Semantic Access Control for Privacy Management of Personal Sensing in Smart Cities

Michał Drozdowicz, Maria Ganzha, Marcin Paprzycki

Abstract—Personal and home sensors generate valuable information that could be used in Smart Cities. Unfortunately, typically, this data is locked out and used only by application/system developers. While vendors are partially to blame, one should consider also the "binary nature" of data access. Specifically, either owner has full control over her data (e.g. in a "closed system"), or she completely looses control, when the data is "opened". In this context, we propose, a semantic technologies-based, authorization and privacy control framework that enables user to maintain flexible, yet manageable data access control policies. The proposed approach is described in detail, including implementation and testing.

Index Terms—Access control, semantic technologies, Smart City, data privacy, sensors, XACML, Attribute Based Access Control

1 Introduction

Many Smart City projects rely on information collected from public sensor networks monitoring, among others: traffic, parking availability, pollution, noise, etc. (see, for instance, [1]). Examples, such as the city of Barcelona [2], show benefits of such knowledge in governing the city, optimizing operational expenditures, and improving citizens' quality of life. However, the initial costs of such initiatives are very high, with major contributing factors including the purchase and installation of sensors, infrastructure cost (e.g. high-throughput network), or development and integration of software. Moreover, introducing new "data sources" often requires deploying new sensor networks, or upgrading existing ones (both generating substantial costs). Finally, to achieve the expected benefits, the ecosystem must be maintained and adapted to follow changes in technology and development and growth of the city.

Some of those shortcomings can be addressed by taking advantage of the rapidly growing number of personal, and home-based, IoT devices, therefore reducing the costs of hardware infrastructure needed to gather data. Moreover, the variety of citizen-owned sensing devices is systematically increasing, generating new dimensions of useful knowledge. For instance, the popularity of fitness tracking solutions has lead to a massive growth of health and lifestyle related data, which could be used to improve the medical and living conditions of the society.

Obviously, motivating the citizens to share their data "with the city" is a serious challenge, but it has been shown [3] that one of the key obstacles to achieving this is "privacy management". Specifically, how to facilitate adequate control over personal data, and thus convincingly assure the protection of privacy. Therefore, a successful solution, gathering personal sensor information for public use, should provide solid means of managing privacy preferences and fine-grained access control. Here, let us note that while there exist ways of anonymizing data to make it less sensitive, research on removing user/data anonymity limits the relevance of such approaches [4], [5].

Furthermore, as discussed in [6], when dealing with health-related data and/or movement patterns extracted by fitness trackers, the actual challenge concerns *relative perception of privacy*. Specifically, it involves not only *which data is to be shared*, but also *with whom* and *for what purpose*.

Finally, let us consider access to personal data by government agencies, e.g. related to criminal investigations, or national security. Analysis by Nojeim et al. [7] shows that, from legal and practical perspective, existing regulations and tools fail to reconcile public security with basic human rights and legal regulations (e.g. the GDPR). Therefore, when facing access requests, businesses storing the personal information rely on own judgment and/or interests, while agencies resort to broad, uncontrolled, and often unnecessarily detailed, surveillance.

The described problem is an instance of a more general topic of *access control*. It requires defining rules of who is allowed to access data, the same way as a company defines who can access a specific area in a building. Access control is well studied, and many approaches to solving it have been suggested and implemented. Here, Access Control Lists, Role Based Access Control or Attribute Based Access Control mechanisms have been created to tackle generalization of user roles and resource groups, static and dynamic Separation of Duties, spatiotemporal authorization, etc. Acknowledging this, note that several specific aspects of privacy management in Smart Cities need to be addressed:

- Access requester is, likely, an organization. Moreover, the structure of the organization is, often, not known up-front. Hence, representation of (hierarchical) orga-

- *M. Drozdowicz, M. Ganzha and M. Paprzycki were with Systems Research Institute, Polish Academy of Sciences, Warsaw, Poland. Corresponding author: michal.drozdowicz@ibspan.waw.pl*
- *M. Ganzha was with Department of Mathematics and Information Science, Warsaw University of Technology, Warsaw, Poland.*

nizational structure is needed. Obviously, question of identity verification arises, but it is out of scope of this contribution.

- Data request is completed on behalf of an external organization, e.g. local government, which should be allowed to access only some data. Hence, token-based authorization (such as OAuth) is not feasible.
- Considered data is often a series/stream of observations that should be abstracted to types/categories, to avoid authorizing them individually. However, due to differences between devices/services, enforcement of a common observation vocabulary is not likely. Hence, use of access control mechanisms that depend on a fixed set of "scopes" (e.g. OAuth), may be challenging.
- Potential (large) scale of social participation, coupled with heterogeneity of data gathering applications, brings interoperability challenges that also materialize in access authorization. Differences in data representation necessitate either (i) conversion of data to some common format, which may not be feasible from closed data ecosystems (e.g. commercial fitness trackers), or (ii) introduction of a mapping layer. This would also result in authorization decisions involving centralized rules and policies.
- Time and location of issuing the request may not be essential, however certain spatiotemporal data related to the accessed information may be of use.
- Legal access to data (e.g. governed by GDPR) must be allowed, while rigorously controlled. Corresponding policies should consider purpose of use, retention routines, type of information, etc.

In this context, in [8] we have proposed a semantically-enriched authorization system for fine-grained control of data access. Here, we expand on the idea, focusing on *if* and *in what way* ontological modeling, and semantic reasoning, can help manage privacy preferences in participatory sensing, within Smart Cities. The proposed solution recognizes that different information may be perceived as more or less private, depending not only on the nature of data, purpose of collection, and requesting entity, but also on purely subjective criteria. This, coupled with semantic representation and processing of pertinent meta-data, enables individuals to precisely manage their data access permissions. Finally, we recognize that certain legal regulations should be enforced and prioritized over the individuals' personal preferences. In this context, let us describe the use case scenarios, which guides the remaining parts of the paper.

2 Overview of the use case scenario

As discussed in [9], fitness data, collected by users for health tracking, could be useful for Smart Cities' agencies (e.g. public health organizations). It may not encounter known problems in adopting participatory sensing (also known as crowdsensing [10]), such as the need for incentives [11] and/or change of behavior. Therefore, the proposed use cases involve tracking of an individual's movement habits, when the data is generated either by a GPS, or a pedometer (e.g. in a smartphone, smartwatch, or smart-shoes).

The general use case, is that of Sally, who uses a smartphone and a fitness application for tracking her running and cycling workouts. Thus far, she has collected data for personal benefits and shared it with friends (using some application). However, she is considering participation in a program analyzing sport activities in her home city. The primary use case (UC1) concerns the local Health Center, wishing to investigate the workout and training habits of the citizens. The second entity interested in her data (UC2) is the Police, investigating a crime in a certain area, searching for potential witnesses. For UC1 Sally would like to specify (independently) what information, and at what level of detail, she will share. UC2 illustrates how the proposed system handles legal obligations, while providing sufficient control over what data may be accessed under what conditions.

To deliver the needed functionality, we will build upon the semantically enriched Attribute Based Access Control system, introduced in [8], [12], [13], and expand it with a more detailed model of privacy preferences, as well as means of handling legal access control policies.

Thus, in Section 3, we give an overview of the state-of-the-art of solutions for enforcing privacy in Smart Cities as well as access control solutions making use of semantic technologies. Further, in Section 4, we briefly describe the SXACML access control system and discuss how to design an ontology that can be used to manage privacy and trust in Smart City. Finally, in Section 5 we revisit our guiding scenarios, to show in detail how proposed system enables citizens to manage access to their personal data, using the proposed system.

3 Related work

3.1 Privacy and access control

Let us first look into the methods of general access control, in which authorization policies and rules are used to validate if a *Subject* is permitted to perform an *Action* on a *Resource* in a certain request *Context*. We purposefully focus on the decision process and leave out consideration of "orthogonal aspects", such as identification, authentication, or action tracing.

Attribute Based Access Control (ABAC) provides the most flexible, and context aware, approach to authorization. Here, *Subject*, *Action*, *Resource*, and *Context* are described with sets of attribute values. Authorization decision is based on evaluation of policies, that specify conditions on the attributes. The most common implementation of ABAC is the eXtensible Access Control Markup Language (XACML[1]; [14]).

In Figure 1 we depict a typical sequence of actions undertaken during request evaluation, which follows the XACML standard.

When a system using XACML is configured, an administrator defines and manages policies within the Policy Administration Point (*PAP*) and supplies them to the Policy Decision Point (*PDP*). Once an access request is sent, by the *Subject*, to the Policy Enforcement Point (*PEP*), it is forwarded to the *Context Handler* which, in turn, notifies the *PDP*. Here, the *PDP* verifies the values of all attributes used in the policy definitions. Such values may be found in the

1. http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html

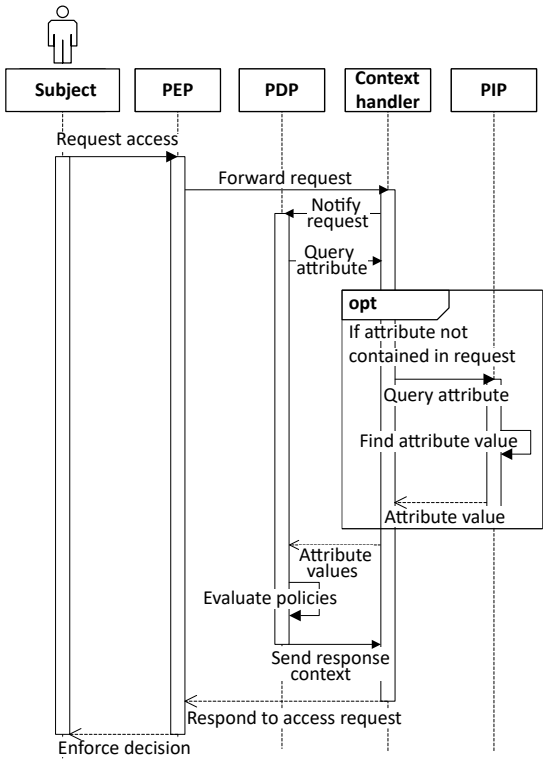

Fig. 1. Evaluation of request in XACML

request itself, or need to be retrieved from one, or more, Policy Information Point (*PIP*) components. Next, the *PDP* evaluates policy rules, combines results (if multiple policies are involved) and builds the response context, which is returned to the *Context Handler*. Results are then sent to the *PEP*, which enforces the decision.

ABAC, however, omits (a) implicit structure of entities requesting access, (b) nature of data, or (c) relations between *Subject(s)* and *Resource(s)*. This leads to significant complexity in defining and managing the policies. Regardless of improvements, like ALFA (Abbreviated Language For Authorization[2]),maintaining policies storing organizational and resource hierarchies, or more sophisticated relationships between data, requires major effort.

Additionally, authoring and administration of policies requires understanding of the policy definition language and the attribute space of the system under control. Therefore, "policy management" requires skilled specialists. As a result, ABAC has been adopted mostly in large organizations, e.g. in military, government, healthcare, or finance, where potential negative implications of unauthorized access justify expenditures related to managing policies.

Note that our first use case assumes privacy control managed by the user who is not an access management expert. Therefore, it is crucial to provide tools and methods to easily express users' attitude towards data access.

Turning our attention back to the personal sensing, [3] explores the objections to participating in such programs, based on studying a group of people who used different sensors. It was observed that raised objections depend on what was recorded, in what circumstances, and on the created value. Moreover, giving users more knowledge and control over data increased the potential for adoption of crowdsensing-type approaches.

In this context, [6] proposes a framework for classifying citizen-related data, which also considers subjective *feelings* about how personal the data is. For instance, location data could be seen as *personal*, while "generic traffic data", not tied to an individual, would likely be *impersonal*. Secondly, the *purpose* of collecting data, ranging from "service" to "surveillance" is examined. For instance, use of location data for traffic management would, most probably, be seen as a "service", while using such information for predictive policing could be considered as "surveillance". When applying this framework to access control and privacy enforcement, two points are worth noticing. First, how "personal" a given "piece of data" appears to the user is rather complex. The data collected by any given sensor can be very personal if connected with personally identifiable information, or collected with high granularity (e.g. exact jogging location data). However, when such data is aggregated over time (e.g. distance run daily) and/or using strong anonymization (e.g. spatial k-anonymity; [15]), it may be considered as not sensitive at all. Furthermore, different persons may care more, or less, about "releasing" personal data (see, information shared within social networks). Overall, "level of sensitivity" cannot be connected to a specific sensor. Instead, it is related to (a) type of observation, (b) its aggregation, (c) anonymization, and (d) "personality" of the user. Second, computer that evaluates access request has limited reasoning capacity. Therefore, it is important to let the user define permission rules, based (i) on the specific organization requesting the information and/or (ii) purpose of use, as specified in the request. For instance, users may assume that requests by the police department represent "surveillance". However, request for data to be used for traffic control, shifts the assessment towards "service" (as long as the user is willing to trust the police).

In [9] and [16], authors discuss the possibility of using personal health and fitness information. They also propose a privacy preserving architecture for data collection. Focus of their work is on aggregating and anonymizing the data, to become "unbreakable" by data mining. Here, let us note that numerous papers describe different solutions to data anonymization, especially for the location data; see, for instance [17], [15], [18] and [19]. However, these papers solve an issue that is, in a way, orthogonal to our concerns. We are interested in designing a system that gives its users the best possible control over their information, regardless of data anonymization. Hence, we accept that, in some cases, sharing fine-grained, personal information may be necessary. Moreover, we assume that the user might not want certain parties to acquire even anonymized information. Finally, we recall that some users may be willing to share "very personal data" regardless if it is anonymized or not.

Authors of [20] propose a framework for managing and enforcing privacy policies in the context of data collection, as well as consent regulations, such as GDPR. The information model used in policies includes devices, entities (agents or organizations), data items, and purposes of use. Semantics of the policy elements can be expressed using subsumption

2. http://docs.oasis-open.org/xacml/alfa-for-xacml/v1.0/alfa-for-xacml-v1.0.doc

of organizations and a partial order of attributes, which represents data item composition (e.g. street name as part of an address attribute). Each policy is a defined as a set of rules governing:

- what attribute can be accessed, by which entity, under what condition (Data Communication Rule),
- for what purposes it can be used, and how long it can be retained (Data Usage Rule), and
- what are the rules for transferring the data to other entities, each governed by separate Data Communication and Usage rules.

Conditions within Data Communication Rules are defined using a formal, logics-based language, including negation and conjunction of predicates. Its relative simplicity brings the possibility of formal verification and provability of policies. The framework, however, offers limited possibilities to express semantics and relationships between the data attributes, which could lead to issues when working with more complex domains and policies. Additionally, given the scope of the framework, covering data collection, usage, and transfer, should the policies be fully enforced, the solution would need to be applied throughout the process of data collection, storage, and handling. This in turn may prove hard to implement in practice, given the wide variety of systems used by the data controllers and processors.

The Personalized Privacy Assistant Project[3], by the Carnegie Melon University, seeks to provide individuals with tools enabling them to control how their data is collected and processed. Related publications such as [21] and [22] describe a mobile solution monitoring the device permissions (e.g. location, camera access) granted to various applications installed on the user's smartphone. It uses machine learning algorithms for clustering and classification, to group these programs into categories based on their functionality profile and purpose of data collection, finally assisting the user in making decisions about their privacy preferences. Additionally, to make configuration of preferences easier, the tool includes a number of privacy profiles and attempts to semi-automatically assign a user to one of them, based on answers to a simple survey. Compared to our research goals, the solution limits the scope of data under control to only the permissions recognized by the Android operating system and, therefore, does not address scenarios where the number of data items, or resources, is large or ever-changing. In a similar way, the set of applications installed on a mobile device do not change as often as the potential consumers of user data, in IoT and personal sensing scenarios.

Recent results of the project, described in [23], expand the solution to cover use cases dealing with a proliferation of IoT devices around the individuals. Here the authors address the problem that a typical person is continuously monitored by various devices, collecting personal information, and has little knowledge or control over the purpose and practices of data processing. They propose a distributed system, in which the personal assistant, residing on the user's smartphone, interacts with registries of surrounding IoT devices and uses its privacy preference policies to control what kind of consent it should give to device owners, i.e. data controllers. Finally, for certain types of more private data, external Policy Enforcement Points can be deployed that control what information can be provided to each data collector, depending on the privacy preferences of the data subject. The approach is sound, but tackles a somewhat different problem than the one we attempt to solve – as in the case of [20], it deals primarily with user consent to collect and process data originating from devices owned and operated by third parties. In our context, we are more concerned with data generated by user-owned devices.

3.2 Ontologies in Access Control

In this context, let us note that knowledge representation and automatic reasoning, based on the structure and semantics of data, are dealt by ontology engineering. Over the years, it developed mature methods for formally representing concepts and their relationships. Here, an ontology is understood as a specification of a vocabulary for a domain, including classes of objects, relations, functions, and other concepts [24]. Ontology-based models have been successfully applied in various areas, e.g. for genome modeling [25]; in healthcare (SAPPHIRE project; [26]), or in the Internet of Things [27].Overall, as shown in [12], [28], [29], by introducing semantic extensions to an ABAC system, it is possible to:

- Define the structure of *Subjects* and *Resources* to closely model actual organizations and domains. Semantics also provides a consistent way of implementing RBAC and hierarchical resources.
- Represent additional relationships and reason over the model, to uncover implicit knowledge, to be used for checking request consistency, applicability of rules, and making decisions.
- Define mapping ontologies, making it straightforward to employ an ontological model of the domain that is reusable in other ways than just authorization.
- Flexibly and efficiently deal with heterogeneity, by utilizing ontology alignment and mapping.
- Define additional attributes that are automatically inferred from the information contained in ontologies.
- Delegate permissions from one *Subject* to others in a hierarchy, by utilizing property transitivity.
- Infer conflicts between roles, by defining disjointness axioms in an ontology (thus, satisfying Separation of Duty). By defining rules, it is possible to tie the role assignment to dynamic conditions, such as time or the *Action* being performed, to handle also Dynamic Separation of Duty.

Another example of combining semantics and access control is [30], where authors propose SenTry – a language and framework for personal privacy control. The solution is based on an OWL ontology modeling policies, and the Semantic Web Rule Language based (SWRL[4]) rules for context-specific predicates used for decision making. Specifically, semantic reasoner evaluates applicable predicates, grouped into: filter, static and dynamic categories. This solution implements the ABAC model, but dismisses the de-facto standard of access control – XACML. By building

3. https://www.privacyassistant.org/

4. http://www.w3.org/Submission/2004/SUBM-SWRL-20040521/

the solution completely from ground up, it misses the opportunity to benefit from the large number of existing (and, sometimes, very mature) tools built around XACML, dealing with handling of rule conflict resolution, request processing, geospatial functions, etc.

Another solution, combining XACML with semantics is reported in [31]. Here, the LAPAR engine uses XML transformations (implemented as XSLT templates) to convert XACML policies to SWRL rules and further transform OWL ontologies and SWRL rules into Jess[5] inference engine statements. Proposed system reasons over combined knowledge, and computes authorization decisions. While presented results concern access to documents in a university, it is not clear (and somewhat doubtful) if the solution is capable of transforming the entire grammars of XACML, OWL, and SWRL into Jess rules solely by processing their XML representations. In absence of other use cases proving the concept also for other domains, we are not convinced the approach is applicable within the context of privacy control.

Summarizing, while the ABAC approach forms the best base for authorization systems in large-scale IoT applications, it lacks flexibility and meta-data modelling capabilities necessary in dynamic, heterogenous environments. Combining ABAC with semantic reasoning on ontological knowledge bases addresses these shortcomings. Finally, extending an established standard like XACML, instead of developing a purely semantic solution, brings numerous benefits, as well as giving the possibility to mix ontological and traditional ABAC-based rules in the same policy set. Let us now pursue this line of reasoning further.

4 Ontologies for privacy management in Smart City

4.1 Semantic XACML

In this context, in [8], [12], we have introduced a semantics-driven implementation of the *PIP*, thus defining the Semantic XACML (SXACML[6]) approach. The complete solution extends the XACML architecture in the following ways:

- In addition to managing XACML policies, the *PAP* module has been complemented with means of administering the ontologies used in the system. It includes a graphical front-end allowing one to define class mappings, expressions, and instances that are then added to the ontology and used during policy processing.
- The *PDP* loads an additional resource finder module that handles multi-resource request scenarios, in which the access request does not specify a concrete resource, but rather a category that needs to be resolved to a set of individuals. The functionality of the semantic resource finder is depicted in Algorithm 1 and also handles resource class hierarchies, i.e. it can traverse an entire class-subclass structure defined in the ontology. Note that, in this scenario, the pure-XACML way of specifying the hierarchy relationships between resources, in policies, is very complex (see the XACML Hierarchical Resource Profile[7]).
- A semantic *PIP* module has been implemented, to enrich the set of attributes provided in the request with new information retrieved from the ontology. Thanks to the use of a semantic reasoner, the attribute values need not be explicitly defined in the knowledge base, but can also be inferred from other known facts. Additionally, we have introduced special attributes denoting the class identifier of each XACML attribute category (subject, resource, action, environment) that can be used in the policies for rules involving the type of a resource, or the role of the subject. As in the case of other attribute values, such classification of request categories can be deduced by means of automatic reasoning. The procedure of retrieving attribute values (labeled in Figure 1 as "Find attribute value") is performed according to Algorithm 2)

input : URI of resource class $class$
output: set of permision decisions

$O_d \leftarrow$ load domain ontology;
$O_m \leftarrow$ load mapping ontology();
$O_r \leftarrow$ new ontology ; `// temp request ontology`
O_r imports {O_d, O_m};
run semantic reasoning on O_r;

$results \leftarrow$ empty set of permission decisions;
$sol \leftarrow$ query ontology for instances of $class$;
`// takes into account entire subclass hierarchy`
foreach *resource individual* I_r *in* sol **do**
  $decision \leftarrow$ evaluate policies for I_r add $decision$ to $results$
end
return $result$;

Algorithm 1: Evaluation for multiple resource class instances

The advantages of the SXACML approach include, but are not limited to:

1) Simplified policies – information common to multiple policies can be "extracted into the ontology", resulting in the policies being represented in a "more compact" form.
2) Better support for RBAC – role hierarchies can be modeled as ontology classes, user membership in a role can be inferred from attributes, and Separation of Duty can be verified by semantic reasoning.
3) More flexibility in defining relationships between concepts – an attribute value may be inferred from a complex graph of linked data, utilizing properties of *Subject*, *Resource*, *Action*, and *Environment*.
4) Improved interoperability, by semantic mapping of disparate concepts in requests and policies – allows decisions even in the case of different vocabularies.

In prior work, we have used the semantic *PIP* only as a provider of attribute values to policies (specified in XACML). The goal, in considered use cases, was to simplify XACML policies, and move domain models to OWL, assuming that the policy administrator has knowledge of XACML but not of OWL. We have also employed the Onto-

5. https://www.jessrules.com/
6. https://github.com/mdrozdo/SXACML
7. http://docs.oasis-open.org/xacml/3.0/rbac/v1.0/xacml-3.0-rbac-v1.0.html

input : evaluation context ctx, id of attribute to find $attrId$
output: bag of values of attribute

$O_d \leftarrow$ load domain ontology;
$O_m \leftarrow$ load mapping ontology();
$O_r \leftarrow$ new ontology ; `// temp request ontology`
O_r imports $\{O_d, O_m\}$;
foreach *category from* ctx **do**
  $I_c \leftarrow$ new OWL individual;
  foreach *attribute in category* **do**
    add property assertion to I_c;
  end
  add I_c to O_r;
end
run semantic reasoning on O_r;

$result \leftarrow$ empty bag of attribute values;
$sol \leftarrow$ query ontology for attribute value;
foreach *result in* sol **do**
  $val \leftarrow$ convert $result$ to an XACML attribute value;
  add val to $result$;
end
return $result$;

Algorithm 2: Finding attribute values

Play[8] ontology editor [13], [32] to assist the administrator in managing ontological concepts, such as *Resource* categories, or *Subject* roles. OntoPlay has proven to be a valuable tool enabling users not accustomed to semantic technologies to create complex class expressions and individual definitions, with applications not only in access control, but also in querying grid computational nodes at the University of Aizu, Japan [33]. Furthermore, it has been utilized in student projects during the Semantic Technologies seminar at the Warsaw University of Technology, as well as in several master's theses defended at that institution, some of them resulting in publications, e.g. [34], [35].

However, in participatory sensing and personal privacy control, there is no system administrator – the solution must be simple enough that the users are able to easily manage their own preferences and policies. We have, therefore, considered how to draw the boundary between OWL and XACML, taking into account that managing XACML policies manually is far beyond the capabilities of a casual user. Hence, we moved most responsibilities for the decision to the the semantic part, by defining the `PermittedRequest` class as a subclass of `Request` and providing the user with an OntoPlay-based interface that lets them define the relevant class expression in a point&click manner, without knowledge of the ontology, and being fully agnostic of the XACML back-end.

On the other hand, in the UC2 use case (i.e. the police investigation), access to the personal data should be rigorously controlled, considering legal conditions and obligations. Here, while it is possible to realize the ABAC model using semantics, it would introduce unnecessary complexity to the user. Moreover, implementing OWL concepts, capturing policy sets, combining multiple policies, obligations, etc., would change the policy processing engine. However, in comparison to subjective personal preferences, legal rules are likely to be relatively static, long-lived, and independent of the user. Therefore, legal policies can be implemented as standard XACML policies (by an access control expert).

In summary, we have separated (1) the subjective, dynamical personal privacy preferences – defined in OWL – and (2) the, potentially complex, static, legal rules defined as XACML policies and the policy sets. In the latter case, the seam between XACML and OWL follows earlier research – the semantic *PIP* infers and provides the *PDP* values of certain attributes, and the remaining parts of policy processing are performed by the *PDP* component. The final solution uses a policy combining algorithm to reconcile the privacy preferences with hard legal rules in the same policy set, giving higher priority to the regulatory requirements.

Note that, we only consider and describe the context of the access request permission and thus focus on the *PAP*, *PDP* and *PIP*. Therefore, we purposefully ignore collecting and storing sensor/activity tracker data. We assume that, in a working system, another layer, responsible for data collection, would be instantiated. Examples of modules, realizing such functionality, can be found in [16], [36], [37]. Likewise, as mentioned earlier, we omit issues related to data anonymization. We assume that data requiring anonymization has already been processed, using techniques stated in Section 3). Nevertheless, these simplifications do not influence the way that the proposed approach works. Finally, we leave out the details of the *PEP* implementation, which would need to be tightly related to the way of storing the data as well as means of requesting the information by third parties. In the prototype implementation we have used WSO2 Identity Server[9] as the gateway and *PEP*.

Let us now turn our attention to another aspect of adapting SXACML to the Smart City use cases. Obviously, the crucial aspect of applying the system to a new domain is the selection or design of ontologies. While the term ontology has many definitions, here we understand it as formal representation of pertinent (application-specific) aspects of knowledge about a domain. Moreover, we have decided to use the Web Ontology Language (OWL[10]; [38]) to formally represent ontologies.

One of the key features of OWL ontologies is that they can be reused by other ontologies, composed and adapted for more specific purposes. Hence, it is important to, first, search for existing resources in ontology catalogues such as the Linked Open Vocabularies[11]. However, we were unable to find a complete ontology covering the privacy management of data acquired from sensing devices. Therefore we have split the domain of interest into several parts, which are then combined into the final representation of the domain of interest.

To this effect, we discuss the following components of the ontological structure, used in the proposed solution:

8. https://github.com/mdrozdo/OntoPlay

9. https://wso2.com/identity-and-access-management
10. https://www.w3.org/TR/owl2-overview/
11. http://lov.okfn.org/dataset/lov

- Access Control ontology – generic representation of ABAC concepts,
- Internet of Things ontology and Fitness Tracking ontology – jointly representing *Resources*,
- Privacy ontology – providing additional privacy-related concepts.

Note that, following principles of ontology engineering, we have been re-using existing ontologies whenever possible, while modifying them only when necessary.

4.2 Access Control ontology

Let us start from the ontology describing concepts related to ABAC and XACML. This ontology has to be generic and domain independent. In [8], [12] this purpose was fulfilled by a simplistic Request Ontology. Here, we introduce a more complete Access Control Ontology (ACO), as an ontological representation of the XACML request elements, but also providing core concepts related to data access. Basic elements of ACO are taken directly from the ABAC model, and reflect the same attribute categories as in XACML:

- `Subject`
- `Resource`
- `Action`
- `Environment`

For the *Subject* part of the ABAC model, an ontology covering relationships between different organizational entities that can be authorized to access the personal data was needed. Hence, we have decided to directly use the W3C Organization Ontology[12]. The XACML *Subject* has been mapped to the `foaf:Agent` class, which may represent a person, group or organization. The ontology also contains concepts and relations needed to define complex organizational structures, and membership in them. Finally, we have reused the `org:Role` class, to be used in policies that assume decisions based on roles of the *Subject* (in the RBAC approach). Moreover, we have defined classes related to the *Resource*:

- `Sensitivity` – capturing how personal the information is, and under what conditions it may be disclosed.
- `Confidentiality` – describing the level of legal restrictions associated with the information.
- `Owner` – specifying the entity (person or organization) owning the *Resource* or being the main object described by the *Resource*.

Another element is the `Trust` class, describing level of confidence of resource owner in given *Subject*. While trust modeling is an interesting topic on its own, here, it is only a class, with subclasses corresponding to different degrees of confidence. Obviously, if needed, this class can be replaced by a more comprehensive ontology (fragment).

Finally, the ontology includes the `PurposeOfUse` class, describing the reason for requesting the *Action* (how the obtained *Resource* will be used).

Figure 2 summarizes the Access Control Ontology.

4.3 Domain ontologies

First of all, considering that the guiding use cases deal with information collected using IoT devices, a "sensor ontology" is needed . This domain is well covered by the W3C Semantic Sensor Network Ontology (SSN[13]), which contains vocabulary describing sensors, observations, and actuators, as well as observed properties, features of interest, etc.

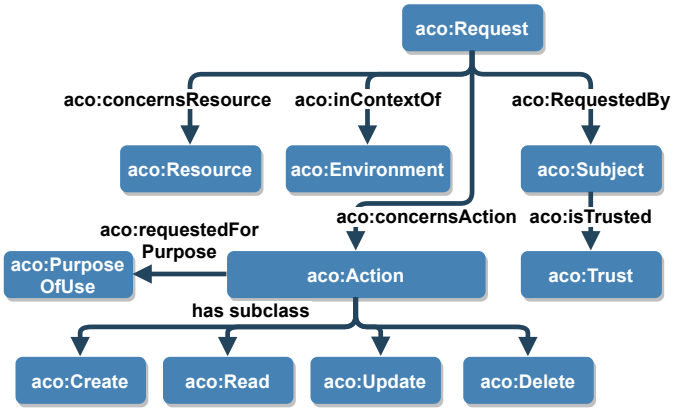

Fig. 2. Access Control Ontology

In SSN the `sosa:Observation` class represents a single act of measurement (`sosa:ObservableProperty`, e.g. heart rate) of a certain "feature of interest" (`sosa:FeatureOfInterest`, e.g. a specific person). It is described with (a) properties related to the sensor (`sosa:Sensor`, e.g. the heart rate monitor), (b) the feature/object that was measured, (c) the measurement procedure (`sosa:Procedure`, e.g. the method of measuring heart rate), and (d) detailed information about the result (`sosa:Result`). When accessing observations, no special requirements on what kind of actions can be performed are present (create, read, update or delete). We have extended the SSN ontology with an `AnonymizationProcedure` class (subclass of textttsosa:Procedure) that describes the method used for removing personal information.

To represent privacy and access control in IoT scenarios, we needed to extend and adapt the SSN ontology to deal with accessing sensor generated observations, hence we have defined a mapping of *Resource* and *Action* from the ACO ontology to appropriate classes in SSN, as described in detail in Section 4.4.

Second, to provide the needed vocabulary we have investigated several ontologies representing training activities and fitness tracking data. Here, the authors of [39] propose ontologies and a rule-based reasoner, for supporting people in following a healthy lifestyle. The presented research focuses on eating habits and omits fitness activities, as well as data tracking, and thus is not a good fit for our needs.

In [40], a framework for inferring person's physical state and activity, based on contextual information obtained from sensor network surrounding user, is proposed. Here, the *Context Modelling Ontology* transforms external information into knowledge about user activity. Unfortunately, the vocabulary is rather high-level and is of limited use for our needs.

The knowledge base in the Physical Activity, Health and Fitness Knowledge Model [41] contains comprehensive

12. http://www.w3.org/TR/2014/REC-vocab-org-20140116/

13. https://www.w3.org/TR/2017/WD-vocab-ssn-20170105/

vocabulary of physical activities. Furthermore, it includes concepts such as *activity frequency*, *activity intensity*, *activity duration* and *activity condition* (in terms of natural, social and legal environment). While some elements of this ontology, especially sports related, or describing workout intensity, could be reused, it is much too broad for our current needs.

As a result of not finding an appropriate solution, we have decided to create a fitness tracking ontology, containing the most important elements relevant to collecting information to workouts and physical activities. To that effect, our Fitness Ontology imports the SSN ontology, and extends it with the following elements (also depicted in Diagrams 3 and 4).

- `Training` and its subclasses represent various workouts, e.g. running or cycling. This element is meant to be extended with a larger set of activities (perhaps using concepts from the Physical Activity Knowledge Base), depending on the application requirements.
- Several subclasses of the `sosa:Observation` class, representing the measurements related to training: (`BloodPressure`, `HeartRate`), or general physical metrics (`Height`, `Weight`).
- `TrainingMetric` representing workout attributes, such as calories burned, distance, or step count.
- `GeospatialMeasurement`, with subclasses `Location` and `Route`, capturing the training location.

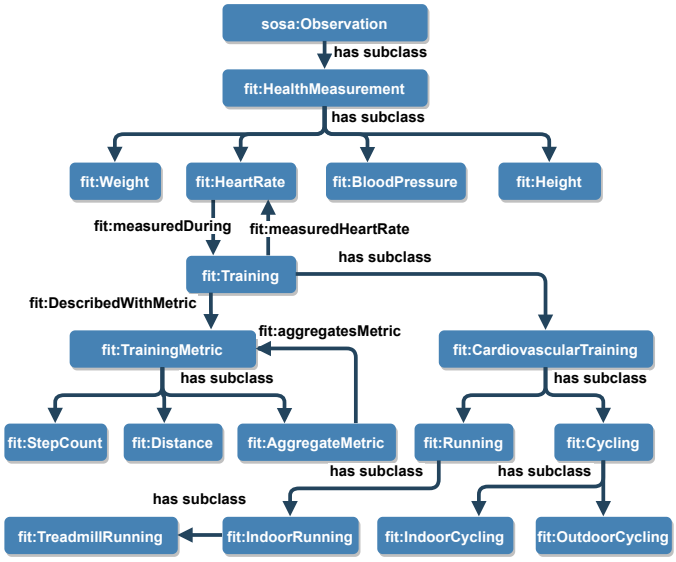

Fig. 3. Fitness Tracking Ontology – classes related to training types

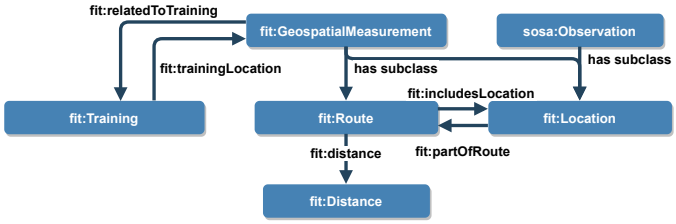

Fig. 4. Fitness Tracking Ontology – classes related to routes and locations

The SSN and Fitness ontologies represent data that is to be subject to access control, i.e. the XACML *Resource* category. The Access Control Ontology allows the *Subject* category to be described as a person, machine agent, or an organization. Let us now consider how to describe the *Action* category, taking into account attributes such as: purpose of use, retention policy, etc.

When it comes to privacy ontologies, authors of [42] have introduced a lightweight ontology for privacy preferences, in the context of Semantic Web and linked data. Moreover, creators of Semantic Cyber Information Modeling Initiative (SCIMI[14]) have proposed a Domain Specific Language for describing a privacy meta-model. However, since 2015, there was no recognizable progress of this work.

The Platform for Privacy Preferences (P3P[15]) is a specification, and an ontology, that allows web site authors to describe privacy practices in a machine readable format, enabling browsers to make semi-autonomous privacy related decisions. Even though the use case of P3P is different, developed ontology contains concepts useful in modeling privacy preferences in the participatory sensing, such as:

- Classes capturing information describing web site visitors: name, email address, IP address, etc.
- Categories to classify information about a person: demographic, financial, health, location, etc.
- Purpose of use categories, such as: administration, contact, telemarketing, etc.
- Retention policies for the collected data.

The P3P specification has been retired in 2018. However, it is a good base for a privacy preference ontology.

In [43], an ontology, describing various aspects of privacy and their interrelationships is presented. The main goal was to categorize data privacy in certain situations (e.g. medical data of a patient admitted to a hospital). Rating is based on aggregating atomic scores, such as: *Data Quality*, *Security*, *Data Subject's Rights*, *Legitimate Grounds of Processing*, *Transparency*, *Consent*, *Anonymity*. While this approach is quite interesting, it is not applicable to our use case as it does not take into account subjectivity of privacy. Specifically, even if policies and procedures are the same, some people may hesitate to expose personal data, while others do so without second thoughts.

The Privacy Preference Ontology (PPO), described in [44], is aimed at providing vocabulary and means of specifying privacy policies using RDF and SPARQL queries. The example uses FOAF to represent resources under control. Unfortunately, it captures only generic concepts (e.g. `PrivacyPreference`, `AccessSpace`, `Resource`, `Condition` etc.), which are already defined in our Access Control Ontology.

Finally, authors of PrOnto [45] decided to model privacy and data protection concepts, in the context of GDPR, to enable legal reasoning and compliance verification. Therefore, it does not contain elements describing privacy preferences or data protection policies, but focuses on legal rules, rights, obligations, purpose of use, etc. Moreover, at the time of writing, no complete ontology could be found. Therefore it was hard to fully evaluate its suitability to our needs.

14. http://privacyontology.org
15. https://www.w3.org/TR/P3P/

Eventually, we have decided that the P3P ontology can best serve as a base, however due to its size, for this report we have used only the elements relevant to our requirements, namely:

- The hierarchy of subclasses of the `Data` class, representing various data categories.
- The `Purpose` class as a representation of purpose of collection or use.
- The `Retention` class.

4.4 Mapping ontology

Having described the ontological components representing the domain under consideration, let us consider in more details how they relate to each other. In order for the *PIP* to access the knowledge base, it must first be consolidated in what we call the Mapping Ontology. It is built of several predefined mapping axioms, complemented with class expressions and / or individuals defined by the user as part of configuring their privacy preferences.

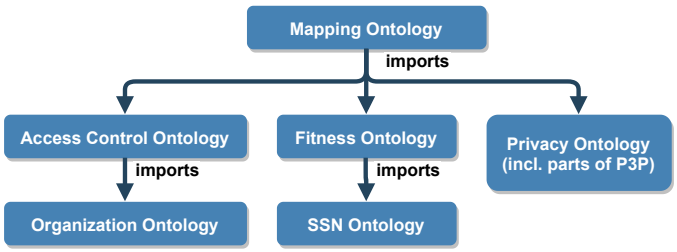

Fig. 5. Imports hierarchy of ontologies

Figure 5 depicts the high-level import hierarchy of the ontologies introduced in the previous sections. Specifically, the predefined mappings state that the `aco:Resource` class is a superclass of: `fit:Training`, `fit:TrainingMetric`, and `fit:GeospatialMeasurement`, which reflects what information could be requested.

The mapping also joins the Fitness and Personal Privacy ontologies – the `fit:HealthMeasurement` and `fit:GeospatialMeasurement` classes are marked as subclasses of `ppo:Health-data-category` and `ppo:Location-data-category` respectively. Moreover, we have added a custom category `ppo:FitnessData`, as a subclass of `ppo:OtherCategories`, that became a superclass for the `fit:Training` and `fit:TrainingMetric` classes. Finally, the `ppo:Purpose` class has been used as the range of the `ppo:hasPurposeOfUse` attribute, describing the `aco:Action` class (representation of the XACML *Action* category). Analogously, we have added the `p3p:Retention` class as an attribute of *aco:Action* (`ppo:hasRetentionPolicy`).

5 Experimental Evaluation

With all elements in place, let us now illustrate how the proposed approach can be used in our two use case scenarios, introduced in Section 2:

- UC1: Health Center requesting aggregated (monthly) information about training metrics.
- UC2: Police department requesting Sally's locations during a specific time period.

5.1 Health center

We start with UC1, where the Health Center requests access to information about Sally's training metrics. First, Sally has to define her privacy preferences. Here, she specifies a permission stating that the requester, belonging to the Health Center organization, may access aggregated monthly distance observations. This policy can be easily created using OntoPlay, as depicted in Figure 6. The result is a class expression describing a subclass of `PermittedRequest` called `HealthCenterPermission` that is subsequently added to the ontology.

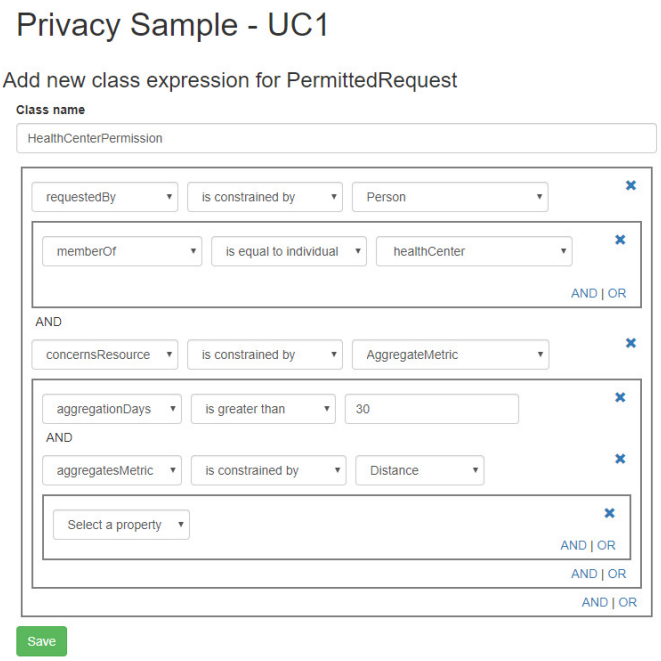

Fig. 6. OntoPlay interface for Health Center permission

Figure 7 presents the XACML request for the `Read` action (line 20), on resources of the `TrainingMetric` class (line 15), made by the Health Centre (line 9). Here, the semantic *PIP* component adds a new individual of type `Request` to the temporary ontology. As the request does not

```
1  <?xml version="1.0" encoding="UTF-8"?>
2  <Request xmlns="urn:oasis:names:tc:xacml:3.0:core:schema:wd-17"
3    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
4    xsi:schemaLocation="urn:oasis:names:tc:xacml:3.0:core:schema:wd-17
5      http://docs.oasis-open.org/xacml/3.0/xacml-core-v3-schema-wd-17.xsd"
     ↪ ReturnPolicyIdList="false">
6    <Attributes
     ↪ Category="urn:oasis:names:tc:xacml:1.0:subject-category:access-subject">
7      <Attribute IncludeInResult="false"
     ↪ AttributeId="urn:oasis:names:tc:xacml:1.0:subject:subject-id">
8        <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#string">
9          healthCentre </AttributeValue>
10     </Attribute>
11   </Attributes>
12   <Attributes Category="urn:oasis:names:tc:xacml:3.0:attribute-category:resource">
13     <Attribute IncludeInResult="false"
     ↪ AttributeId="sxacml:resource:resource-class-id">
14       <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#anyURI">
15         TrainingMetric</AttributeValue>
16     </Attribute>
17   </Attributes>
18   <Attributes Category="urn:oasis:names:tc:xacml:3.0:attribute-category:action">
19     <Attribute IncludeInResult="false"
     ↪ AttributeId="urn:oasis:names:tc:xacml:1.0:action:action-id">
20       <AttributeValue
     ↪ DataType="http://www.w3.org/2001/XMLSchema#string">read</AttributeValue>
21     </Attribute>
22   </Attributes>
23 </Request>
```

Fig. 7. XACML request for UC1

refer to a specific resource, but rather to a resource category – the *Resource* is only described as the `TrainingMetric` class. This is because the organization would like to collect as much relevant data as possible. Therefore, the initial step is to retrieve the appropriate resource individuals from the ontology. Considering the hierarchy of metrics shown in Figure 3, there exist several types of metrics: `StepCount`, `Distance`, and `AggregateMetrics`. Applying our semantic implementation of the XACML Hierarchical Resource Profile, the system translates this request into multiple decisions, based on the results of inference, which individuals in the ontology are instances of the `OutdoorTraining` class or its subclasses. The following evaluation steps are subsequently repeated for each resource.

The *Subject* is specified with the class `HealthCentre`. Hence, an individual of that class is added to the request ontology, and linked to the request individual. The request individual is also connected to the *Resource*. Next, the semantic reasoner can infer if the request individual satisfies all constraints specified by Sally, and classify it as `HealthCenterPermission`, and as `PermittedRequest`. Note that, the XACML policy includes a condition on the request class id attribute. Therefore the *PIP* returns a collection of classes describing the request individual: `Request`, `PermittedRequest` and `HealthCenterPermission`. The rule evaluates to a *true* value and the request is permitted by the *PDP*.

5.2 Police department

Next, let us consider the case of the Police Department, requesting location records from the night a crime took place (case UC2). Here, it is not up to the individual to define the rules of data access, as they represent en existing legal framework. In our, somewhat artificial, example we assume that the policy permits the Police to access information about the location of an individual as related to a committed crime. In the real-world, such a request would need to be accompanied by a warrant, i.e. specifying event location and time. Such warrant would need to be verified and digitally signed before being included in the data access request. Here, we will omit details as to how a warrant issuer can secure the request, and protect it from tampering. Nevertheless, let us stress that existing specifications such as the XACML XML Digital Signature Profile[16], provide appropriate solutions. The example request, depicted in Figure 8, contains the following attribute values:

- The `Subject` is Police Department (line 9).
- The `Resource` to be accessed is `Location` – which should be understood not as a single, specific, position, but rather as all permitted locations (line 14).
- The `Action` is Read (line 18).
- `Environment` encloses attributes related to the crime event location and time (lines 22-29).

The associated policy is presented in Figure 9 (encoded in ALFA for brevity). It contains a number of conditions on different attributes. The *Subject* is limited to the Police Department, and only `Read` actions are permitted. Moreover, the

16. http://docs.oasis-open.org/xacml/3.0/dsig/v1.0/xacml-3.0-dsig-v1.0.html

```
1  <?xml version="1.0" encoding="UTF-8"?>
2  <Request xmlns="urn:oasis:names:tc:xacml:3.0:core:schema:wd-17"
3    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
4    xsi:schemaLocation="urn:oasis:names:tc:xacml:3.0:core:schema:wd-17
5      http://docs.oasis-open.org/xacml/3.0/xacml-core-v3-schema-wd-17.xsd"
     ↪ ReturnPolicyIdList="false">
6    <Attributes
     ↪ Category="urn:oasis:names:tc:xacml:1.0:subject-category:access-subject">
7      <Attribute IncludeInResult="false"
     ↪ AttributeId="urn:oasis:names:tc:xacml:1.0:subject:subject-id">
8        <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#string">
9          policeDepartment</AttributeValue></Attribute>
10   </Attributes>
11   <Attributes Category="urn:oasis:names:tc:xacml:3.0:attribute-category:resource">
12     <Attribute IncludeInResult="false"
     ↪ AttributeId="sxacml:resource:resource-class-id">
13       <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#anyURI">
14         Location</AttributeValue></Attribute>
15   </Attributes>
16   <Attributes Category="urn:oasis:names:tc:xacml:3.0:attribute-category:action">
17     <Attribute IncludeInResult="false"
     ↪ AttributeId="urn:oasis:names:tc:xacml:1.0:action:action-id">
18       <AttributeValue
     ↪ DataType="http://www.w3.org/2001/XMLSchema#string">read</AttributeValue>
19     </Attribute>
20   </Attributes>
21   <Attributes
     ↪ Category="urn:oasis:names:tc:xacml:3.0:attribute-category:environment">
22     <Attribute IncludeInResult="true" AttributeId="EventLocation">
23       <AttributeValue DataType="urn:ogc:def:dataType:geoxacml:1.0:geometry">
24         <gml:Point xmlns:gml="http://www.opengis.net/gml/3.2"
     ↪ srsName="EPSG:4326">
25           <gml:pos srsDimension="2">38.889444 -77.035278</gml:pos>
26         </gml:Point></AttributeValue></Attribute>
27     <Attribute IncludeInResult="true" AttributeId="EventTime">
28       <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#dateTime">
29         2019-05-30T09:30:10-06:00</AttributeValue></Attribute>
30   </Attributes>
31 </Request>
```

Fig. 8. XACML request for UC2

request is only permitted if the *Resource* location attribute is within a one kilometer radius from the event location, and the position record time is in the period of one hour before and after the event. To apply geospatial comparison, the policy makes use of an extension to the XACML standard – the Geospatial eXtensible Access Control Markup Language (GeoXACML[17]). Such spatiotemporal conditions would be hard to define in a typical OWL reasoner and would require specific geospatial extensions to the ontology language, as well as the inference engine. The policy could also be expanded to feature conditions related to the purpose of use of the information, warrant chain, etc. Additionally, not illustrated in the listing, the policy is also part of a Policy Set together with other policies (e.g. the one used in UC1), configured with a combining algorithm which secures that the law enforcement regulations override any user-defined preferences.

```
namespace policePolicy{
    policy checkLocation {
        apply firstApplicable
        rule checkLocation {
            target clause subjectId == "policeDepartment"
                and action = "read"
                and resourceClass = "Location"
            condition timeInRange(locationTime,
                    dateTimeAddDayTimeDuration(eventTime, "P1H"),
                    dateTimeSubtractDayTimeDuration(eventTime, "P1H"))
                and isWithinDistance(eventLocation, locationPoint, 1000)
            permit
        }
    }
}
```

Fig. 9. Policy for UC2

In this case, the request defines the resource as an instance of the `Location` class and therefore the *PDP*, again,

17. https://www.opengeospatial.org/standards/geoxacml

needs to resolve it to a number of individual resources. The semantic reasoner is used by the resource finder to fetch concrete instances of the `Location` class from the ontology. While evaluating the policy for each of the locations, the *PDP* does not encounter the `locationTime` and `locationPoint` attributes from the `Resource` category and therefore it requests their values from the (semantic) *PIP*. Having acquired the attribute values, the policy condition is evaluated, by means of date-time XACML functions and the GeoXACML engine. The result is a multi-resource decision consisting of a number of individual responses, one for each location contained in the ontology.

6 Concluding remarks

In this paper we have considered how a semantically enriched Attribute Based Access Control system can be applied to (self-)management of user data privacy in Smart Cities. We have shown that practical application of semantic technologies brings important advantages for development of flexible, though robust, privacy controlling environments. In this context, our main contributions are as follows.

- We have reviewed existing ontologies, covering different aspects of the domain of interest, including sensors, fitness tracking and personal privacy, finally composing well established vocabularies into a complete ontology. This outlines the path that should be followed when systems similar to ours are to be developed for other domains.
- On the basis of our earlier work, we have presented a more refined approach to combining the XACML policies with semantical reasoning. Here, among others, we have taken into account observation that certain rules may need to be more rigid than others. The proposed approach allows "mixing and matching" (depending on specific circumstances of the developed system) XACML rules with attributes resulting from semantic reasoning. In other words, the boundary between XACML and semantics can be instantiated as needed.
- We managed to separate the (fixed) "rules of the system" that are to be formulated by specialists, from user-preferences. Here, each user can ("dynamically") formulate her/his rules, representing personal attitude towards privacy. User preferences will be captured within the system, without the need to change the rules.
- Expression of personal preferences does not require knowledge of semantic technologies. Rather, it is realized using OntoPlay, a novel interface to ontology-driven systems.
- Moreover, use of OntoPlay allows easy modification of system ontology. After ontology is modified, it will automatically materialize in the user interface, without the need of changing the system code.

One area that we have not included in the research, but intend to investigate in the future, is the specification of obligations, i.e. additional actions that should be performed by the PEP, following the enforcement of the decision. The considered solution will fully support the default XACML obligation specification format, but taking advantage of the rich body of knowledge regarding semantic web services could improve the possibilities of describing mandatory data storage and processing regulations.

References

[1] L. Sanchez, L. Muñoz, J. A. Galache, P. Sotres, J. R. Santana, V. Gutierrez, R. Ramdhany, A. Gluhak, S. Krco, E. Theodoridis *et al.*, "SmartSantander: IoT experimentation over a smart city testbed," *Computer Networks*, vol. 61, pp. 217–238, 2014.

[2] T. Bakıcı, E. Almirall, and J. Wareham, "A smart city initiative: the case of Barcelona," *Journal of the Knowledge Economy*, vol. 4, no. 2, pp. 135–148, 2013.

[3] P. Klasnja, S. Consolvo, T. Choudhury, R. Beckwith, and J. Hightower, "Exploring privacy concerns about personal sensing," *Pervasive Computing*, pp. 176–183, 2009.

[4] K. El Emam, E. Jonker, L. Arbuckle, and B. Malin, "A systematic review of re-identification attacks on health data," *PLOS ONE*, vol. 6, no. 12, pp. 1–12, 12 2011.

[5] C. C. Porter *et al.*, "Constitutional and regulatory: De-identified data and third party data mining: The risk of re-identification of personal information," *Shidler JL Com. & Tech.*, vol. 5, pp. 3–24, 2008.

[6] L. van Zoonen, "Privacy concerns in smart cities," *Government Information Quarterly*, vol. 33, no. 3, pp. 472 – 480, 2016.

[7] G. T. Nojeim, R. D. Lee, and I. S. Rubinstein, "Systematic government access to personal data: a comparative analysis†," *International Data Privacy Law*, vol. 4, no. 2, pp. 96–119, 05 2014.

[8] M. Drozdowicz, M. Ganzha, and M. Paprzycki, "Semantically enriched data access policies in eHealth," *Journal of Medical Systems*, vol. 40, no. 11, p. 238, 2016.

[9] A. Clarke and R. Steele, "How personal fitness data can be reused by smart cities," in *2011 Seventh International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, Dec 2011, pp. 395–400.

[10] Wikipedia, "Crowdsensing — wikipedia, the free encyclopedia," 2017, online, Accessed: 2017-05-25. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Crowdsensing&oldid=781041523

[11] H. Gao, C. H. Liu, W. Wang, J. Zhao, Z. Song, X. Su, J. Crowcroft, and K. K. Leung, "A survey of incentive mechanisms for participatory sensing," *IEEE Communications Surveys Tutorials*, vol. 17, no. 2, pp. 918–943, Secondquarter 2015.

[12] M. Drozdowicz, M. Ganzha, and M. Paprzycki, "Semantic Policy Information Point – preliminary considerations," in *ICT Innovations 2015*, S. Loshkovska and S. Koceski, Eds. Cham: Springer International Publishing, 2016, pp. 11–19.

[13] M. Drozdowicz, M. Alwazir, M. Ganzha, and M. Paprzycki, "Graphical interface for ontology mapping with application to access control," in *Intelligent Information and Database Systems*, N. T. Nguyen, S. Tojo, L. M. Nguyen, and B. Trawiński, Eds. Cham: Springer International Publishing, 2017, pp. 46–55.

[14] V. C. Hu, D. Ferraiolo, R. Kuhn, A. Schnitzer, K. Sandlin, R. Miller, and K. Scarfone, "Guide to attribute based access control (ABAC) definition and considerations," *NIST Special Publication*, vol. 800, p. 162, 2014.

[15] "Spatial k-anonymity," in *Encyclopedia of Database Systems*, L. Liu and M. T. Özsu, Eds. Boston, MA: Springer US, 2009, pp. 2714–2714.

[16] A. Clarke and R. Steele, "Smartphone-based public health information systems: Anonymity, privacy and intervention," *Journal of the Association for Information Science and Technology*, vol. 66, no. 12, pp. 2596–2608, 2015.

[17] K. L. Huang, S. S. Kanhere, and W. Hu, "Towards privacy-sensitive participatory sensing," in *2009 IEEE International Conference on Pervasive Computing and Communications*, March 2009, pp. 1–6.

[18] C. Liu, S. Chen, S. Zhou, J. Guan, and Y. Ma, "A novel privacy preserving method for data publication," *Information Sciences*, vol. 501, pp. 421 – 435, 2019.

[19] M. Li, L. Zhu, Z. Zhang, and R. Xu, "Achieving differential privacy of trajectory data publishing in participatory sensing," *Information Sciences*, vol. 400-401, pp. 1 – 13, 2017.

[20] R. Pardo and D. Le Métayer, "Analysis of privacy policies to enhance informed consent," in *Data and Applications Security and Privacy XXXIII*, S. N. Foley, Ed. Cham: Springer International Publishing, 2019, pp. 177–198.

[21] 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 *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security*, ser. SOUPS '16. USA: USENIX Association, 2016, p. 27–41.

[22] B. Liu, J. Lin, and N. Sadeh, "Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help?" in *Proceedings of the 23rd International Conference on World Wide Web*, ser. WWW '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 201–212. [Online]. Available: https://doi.org/10.1145/2566486.2568035

[23] A. Das, M. Degeling, D. Smullen, and N. M. Sadeh, "Personalized privacy assistants for the internet of things: Providing users with notice and choice," *IEEE Pervasive Computing*, vol. 17, pp. 35–46, 2018.

[24] T. R. Gruber, "A translation approach to portable ontology specifications," *Knowledge Acquisition*, vol. 5, no. 2, pp. 199 – 220, 1993.

[25] M. Ashburner, C. A. Ball, J. Blake, D. Botstein, H. Butler, J. Michael Cherry, A. P. Davis, K. Dolinski, S. Dwight, J. Eppig, M. Harris, D. P. Hill, L. Issel-Tarver, A. Kasarskis, S. Lewis, J. Matese, J. E. Richardson, M. Ringwald, G. M. Rubin, and G. Sherlock, "Gene ontology: tool for the unification of biology. the gene ontology consortium," *Nature genetics*, vol. 25, pp. 25–9, 06 2000.

[26] L. Feigenbaum, I. Herman, T. Hongsermeier, E. Neumann, and S. Stephens, "The semantic web in action," *Scientific American*, vol. 297, pp. 64–71, 01 2008.

[27] M. Ganzha, M. Paprzycki, W. Pawłowski, P. Szmeja, and K. Wasielewska, "Semantic interoperability in the Internet of Things: An overview from the INTER-IoT perspective," *Journal of Network and Computer Applications*, 2016.

[28] T. Priebe, W. Dobmeier, and N. Kamprath, "Supporting attribute-based access control with ontologies," in *Availability, Reliability and Security, 2006. ARES 2006. The First International Conference on*. IEEE, 2006, pp. 8–pp.

[29] R. Ferrini and E. Bertino, "Supporting RBAC with XACML+OWL," in *Proceedings of the 14th ACM symposium on Access control models and technologies*. ACM, 2009, pp. 145–154.

[30] S. Alcalde Bagüés, A. Zeidler, C. Fernandez Valdivielso, and I. Matias, "Towards personal privacy control," in *On the Move to Meaningful Internet Systems 2007: OTM 2007 Workshops*. Springer, 2007, pp. 886–895.

[31] I. C. Hsu, "Extensible access control markup language integrated with semantic web technologies," *Information Sciences*, vol. 238, pp. 33 – 51, 2013.

[32] M. Drozdowicz, M. Ganzha, M. Paprzycki, P. Szmeja, and K. Wasielewska, "OntoPlay – a flexible user-interface for ontology-based systems." in *AT*, 2012, pp. 86–100.

[33] A. Vazhenin, Y. Watanobe, K. Hayashi, M. Drozdowicz, M. Ganzha, M. Paprzycki, K. Wasielewska, and P. Gepner, "Agent-based resource management in tsunami modeling," in *2013 Federated Conference on Computer Science and Information Systems*. IEEE, 2013, pp. 1047–1052.

[34] P. Chmiel, M. Ganzha, T. Jaworska, and M. Paprzycki, "Combining semantic technologies with a content-based image retrieval system – preliminary considerations," *AIP Conference Proceedings*, vol. 1895, no. 1, p. 100001, 2017.

[35] R. Szczekutek, M. Ganzha, M. Paprzycki, S. Fidanova, I. Lirkov, C. Badica, and M. Ivanovic, "System for semantic technology-based access management in a port terminal," *AIP Conference Proceedings*, vol. 2025, no. 1, p. 090002, 2018.

[36] M. Mun, S. Hao, N. Mishra, K. Shilton, J. Burke, D. Estrin, M. Hansen, and R. Govindan, "Personal data vaults: A locus of control for personal data streams," in *Proceedings of the 6th International COnference*, ser. Co-NEXT '10. New York, NY, USA: ACM, 2010, pp. 17:1–17:12.

[37] V. Gay and P. Leijdekkers, "Bringing health and fitness data together for connected health care: Mobile apps as enablers of interoperability," *J Med Internet Res*, vol. 17, no. 11, Nov 2015.

[38] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph, "OWL 2 web ontology language primer," *W3C recommendation*, vol. 27, no. 1, p. 123, 2009.

[39] M. Dragoni, M. Rospocher, T. Bailoni, R. Maimone, and C. Eccher, "Semantic technologies for healthy lifestyle monitoring," in *International Semantic Web Conference*. Springer, 2018, pp. 307–324.

[40] A. Dewabharata, D. M.-H. Wen, and S.-Y. Chou, "An activity ontology for context-aware health promotion application," in *Computer Software and Applications Conference Workshops (COMPSACW), 2013 IEEE 37th Annual*. IEEE, 2013, pp. 421–426.

[41] "Physical activity, health and fitness knowledge model," https://github.com/IC-FOODS/physical-activity-health-fitness, 2019.

[42] O. Sacco and A. Passant, "A privacy preference ontology (PPO) for linked data," in *Linked Data on the Web Workshop at WWW2011*, 2011.

[43] M. Hecker, "A generic privacy ontology and its applications to different domains," Ph.D. dissertation, Curtin University, 2009.

[44] O. Sacco and R. Passant, "A privacy preference ontology (ppo) for linked data," in *Linked Data on the Web Workshop at the World Wide Web Conference*, 2011.

[45] M. Palmirani, M. Martoni, A. Rossi, C. Bartolini, and L. Robaldo, "Pronto: Privacy ontology for legal reasoning," in *Electronic Government and the Information Systems Perspective*, A. Kő and E. Francesconi, Eds. Cham: Springer International Publishing, 2018, pp. 139–152.

Michał Drozdowicz obtained M.Eng. degree in Applied Computer Science from Warsaw University of Technology, Warsaw, Poland in 2007 and is currently pursuing Ph.D. at the Systems Research Institute, Polish Academy of Sciences. His research interests include semantic data processing, information privacy and distributed computing.

Maria Ganzha is an Associate Professor in the Warsaw University of Technology (Warsaw, Poland). She has an MS and a PhD degree in mathematics from the Moscow State University, Russia, and a Doctor of Science degree in Computer Science from the Polish Academy of Sciences. Maria has published close to 200 research papers, is on editorial boards of 5 journals and a book series, and was invited to program committees of more than 250 conferences.

Marcin Paprzycki is an associate professor at the Systems Research Institute, Polish Academy of Sciences. He has an MS from Adam Mickiewicz University in Poznań, Poland, a PhD from Southern Methodist University in Dallas, Texas, and a Doctor of Science from the Bulgarian Academy of Sciences. He is a Senior Member of IEEE, a Senior Member of ACM, a Senior Fulbright Lecturer, and an IEEE Computer Society Distinguished Visitor. He has contributed to more than 500 publications and was invited to the program committees of over 800 international conferences.