

Applying Formal Methods to Detect and Resolve Ambiguities in Privacy Requirements

Ioannis Agrafiotis, Sadie Creese, Michael Goldsmith, Nick Papanikolaou

International Digital Laboratory
University of Warwick, Coventry, UK

{I.Agrafiotis,S.Creese,
M.H.Goldsmith,N.Papanikolaou}@warwick.ac.uk

Abstract. In this paper, we demonstrate how formal methods can be used to unambiguously express privacy requirements. We focus on requirements for consent and revocation controls in a real world case study that has emerged within the EnCoRe project. We analyse the ambiguities and issues that arise when requirements expressed in natural language are transformed into a formal notation, and propose solutions to address these issues. These ambiguities were brought to our attention only through the use of a formal notation, which we have designed specifically for this purpose.

1 INTRODUCTION

It is common practice for individuals to disclose personal information via the Internet in order to acquire access to services, products and benefits of today's society. Thus, enterprises, organisations and government institutions alike have developed facilities to collect, process and share personal data with third parties. However, concerns about invasion of privacy are growing, mainly because of the way individuals' personal data is handled by these parties. In this paper we use the term "data controllers" to describe all the parties that handle and process personal data. That these concerns are based on solid ground is illustrated by the increasing number of incidents where data has been lost, mistreated, or shared without authority [4], making the use of privacy-enhancing technologies essential for every Internet user.

Although the right to privacy has been fundamental to all democratic societies and its importance is highlighted throughout the published literature [1], the term *privacy* has no inherent definition [3]. It is difficult to define privacy because it is a complex, multidimensional and highly context-dependent notion. People feel differently about what privacy means to them and have developed different meanings and interpretations according to their culture and experiences. The volatility of the notion of privacy, its highly contextual nature, and the widespread availability of powerful technologies for the collection, processing and sharing of personal data, all justify the need to carefully study, develop and enforce suitable privacy controls for users in modern information systems.

Definitions in the information-privacy literature describe an “implicit and limited view” of controls that an individual can invoke [3]. Westin [1] defines privacy as “the claim of individuals, groups or institutions to determine for themselves when, how and to what extent information about them is communicated to others.” Inspired by this view, the EnCoRe¹ project [5] is working to make available a virtual smörgåsbord of consent and revocation controls that an individual may choose to use in order to manage her/his data flow.

In the EnCoRe project we perceive “controls” as means enabling people to manage the flow of their personal data, by expressing consent and revocation preferences that can be implemented through non-interference and privacy policies. The overall vision of the project is “to make giving consent as reliable and easy as turning on a tap and revoking that consent as reliable and easy as turning it off again” [5]. To this end, we are taking into account a variety of perspectives, including social, legal, regulatory and technological aspects.

We have devised a model of consent and revocation, based on the published literature and on workshops held within the scope of the project. We have developed an accompanying logic of consent and revocation (C&R), which we use to formalise specific contextual requirements, enabling us to translate natural language expressions of C&R needs into an unambiguous form suitable for checking implementations against. In this effort we used a real world case study to validate our Logic.

This paper describes the ambiguities and the problems that came to light when we applied our consent and revocation logic in a real world case study in order to represent formally the specification requirements of the system. When formal methods are applied to privacy problems, “the nature of privacy offers new challenges and thus new opportunities for the formal methods” [14]. The new challenge that we describe in this paper is the ambiguities, which were not evident at first consideration of the consent and revocation models, but derive from the challenges and gaps created when the details of law, regulation, policy and social factors are combined and applied by computer scientists [8]. We extend this view and argue that these ambiguities, as well as highlighting the gap between high-level and low-level methods, unveil the complexity of the privacy problem and also arise from the gap between peoples’ desire for privacy and the data controllers’ will for security.

In the second section we describe the different controls envisaged in the information-privacy literature and how we extend these controls within the project to develop a consent and revocation logic, and present a brief explanation of the logic’s semantics. In the third section we present the real world case study and in the fourth section we categorise the ambiguities that emerge during the process of formalising this scenario and we illustrate with examples of formal descriptions of some of the scenario’s use cases. We also, propose solutions to the emerging ambiguities. Finally, we propose opportunities for future work.

¹ The EnCoRe project [5] is an interdisciplinary research project, undertaken collaboratively by UK industry and academia, and partially funded by the Technology Strategy Board (TP/12/NS/P0501A), the Engineering and Physical Sciences Research Council and the Economic and Social Research Council (EP/G002541/1).

2 MODELLING CONSENT AND REVOCATION

In the literature of information privacy, controls have been conceptualised mainly during the process of consent [3]. Researchers identify controls that are applied at the start of a disclosure, during the processing of data and by providing the choice for the individual to be notified. Furthermore, controls could be exercised on what personal data is made available to others and with whom this data is shared [3]. There are limited references to revocation controls and these are only focused on opt-out choices.

We have developed a model of consent and revocation to provide a more holistic view and offer richer control mechanisms to the individuals whose personal data is held by data controllers [6]. In this paper we refer to this category of individuals with the term “data subjects”. From the literature, we have identified the different consent controls highlighted above and we have conducted workshops in order to identify different types of revocation controls. We have concluded that there exist at least eight different types of revocation [6]. These are:

- No Revocation At All
- Deletion
- Revocation of Permissions to Process Data
- Revocation of Permissions for Third Party Dissemination
- Cascading Revocation
- Consentless Revocation
- Delegated Revocation
- Revocation of Identity (Anonymisation)

We have applied our model to a real world case study, in order to validate it and elicit requirements for the EnCoRe system [7]. Our logic is designed to provide a formal verification framework for privacy and identity management systems. It fills the gap between data-privacy policy languages and high-level requirements by focusing on the semantics of the process of consent and revocation when applied to the handling and use of personal data [7].

The application of formal methods to privacy mainly focuses on translating privacy policies [14] which are mostly written in natural language, into machine readable formats. Languages like P3P [12] and EPAL [13] are examples of these. Barth et al [11] have formed a logic of “contextual integrity” based on Nissenbaum’s theory about dissemination of information [10]. The logic describes how different roles are allocated to people according to context and allows or set constrains on how people of these roles transmit data between them. They applied this logic to privacy policies such as HIPAA [11] and the Children’s Online Privacy Protection Act [11]. However none of these methods handle consent while they completely neglect the notion of revocation.

The logic consists of two novel models of consent and revocation, namely an access control model, described in Section 3, and a Hoare logic, described in Section 4. The access control model and the Hoare logic have been developed so as to be complementary to one another. The first model immediately supports policy enforcement architectures such as the one being developed within EnCoRe, but it

does not provide an intuitive language for data subjects to express their consent and revocation behaviour. The second model provides a core set of consent and revocation actions axiomatised in their effect on rights and permissions in a way that is more familiar to data subjects.

2.1 AN ACCESS CONTROL MODEL FOR CONSENT AND REVOCATION

In the access control model we formalise the semantics of consent and revocation processes using labelled transition systems [7]. The objective of this part is to express the requirements of such processes effectively. In this model, suitable for expressing privacy preferences, consent and revocation are perceived as dynamic modifications of those preferences.

There are three main tasks for which a data collector requires consent from an individual:

- collection of personal data (for storage in a database)
- use of personal data (for analysis, processing marketing or one of many other purposes)
- sharing or dissemination of personal data (to the public domain, or to another data collector)

We identify these three cases as permissions in the access control model that describe the state of the system. These permissions may be shared or revoked from the data subject. Every action in the model is of the form $r, action(\sigma, \delta, \Phi, q), v \rightarrow v'$, where r is the data subject who gives the permission, δ is the data that the action refers to and q is the data controller to whom the permission is shared. The letter σ describes which permissions are shared or revoked from the action and tracks the changes in the state of the system. The actions described in the model are these of consent, revocation, deletion, update and notification. Furthermore, we set guards or preferences on these actions, captured in the condition Φ , which contain the data subject's options that change according to context.

The variables contained in the Φ condition that set guards on the actions and describe the data subjects' consent and revocation preferences are depicted in Table 1.

Variables	Meaning
t	duration of consent (time-out)
v	volume of data held
s	sensitivity of data held
Π	parties that may access the data
a	persistence: how data is treated after consent has lapsed

Table 1. The variables in the Access Control Model

For example the operation $grant(\sigma, \delta, \Phi, q)$ simply updates the rights matrix with a new permission on datum d for a principal q and ensures the resulting consent and revocation state satisfies the condition Φ .

2.2 DESCRIPTION OF THE HOARE LOGIC

In the Hoare Logic we define consent and revocation processes with a set of rights for principals. We identify how the rights and actions are combined to affect permissions and create obligations. This logic differs to the access control model in its treatment, as it effectively models C&R as the application of rights that allow certain permissions. Consequently, one action in the access control model may be described with a combination of actions in the Hoare Logic. This is because we believe, the assignment of “rights” in this manner, to be much more intuitive to the way data subjects express themselves. Furthermore, in the Hoare Logic the conditions that guard each action are not expressed. We deliberately abstract the Φ conditions in the Hoare logic as the model focuses only on permissions in order to be more familiar to the data subjects.

We use the notation $\langle \text{cond1} \rangle t \langle \text{cond2} \rangle$ to express obligations, with the following intuitive meaning: from a state satisfying $\langle \text{cond1} \rangle$ there is a requirement to apply term t to produce a new state satisfying $\langle \text{cond2} \rangle$. The rules for the logic will be given in the form of Hoare triples, as follows:

$$\{precondition\} t \{postcondition\}$$

The precondition is a combination of rights and obligations. Every time the t action is triggered and the precondition is true then the result is that the post condition, which is a combination of new rights and obligations as well, is true. There is also the case where more than one action t could be executed at the same time. We capture the concurrency of actions by using the symbol “||”. Thus the triple will be formalised as

$$\{pre-condition\} t_1 || t_2 \{post-condition\}$$

All the permitted actions and rights are presented in the figures below.

We identified six different rights in the Hoare Logic. These rights are presented in Table 2. The actions that allow data subject to share and revoke rights, delete, anonymise and aggregate data are illustrated in Table 3.

Right	Meaning
$aO\delta$	a owns (originates) δ
$aL\delta$	a knows (where to locate) δ
$aP\delta$	a may process δ
$aA\delta$	a may aggregate δ
$aS\delta$	a may share δ (one-step further)
$aS^*\delta$	a may share δ transitively

Table 2. The rights in the Hoare Logic

Action	Meaning
grant (a, b, δ)	grant consent for b to process δ
grant ¹ (a, b, δ)	grant consent for b to share onward δ one step further
grant ⁺ (a, b, δ)	grant consent for b to share onward δ transitively
release (a, b, δ)	release δ for anonymous aggregation to b
revoke (a, b, δ)	revoke permission from b to process δ (personally identifiable)
revoke ⁺ (a, b, δ)	cascade revoke permission from b and friends to process δ
delete (a, b, δ)	delete δ at b
delete ⁺ (a, b, δ)	cascade delete δ

Table 3. The actions in the Hoare Logic

For example a principal a may grant consent for processing of a datum δ to a principal b only if a owns δ or is able to share it. Once consent has been granted, b will know where to find and process δ . Thus, the first rule for consent is as follows:

$$\frac{\{aO\delta \vee aS\delta\}}{\mathbf{grant}(a, b, \delta)} \frac{}{\{bL\delta \wedge bP\delta\}}$$

3 DESCRIPTION OF THE CASE STUDY

The case study selected for the validation of the models and the logic is the Enhanced Employee Data Scenario [2]. Our choice was informed by the fact that the management of employee data in organisations is a well-understood problem, and employees' privacy offers interesting issues in terms of managing consent and revocation controls in a context where different business, legal and personal requirements need to be taken into account [2].

The case study describes a number of use case scenarios and elicits from these a list of requirements. We explore the implications of invoking consent and revocation controls. These use cases are meant to illustrate key points affecting the management of consent and revocation such as: provision or revocation of consent by a data subject; enforcing consent and revocation preferences; dealing with the overall consent and revocation controls and its impact on data including notifications, updates, auditing, assurance aspects, etc.

In the employee data scenario, we focus on PII and sensitive information, such as trade union membership, financial/payment detail, home address details, family details, etc. Personal data can be gathered by different sub-organisations within the

enterprise (e.g. HR department, Payroll, Occupational Health department, pensions, customer care department, help-desk and customer-relationship management services, web services to sell products, etc.).

Although, this scenario is far from what can be achieved today in terms of consent and revocation controls, we believe that the emerging requirements, with reality checks, contribute to our understanding of these controls and indicate problems and ambiguities when implementing them.

4 AMBIGUITIES IN REQUIREMENTS

We categorise the ambiguities that emerge from our formalisation of the requirements into two classes. The first class comprises of the ambiguities created from the application of the details of law, regulation, policy and social factors by computer scientists [8]. The second class consists of ambiguities that emerge from the complexity of the notion of privacy and the gap that exists between the data subjects demand to control the flow of their data and the data controllers' desire to reduce data subjects' interference.

We argue that to detect and solve these ambiguities the use of formal methods is necessary. Initially, we identified these issues when we applied a simple version of the Logic on this environment (where the Logic was designed in response to informal requirements expressions). We quickly identified that once we had resolved the ambiguities in requirement we could not express them using our simple Logic. In order to address these issues we need to represent significantly larger rule sets than we had done before [8].

4.1 AMBIGUITIES OF THE FIRST KIND

In the first class ambiguities occur when the data subject performs controls in order to update, delete, revoke or change his or her given consent. More specifically, in the case where the data subject wishes to update his or her personal data, there can be ambiguities emerging as to whether previous data should be deleted or linked with the new data. Furthermore, in the case where the organisation has shared his data transitively it is not clarified whether the changes should affect the third parties as well.

Consider the example below which we draw from the EnCoRe project's case study on Employee Data Handling. This example highlights the aforementioned ambiguity and we propose solutions described in consent and revocation logic. Mary (an employee of the Company X) is getting married. She has to update her personal profiles and data within Company X and some of the third party services (including change address, financial details, indication of next of kin, etc.). It is clear that we need additional actions in the Hoare Logic which will enable us to explicitly model the "update process" and a rule to define the pre and post conditions. We introduce two actions, the "**update***" and "**update**", which allow data subjects to determine whether to delete or link the old data with the new respectively. Furthermore, we

define the right $aO\delta$ for someone to update her/his data and to express preferences whether his updates will be transitively shared onwards or not. This is formalised below:

$$\begin{aligned} & \{ mU\delta \wedge hL\delta \} \\ & \mathbf{update}^*(m, h, \delta, \delta') \\ & \{ (\neg hL\delta \wedge hL\delta') \wedge \langle \langle hP\delta \parallel hS\delta \parallel hS^*\delta \rangle \rangle \\ & \mathbf{grant}(m, h, \Phi, \delta') \parallel \mathbf{grant}^1(m, h, \Phi, \delta') \parallel \mathbf{grant}^*(m, h, \Phi, \delta') \\ & \langle \langle hP\delta \parallel hS\delta' \parallel hS^*\delta' \rangle \rangle \} \end{aligned}$$

With the access control model we express the same actions and permissions but we also set the preferences of the data subject that guard every action.

$$(r, \mathbf{update}(\sigma, \delta, \Phi, q), v) \rightarrow v'$$

where $v' = v [\rho \mapsto \rho']$ such that $v' \models \Phi$

where $\Phi := (\varphi_1 \wedge \varphi_2)$

$$\varphi_1 := \psi_t, \varphi_2 := \psi_s, \psi_t := t < 30, \psi_s := s = \{ \text{sensitive} \}$$

From the above formalisation, Mary decides to update her data by linking the new with the old ones. She doesn't choose to disseminate these updates to third parties and she is able to set time and sensitivity preferences. This means that she controls what information is updated, how it is updated, who will process the updated information and she also sets guards on these controls, expressed via the values of the variables that she chooses. Similar to the process of update when an individual revokes or changes his initial consent it is not clarified whether these changes will affect third parties.

Revocation of consent, even though it is analysed in detail in our model, obtains different meanings depending on the circumstances and purposes that the data is being held for. In the case of deletion, ambiguities emerge from differences on how higher level people perceive the notion of deletion and how it could be technically implemented. For example deleting data could have multiple meanings. We could render the data useless, scramble data, delete it from the back-up system or physically destroy the hard discs. Consider the case below that combines both types of these ambiguities:

Company X outsources the provisioning of a few mandatory enterprise services (including travel agency services, pension fund management) and voluntarily services (Sport and Social Club - SSC). Part of Mary's data (financial details, address, employee references, etc.) needs to be disclosed to these third parties. Mary decides to withdraw from the voluntary SSC service. She revokes her consent to use her personal data.

The formalisation of this example in the logic is illustrated below.

$$\begin{aligned} & \{ mO\delta \} \\ \mathbf{revoke}(m, t, \delta) \parallel \mathbf{delete}(m, t, \delta) \\ & \{ (\neg tL\delta \wedge \neg tA\delta) \wedge \neg tP\delta \} \end{aligned}$$

With the access control model we express the preferences of the data subject regarding the process of deletion.

$$\begin{aligned} & (r, \mathbf{delete}(\sigma, \delta, \Phi, q), v) \rightarrow v' \\ & \text{where } v' = v [\rho \mapsto \rho'] \text{ such that } v' \models \Phi \\ & \quad \text{where } \Phi := \varphi_1 \\ & \quad \quad \varphi_1 := \alpha \\ & \psi_\alpha := \alpha = \{ \text{delete from back-up system} \} \end{aligned}$$

In this example Mary chooses to revoke the ability from the third parties to process her data. Additionally, she expresses her preference that the data stored should also be deleted. It is important to highlight the difference illustrated by this example, between revocation of consent and deletion. This is a key misunderstanding between data subjects. The logic describes four different types of revocation that allows us to express formally all the different meanings that the term has, depending on the context. Another issue addressed with this formalisation is that of deletion. We include the variable α where Mary expresses what she perceives as deletion in the specific context.

4.2 AMBIGUITIES OF THE SECOND KIND

The second class highlights the complexity of the privacy problem and underlines the conflicts that emerge between a data subject and a data controller. The most interesting issues but at the same time most difficult to address is that of aggregation and anonymity. The complex nature of these issues could lead to a situation where it may be technically infeasible to express data subjects' preferences or the privacy regulations in place.

Aggregation unveils more information about the data subject, by combining pieces of information already available. Data could be processed and shared with the proper way by the data controller, but when aggregated, this data create new information that may compromise data subject's privacy. In the Hoare Logic we assume that every time when data is shared that could be aggregated by the data controller, which as long as he has permission to collect data has inherently the right to aggregate that. A possible solution to the problem of aggregation is for the data subject to define the purpose for which the data is shared and also control what further personal data the data controller collects.

The ambiguities that arise in the case of anonymity concern the way which data is anonymised. These ambiguities were unveiled when we tried to formalise the case where Mary requested her medical records to be anonymised if shared with another third party. Although Mary may consent to share her anonymised medical records, the danger of her identity to be unveiled always lurks, as "data can either be useful or

perfectly anonymous but never both” [9]. “There is growing evidence that data anonymisation is not a reliable technical measure to protect privacy. The development of more powerful re-linking technology capabilities and the wider access to increasing amounts of data from which to drive these are reducing its effectiveness” [15].

Even if methods such as k-anonymity [9] become efficient the link between the data controller and the third party captured by the logic, potentially could lead to de-anonymisation of the data. In our logic we capture data subjects request to anonymise data first and then disseminate to a data controller but we consider the anonymised data as new data where the data subject has no controls on that data and we forbid any share of the old data between the data subject the data controller and the third parties that have access to the anonymised data.

Ambiguities also emerge when the data subject exercises her/his right to revoke consent but at the same time the data controller is unable to perform such an action. For example a data subject may request his data to be deleted but the organisation is still processing his data.

To address these issues we insert a new action, which allows the data subjects to express transparency in their decisions. Additionally, we tackle the conflicts that occur between the data subjects and the data controllers by introducing a combination of permissions and obligations under certain conditions. For example in order to revoke a data subject his consent to process data there is a condition that the data controller does not process the specific data at that time.

Consider the example where Mary decides to withdraw from the voluntary SSC service. She revokes her consent to use and share her personal data. The formalisation of the example is shown below:

$$\begin{aligned} & \{ mO\delta \wedge hP\delta \wedge hS\delta \} \\ & \mathbf{revoke}(m, t, \delta) \parallel \mathbf{revoke}^1(m, t, \delta) \\ & \{ (\neg tP\delta \wedge \neg tA\delta) \wedge (\neg tS\delta) \} \end{aligned}$$

In the access control model we define a binary variable p which is true only if the data controller does not process the data. We include this variable in the Φ condition to permit the execution of the action of revocation only if the data controller does not process the data.

$$\begin{aligned} & (r, \mathbf{revoke}(\sigma, \delta, \Phi, q), v) \rightarrow v' \\ & \text{where } v' = v [\rho \mapsto \rho'] \text{ such that } v' \models \Phi \\ & \text{where } \Phi := \varphi_1, \varphi_1 := \psi_b, \psi_b := p = \{ \text{true} \} \end{aligned}$$

In this formalisation Mary chooses to revoke the ability of third parties to share her data onwards and their ability to process her data as well. In order to solve the conflicts between data subjects will and data controllers need to process the data these actions can only be fulfilled under certain circumstances controlled by both Mary and

third parties. The p variable ensures that there will be no ambiguities when these actions take place.

Another area of conflict is the notification process. It is ambiguous whether it is an obligation for the data controller to notify the data subject or a right for the data subject to request to be notified. Thus, the question that arises is who will trigger the action. For example consider the case where a data subject wants to be notified when a specific data is processed. There isn't a right to provide the owner of the data the ability to request such an action. But also the enterprise is not allowed to contact the data subject unless she/he has consented to such an action.

In the law, there isn't an explicit reference on whether the data subject could request to be notified or not. However, there is a "data subject's request" right that allows data subjects to request all the data that a data controller possess about them. In alliance with this reference, we will formalise the notification process as an obligation for the data controller.

Furthermore, ambiguities occur with the handling of meta-data. In the logic, meta-data mainly comprises of variables and ranges of values that set the context and describe data subject's consent and revocation preferences. In the high-level models it is not clear what happens with the data subject's control preferences. An interesting example is that of notification. How can an enterprise notify a data subject that their data were deleted completely if they do not keep their email and consent and revocation preferences describing the conditions for the action of notification?

The following example illustrates both ambiguities and the formalisation presented proposes a way to address these issues. Consider the case where Mary is offered the opportunity to express notifications preferences about access/internal usage/disclosure to third parties of this data. The formalisation is:

$$\begin{aligned} & \{ mO\delta \wedge tP\delta \} \\ & \mathbf{notify}(m, t, \delta) \\ & \{ \langle mN\delta \rangle \mathbf{notify}(m, t, \delta) \langle \text{true} \rangle \wedge tLn^\dagger \wedge tLn^* \} \end{aligned}$$

In this example, Mary expresses her preference to be notified when her data are accessed and processed by third parties. She also chooses to be notified by e-mail. We solve the problem of who will trigger the action of notification by introducing an obligation for the company to notify Mary. Furthermore Mary controls the possible channels through which she will be notified and the reason that will trigger such an action. Also by defining the meta-data stored in the company, in every action performed by the data subject she/he could express exactly the same controls that apply on her/his personal data.

5 CONVERGENCE

The formalisation of the case study revealed ambiguities and areas of conflict, enabling us to extend and improve system's requirements. However, their formalisation could not be effectively addressed with the initial state of the Logic. Tackling the ambiguities created both from the formalisation of law, regulation and social factors and from the complexity of the notion of privacy, enhanced the effectiveness of our logic by introducing new actions and rights and enriched its descriptiveness by identifying new variables and options for the data subject to express his preferences.

We will briefly mention the novelties in the Logic that allowed us to formalise effectively and unambiguously all the requirements for the first case study. We introduced four actions for updating data, enabling data subjects to update data either by deleting the previous data or by linking that with the new or propagate the updates to third parties as well. We defined an action for notification and created an obligation for the data controller providing the means to the data subjects to be notified under certain conditions and through certain communication channels (e-mail).

Further to the introduced actions, we identified rights that will determine whether the actions will be completed. The data subject now has the right to be notified, the right to update data and the right to delegate rights to other individuals. Furthermore, the data controller has the right to know the location of every meta-data, enabling the data subject to express preferences on the way that the meta-data will be treated.

Last but not least, the effectiveness of the old and new actions was increased by the new variables. We now created variables to determine when the data subject could revoke permissions from the data controller, the way to delete data and the purpose for which the data will be used in order to address the problem of aggregation.

6 CONCLUSIONS AND FUTURE WORK

Enabling individuals to control the flow of their personal data is a complex issue. In this paper we focused on consent and revocation controls and their practical implications when formalising high-level requirements. We discussed the ambiguities that occurred in a real-case scenario of a large organisation, during that process. We categorised these issues into two kinds, according to the source of their existence. Furthermore we proposed the use of formal methods in order to address these problems and we described parts of the solution.

Future work will focus on validating the extended Logic in a different real-case scenario and identifying new ambiguities. As our aim is to develop a general applicable Logic, the emerging ambiguities should be minor and solved without extending the Logic with more actions and rights.

Implementing consent and revocation controls raises technological, legal and business challenges. Thus, we need to combine effectively diverse scientific fields that are not necessarily complementary. Developing a logic for consent and revocation and applying it to a real-case scenario in order to identify and address the emerging difficulties, is the first step towards that objective.

7 REFERENCES

1. Alan Westin (1967). *Privacy and Freedom*. New York: Atheneum.
2. Marco Casassa Mont, Siani Pearson, Gina Kouna, Yun Shen, and Pete Bramhall (2009). *On the Management of Consent and Revocation in Enterprises: Setting the Context*. Technical Report HPL-2009-49, HP Labs, Bristol.
3. Edgar A. Whitley (2009). *Information privacy consent and the “control” of personal data*, Inform. Secur. Tech. Rep. doi:10.1016/j.istr.2009.10.001
4. Edgar A. Whitley (2009). *Perceptions of government technology, surveillance and privacy: the UK identity cards scheme*, Neyland D, Goold B, editor, new directions in privacy and surveillance, Gullompton: William; p. 133-156.
5. EnCoRe. <http://www.encore-project.info>.
6. Ioannis Agrafiotis, Sadie Creese, Michael Goldsmith, and Nikolaos Papanikolaou (2009). *Reaching for Informed Revocation: Shutting Off the Tap on Personal Data*, Proceedings of Fifth International Summer School on Privacy and Identity Management for Life, Nice, France, 7th – 11th September 2009.
7. Ioannis Agrafiotis, Sadie Creese, Michael Goldsmith, Nick Papanikolaou (2010). *The Logic of Consent and Revocation*. Submitted.
8. K. Krasnow Waterman (2010). *Pre-processing Legal Text: Policy Parsing and Isomorphic Intermediate Representation*, Intelligent information Privacy Management Symposium at the AAAI Spring Symposium.
9. Paul Ohm (2009). *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, University of Colorado Law Legal Studies Research Paper No. 09-12 2009 [<http://ssrn.com/abstract=1450006>].
10. Helen Nissenbaum (2004). *Privacy as contextual integrity*. *Washington Law Review*, 79(1),
11. Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum (2006). *Privacy and contextual integrity: Framework and applications*. In *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 184–198, Washington, DC, USA, 2006. IEEE Computer Society.
12. Lorrie Faith Cranor 2002. *Web Privacy with P3P*. O'Reilly, September.
13. Calvin Powers, Matthias Schunter (2003). *Enterprise privacy authorization language (EPAL 1.2)*. W3C Member Submission.
14. Michael Carl Tschantz, Jeannette M. Wing (2009). *Formal Methods for Privacy*. FM 2009: Formal Methods, Second World Congress Eindhoven, The Netherlands, LNCS 5850.
15. EnCoRe Press Briefing, London School of Economics, June 29th 2010.