

Automating Consent Management Lifecycle for Electronic Healthcare Systems*

Muhammad Rizwan Asghar and Giovanni Russello

Abstract The notion of patient's consent plays a major role in granting access to medical data. In typical healthcare systems, consent is captured by a form that the patient has to fill in and sign. In e-Health systems, the paper-form consent is being replaced by access control mechanisms that regulate access to medical data while taking into account electronic content. This helps in empowering the patient with the capability of granting and revoking consent in a more effective manner. However, the process of granting and revoking consent greatly varies according to the situation in which the patient is. Our main argument is that such a level of detail is very difficult and error-prone to capture as a set of authorisation policies. In this chapter, we present ACTORS (Automatic Creation and lifecycle managementT Of authoRisation policieS), a goal-driven approach to manage consent. The main idea behind ACTORS is to leverage the goal-driven approach of Teleo-Reactive (TR) programming for managing consent that takes into account changes regarding the domains and contexts in which the patient is providing her consent.

1 Introduction

Healthcare information refers to any data containing information about an individual's health conditions. As it contains sensitive personal information, its improper disclosure may influence several aspects of an individual's life. Today, medical data is massively being converted into electronic format. Individuals' medical data can be

Muhammad Rizwan Asghar
Department of Computer Science, The University of Auckland, New Zealand
e-mail: r.asghar@auckland.ac.nz

Giovanni Russello
Department of Computer Science, The University of Auckland, New Zealand
e-mail: g.russello@auckland.ac.nz

* This chapter extends our work that appeared in the Proceedings of POLICY 2012 [1].

now easily accessible to a very large number of health-care professionals. Although this is done with the best of intentions to improve the processing and streamline healthcare delivery, it also poses very concrete threats to the individual's privacy.

Since the medical information of an individual is confidential, the only basis for accessing it is through that individual's consent. In traditional healthcare systems, an individual provided her consent by signing a paper form. In these settings, withdrawing consent was very difficult for an individual because she had to go through complicated bureaucratic processes. Moreover, the granularity of consent was very coarse-grained. The individual agreed in providing consent in advance for all her medical data, thus violating the principle of least privilege [2] – a principle that advocates for providing only legitimate access to requested resources for a limited time necessary to complete the job.

Policy-based authorisation mechanisms have successfully been used in managing access rights given the flexibility and re-usability that they offer. In literature, several approaches have been realised where the notion of consent is integrated with the policy decision mechanism. For instance, Russello *et al.* [3] propose to capture the notion of consent through the use of medical workflow and to integrate it with Ponder2 authorisation policies [4]. Ponder2 authorisation policies are represented as a (S, A, T) tuple, meaning a subject S can take action A on target T. For instance, a nurse (S) can read (A) patients' records (T). Wuyts *et al.* [5] have extended the eXtensible Access Control Markup Language (XACML) [6] authorisation model with the notion of consent. XACML is an eXtensible Markup Language (XML)-based language designed for specifying fine-grained access control. It is a standard ratified by the Organisation for the Advancement of Structured Information Standards (OASIS). XACML policies are expressed as a set of rules for regulating access to the resources. A XACML request, containing necessary information in making authorisation decision, is evaluated against XACML policies.

To specify a set of authorisation policies that capture all the details required to enforce correctly an individual's decisions about consent is very complex. First of all, each authorisation policy has conditions to express when it should be enforced that might be in conflict with other policies. Although work has been done to address the problem of automatically resolving conflicts [7], it is not possible to completely automate the decision since in the specific case of the healthcare scenario, humans are also involved. To complicate matters further, contextual information needs to be captured to identify the purpose of the access being requested. If these details are not captured correctly in the policy specification by the security administrator then there may be serious consequences.

For instance, the way in which an individual wants to provide and revoke her consent differs according to the caregivers that she is interacting with. With her General Practitioner (GP), a patient typically establishes a lasting relationship; therefore, consent can be given for a long time. On the other hand, when she is visiting a specialist in a hospital, she wants to give consent only for the time the treatment will last and only for the data that is required for the specific treatment. Still, another different situation is in the case of an emergency where the paramedics have to provide

first care before reaching the emergency room. In this case, consent can be given to the medical data however for the short period of time required to reach the hospital.

From the above scenario, it emerges that specifying in one single policy set all the requirements for managing consent is a very error-prone task. Moreover, as argued in [8] users are not engaging with their privacy tools and often prefer to ignore them. In the light of this, in this chapter we propose ACTORS (Automatic Creation and lifecycle management Of authoRisation policieS) where a goal-driven approach is used to *glue* together and manage authorisation policies that have a common aim, that is the handling of consent in a specific context (*i.e.*, consent for the GP, for the specialist, and paramedics). In particular, our observation is that we can simplify the specification of authorisation policies when these are treated as a *program sequence* towards a specific goal. The main contribution and novelty of our approach is to propose the idea of using Teleo-Reactive (TR) programs to glue together authorisation policies aiming at a specific goal. The idea of TR programs was initially introduced by Nilsson [9]. The main advantage of TR programs is that the way in which they are specified is very natural for humans. Therefore, a security administrator can capture more naturally the security requirements in a TR sequence.

The rest of this chapter is organised as follows. In Section 2, we discuss the legal aspect related to consent and set some of the terminology that will be used in the rest of this chapter. Section 3 describes an overview of a case study that we use to demonstrate the feasibility of our approach. Next, we provide a brief overview of TR Policies in Section 4. In Section 5, we present our proposed approach. In Section 6, we show how the case study scenarios can be modelled using the proposed approach. Related approaches are reviewed in Section 7. Finally, we conclude and indicate some directions for future work in Section 8.

2 Legal Background

In this section, we will discuss some of legal frameworks related to data privacy and consent. This is not intended to be an exhaustive discussion on all the legal frameworks out there. On the other hand, we feel it is necessary to put within the law perspective the technical discussion that will follow in this chapter.

2.1 Legal Framework for Consent

When dealing with people's data, the most developed countries have established legal frameworks to provide individuals with rights to allow them to make decisions regarding collection, use and disclosure of personal data. As discussed in [8], this approach can be considered as a "privacy self-management" and relies entirely on the user to take decisions and actions to either protect or disclose her data.

In the last 25 years, to deal with the growing demand and new capabilities for data collection and aggregation of digital data, a significant number of new laws have been proposed and passed in the U.S.: these include Organisation for Economic Cooperation and Development (OECD) Privacy Guidelines in 1980, the Asia-Pacific Economic Cooperation (APEC) Privacy Framework in 2004, and more recently in 2012 the Federal Trade Commission (FTC) and the White House issued major new frameworks for protecting privacy. All these efforts have in common a set of principles for protecting privacy that was first proposed in the Fair Information Practice Principles (FIPPs) as a report by the U.S. Department of Health, Education, and Welfare in 1973 [10]. The FIPPs include several guidelines such as 1) transparency of the record system of personal data; 2) the right to notice; 3) the right to prevent the use of personal data for new purposes without consent; 4) the right to correct/amend one's record; and 5) responsibility of the data holder for protecting data from misuse.

The privacy law framework in Canada is also based on the OECD Guidelines proposed in 1980 and relies on consent for collection, use and disclosure of personal data. However, as discussed in [11], Canada's legislation on data handling and processing makes a clear separation on role of consent between public and private sectors. In particular, in the public sector, consent is seen as a *justification* whether in the private sector consent is a *requirement* for collecting, using and disclosing data. The Australian Privacy Act provides general guidelines for collecting and processing personal data that are also based on consent. However, for larger commercial entities (with annual income greater than 3 million AUD), there is also an extra burden to destroy personal data once it has been used as intended at collection time. For instance, if a customer's address was collected because of the delivery of goods, once the goods are delivered then the information should be destroyed. New Zealand has a more relaxed approach when it comes to collection of data for commercial purpose. For instance, an entity can collect information about an individual as long as the individual has been informed about the collection and the purpose for which the data has been collected. More interestingly, an entity does not need to inform again an individual if the same type of data is collected again after the person has been correctly informed the first time. However, in New Zealand, user consent for collecting, using, and disclosing personal data is required in specific sectors related to healthcare, telecommunications and credit records.

Compared to the frameworks of the countries above, the EU directives have a more paternalistic approach when it comes to data processing, as discussed in [12]. According to article 2(h) of the EU Data Protection Directive (DPD) [13], consent is defined as: "*the data subject's consent shall mean any freely given specific and informed indication of his wishes by which a data subject signifies his agreement to personal data relating to him being processed*", where the term *data subject* describes an individual whose data is handled and *data controller* indicates any entity that handles personal data.

The EU law framework also supports the concept of privacy self-management. According to article 7 (a) of the EU DPD [13], a data subject's personal data may only be processed if she has given her consent. However, the way in which a data

handler seeks the data subject's consent is much more regulated. Furthermore, data processing in the EU is always controlled through a legal framework; whereas, in the US, data processing is always granted by default unless explicitly forbidden by the law.

2.2 Consent in Healthcare System

In the context of healthcare systems, consent indicates agreement of the patient on sharing her personal health information [14]. In traditional healthcare systems, a data subject provides her paper-based consent typically once she is enrolled within the system. Generally, the paper-based consent is considered valid once signed by the data subject. Unfortunately, there are two main problems with the paper-based consent. First, it becomes very cumbersome for the data subject to withdraw her paper-based consent. That is, she has to go through complicated bureaucratic processes where she has to call on the responsible authority to withdraw her consent with some considerable effort, waste of time and a huge sense of frustration. Second, a data subject provides her consent in advance for all her medical data at the time of registration with the healthcare system even when it may not be necessarily used, thus violating the principle of least privilege.

With the introduction of electronic healthcare systems, we have moved from the paper-based consent to the *electronic consent*, or *e-consent* in short. e-consent has been established as a new industry standard [15] and aims at replacing the traditional paper-based control, thus providing more control to patients for controlling the way they share their Electronic Health Records (EHR).

In current IT healthcare systems, the notion of e-consent is captured as *authorisation policies* that control the access to the data, such as in [3]. Technically, the creation or editing of these authorisation policies is delegated to an IT security administrator. The security administrator operates on behalf of the data subject to deploy policies in the IT infrastructure of the data controller. In some countries, specific legislation may require the digital consent to be digitally signed by the data subject to be considered equivalent to the manually signed paper-based consent [16].

Using the classification proposed in [17], it is possible to identify the following essential elements of e-consent:

- **Requester:** An entity to whom the authorisation is provided. It could be a person, a role or even an organisation.
- **Actions:** These are the set of rights that are authorised by the consent.
- **Purpose:** A purpose is a reason for which the authorisation is given.
- **Validity:** It is a time period in which the authorisation is applicable.
- **Revocation:** This is a feature of consent using which one can revoke her consent.
- **Delegation:** Delegation is an authority given to someone who can manage consent on behalf of someone else.

In the following discussion, we will use the term consent to indicate in general electronic consent. We can identify two categories of consent [14, 18]: implicit and explicit consent. Implicit (*a.k.a.* implied) consent is one that is inferred from the actions. Explicit consent is one that is given explicitly. There is another type of consent called informed consent. In the context of healthcare systems, informed consent requires patients to be informed about what they are going to agree with. In [18], Coiera and Clarke list the following four forms of consent:

- **General Consent:** This is one time consent given by a patient to any medical professional, for any purpose and valid as long as it is not revoked by the patient. This form ensures ease of use but might hamper protection due to open access.
- **General Consent with Specific Denials:** The patient provides a general consent except some specific conditions that could be based on expressive policies *e.g.*, based on time, purpose and/or validity. From the security point of view, it improves the general consent form but reduces availability.
- **General Denial with Specific Consent(s):** It is opposite of the previous form. In this case, the patient provides a general denial except specific conditions under which she would like to give her consent. This consent type provides reasonable control as well as restricted availability.
- **General Denial:** It is opposite of the first form. In this case, a patient denies access to her personal information.

There are two major factors affecting the evaluation criteria for the consent. First, it is the *ease of use*, meaning how easy it is to access and use the consent. The second one is *privacy*, which means the level of protection offered by consent. If we evaluate the above four forms of consent under this evaluation criteria then the general denial provides better protection of privacy and this level decreases as we choose other forms (from bottom to up) of consent. Thus, the protection of privacy provided by the general consent is at the lowest level among all other forms. However, the general consent ensures the ease of access. The ease of access decreases as we move down from the general consent to the general denial.

It is important to know that consent is required for providing access to medical data that is not anonymised. However, consent might not be required when the patients' data is first anonymised and then shared, given the data anonymisation technique can guarantee privacy of the patients.

2.3 Consent Limitations

In practice, these law frameworks rely on the data subjects to make decisions on whether it is beneficial to them to consent access and usage of their data: consent legitimises any collection, use and disclosure of someone's data.

There are several limitations with this approach. First of all, there are cognitive problems with privacy self-management. As demonstrated by several empirical studies in social science research, people do not engage with privacy self-

management: the main reason is that they do not read the terms and conditions notices [19] or when they read them they do not understand them [20]. However, even when people read and understand the notices they are not able to take rational and informed decisions on the costs and benefits on consenting access to their data [21].

Another issue is more related to the scale of the problem. Assuming that people had a complete understanding of the risks involved in consenting access to their personal data, there are far too many entities collecting data to make self-managing privacy practically possible. One study has estimated that the cost associated with the lost productivity if each of us were to read the terms and conditions notices of each website we visit on a given year to a staggering \$781 billion [22]. To complicate matters even further, each entity quite frequently changes privacy policies, which would require further engagement from the user.

Third, the major harm to one's privacy comes from the aggregation of data collected by different parties over a period of time. It is almost impossible for a user to be able to understand the risks and benefits at the time of the release of a piece of information without a proper knowledge of how the data will be used in aggregation with other information. For instance, several entities that have received consent to use the data subject's data, that in isolation and at a given point in time are not harmful, could decide later to collaborate or aggregate their data resulting in a violation of the data subject's privacy. How can a data subject be able to predict such an event at the time the consent is given? Privacy regulations aim mainly at protecting one's Personally Identifiable Information (PII). However, PII is not a static label that can be associate to a piece of data for its entire life cycle. With the huge amount of facts that we leave in our digital trail online and the advancements of data mining technologies, identifying someone is becoming very easy from data that taken in isolation is pretty harmless [23]. The result is that with data mining and aggregation, it is nearly impossible to be able to manage one's personal information.

Another negative aspect of privacy self-management is that it is always considered as an isolated transaction between an individual and an entity. However, some aspects of the privacy have an effect not only on the individual but to a society as a whole. The decisions taken by an individual for consenting collection, usage and disclosure of her data might not have the most desirable effect on a larger scale. On the other hand, sometime overriding one's privacy can be beneficial for the protection or advantages of our society. For instance, as discussed in [24], the use of data analytics can lead to better medical treatments as well as better responses to data breaches. Privacy self-management fails to address the global outcomes on a social level, focusing only on the single individual and on isolated transactions. Last but not least, data subjects may withdraw their consent at any time.

Considering the criticisms above, one would be tempted to abandon consent as a mean to safeguard one's privacy. However, we argue that controlling consent could be improved by bringing the tools to access and manage consent closer to the data subject, such as to her mobile device. Also, instead of doing consent micro management by presenting endless requests of binary consent decisions, we want to provide to the data subject with an approach that takes into account her goals when it comes

to protect her privacy and that will learn from the decisions the data subject takes in a particular context.

3 A Case Study

In this section, we introduce the case study that we will use throughout the paper to demonstrate the feasibility of our approach. The case study is partially inspired from the European funded projects including ENDORSE [25] and EnCoRe [26]. Both projects focus on developing IT solutions for privacy preserving data management, where *consent* is one of the main point of focus.

In this section, we describe several scenarios based on the IT healthcare system currently deployed in one of the major hospitals in Italy. We assume that each patient has a smartphone that she uses to receive requests for giving her consent when she is interacting with the medical personnel. A patient can review through her smartphone who is requesting the access, the purpose of the request, and which data is requested.

At the time of providing consent, a patient may decide to save her preferences for subsequent consent requests made in the same context and/or by the same entity. Afterwards, a patient may withdraw her saved preferences regarding consent. Furthermore, a patient may activate withdrawn preferences regarding her consent. Last but not least, a patient may intend to delete, forever, her preferences, initially saved for providing consent automatically.

Patient visiting her GP. Let us consider the healthcare scenario where Alice moves to Milan and visits her GP for the first time. The GP requires access to Alice's medical history consisting of several medical tests and reports. For this purpose, the GP requires Alice's consent. Alice receives the consent request on her smartphone and decides to provide her consent also in the future.

Patient visiting a cardiologist. Later, the GP of Alice discovers that she has a heart disorder. In this case, the GP refers Alice to a cardiologist for further testing. For visiting the cardiologist, Alice needs to contact the hospital booking service for getting an appointment. The hospital has several cardiologists; thus, it is not known in advance which one is assigned prior to the actual appointment. On the day of appointment, Alice will know the assigned cardiologist and can consent the cardiologist to access her medical data. However, Alice's consent should be valid for the duration of the treatment and the data accessed should be within the scope of the treatment (*i.e.*, the cardiologist should not have access to Alice's gynaecological reports). Moreover, if Alice is not happy with the assigned cardiologist then she may withdraw her consent and request a new cardiologist.

Patient in an emergency situation. While Alice is driving in her car, she has a car accident and gets injured. The emergency response team reaches the accident location and starts treating Alice. For the treatment, the paramedic requires Alice's consent to access her medical history to get information about her allergies and any serious conditions that she already may have. Alice provides consent to access her medical records so that the paramedic is aware of her heart problem and provides

the appropriate treatment that does not interfere with the treatment prescribed by the cardiologist. Although the paramedic has access to Alice's full medical record, consent should be revoked when the emergency is over.

4 Overview of Teleo-Reactive Policies

From the above scenarios, it is clear that to capture all the details required to express the data subject's consent in different settings is very complex. If these details are not captured correctly by the security administrator in the policy specification then serious consequences might happen. In our experience, capturing all the security requirements through the specification of several independent authorisation policies is a very hard task. In the specific case of capturing a data subject's consent, it becomes even more complicated since there is the involvement of a human (which is the data subject that can grant, hold and withdraw consent) and contextual information expressed in the policies (such as the location and time of the access).

In this chapter, we propose to employ a goal-driven approach to *glue* together and manage authorisation policies that have a common aim, *i.e.*, the handling of consent. In particular, our observation is that we can simplify the specification of authorisation policies when these are treated as a *program sequence* towards a specific goal. In this chapter, we propose to leverage the idea of TR programs to glue together authorisation policies aiming at a specific goal. The idea of TR programs was initially introduced by Nilsson [9]. A TR program is a control sequence directing towards a goal while taking into account changes in environmental circumstances. TR programs were used for automating behavioural robotics where a robot was continuously observing its environmental changes.

In the following, we provide a brief overview of TR policies that is similar to one introduced by Marinovic *et al.* in [27].

```

1 tr-policy name( $P_1, P_2, \dots, P_m$ )
2  $cond_1(V) \rightarrow action_1(V)$ 
3  $cond_{2_a} \wedge (cond_{2_b} \vee \neg cond_{2_c}) \rightarrow action_{2_y} \otimes action_{2_z}$ 
4  $cond_3(P_1) \rightarrow action_{3_a} \parallel action_{3_b}$ 
5 ...
6  $cond_{n_1} \wedge cond_{n_2} \dots \vee cond_{n_x} \rightarrow action_{n_1} \parallel action_{n_2} \dots \otimes action_{n_y}$ 

```

Fig. 1 A layout of TR policies.

4.1 TR Policy Representation

A TR policy is an ordered list of rules as shown in Figure 1, where each rule contains (Line 2) a condition part and an action part. The condition part contains a predicate that is bound with a variable, which is denoted with V . These variables may describe facts or states of the system or environment in which a TR policy is evaluated. A variable starts with a capital letter while a condition or an action starts with a small letter. The action part contains a function that is called by the TR policy. The action part may contain variables. The condition and action parts are separated by \rightarrow . Each TR-policy has a name starting with a small letter and can be instantiated with some parameters (Line 1). The condition part may include parameters, each denoted by P_i (Line 4). The condition part can contain either a single condition or form (Line 2 or Line 6) a conditional expression where multiple conditions can be combined using logical operators \wedge and \vee . Similarly, the action part can contain either a single function or multiple functions that may be executed sequentially and/or concurrently. The sequential and concurrent execution of functions can be represented with \otimes operator (Line 3 and Line 6) and \parallel operator (Line 4 and Line 6), respectively. In a TR policy, rules are specified in the descending order with respect to their priorities. That is, a high priority rule comes first.

```

1 tr-policy superStore(E)
2
3 isStoreCrowded  $\wedge$  isAvailable(CC)  $\rightarrow$  serverAtCheckoutCounter(E,CC)
4
5 askedForHelp(E,C)  $\rightarrow$  helpCustomer(E,C)
6
7 isShelfEmpty(S)  $\rightarrow$  stackShelf(E,S)

```

Fig. 2 An example of a TR policy.

In Figure 2, we have illustrated an example of a TR policy representing job specification of an employee E who works at a superstore. For simplicity, we mainly consider three job responsibilities: serving at checkout counters, helping customers and stacking shelves. The top priority will be given to serving at a checkout counter. An employee E can server at a checkout counter CC when each occupied checkout counter is crowded by a large number of customers. Of course, a checkout counter CC should be available before an employee E can start serving. This job responsibility is specified as the first rule at Line 3. The next priority will be given to helping any customer C , meaning if an employee E is asked for any help then she should help the customer C . The second rule (Line 5) in the TR policy represents this job responsibility. The lowest priority job responsibility is stacking a shelf S if it is empty (or about to empty) – see the last rule (Line 7) in the TR policy.

4.2 TR Policy Evaluation

The runtime of the TR policy monitors changes in facts or states about the system or environment in which evaluation is performed. These changes can result in the condition part of a rule becoming either *true* or *false*. The functions in the action part of a rule will be executed if its condition part is evaluated to *true* by the runtime. In a TR policy, the condition part corresponding to the highest priority rule is evaluated first. If it evaluates to *false*, the condition part of the next high priority rule will be evaluated. In other words, if the action part of any rule is being executed, it means the condition parts of all higher priority rules (as compared to the current rule) are evaluated to *false*. The action part of any rule is executed as long as its condition part evaluates to *true* while condition parts of all higher priority rules (as compared to the current rule) remain *false*.

As an example, let's discuss evaluation of the TR policy illustrated in Figure 2. This TR policy consists of three rules listed in the order of priority (from higher to lower). The runtime of the TR policy checks if the superstore is crowded, *i.e.*, *isStoreCrowded*. Then, it identifies whether any checkout counter is available, *i.e.*, *isAvailable(CC)*. If both conditions are met (*i.e.*, they evaluate to *true*), the first rule is fired and the employee E starts serving, *i.e.*, *serverAtCheckoutCounter(E, CC)*. However, if the superstore is not crowded (*i.e.*, *isStoreCrowded* evaluates to *false*) or it is crowded but no checkout counter is available (*i.e.*, *isAvailable(CC)* evaluates to *false*) then the runtime will start evaluating conditions of next rules in the TR policy. In our case, the next rule is to monitor if the employee E is asked for help by any customer C (*i.e.*, *askedForHelp(E, C)*). If so (*i.e.*, *askedForHelp(E, C)* evaluates to *true*), the employee will start serving the customer (*i.e.*, *helpCustomer(E, C)*). Otherwise, the runtime will evaluate condition of the last rule, *i.e.*, if a shelf is empty. It will evaluate *isShelfEmpty(S)*. If it is *true*, the employee E will start stacking the shelf S (*i.e.*, *stackShelf(E, S)*). As soon as execution of any rule is completed, the runtime will start evaluating condition of the first rule in the TR policy (*i.e.*, serving at a checkout counter if the superstore is crowded).

5 The ACTORS Approach

ACTORS aims at automating creation and management of consent related authorisation policies using a goal-driven approach. Figure 3 illustrates the ACTORS architecture. There are three main system entities:

- **Data Subject:** Data subjects represent end-users and are key entities, responsible for granting, updating or withdrawing their consent.
- **Data Requester:** A data requester is an entity who makes a consent request. In the context of healthcare systems, this entity could be a doctor or a GP.

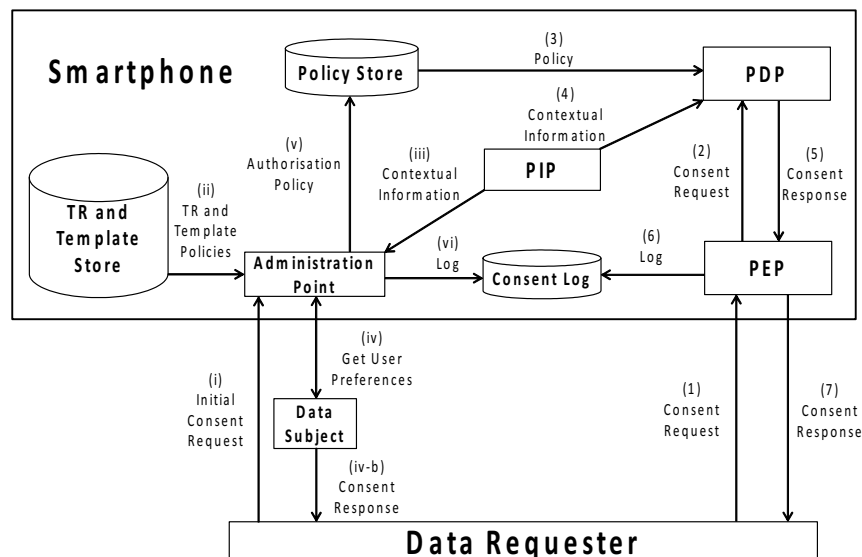


Fig. 3 The ACTORS architecture for managing consent lifecycle.

- **Smartphone:** It is a mobile device that is possibly owned or at least operated by data subjects. It automatically manages lifecycle of consent authorisation policies.

In Figure 3, an initial consent request is issued by the data subject as we can see in Step (i). The Administration Point, managed by the data subject's smartphone, receives the request and fetches corresponding TR and template policies from the TR and template store in Step (ii). The main idea is that each TR policy captures a specific goal, such as managing consent for the GP. TR policies are used for instantiating authorisation policies from a standard set of policy templates. TR policies also manage the lifecycle of instantiated authorisation policies.

Since all the details required in an authorisation policy may not be known in advance (such as, ID of the specific cardiologist assigned on the day of the visit, location where the visit will take place), we use policy templates to define abstract authorisation policies. When all the required information is available, TR policies can instantiate the required authorisation policies from the given templates. This instantiated authorisation policy is stored and enforced by the data subject's smartphone, thus providing greater control to data subjects to manage their consent.

Authorisation policies are created and managed based on the data subject's intent while taking into account contextual information retrieved from the Policy Information Point (PIP) in Step (iii). The contextual information may be information about

facts or states of the environment or the system. For collecting contextual information in an automated manner, we assume that data subjects have smartphones equipped with some sensors for capturing environmental conditions. For instance, a smartphone can detect a fire alarm or an emergency situation such as a road accident. After collecting contextual information, an authorisation policy is populated and a data subject is asked about saving her preferences – in Step (iv) – for saving such an authorisation policy in order to authorise subsequent consent requests in an automated manner. Next, a consent response is sent to the data requester in Step (iv-b). Next, an authorisation policy is stored in the policy store in Step (v). Finally, this instantiation of authorisation policy is logged in Step (vi).

After an authorisation policy has been instantiated, any subsequent consent request will be received by the Policy Enforcement Point (PEP) running on the data subject's smartphone as we can see in Step (1). The PEP forwards this request to the Policy Decision Point (PDP) in Step (2). The PDP is responsible for fetching corresponding authorisation policies and collecting contextual information as we can see in Steps (3) and (4), respectively. Next, the PDP makes the decision by evaluating authorisation policies against the request and contextual information provided. Then, it sends the consent response to the PEP in Step (5). The PEP also logs the decision made by the PDP in Step (6) and finally the consent response is sent to the data requester in Step (7).

We assume that the PEP and the administration point have access to the data subject's signing key that could be used for signing consent responses sent to the data requester. The consent log captures the complete details of actions taken by the data subject and decisions (*e.g.*, signed consent responses) automatically made by the smartphones based on data subjects' intent. A data subject has full access to her consent log.

The data subject has a right to update her consent policy, withdraw her consent by deactivating the consent policy or delete the policy altogether. In all these cases, data subjects have to interact with the administration point for any modification. Our architecture is flexible enough to cope with updates in the workflow of the healthcare providers or even in the law. All the healthcare providers have to do is to update TR and template policies stored in the TR and template store and delete, if any, existing authorisation policies managed by the policy store.

It is important to mention that data requesters can get access to the data in case of emergency using the *break glass* policy. In this case, the healthcare system could expect evidence of being in emergency situation. The healthcare system could defer verification of such evidence for the post-incident investigation.

5.1 Authorisation Policies

An authorisation policy specifies who is permitted (or denied) access to a resource under specific conditions. In ACTORS, an authorisation policy contains the following fields:

- **Data Requester Role:** It is role of the entity who makes the access request. It can contain either a single role or a set of roles.
- **Data Requester ID:** It is ID of the one who makes the access request. Like the above field, this field can contain either a single ID or a list of IDs. This field is optional as permissions can be assigned to roles instead of specific IDs.
- **Data Subject ID:** It refers to the data subject who owns the resources.
- **Data Subject Resource:** It contains data subject resource(s) protected through the authorisation policy.
- **Access Rights:** Access rights define the permission on the data subject resource.
- **provided:** It contains a conditional expression that may contain a set of conditions combined with *and* and *or* logical operators. Each condition is a predicate that is bound to a variable. These variables can come from contextual information that may be facts or states about the system or the environment. The contextual information may include access purpose, access time, access date, data requester location and data subject location.

```

1 DataRequester.Role = { 'Doctor' }
2 DataRequester.ID = { 'Bob' }
3 DataSubject.ID = 'Alice'
4 DataSubject.Resource = { 'Blood Test' }
5 AccessRights = { READ }
6 provided
7   (AccessPurpose = 'Diagnosis' or
8    AccessPurpose = 'Treatment') and
9    AccessTime >= 9:00

```

Fig. 4 An example of an authorisation policy.

Figure 4 illustrates an example of an authorisation policy where Bob in a role doctor is permitted to have read access on Alice's *Blood Test* report provided he makes the access request after 9:00 hrs for the purpose of diagnosis or treatment. The use of the Data Requester ID might seem redundant given the fact that the policy already has a Data Requester Role. However, it might be the case that the data subject might not want a specific requester to access her data. For instance, Alice does not want Eve (another doctor and Bob's colleague) to read her *Blood Test* report. This requirement can be captured by specifying in the Data Requester ID the condition $\neg 'Eve'$. The introduction of both positive and negative conditions could result in conflicting authorisation policies. For resolving conflicts, we can use existing resolution techniques, such as one proposed in [7].

We assume that once authorised as per authorisation policy, a data requester can access medical data of the patient for a certain time (say for the duration of the appointment). Once a time limit is reached, the data requester would not be able to access the data anymore.

5.2 Policy Templates

A policy template provides a structured format for instantiating authorisation policies on-the-fly. It is the authorisation policy specification with placeholders for variables that are assigned a value based on contextual information and a data subject's intent. A data subject's intent is about what a data subject can expect and can be captured based on actions taken by her. A policy template contains almost the same fields as an authorisation policy does. The fields of a very generic policy template are left blank so that they can be assigned a value based on contextual information. However, a list of options can be provided for each field. It means that a template field can only be filled, at the time of policy instantiation, with a value out of the list of options.

```
1 DataRequester.Role = {'Dentist'}
2 DataRequester.ID
3 DataSubject.ID
4 DataSubject.Resource = {'Dental Report'}
5 AccessRights = {READ, WRITE}
6 provided
7   AccessPurpose is 'Diagnosis' or 'Treatment'
```

Fig. 5 An example of a policy template.

Figure 5 illustrates an example of a policy template. This policy template can be applied when a data requester is in role *Dentist* and the requested resource is *Dental Report* with access rights either *READ* or *WRITE* access and access purpose is either *Diagnosis* or *Treatment*. For rest of the fields, any value can be assigned based on contextual information and the data subject's intent.

Generally, specifying an authorisation policy is difficult. However, policy templates, which could be provided by healthcare providers, make it easy for patients to instantiate required authorisation policies. For instantiation of authorisation policies from policy templates, a patient should be provided with usable interfaces with simple privacy controls. These privacy controls will lead to automatic generation of authorisation policies. Without loss of generality, our proposed architecture enables data subjects to update existing authorisation policies or create new ones.

Policy templates are associated with TR policies and goals that the TR policy is trying to achieve. For instance, the policy template in Figure 5 can be applied when the goal of the patient is to visit a dentist. Therefore, such a template is associated with the TR policy managing that specific goal. Each TR policy can be associated with several templates. Based on contextual information and a data subject's intent, the TR policy can identify which policy template fulfils the criteria and then instantiates the required authorisation policy.

5.3 TR Policies

As already explained in Section 4, TR programs were introduced for continuously monitoring the behaviour of a robot while taking into account environmental changes. In ACTORS, we use TR policies for controlling the lifecycle of authorisation policies towards a specific goal, which is the management of data subject's consent in a given situation. Each TR policy might be associated with several policy templates from which authorisation policies can be instantiated. Several TR policies might be present on the data subject's smartphone. The selection of the appropriate TR policy is based on contextual information. The main advantage in using TR policies is that they provide a built-in prioritisation of actions needed for controlling the granting and revocation of data subjects' consent that reacts to the changes in the context in which the data subjects are interacting.

In the following section, we are going to provide details of how ACTORS can be used for the case study presented in Section 3.

6 Managing Consent in Healthcare Scenarios

ACTORS can be applied to any domain; however, we focus on healthcare scenarios as already described in Section 3, where consent needs to be captured and saved based on contextual information and the patient's intent

Note 1. in this context we assume that the patient is the data subject. For automatically instantiating authorisation policies regarding consent and managing lifecycle of those policies, we assume that each patient is provided a set of TR policies and policy templates at the time of registration with her healthcare provider. In fact, TR policies and policy templates are deployed on patients' smartphone together with an application. Each TR policy can be associated with multiple policy templates. The smartphone application automatically selects the most appropriate TR policy and the policy template based on the consent request and contextual information. After instantiation of authorisation policies regarding consent, they are stored and enforced by the patient's smartphone. It should be noted here that only policies and patient's decisions are stored in the smartphone while the medical data is stored in the caregiver IT infrastructure. In this section, we explain in detail how we exploit the proposed approach, described in Section 5, for providing solutions for each scenario described in Section 3.

Patient visiting her GP. In the scenario when a GP needs the patient consent, a consent request is sent to the patient for providing access to a GP to requested resources. This consent request may be directly sent by the healthcare system to the patient when a GP makes an access request to the patient resources. This consent request may include information about the GP and the patient, the patient resources, an access purpose and access duration details. Based on the consent request together

with contextual information, the most appropriate applicable TR policy and policy template are selected.

```

1 tr-policy consentAtGPClinic(Patient)
2
3 consentAvailable(Patient,GP) ^ saveCurrentPreferences → instantiatePolicy(Patient) ⊗
   activate(Patient.Policy) || sendConsent(Patient,GP)
4
5 consentAvailable(Patient,GP) → sendConsent(Patient,GP)
6
7 needsConsent(Patient,GP) ^ instantiatedPolicy(Patient) ^ ¬withdrawn(Patient.Policy) →
   evaluatePolicy(Patient)
8
9 needsConsent(Patient,GP) → waitPatientDecision(Patient,GP)
10
11 deleteSavedPreferences(Patient) → remove(Patient.Policy)
12
13 activatePolicyRequest(Patient) → activate(Patient.Policy)
14
15 withdrawPolicyRequest(Patient) → withdraw(Patient.Policy)

```

Fig. 6 A TR policy for managing authorisation policy for providing consent to a GP.

Figure 6 describes a TR policy that is applied when a GP needs a patient's consent for accessing her data from his clinic. The name of this TR policy is *consentAtGPClinic* and *Patient* is the parameter. When the first consent request is made, consent is not available and the condition parts of rules at Line 3 and Line 5 evaluate to *false*. The condition part of rule at Line 7 also evaluates to *false* as no authorisation policy is instantiated yet, *i.e.*, *instantiatedPolicy(Patient)* is *false*. However, the condition part of rule at Line 9 evaluates to *true*, so the action part of this rule is executed and the system waits for the patient decision for providing consent to her GP, *i.e.*, *waitPatientDecision(Patient,GP)* is executed.

Once the patient provides consent for granting access to her GP on her resources, then *consentAvailable(Patient,GP)* becomes *true*. At the time of providing consent, a patient can be given an option to save her current preferences for providing her consent for similar consent requests when made in the same environment. If a patient does so, the condition part of rule at Line 3 becomes *true*; therefore, the authorisation policy regarding consent is instantiated from the policy template and then it is activated while at the same time, consent is sent.

Figure 7 illustrates a policy template that is applied when a patient visits her GP, as is evident from the data requester role that is GP only. The empty fields including data requester name, data subject name, data subject resource and access rights can be filled with values based on the consent request. However, there are certain conditions in the *provided* part of the policy template that are formulated at the time of instantiating an authorisation policy. These conditions include: the access purpose must be either *diagnosis* or *treatment*; access time must be in office hours; and both the patient and the GP must be present in the GP's clinic. These conditions are formulated based on contextual information that is collected from either patient's smartphone or the external information point, such as made available by the health-

```

1 DataRequester.Role = {'GP'}
2 DataRequester.Name
3 DataSubject.Name
4 DataSubject.Resource
5 AccessRights
6 provided
7   AccessPurpose is 'Diagnosis' or 'Treatment'
8   AccessTime is within DutyHours
9   DataRequester.CurrentLocation = DataSubject.CurrentLocation
10  DataRequester.CurrentLocation = DataRequester.Clinic.Location

```

Fig. 7 A policy template for generating an authorisation policy for providing consent to a GP.

care provider. The contextual information from a patient's smartphone may include information like patient's current location, while contextual information from the external information point may include information about location of GP's clinic and GP's duty hours. Once all the required information for the applicable policy template is retrieved, the authorisation policy is instantiated and activated.

```

1 DataRequester.Role = {'GP'}
2 DataRequester.ID = {'Bob'}
3 DataSubject.ID = 'Alice'
4 DataSubject.Resource = {'Blood Test'}
5 AccessRights = {READ}
6 provided
7   AccessPurpose = 'Diagnosis' and
8   (AccessTime ≥ 9:00 and AccessTime ≤ 17:00) and
9   DataSubject.CurrentLocation = 'Milan' and
10  DataRequester.CurrentLocation = 'Milan'

```

Fig. 8 An authorisation policy for providing consent to a GP.

Figure 8 shows the instantiated authorisation policy regarding consent, expressing that a GP Bob can get patient Alice's consent for *READ* access on Alice's *Blood Test* when accessed for the *Diagnosis* purpose during the duty hours (that is, between 9:00 and 17:00 hrs) from Bob's clinic located in *Milan*.

A patient may decide to withdraw her consent. In this case, the condition part of rule at Line 15, *i.e.*, condition *withdrawPolicyRequest(Patient)*, becomes *true* and the authorisation policy is withdrawn by invoking *withdraw(Patient.Policy)* function. Furthermore, a patient can decide to activate her withdrawn consent. In this case, condition *activatePolicyRequest(Patient)* becomes *true* and *activate(Patient.Policy)* function is invoked for activating the authorisation policy. Last but not least, a patient may also choose to delete forever her saved preferences for automatically providing consent. In this case, *deleteSavedPreferences(Patient)* becomes *true* and *remove(Patient.Policy)* function is invoked for deleting the instantiated authorisation policy.

In case if a GP needs the patient consent when the patient has already saved preferences for providing consent automatically to her GP and consent is not withdrawn yet then consent will be provided after evaluating the consent request and contextual

information against the instantiated authorisation policy, see rule in Figure 6 at Line 7. We assume that the consent request is same as already described above. However, we have to collect contextual information in order to evaluate the authorisation policy for providing consent. The patient's smartphone may provide information about her location and the current time while the information about the GP's location can be collected from the external information point. This may be the healthcare system or the GP's smartphone which may provide GP's location information to the patient's smartphone. Based on the consent request and contextual information, the authorisation policy is evaluated (see rule in Figure 6 at Line 7). After the evaluation of the authorisation policy, consent becomes available and the consent response is automatically sent by the patient's smartphone (see rule in Figure 6 at Line 5). The consent response contains patient consent if the authorisation policy evaluates to *true*, otherwise it may contain an error message.

A patient may decide not to save her current preferences for providing consent automatically to her GP. In such a case, the patient will be explicitly asked each time (see rule in Figure 6 at Line 9) and consent will be provided once the patient takes her decision (see rule in Figure 6 at Line 5).

Patient visiting a cardiologist. A cardiologist may also need the patient consent while accessing the patient resources. Like the above scenario, a patient receives the consent request. This consent request may include information about the cardiologist and the patient, the patient resources, an access purpose and access duration details. The additional point in this scenario as compared to the previous scenario is that a cardiologist is provided consent for getting access on the patient resources as long as the treatment may last. In other words, the saved preferences for providing consent are deleted automatically right after the treatment.

```

1 tr-policy consentAtSpecialistClinic(Patient)
2
3 consentAvailable(Patient,Specialist) ∧ saveCurrentPreferences → instantiatePolicy(Patient) ⊗
   activate(Patient.Policy) || sendConsent(Patient,Specialist)
4
5 consentAvailable(Patient,Specialist) → sendConsent(Patient,Specialist)
6
7 needsConsent(Patient,Specialist) ∧ instantiatedPolicy(Patient) ∧ ¬withdrawn(Patient.Policy) →
   evaluatePolicy(Patient)
8
9 needsConsent(Patient,Specialist) → waitPatientDecision(Patient,Specialist)
10
11 timeout(Patient.Policy) ∨ deleteSavedPreferences(Patient) → remove(Patient.Policy)
12
13 activatePolicyRequest(Patient) → activate(Patient.Policy)
14
15 withdrawPolicyRequest(Patient) → withdraw(Patient.Policy)

```

Fig. 9 A TR policy for providing consent to a specialist.

Figure 9 shows the TR policy for managing authorisation policy in order to provide consent to a specialist. The name of this TR policy is *consentAtSpecialistClinic*. The TR policy is similar to one already described in Figure 6. In case of a cardiolo-

```

1 DataRequester.Role = {'Cardiologist'}
2 DataRequester.ID
3 DataSubject.ID
4 DataSubject.Resource = {'ECG Report', 'Cardiography', 'Engyography'}
5 AccessRights = {READ, WRITE}
6 provided
7     AccessPurpose is 'Diagnosis' or 'Treatment'
8     AccessTime is within DutyHours
9     DataRequester.CurrentLocation = DataSubject.CurrentLocation
10    DataRequester.CurrentLocation = DataRequester.Clinic.Location

```

Fig. 10 A policy template for generating an authorisation policy for providing consent to a cardiologist.

gist, the TR policy of specialist is selected. As we can observe that the TR policy of a specialist is very generic, it can be applied to other specialists such as a dentist and a gynaecologist. However, there is a specific policy template for each specialist. The policy template for cardiologist is shown in Figure 10. The policy template is restricted to only resources that could be accessed by a cardiologist. These resources include *ECG Report*, *Cardiography* and *Engyography*. This is different from the policy template of above scenario as resource field in Figure 7 is left empty, indicating that a GP can obtain consent to access any resource.

```

1 DataRequester.Role = {'Cardiologist'}
2 DataRequester.ID = {'David'}
3 DataSubject.ID = 'Alice'
4 DataSubject.Resource = {'ECG Report'}
5 AccessRights = {READ, WRITE}
6 provided
7     AccessPurpose = 'Diagnosis' and
8     (AccessTime ≥ 9:00 and AccessTime ≤ 17:00) and
9     DataSubject.CurrentLocation = 'Como' and
10    DataRequester.CurrentLocation = 'Como'

```

Fig. 11 An authorisation policy for providing consent to a cardiologist.

Figure 11 shows the authorisation policy regarding consent for a cardiologist when a patient intends to save her preferences until she is treated. The authorisation policy expresses that a cardiologist David can get patient Alice's consent for *READ* and *WRITE* access on Alice's *ECG Report* when accessed for *Diagnosis* purpose during the duty hours (that is, between 9:00 and 17:00 hrs) from David's clinic located in *Como*.

The authorisation policy regarding consent for a cardiologist may automatically be deleted once the treatment completes. This information about treatment duration can be collected by the patient at the time of saving her preferences. For instance, it may be included in the consent request or can be collected as contextual information from the information point made available by the service provider. Once the treatment duration expires (starting from when the first consent request is made), condition *timeout(Patient.Policy)* becomes automatically *true* and

remove(Patient.Policy) function is invoked for deleting the instantiated authorisation policy according to the rule at Line 11 in Figure 9. Alternatively, a patient may decide to delete her saved preferences during the treatment duration as already considered in above scenario.

Patient in an emergency situation. In an emergency situation, the emergency response team may need a patient's consent in order to get an access to her medical data for the treatment purpose. Similar to above scenarios, the patient receives the consent request, which may include information about the emergency response team, the patient resources, an access purpose and access duration details. Similar to the cardiologist scenario, we consider that the patient intends to provide her consent as long as the treatment may last. Technically, the saved preferences for providing consent are deleted automatically right after the treatment. The TR policy for specialist, shown in Figure 9, can also be applied for this scenario.

```

1 DataRequester.Role = {'EmergencyResponseTeam'}
2 DataRequester.Name
3 DataSubject.Name
4 DataSubject.Resource = {'Allergy Report', 'Blood Test'}
5 AccessRights = {READ}
6 provided
7     There is an Emergency situation
8     AccessPurpose is 'Diagnosis' or 'Treatment'
9     DataRequester.CurrentLocation = DataSubject.CurrentLocation

```

Fig. 12 A policy template for generating an authorisation policy for providing consent to the emergency response team.

The policy template applied in emergency situation is shown in Figure 12. In the *provided* part of the policy template for emergency situations, we include the condition for capturing the notion of emergency situation, *i.e.*, *There is an Emergency situation*. Furthermore, we omit also the condition *AccessTime is within DutyHours*, in contrast to the policy template for a GP shown in Figure 7, considering the fact that the emergency can happen at any time. For restraining access in emergency situations, the resource field of the policy template is set to *Allergy Report* and *Blood Test*. Moreover, we consider *READ* only access in emergency situations.

```

1 DataRequester.Role = {'EmergencyResponseTeam'}
2 DataRequester.ID = {'Payne'}
3 DataSubject.ID = 'Alice'
4 DataSubject.Resource = {'Allergy Report'}
5 AccessRights = {READ}
6 provided
7     Emergency = TRUE and
8     AccessPurpose = 'Diagnosis' and
9     DataSubject.CurrentLocation = 'Aachen' and
10    DataRequester.CurrentLocation = 'Aachen'

```

Fig. 13 An authorisation policy for providing consent to the emergency response team.

Figure 13 shows the authorisation policy for providing consent to the emergency response team. This authorisation policy is instantiated when an emergency happens in *Aachen* and *Fayne*, a member of emergency response team, requests *READ* access on (a patient) *Alice's Allergy Report* for *diagnosis* while *Alice* provides her consent and also saves her preferences for subsequent requests in the same environment. The occurrence of emergency situation may be detected using a patient's smartphone.

There are few important points to be considered. First, we are instantiating one authorisation policy per instance of the emergency response team. Alternatively, it may also be possible to instantiate the authorisation policy at the role level (*i.e.*, *EmergencyResponseTeam*) instead of at the instance level (*i.e.*, *Fayne*). Second, the patient may be in the unconscious state and may not be able to provide her consent. In such situations, authorisation policies can be instantiated from break-the-glass policy templates without asking patients. In other words, the emergency response team may provide consent on patient's behalf when patients are in the unconscious state. Here, the unconsciousness state can be incorporated at the time of sending consent request by members of the emergency response team to the patient's smartphone. Finally, as ultimate break-the-glass in case the smartphone is not reachable or not functioning, the emergency team can specify the current circumstances together with the request for accessing the patient's medical data. This information then can be checked in a post-incident analysis to make sure that such access mode is not abused. Again, it should be noted here that the medical data are not stored in the smartphone.

7 Related Work

In [28], *Aboelfotoh et al.* propose a mobile-based architecture for integrating Personal Health Record (PHR), where allows patients to control their data through their mobile devices. To manage the lifecycle of consent, they use the goal-driven approach, proposed by *Asghar and Russello* in [1]. In addition, they comply with the privacy consent direction [29] of Health Level 7 (HL7) [30], a reference guide for exchanging healthcare data.

Curren and Kaye [31] provide a legal background that signifies importance of consent withdrawal and revocation. According to their analysis, implementing consent withdrawal and revocation is not straightforward in practice, in particular when we address all the related legal complications. In [14], *Pruski* introduces e-CRL, a language for expressing patients' consent to regulate access to their health information.

Russello et al. [3] propose a consent-based framework that enables patients to control disclosure of their medical data, where the mechanism of capturing consent is integrated with workflows. The idea is to automatically generate Ponder2 style of authorisation policies [32] that depend on workflows. However, there is no automatic mechanism for managing the lifecycle of consent, such as consent withdrawal, activation or deletion. *Asghar and Russello* [33] suggest a mechanism for

managing the consent lifecycle. They introduce a notion of very expressive consent represented as a consent policy. However, they assume that a data subject defines his/her consent policies; unfortunately, such a solution may not be acceptable because data subjects may not be able to understand low-level policy details.

Wuyts *et al.* [5] incorporate patient consent with healthcare systems. They use the XACML policy language for defining access control on medical data and retrieve consent from the Policy Information Point (PIP). They express consent as a set of pre-defined attributes and store it in the database. The similar approach is used by Jin *et al.* in [34], which is an authorisation framework for sharing EHR. The main issue with both approaches is that the set of pre-defined attributes may not be sufficient to capture consent as it may involve certain conditions. In order to overcome this issue, there are approaches [33, 35] in which consent is treated as an authorisation policy; however, it raises some other problems. First, this approach requires users to specify low-level details, which a normal user may not be aware of, at the time of policy creation. Second, there is no automatic mechanism for managing the consent lifecycle.

EnCoRe [26, 36] aims at managing consent of users in order to regulate access to their personal data. In EnCoRe, a user is expected to define her preferences regarding consent, which are stored by enterprises. Once any piece of personal data is requested, these preferences are checked by the enterprises before granting access to the requested data. However, it may be cumbersome for users to define such complicated preferences. In our proposed solution, users' consent can be captured and managed dynamically by taking into account contextual information. Furthermore, our proposed approach offers more control and access to users as consent is stored and managed on their smartphone.

Luger and Rodden [37, 38] advocate how consent is a critical concern in pervasive computing. They describe issues with existing consent systems and highlight challenges and recommendations for a consent management system.

Marinovic *et al.* [27] employ TR policies for continuously monitoring the nursing home, where caregivers (including nurses, head-nurses, patients and students) are equipped with mobile devices for running their corresponding TR policies. They use TR policies to manage all activities of a caregiver using one workflow specification while we use TR policies with the goal of capturing consent that may involve instantiation of authorisation policies regarding consent and management of their lifecycle, consisting withdrawal and activation of consent.

Illner *et al.* [39, 40] suggest an automated approach for managing services related to distributed and embedded systems in dynamic environments. In their approach, various configurations for the services are generated and mapped to specific environmental conditions only once at the design time when system is setup while appropriate configurations for the services are activated at runtime when certain environmental conditions hold. The shortcoming of this approach is that the configurations are defined statically while our goal-based approach is dynamic in a sense that authorisation policies do not need to be specified in advance and are instantiated automatically while taking into account environmental conditions.

Johnson *et al.* [41] suggest a general approach for creating policy templates. A policy template provides users with a structured format for authoring policies. In our proposed solution, a healthcare provider may consider this work for generating policy templates. Chan and Kwok [42] describe a method to create policies automatically based on observed events. They use the Singular Value Decomposition (SVD) technique for modelling correlation between events and policies and then create new policies or select recommended policies based on the correlation. Unfortunately, the SVD technique may not always choose the fine-grained policies while our proposed approach always generates the fine-grained authorisation policies based on environmental conditions.

Fu *et al.* [43, 44] propose how to automatically generate required IPsec policies without manual configuration. The idea is to define high-level security requirements and then automatically generate a set of IPsec policies that can satisfy all security requirements. The main problem is that this approach incurs high performance overhead for finding the required set of policies as the proposed algorithm needs to go through a large number of possibilities before halting. Instead of generating a set of authorisation policies, our proposed approach generates only a single authorisation policy while taking into account contextual information and user intent.

8 Conclusions and Future Work

With the increasing attention towards the notion of data subjects consent to be integrated in access control mechanisms, the task of properly capturing security requirements in policy specification is becoming very daunting. This increases the risk of introducing errors in the policy specification that might compromise the privacy of the medical data. In the light of this, in this chapter we have proposed ACTORS a goal-driven approach, where authorisation policies are managed by TR policies that have goal of capturing the consent preferences of a data subject. As we have shown in our scenario, data subjects might want to handle consent in accordance with the actual situation and context. TR policies are structured in such a way that rules at the top are closer to the goal of the policy while rules at the bottom are more relevant when the goal is not close to be achieved. This is very natural for humans to grasp; therefore, a security administrator can capture more naturally the security requirements.

As future work, we are planning to focus on securing the mobile device where the consent application is installed. As mobile device could be stolen or misplaced, there is a need to make sure that the data subject does not lose control over her data and consent. We are exploring several approaches including dynamic authentication mechanisms to authenticate users in a seamless manner, *i.e.*, without requiring unnecessary interactions with the device.

Another area that requires investigation is enforcement of cross-domain policies. In this setting, it is difficult for the security administrator to have all the details of the different domains in which the data of the user might end up. Our idea is to

have mapping of the policy templates from one domain to the other, say by means of ontologies.

We are planning to perform a thorough evaluation with the medical school at our university. Our current experience so far with the capturing of security requirements with TR policies is very promising. ACTOR has already captured the requirements of one of the testbeds in the ENDORSE project. Another interesting one would be capturing consent for handling personal data of customers of a commercial entity. We can apply the same concept to solve other real-world problems. Terms and conditions at signing up or installing a software and granting app permissions in Android (or approving app restrictions in iOS) are few interesting problems among many others.

References

1. M. R. Asghar and G. Russello, "ACTORS: A goal-driven approach for capturing and managing consent in e-health systems," in *2012 IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY)*, July 2012, pp. 61–69.
2. J. Saltzer and M. Schroeder, "The protection of information in computer systems," *Proceedings of the IEEE*, vol. 63, no. 9, pp. 1278–1308, September 1975.
3. G. Russello, C. Dong, and N. Dulay, "Consent-based workflows for healthcare management," in *Policies for Distributed Systems and Networks, 2008. POLICY 2008. IEEE Workshop on*, June 2008, pp. 153–161.
4. K. Twidle, N. Dulay, E. Lupu, and M. Sloman, "Ponder2: A policy system for autonomous pervasive environments," *Autonomic and Autonomous Systems, International Conference on*, pp. 330–335, 2009.
5. K. Wuyts, R. Scandariato, G. Verhenneman, and W. Joosen, "Integrating Patient Consent in e-Health Access Control," *International Journal of Secure Software Engineering, IGI Global*, vol. 2, no. 2, June 2011, partner: KUL; project: NESSoS.
6. OASIS Standard, "eXtensible Access Control Markup Language (XACML) Version 3.0," <http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.pdf>, January 2013.
7. G. Russello, C. Dong, and N. Dulay, "Authorisation and conflict resolution for hierarchical domains," in *Policies for Distributed Systems and Networks, 2007. POLICY '07. Eighth IEEE International Workshop on*, June 2007, pp. 201–210.
8. D. J. Solove, "Introduction: Privacy self-management and the consent dilemma," *Harv. L. Rev.*, vol. 126, p. 1880, 2012.
9. N. J. Nilsson, "Teleo-reactive programs for agent control," *J. Artif. Int. Res.*, vol. 1, pp. 139–158, January 1994.
10. "Secys advisory comm. on automated pers. data sys., u.s. dept. of health, educ. & welfare, records, computers, and the rights of citizens," pp. 48–50, 1973.
11. P. Lawson and M. ODonoghue, "Approaches to consent in canadian data protection law," *Lessons from the Identity Trail: Anonymity, Privacy and Identity in a Networked Society*, pp. 23–42, 2009.
12. P. M. Schwartz, "The eu-us privacy collision: A turn to institutions and procedures," 2013.
13. European Communities, "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, year=1995, howpublished = http://ec.europa.eu/justice/policies/privacy/docs/95-46-ce/dir1995-46_part1_en.pdf."
14. C. Pruski, "e-CRL: A rule-based language for expressing patient electronic consent," in *eHealth, Telemedicine, and Social Medicine, 2010. ETELEMED '10. Second International Conference on*, February 2010, pp. 141–146.

15. L. McNair and A. Costello, "Electronic informed consent: A new industry standard," http://www.wcgclinical.com/wp-content/uploads/2014/03/eConsent-White-Paper_FINAL.pdf, 2014.
16. E. Communities, "Directive 1999/93/EC of the european parliament and of the council of 13 december 1999 on a community framework for electronic signatures," <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CONSLEG:1999L0093:20081211:EN:PDF>, December 1999.
17. R. Clarke, "econsent: A critical element of trust in ebusiness," *BLLED 2002 Proceedings*, p. 12, 2002.
18. E. Coiera and R. Clarke, "e-consent: The design and implementation of consumer consent mechanisms in an electronic environment," *Journal of the American Medical Informatics Association*, vol. 11, no. 2, pp. 129–140, 2004.
19. H. Nissenbaum, *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
20. J. B. Earp, Q. He, W. Stufflebeam, D. Bolchini, C. Jensen *et al.*, "Financial privacy policies and the need for standardization," *IEEE Security & Privacy*, vol. 2, no. 2, pp. 36–45, 2004.
21. J. Turov, L. Feldman, and K. Meltzer, "Open to exploitation: America's shoppers online and offline," 2005.
22. A. M. McDonald and L. F. Cranor, "Cost of reading privacy policies, the," *ISJLP*, vol. 4, p. 543, 2008.
23. X. Zhou, S. Demetriou, D. He, M. Naveed, X. Pan, X. Wang, C. A. Gunter, and K. Nahrstedt, "Identity, location, disease and more: Inferring your secrets from android public resources," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 1017–1028.
24. P. M. Schwartz and D. J. Solove, "Pii problem: Privacy and a new concept of personally identifiable information, the," *NYUL Rev.*, vol. 86, p. 1814, 2011.
25. P. Malone, M. McLaughlin, R. Leenes, P. Ferronato, N. Lockett, P. B. Guillen, T. Heistracher, and G. Russello, "ENDORSE: a legal technical framework for privacy preserving data management," in *Proceedings of the 2010 Workshop on Governance of Technology, Information and Policies*. ACM, 2010, pp. 27–34.
26. E. A. Whitley, "Informational privacy, consent and the "control" of personal data," *Inf. Secur. Tech. Rep.*, vol. 14, no. 3, pp. 154–159, August 2009.
27. S. Marinovic, K. Twidle, N. Dulay, and M. Sloman, "Teleo-reactive policies for managing human-centric pervasive services," in *Network and Service Management (CNSM), 2010 International Conference on*, October 2010, pp. 80–87.
28. M. Aboelfotoh, P. Martin, and H. Hassanein, "A mobile-based architecture for integrating personal health record data," in *e-Health Networking, Applications and Services (Healthcom), 2014 IEEE 16th International Conference on*, October 2014, pp. 269–274.
29. Health Level Seven International, "H17 implementation guide for cda release 2: Privacy consent directives, release 1," http://gforge.hl7.org/gf/download/frsrelease/977/10295/CDAR2_IG_CONSENTDIR_R1_N1_2013MAY.pdf, May 2013.
30. R. H. Dolin, L. Alschuler, S. Boyer, C. Beebe, F. M. Behlen, P. V. Biron, and A. S. Shvo, "H17 clinical document architecture, release 2," *Journal of the American Medical Informatics Association*, vol. 13, no. 1, pp. 30–39, 2006.
31. L. Curren and J. Kaye, "Revoking consent: A blind spot in data protection law?" *Computer Law & Security Review*, vol. 26, no. 3, pp. 273 – 283, 2010.
32. N. Damianou, N. Dulay, E. Lupu, and M. Sloman, "The ponder policy specification language," in *Policies for Distributed Systems and Networks*, ser. Lecture Notes in Computer Science, M. Sloman, E. Lupu, and J. Lobo, Eds. Springer Berlin Heidelberg, 2001, vol. 1995, pp. 18–38.
33. M. Asghar and G. Russello, "Flexible and dynamic consent-capturing," in *Open Problems in Network Security*, ser. Lecture Notes in Computer Science, J. Camenisch and D. Kesdogan, Eds. Springer Berlin Heidelberg, 2012, vol. 7039, pp. 119–131.
34. J. Jin, G.-J. Ahn, H. Hu, M. J. Covington, and X. Zhang, "Patient-centric authorization framework for sharing electronic health records," in *Proceedings of the 14th ACM Symposium on*

- Access Control Models and Technologies*, ser. SACMAT '09. New York, NY, USA: ACM, 2009, pp. 125–134.
35. C. M. O'Keefe, P. Greenfield, and A. Goodchild, "A decentralised approach to electronic consent and health information access control," *Journal of Research and Practice in Information Technology*, vol. 37, no. 2, pp. 161–178, 2005.
 36. M. C. Mont, S. Pearson, G. Kounga, Y. Shen, and P. Bramhall, "On the management of consent and revocation in enterprises: Setting the context," *HP Laboratories, Technical Report HPL-2009-49*, 2009.
 37. E. Luger and T. Rodden, "Terms of agreement: Rethinking consent for pervasive computing," *Interacting with Computers*, p. iws017, 2013.
 38. —, "An informed view on consent for ubicomp," in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '13. New York, NY, USA: ACM, 2013, pp. 529–538.
 39. S. Illner, H. Krumm, A. Pohl, I. Lück, D. Manka, and T. Sparenberg, "Policy controlled automated management of distributed and embedded service systems," in *Parallel and Distributed Computing and Networks*, 2005, pp. 710–715.
 40. S. Illner, A. Pohl, H. Krumm, I. Luck, D. Manka, and T. Sparenberg, "Automated runtime management of embedded service systems based on design-time modeling and model transformation," in *Industrial Informatics, 2005. INDIN '05. 2005 3rd IEEE International Conference on*, August 2005, pp. 134–139.
 41. M. Johnson, J. Karat, C. Karat, and K. Grueneberg, "Usable policy template authoring for iterative policy refinement," in *Policies for Distributed Systems and Networks (POLICY), 2010 IEEE International Symposium on*, July 2010, pp. 18–21.
 42. H. Chan and T. Kwok, "A policy-based management system with automatic policy selection and creation capabilities by using a singular value decomposition technique," in *Policies for Distributed Systems and Networks, 2006. Policy 2006. Seventh IEEE International Workshop on*, June 2006, pp. 4 pp.–99.
 43. Z. J. Fu and S. F. Wu, "Automatic generation of IPSec/VPN security policies in an intra-domain environment," 2001.
 44. Z. Fu, "Network management and intrusion detection for quality of network services," PhD in Computer Science, North Carolina State University, 2001.

About the Authors

Muhammad Rizwan Asghar

The University of Auckland,
New Zealand
r.asghar@auckland.ac.nz



Muhammad Rizwan Asghar is a Lecturer at The University of Auckland in New Zealand. Prior to joining this tenure-track faculty position, he was a Post-Doctoral Researcher at international research institutes including Saarland University in Germany and CREATE-NET in Trento Italy, where he also served as a Researcher. He received his Ph.D. degree from the University of Trento, Italy in 2013. As a part of his Ph.D. programme, he was a Visiting Fellow at the Stanford Research Institute (SRI), California, USA. He obtained his M.Sc. degree in Information Security Technology from the Eindhoven University of Technology (TU/e), The Netherlands in 2009. His research interests include access control, applied cryptography, security, privacy, cloud computing and distributed systems.

Giovanni Russello

The University of Auckland,
New Zealand
g.russello@auckland.ac.nz



Giovanni Russello is a Senior Lecturer and Leader of the Digital Security Programme at the University of Auckland, New Zealand. He received his Ph.D. from the Eindhoven University of Technology, The Netherlands. After obtaining his Ph.D., he was a Research Associate in the Department of Computing at Imperial College London, UK. His research interests include policy-based security systems, privacy and confidentiality in cloud computing, smartphone security and applied cryptography.