define* operation) using the ATT protocol.

Direct communications have several benefits: first, they do not require Internet connectivity. Also, the locality of the communications reduces the risk of further tracking by a remote entity. From the point of view of the DS, the information part is collected passively by collecting the data transmitted by the *BLE Privacy Beacon*; this means that, in order to be informed, the DS does not expose his presence. They also raise several challenges. First, all devices should be able to declare themselves. Tracking systems involving passive devices thus need to be enhanced (for example with a *Privacy Beacon*) to enable these declarations. Also, the communication protocol should support the communication of the parameters of the *declare* and *define* operations. In addition, the coverage of the declaration mechanism should match the area in which the data collection is taking place. Finally, all the above features should be possible at reasonable cost and without disrupting existing services.

B. Indirect communications

Another option to implement information and consent is to use a registry (hereinafter “indirect communication”). Registries can be used both by DC (implementation of the *declare* operation) and by DS (implementation of the *define* operation). A DC registry is a database freely accessible through the Internet, storing all relevant information about DC devices, including the parameters of the *declare* operation. The DC registry declaring a DC device δ must be accessible to any DS device σ (or the paired Gateway Device) before it enters the range of this DC device (*i.e.* when $Within(\lambda_\sigma, \lambda_\delta, \rho_\delta)$) if λ_σ is the location of σ and λ_δ and ρ_δ are respectively the location and range of δ , for example via a Web site or through an application. They must provide the required information in machine-readable format, for instance a structured format such as JSON provided through an API. They should also include a human-readable version that can be consulted directly by DS.

Indirect communications through DC registries have several advantages compared to direct communications: (1) they enable the visualization of DC policies regardless of the location of DS, which means that DS can be informed about the collection of data before visiting an area and (2) they provide a flexible management approach for DC policies – they do not require a specific infrastructure or particular capabilities of the devices except for an Internet connection.⁴

However, their implementation raises several challenges: (1) DS devices must always be aware of all surrounding devices; therefore, registries should be easily accessible; (2) registries

⁴Therefore, they can be well-suited to passive devices such as cameras.

must be properly managed, up-to-date and accurate in order to meet the requirements defined in the previous section. Managing a registry can be achieved in different ways: it can be centralized or distributed, and contributions can be restricted to authenticated parties.

In the same spirit, DS registries can be used by DC to retrieve the consents provided by DS. DC should be able to prove that the retrieved consents have effectively been provided by the right DS. This proof of identity could be implemented by using an authentication mechanism (*e.g.* with tokens). A drawback of DS registries is the fact that they may represent a weakness in terms of privacy since DS have to disclose their privacy policies. These registries should be secure to ensure that only the concerned DC can get access to them.

C. Personal Data Custodian

The interactions described in the previous sections involve a variety of devices. However, the ultimate recipients of DC policies and original sources of DS policies are the DS themselves. In order to describe the interactions between DS and their devices, we assume that a software tool, called the Personal Data Custodian (hereinafter “PDC”), is installed on DS Gateway Devices. The roles of the PDC are the following:

- Interact with the DS to allow him to consult the information received from DC devices (pursuant to the *declare* operation).
- Interact with the DS to allow him to express his privacy choices (implementation of the *define* operation).
- Interact with DC devices to communicate personal data with their DS policies (implementation of the *collect* operation) or to reject collection requests from DC devices when the associated DC privacy policy does not comply with the DS privacy policy. Instead of sending the DS policy, the PDC can also send an explicit consent message to confirm the acceptance of the policy previously sent by the DC (*e.g.* by signing a hash of this policy, as discussed in Section VI).

As discussed in the introduction, consent is valid only if it is freely given, specific, informed and unambiguous. Each of these conditions brings forward strong requirements on the PDC and the language used to express privacy policies:

- *Consent must be freely given:* any personal data and privacy policy communicated by the PDC should reflect the genuine choices of the DS.
- *Consent must be specific:* the privacy policy language must be rich enough to allow DS to express granular choices, for example about types of data, data controllers or authorized purposes.
- *Consent must be informed:* the PDC must not disclose personal data to a DC device that has not communicated its privacy policy.
- *Consent must be unambiguous:* in order to avoid any ambiguity, the privacy policy language should be endowed with a formal semantics and the interface used to interact with the DS should not give rise to any misunderstanding.

A privacy policy language meeting these requirements is described in a companion paper [12]. In the present paper, we focus on the PDC itself and its interactions with the DS. The main challenge is to find the appropriate level of automation and type of interaction to meet the GDPR requirements while avoiding information fatigue. If the level of automation is low and interactions too frequent, consents may apparently meet all the GDPR requirements but in fact result from routine, mechanical, acceptance from DS, as already observed on the internet. If the level of automation is high, the reason may be that privacy policy rules are defined in a very coarse way (for example, “always accept the disclosure of my *MAC address*”) which would not meet the GDPR requirements. Careful PDC design choices can help resolve this tension. For example, DS should be able to express positive rules (conditions in which they agree to communicate a type of personal data) and negative rules (conditions in which they refuse to communicate a type of personal data). The PDC would then interact with the DS only in situations for which it has not received any instruction (for example when an unknown category of DC device requests personal data).

IV. DESIGN SPACE

As illustrated by the previous section, a variety of technical solutions can be implemented to make information and consent more effective in IoT environments. However, depending on the precise context, in particular the features of the devices, not all options are always possible. In this section, we provide some guidance to designers with an outline of the main parameters to be considered and their impact on the available options.

Table I and Table II show, for each feature in the first column, the technical options that are possible or not for the implementation of information and consent respectively. In Table I, the first column refers to the DC device (collecting device) whereas in Table II it refers to the DS device (source of the personal data). For the sake of readability, we show only negative answers in the tables and take the convention that empty boxes are interpreted as the fact that the feature does not prevent the technical option. In order to decide if a technical option is possible in a given context, the designer must check that none of the features of this context corresponds to a “**X**” in the column representing this option. Occurrences of “(**X**)” denote situations in which the feature does not prevent the technical option but the technical option is likely to be either unnecessary (for example, in Table I, the use of an additional beacon is probably not necessary for sensors endowed with an extensible protocol) or insufficient (for example, in Table I, indirect communications will probably not be sufficient for moving sensors, such as cameras mounted on vehicles, unless the registers can be updated in real time).

A passive sensor in Table I is a sensor, such as a camera, able to collect data but not to communicate a privacy policy. An extensible protocol is defined as a protocol, such as Wi-Fi or Bluetooth, which can be configured to communicate a privacy policy. The beacons considered in Table I are beacons

that can be added to a device to allow it to communicate a privacy policy.

In Table II, a device with substantial resources is assumed to be a device that can be used to define and manage privacy policies without significant drawbacks. The required resources include memory, computing power, communication means and a user-friendly interface (for example, the screen of a smartwatch is not appropriate for the definition of a privacy policy). The difference between a systematic collection process and a selective collection process is the fact that in the first case it is not possible to prevent the collection of certain data while this is possible in the latter. For example, video recording is a systematic collection process. The collection process can be systematic for certain data (such as MAC addresses for Wi-Fi access points) and selective for other data (for example, payload data in a Wi-Fi protocol). When the collection process is systematic, it cannot filter out data for which consent has not been granted: the only solution in this case is to implement consent a posteriori, by deleting or anonymizing the data as soon as it is collected. Pre-existing relationships corresponds to the situation where the DS and the DC know each other through any type of identifier that can be used to declare their respective privacy policies.

V. CASE STUDIES

In order to illustrate the versatility of our framework and the technical options introduced in Section III, we present their application to a challenging Bluetooth-based tracking case study in Section V-A. We discuss more briefly another application of vehicle tracking in Section V-B. We emphasize that all the solutions proposed here apply without any registration phase or assumption about prior contacts between DS and DC.

A. Bluetooth-based tracking

Systems that monitor individuals based on the Bluetooth MAC address of their device are becoming commonplace and are deployed in venues such as shopping malls [10] and music festival events [8]. Those systems passively collect the MAC addresses [4] found in messages broadcast by portable and wearable devices such as smartphones, smartwatches, wristbands, etc.

Any person entering the operation area of a such system should be informed and should be able to provide his consent. This is particularly challenging in open venues like shopping malls where there is generally no existing link between visitors and the entity operating the system; as a result, visitors are currently informed of those tracking systems via posters and consent requirement is simply ignored [4]. This use case is challenging but it can be addressed in our framework either in the direct mode or in the indirect mode. In both cases, DS can be informed that Bluetooth data collection is taking place and can in turn send their consent, including their radio identifiers, through their PDC (if their own privacy policy allows for this collection).

Let us consider an area (shopping mall, museum, music festival ...) in which a Bluetooth tracking system is operating,

i.e. all active Bluetooth devices in this area can have their MAC address recorded and associated with other data (time, location, etc.). Let us further assume that the DC in charge of the tracking system has no prior link with the DS. This means that, before the DS enters the area, the DC and the DS have not been able to communicate and that any exchange of information must be done on the spot. For what concerns the information of the DS, written signs could be displayed at the entrances of the area but they might remain unnoticed and would not convey the level of information required in this context. In this challenging situation, the technical solutions presented in Section III can be used to implement mechanisms for seamless information and consent.

We consider, as an illustration, a DS equipped with a Bluetooth wristband visiting a location where Bluetooth tracking is deployed. This is a situation in which the DC devices (Bluetooth sensors) are *fixed* and *passive* and they implement a *systematic collection process*. Furthermore, there is *no pre-existing relationship* between the DS and the DC. The design guidelines of Section IV show that enforcement must be done a posteriori and the applicable options include: *Direct communications with beacon* and *Indirect communications*. The direct communication option can be further refined based on the features of the DS device (support of extensible protocol and device resources). As of today, wearable devices such as wristband are supporting an *extensible protocol* (Bluetooth – see section III-A), but they have *reduced user interface* and *limited energy resources*. These additional constraints imply that the direct communications must be done with pairing.

The first option to address this case is to use direct communications via *BLE Privacy Beacons* as presented in Section III-A. Such beacons can be deployed in the monitored area and around it in order to inform DS as soon as they enter the area. Furthermore, if the DS has configured his PDC with the relevant identifiers (in this case the Bluetooth MAC addresses of his devices), the PDC will also be able to use direct communications with the beacon in order to send to the DC the potential consent of the DS (if the privacy policy of the DC complies with his own policy). The scenario is the following (see Figure 1): when entering the area, the PDC of the DS will automatically detect the *BLE Privacy Beacon* and retrieve the privacy policy of the tracking system thanks to the direct communication mechanism. If this DC privacy policy complies with the DS privacy policy, the PDC of the DS automatically sends the consent through the BLE direct communication channel. This consent contains the MAC address, which is the identifier of the DS in the context of this data collection. Once this consent is received and securely stored, the tracking system is allowed to collect data on the subject identified by this address. By default, the system discards any data for which it has not obtained consent.

In the case where there is a *pre-existing relationship* between the DS and the DC, another option to deal with Bluetooth tracking systems would be to rely on indirect communications as described in section III-B. In this case, the DS is informed of the presence of the system through a DC

TABLE I
TECHNICAL OPTIONS FOR INFORMATION AS A FUNCTION OF THE DC DEVICE

Features of DC device	Direct communications without beacon	Direct communications with beacon	Indirect communications
Passive sensor	(X)		
Active sensor with extensible protocol		(X)	
Active sensor without extensible protocol	(X)		
Fixed sensor			
Moving sensor			(X)

TABLE II
TECHNICAL OPTIONS FOR CONSENT AS A FUNCTION OF THE DS DEVICE

Features of the DS device	Direct communications without pairing	Direct communications with pairing	Indirect communications	A priori enforcement	A posteriori enforcement
Device with extensible protocol		(X)	(X)		
Device without extensible protocol	(X)				
Device with substantial resources		(X)	(X)		
Device with scarce resources	(X)				
Systematic collection process				(X)	
Selective collection process					(X)
Pre-existing relationship					
No pre-existing relationship			(X)		

registry: the DS Gateway Device regularly sends requests with its current localization to a dedicated server in order to get the list of nearby data collection systems and the associated DC privacy policies. If the DC privacy policy complies with the DS privacy policy, the PDC sends the DS consent automatically to the dedicated online DS registry.

B. Vehicle Tracking

Vehicles may be subject to passive data collection by systems that detect and record their presence. One of the main technologies allowing this data collection is Automatic Number Plate Recognition (ANPR) [5] based on images captured by CCTV cameras. Given the nature of the location where those systems are deployed (road sections, parking lots), DS and DC have usually no pre-existing relationship. Furthermore, the driving activity requires full attention of the driver, which cannot be disturbed by tasks related to information and consent. This scenario describes *fixed* and *passive* sensors, with *no pre-existing relationship* and a *systematic collection process*. The situation is therefore similar to Section V-A except that the identification of DS relies on number plates rather than MAC addresses. Before entering a monitored area, the Gateway Device of the driver (*e.g.* his smartphone) communicates with the *BLE Privacy Beacons* to retrieve DC privacy policies and to return a potential consent including his car plate number. Furthermore, the Personal Data Custodian (see Section III-C) running on the DS Gateway Device implements the consent decision without requiring the driver interaction, preserving its attention for the driving task.

VI. PROTOTYPE IMPLEMENTATION AND EVALUATION

In this section, we briefly describe our prototype implementation of the direct communication solution introduced in

Section III-A and evaluate it against the objectives set forth in Section II.

A. Implementation

We use the case study “Bluetooth-based tracking” described in Section V-A as an illustration of the prototype. In this scenario, DC devices are *fixed* and *passive*, and collect data *systematically*. We also consider that there is *no pre-existing relationship* between the DS and the DC. We use a BLE Privacy Beacon combined with a mobile application running on an Android phone implementing the PDC presented in Section III-C. The BLE Privacy Beacon is based on a low cost (less than \$6) hardware (*Espressif ESP32*⁵) that implements the information and consent mechanisms (the code of the BLE Privacy Beacon and PDC are available online⁶). We have implemented a prototype tracking system monitoring Bluetooth signals and storing MAC addresses and timestamps. The tracking system is augmented by Privacy Beacons and a consent management mechanism that discards data for which consent has not been obtained. The DS device is a Garmin forerunner 235 smartwatch. A sketch of this implementation is pictured in Figure 1.

The mobile application acts as a PDC and enables the definition of DS privacy policies in a user-friendly manner. DS can add, update and delete rules through a scroll-down menu. The PDC can also instantiate generic consents as described in Section III-C. The Privacy Beacons broadcast their 86 bytes long DC privacy policy, taking advantage of the *Advertising* features of the BLE protocol. Upon reception, the DC policy is compared with the current DS policy, and the PDC issues a consent message in case of compliance. The consent is

⁵<https://www.espressif.com/en/products/hardware/esp32/overview>

⁶Respectively at https://github.com/cunchem/BLE_Privacy_Beacon.git and <https://gitlab.inria.fr/vmorel/coiot>

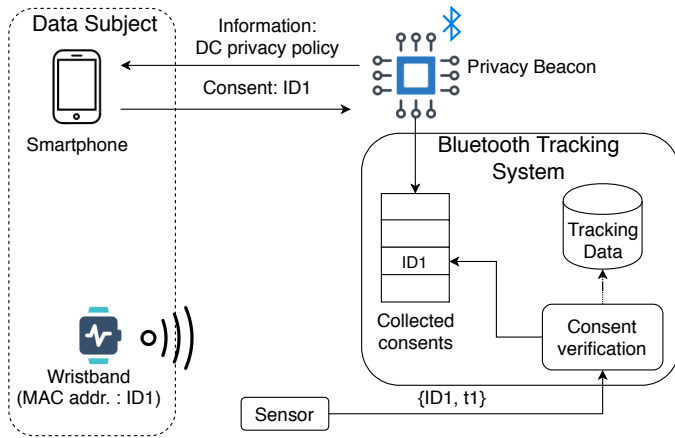


Fig. 1. Illustration of the Bluetooth tracking scenario. The DS is informed of the tracking system via direct communication between his Gateway Device (smartphone) and the DC Privacy Beacons. Potential consent is then sent by the DS Gateway Device to the Privacy Beacons, to be stored by the tracking system. Communication between the smartphone and the Privacy Beacons happens over BLE. When in range of sensors, the Bluetooth identifier of the wristband is passively collected and only stored if consent has been retrieved.

sent through the Attribute Protocol. The consent message comprises the identifiers of DS devices — including the MAC address of the smartwatch — as well as a hash of the DC privacy policy to which the DS consents. Once the consent has been retrieved by the DC Privacy Beacon, it is stored by the tracking system. For each collected data item, the system checks whether a consent has been collected for this device identifier (in our case, a Bluetooth MAC address). Data is stored only if it is the case.

B. Evaluation

Even if the prototype presented here does not intend to be production-ready, it is fully functional and it meets the essential requirements set forth in Section II. In the following, we assess the strengths and limitations of the prototype with respect to the six facilities identified in Section II:

Information:

- The declaration by DC of their devices is performed by the Privacy Beacons. In the current version of the prototype, the declaration does not include the position and range of the device but the extension is straightforward and the next version will include it.
- The range of the collecting devices can be tuned to fit with the range of Privacy Beacons to ensure that any DS about whom personal data can be collected receives the declaration sent by the DC. In practice, DC privacy policies are retrieved between one and five seconds after the PDC enters the area.
- This information is presented to the DS on the PDC. The presentation highlights information of interest to the DS, such as the type of data collected and the retention time. In the next version of the prototype, the presentation will also include a link to the full DC privacy policy.

Consent:

- The PDC makes it possible for DS to define, modify and delete privacy policy rules. These rules express the conditions under which DS consent to the collection of their data. In the future, consents will be enhanced with secure authentication to ensure their integrity and authenticity.
- DC receive the consents through Privacy Beacons. When consent is not received, the personal data is immediately deleted.
- DC can store the consents on a central server through Privacy Beacons. In the next version of the prototype, they will be stored on a secure ledger to ensure their integrity.

As far as costs are concerned, the prototype demonstrates that the framework can be instantiated in real life use cases with a low-cost implementation.

VII. RELATED WORK

The Privacy Assistant project led by the Carnegie-Mellon University (CMU) is an example of use of registries to declare and retrieve privacy policies of IoT devices [2]. A prototype has been deployed on the CMU campus, where DS are able to locate cameras. Combined with an assistant on a mobile phone, subjects are warned about personal data collection in their vicinity. Our contributions with respect to this work are threefold: first, we present in Section II a generic framework that can be instantiated in different ways as shown in Section III. The registries proposed in [2] represent one of the possible technical options. Also, we emphasize the control of DS over the disclosure of their personal data through their Gateway Devices (rather than through external privacy enforcement points). Last but not least, our first motivation is the implementation of the GDPR requirements in the context of IoT. As argued in Section II, the design choices of our framework are driven by this objective, especially the need to ensure that the criteria for valid consent are met.

Another example of registry dedicated to the declaration of privacy choices is the Smart Places⁷ service proposed by the Future of Privacy Forum (FPF). With this service a data subject can provide his Wi-Fi or Bluetooth MAC address in order to opt-out from tracking by participating companies.

In [11] the authors introduced the concept of *IoT Resource Registries (IRRs)* that broadcast data collection policies and sharing practices of the local IoT systems. Our direct communication mode has some similarities with this proposal. However, in contrast with [11], we provide technical solutions for its implementation and demonstrate an actual prototype. Furthermore, our approach is not limited to information since it also includes consent.

Previous work has also been done on privacy assistants [2], including the PawS architecture [7] and IoTA [2], [3]. In contrast with previous work in this area, our framework does not make any assumption about policy enforcement points and prior contact (*e.g.* through registration) between DS and DC:

⁷<https://smart-places.org/>

beacons announce themselves for direct communications, and registries are automatically retrieved by the PDC for indirect communications. Our framework enables local communications: DS policies do not have to be disclosed. In addition, our framework and privacy policies rely on a formal semantics which makes it possible to reason about privacy policies and provide further guidance to DS (for example about the risks related to their policies) [12].

In a nutshell, the main contributions of this paper with respect to previous work are the following:

- Our framework is generic, including both direct and indirect communication modes, for both information and consent.
- It is able to deal with the situation, which is common in the IoT, where DS do not have any prior contact with DC.
- It relies on a formal semantics which makes it possible to avoid ambiguities and to provide formal guarantees.
- It has been devised to meet the requirements for information and consent in the GDPR and to enhance local control of the DS over the collection of their data.

VIII. CONCLUSION

Beyond GDPR compliance, we believe that the adoption of the measures suggested in this paper would contribute to reduce the imbalance of powers between DC and DS without introducing prohibitive costs or unacceptable constraints for DC. The effectiveness of these solutions also depends on organizational and regulatory measures. For example, DC deploying or using IoT devices must have the legal obligation to declare their devices (with the required information) using electronic means.⁸ These solutions also require a standardization effort (e.g. about the declaration protocol and the privacy policy language).

On the technical side, further work is required to improve the user-friendliness of the interface of our PDC and also to make it easier for DC to declare their devices. The fact that our framework and privacy policy languages are endowed with formal semantics also paves the way for richer user interfaces. For example, we suggest in a companion paper [12] the verification of properties based on different risk assumptions. This facility could be useful to enhance DS awareness and to allow them to make better informed decisions.

The proposals made in this paper are also very relevant to the ongoing discussions about the future ePrivacy Regulation [6]. As stated by the WP29 [17], the current draft “gives the impression that organisations may collect information emitted by terminal equipment to track the physical movements of individuals (such as Wi-Fi-tracking or Bluetooth-tracking) without the consent of the individual concerned.” The text is still evolving but this would be all the more unacceptable that, as discussed in this paper, solutions can be developed to make information and consent more effective,

⁸The use of electronic means is not required by the GDPR and, so far, the WP29 seems to consider signposts as an acceptable information means.

without introducing excessive constraints neither for data controllers nor for data subjects.

ACKNOWLEDGMENTS

This work has been partially funded by the CHIST-ERA project UPRISE-IoT (User-centric Privacy and Security in the IoT) and the ANR project CISC (Certification of IoT Secure Compilation).

REFERENCES

- [1] M. Cunche, D. Le Métayer, and V. Morel. A Generic Information and Consent Framework for the IoT (Extended Version). Research Report 9234, Inria. <https://hal.inria.fr/hal-01953052>.
- [2] A. Das, M. Degeling, D. Smullen, and N. Sadeh. Personalized privacy assistants for the internet of things: An infrastructure for notice and choice in the internet of things. *IEEE Pervasive Computing*, 17(3):35–46, Jul 2018.
- [3] A. Das, M. Degeling, X. Wang, J. Wang, N. Sadeh, and M. Satyanarayanan. Assisting users in a world full of cameras: A privacy-aware infrastructure for computer vision applications. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1387–1396, July 2017.
- [4] L. Demir, M. Cunche, and C. Lauradoux. Analysing the privacy policies of Wi-Fi trackers. In *Workshop on Physical Analytics*, Bretton Woods, United States, June 2014. ACM.
- [5] S. Du, M. Ibrahim, M. Shehata, and W. Badawy. Automatic License Plate Recognition (ALPR): A State-of-the-Art Review. *IEEE Transactions on Circuits and Systems for Video Technology*, 23(2):311–325, February 2013.
- [6] European Commission. Proposal for a Regulation of the European Parliament and of the Council concerning the respect for private life and the protection of personal data in electronic communications and repealing Directive 2002/58/EC (Regulation on Privacy and Electronic Communications), January 2017.
- [7] M. Langheinrich. A privacy awareness system for ubiquitous computing environments. In G. Borriello and L. E. Holmquist, editors, *UbiComp 2002: Ubiquitous Computing*, pages 237–245, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [8] Jakob Eg Larsen, Piotr Sapiezynski, Arkadiusz Stopczynski, Morten Mørup, and Rasmus Theodorsen. Crowds, bluetooth, and rock’n’roll: understanding music festival participant behavior. In *Proceedings of the 1st ACM international workshop on Personal data meets distributed multimedia*, pages 11–18. ACM, 2013.
- [9] Official Journal of the European Union. Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), May 2016.
- [10] Dieter Oosterlinck, Dries F Benoit, Philippe Baecke, and Nico Van de Weghe. Bluetooth tracking of humans in an indoor environment: An application to shopping mall visits. *Applied geography*, 78:55–65, 2017.
- [11] P. Pappachan, M. Degeling, R. Yus, A. Das, S. Bhagavatula, W. Melicher, P. E. Naeini, S. Zhang, L. Bauer, and A. Kobsa. Towards Privacy-Aware Smart Buildings: Capturing, Communicating, and Enforcing Privacy Policies and Preferences. In *Distributed Computing Systems Workshops (ICDCSW), 2017 IEEE 37th International Conference on*, pages 193–198. IEEE, 2017.
- [12] R. Pardo and D. Le Métayer. Analysis of Privacy Policies to Enhance Informed Consent. In *Data and Applications Security and Privacy XXXII - 32nd Annual IFIP WG 11.3 Conference, DBSec 2019, Charleston, USA, July 15-17, 2019, Proceedings*, 2019.
- [13] Bluetooth SIG. *Specification of the Bluetooth System v5.0*. December 2016.
- [14] WP29. Opinion 8/2014 on Recent Developments on the Internet of Things, 2014.
- [15] WP29. Guidelines on Consent under Regulation 2016/679, November 2017.
- [16] WP29. Guidelines on transparency under Regulation 2016/679, December 2017.
- [17] WP29. Opinion 01/2017 on the Proposed Regulation for the ePrivacy Regulation (2002/58/EC), December 2017.