

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/103869>

Please be advised that this information was generated on 2020-10-29 and may be subject to change.

User Privacy in Applications for Well-being and Well-working

Requirements and approaches for user controlled privacy

Wouter Bokhove, Bob Hulsebosch

Novay

Enschede, The Netherlands

wouter.bokhove@novay.nl, bob.hulsebosch@novay.nl

Bas van Schoonhoven, Maya Sappelli

TNO

Delft, The Netherlands

bas.vanschoonhoven@tno.nl, maya.sappelli@tno.nl

Kees Wouters

Philips Research

Eindhoven, The Netherlands

kees.c.b.a.wouters@philips.com

Abstract—Well-being applications at work and at home are expected to help people to continue contributing to society, the marketplace and the economy. To make them adaptive and intuitive, and allow them to provide personalized information and coaching to the user at the right time requires the availability of context information. The use of sensory devices for this purpose gives rise to an increased information level about users but also poses an increased privacy risk, especially when ubiquitous sensors and devices are networked and connected to on-line services. This paper describes user-centric approaches for protecting the privacy of users when applications use sensor data. Moreover, it assesses the compliance of these approaches with requirements for user controlled privacy and their suitability for well-being and well-working applications. Based on this assessment a number of privacy control approaches have been selected that are suitable for well-being and well-working applications.

Keywords - *privacy; context; user-centric; control*

I. INTRODUCTION

Well-being applications at work and at home assist individuals to continue their participation in society. In this paper, we advocate that user-centric sensing and reasoning techniques improve the efficiency and acceptability of applications for physical and mental well-being (mostly in a private context) and well-working (in a work context). These applications exploit information that describes the current context of the user. The availability of up-to-date contextual information enables application-developers to develop adaptive personalized apps.

However, this also results in privacy risks. The ubiquity of context information, and the relatively easy way of sharing this context information, increases the privacy risks for users. Data collection and data processing with respect for privacy [1] and data protection [2, 3], especially with regard to user awareness and control, are essential for privacy preservation and ultimately for the acceptance of well-being and well-working services. It is, therefore, important that the user remains in control of the collection, processing and distribution of data that is related to this user. In situations where user control is not feasible (e.g., in some

professional or medical applications where context data is required to properly perform a job) users should still be informed properly about the way their personal information is handled.

These context-based reasoning systems make it virtually impossible for individuals to control access to privacy-sensitive information. Often, the user may not even know about the contextual information that has been sensed or inferred about her. In this paper, we argue that software frameworks that support context-management must obey two design considerations: *user control* and *usability*.

User control means that the user is empowered to decide what fragments of context information they consider sensitive, and in what situations they are prepared to share it with other parties for what purposes and under which conditions.

From the *usability* perspective, it is required that privacy control is user friendly and intuitive so users understand what they have specified, and are encouraged to actively use these controls instead of relying on default privacy settings.

Obviously, the trade-off triangle will play an important role in these considerations [4]: a balance should be found between ease of use, amount of control, and intrusiveness towards the user. In many cases, this results in conflicts.

Privacy control should be considered as a process of continuous adaption and conformation of preferences to the situational context and social practices. Key elements in this dynamic privacy control process are the ability for users to gain insight in and control over their current privacy settings and to get feedback on the impact of these settings given a situational context. These elements are fundamental to successful deployment of privacy-preserving well-being and well-working applications.

Improved privacy control and awareness might result in modifications of the privacy policies. These modifications can be considered as ad-hoc and executed manually by the user or automatically by the privacy control system and can be considered as part of a single loop learning cycle. As part of a so-called second or double loop, awareness could also result in a modification of the mindset of the user regarding her privacy policy settings, e.g., make them less tight.

This paper discusses requirements for privacy control in context-aware services architectures. In addition, it presents the different functionalities needed to facilitate this control. The main objective of this control is to assist end-users in making decisions regarding privacy-sensitive information used by well-being and well-working applications. Although the actual privacy settings might be different for each application, the requirements for the controls are the same.

The structure of the paper is as follows. It starts with an overview of privacy control requirements followed by an inventory of known privacy control approaches is presented. Subsequently, each approach is assessed against the requirements. Finally, we summarize the outcome of this assessment and draw conclusions on the most suitable privacy control approaches for well-being and well-working applications.

II. PRIVACY CONTROL REQUIREMENTS

Privacy architectures try to meet the fair information practices principles developed since the 1970s [5]. Since then, a lot of organizations have come up with privacy guidelines, directives, frameworks and/or principles to further specify or explain the privacy issues at hand and how these should be handled [1, 2, 3, 6, 7, 8, 9, 10, 11].

It has been recognized that implementation of privacy principles is especially difficult in ubiquitous systems involving (large) sensor systems which typically collect a lot of context information. Langheinrich [12] has tried to develop a comprehensive set of guidelines for designing privacy-aware ubiquitous systems based on a number of the aforementioned guidelines. Inspired by these privacy principles and focussing on user-centric privacy control, overview and usability, the following requirements can be distinguished.

A. User-centric privacy control requirements

1) *Users must have privacy control over context information.*

Many users don't mind sharing personal information as long as they control how, where, when and with whom information is shared [13]. This is not only limited to static information like a user's name, birth date or more dynamic information like health records, status updates on social networks or the contents of emails, but is also applicable to a user's context. This results in the prime requirement that users must have privacy control over their context information.

2) *Users must be asked for permission at the time the context is requested.*

Users will not be able or willing to configure their privacy policies (completely) in advance. When context-information is requested, the user should be able to give or withhold an (informed) consent. Any solution for privacy control should thus allow for just-in-time (JIT) context requests. By allowing JIT consent requests, privacy policies must be applied in real-time.

This requirement is in direct contradiction with the usability requirement of unobtrusiveness (see also C1). A

solution to reduce the invasiveness of JIT consent requests is to let the control system learn from responses and thus increasingly develop its privacy policies.

3) *Users must be able to modify and revoke their consents.*

When a user gives consent for accessing context-related information, the user should be able to revoke or modify this consent at any time. If consent is revocable, research [14] shows that this can reduce risk perception. In contrast to the current practice, where consents are mostly permanent (until revoked, if the user is able to find this option), it would be better to use access tokens with a limited life span or a limited number of uses.

4) *Users must have fine grained privacy control.*

The different attitude towards privacy can also be translated to a requirement with respect to the level of control that a user wants to exert with respect to privacy settings.

5) *Users must be able to define the granularity of the context information.*

Besides the level of detail of the privacy control settings, also the context information itself can be more or less detailed. Users should have control over this granularity. With respect to location, users may want to provide their exact GPS location or maybe just a (descriptive) derivative: home or work, neighborhood or city or region or country. Similar granularity levels (Quality of Context) can also be defined for other types of context information. This implies that the user is able to specify the granularity of the context information she is willing to share with a service provider.

B. Overview requirements

1) *Users must be able to get an overview of all their privacy control settings.*

As users might be confronted with context-related privacy control issues throughout a long period of time, the user should have some way to get an overview of all their settings and consents, preferably in a single overview. Such an overview must provide insight into the users, applications and services that have access to a user's (aggregated) context information (and preferably also into the times and frequency this context information is accessed).

2) *Users must be able to get an overview of personal data provided to or accessed by a service.*

When users give consent to a service to access some of their personal data they may not be aware of the frequency this data will be used or the quality of context of the information. Therefore, it is required to have the possibility of an overview showing what data is used by which specific service. A step further would be to get insight in what is *derived* from the collected data by the context information consuming and/or aggregating parties.

C. Usability requirements

1) *Users must not perceive privacy control as annoying or interruptive.*

Applications for well-being or well-working may need access to several different types of context information at

different times and at different frequencies. The user should not need to grant or refuse access each and every time access is requested as this will make the control over context-related privacy settings a full-time job. Each time the user is asked to give permission this should be done in a manner which is neither interruptive nor annoying to the user. Therefore it should be done using a user friendly interface that enables an unobtrusive control of the privacy settings.

2) *Users must be able to understand the provided privacy controls*

Users must be able to understand what they give consent to, or put differently, the consent should be an informed consent. Informed consent is one of the requirements of the European Directive [2]. Since users have different levels of understanding and background knowledge, upholding this requirement is far from trivial.

3) *Privacy policies must be personalised.*

Many studies have investigated the attitude of users towards privacy issues. It is generally accepted to classify a person as being a privacy unconcerned (approximately 25%), pragmatist (approx. 50%) or fundamentalist (approx. 25%) [15]. Although this classification should not be used as a predictor for disclosing location information [16], it is found that users have a different attitude towards privacy.

III. PRIVACY CONTROL APPROACHES

Multiple approaches to privacy control can be found both in literature and in the current practice of social networking sites. This section describes these approaches.

A. *Quality of Context*

A form of obfuscation of context information is to alter its quality [17]. The assumption here is that detailed and specific context information is more privacy sensitive. From a privacy viewpoint, a user might want to restrict certain requesters from accessing very precise information.

B. *Symmetry*

An important approach to maintain privacy in context-aware environments is the principle of minimal asymmetry, which in short states that the ability to obtain information should be coupled with the sharing of information between the data owner and consumer [18].

Balancing the amount of information flowing between peers is important to maintain the balance in any relationship. This is particularly the case for social relationships. Social systems often approach the symmetry principle by allowing the user to see the status of the other users she is connected with.

C. *Lying about yourself*

Adapting data is a method for controlling what information that is sent out. A user can plan to lie or adapt the data after it is recorded and checked. This method reduces the tractability of the user's actions. Another method for lying is by adding fake data to obfuscate the actual information.

D. *K-anonymisation / hiding in the crowd*

In essence, the concept of k-anonymity relies on a simple protection mechanism: obfuscation. It then measures the provided privacy with a single parameter k. The value k determines the privacy protection in place: the larger the k is, the higher the privacy protection is. The k-anonymity scheme for location privacy has become popular, mainly due to its simplicity.

Another method for obfuscation of data is by hiding it in the crowd [19, 20]. This is a method based on k-anonymisation. By adding more or less random data (noise) to the signal it becomes more difficult to track down the original data. It can be seen as artificially creating other persons in the user's region such that the conditions for k-anonymity are met automatically.

E. *Anonymisation and pseudonymisation*

Pseudonymity is the ability to prove a consistent identity without revealing one's actual name, instead using an alias or pseudonym. Pseudonymity combines many of the advantages of both a known identity and anonymity. In anonymity, one's identity is unknown, but pseudonymity creates a separate, persistent "virtual" identity that cannot be linked to a specific person, group or organization. The purpose is to render the data record less identifying resulting in less customer or patient objections to its use. Data in this form is suitable for extensive analytics and processing.

Anonymity is often used as an underlying building block when implementing pseudonymity. In case of anonymity, no persistent name is used. It conceals the relationship between a particular user and the data about him. User model entries can no longer be assigned to a particular user, thus ensuring that they will remain secret. As a consequence, an anonymous communicating party cannot be remembered. It is also known as unlinkable anonymity.

F. *Consent*

Consent is often required by legislation and is part of many fair information practices. By asking users for consent before sharing or accessing personally identifying information (PII) the user has great control over his privacy. In practice, however, most consents are based on 'take-it-or-leave-it' and thus leave little choice and control to the user with respect to his or her privacy. Several extensions to the 'simple' user consent questions can be defined, such as the option to decide which attributes will be released or the option to give only consent for a limited amount of time.

G. *Privacy control layers*

When more control options become available it is likely to divide these options into layers where three layers are most common. Every layer contains more detailed settings. For controlling privacy settings, the top level is roughly suitable for users which are unconcerned about their privacy while the privacy fundamentalists can use the lowest level to configure their settings (almost) on policy level.

H. Fine-grained control

Research demonstrates that users have nuanced privacy preferences and that providing them with the ability to control personal information sharing based on more fine-grained and expressive privacy controls offers substantial benefit over simpler privacy controls.

There clearly is a need for greater expressiveness in privacy mechanisms, which control the conditions under which private information is shared on the Web [21]. Any increase in allowed expressiveness for privacy mechanisms leads to a strict improvement in their efficiency (i.e., the ability of individuals to share information without violating their privacy constraints), but comes at the cost of user friendliness as most privacy preferences will become relatively complex privacy.

I. Grouping

Grouping attribute, people or service providers can help the user defining privacy policies.

1) Grouping attributes

The clustering of several personal data attributes for which the same privacy policy will hold is a common way for current online services to organize consent of users. Clustering of attributes offers users a clear overview of which attributes will be shared and it provides a fast and easy way to give consent. However, in many current services users lack the possibility to cluster attributes themselves, or to alter the predefined clustering.

2) Grouping people

Another way of clustering in privacy settings is to cluster people that have access to a particular attribute or several attributes. In the EU project PrimeLife [14], the social network Clique was developed which was based on this idea of clustering. Clustering makes the audience for users who see their information more transparent and it allows users to keep different parts of their identity separate (for example professional and personal life). This type of clustering thus allows for audience segregation [22]. Currently several social network sites such as Google+ (which named their groups 'Circles'), Facebook (at which you can define multiple Lists of friends) have incorporated this clustering into their privacy settings.

3) Grouping service providers

The third way in which privacy settings may be clustered is by arranging service providers in groups that may receive certain data based on certain characteristics of the service provider. A potential issue with this approach is the question who is determining in what cluster a service provider fits in.

J. Removing Policies

Another privacy control option is to remove existing policies. Policies can be rules that the system has learned regarding consents the user have given. Kill switches exist that revoke all privacy settings at once and can be considered to be a batch version of the possibility to remove policies.

K. Overview

Awareness starts with having an overview that captures the kind of information that is being shared with consuming

parties (other users or service providers) under what conditions. At the moment this sharing information is far too scattered. Typically, consent is given once during installation of the application and forgotten afterwards. Having an overview of all consents given in the past to service providers that control certain personal data attributes would be an ideal starting point for privacy control. The size and complexity of the overview will strongly depend on the user's privacy attitude: unconcerned, pragmatic, or concerned [15]. Overviews could include all given consents, which information is available to specific others or when or how often a service retrieves specific context information and exactly which information is retrieved. The overview may lead to an increased user awareness concerning her privacy settings and prevention of inadvertent invasions of privacy.

L. Privacy Mirror

Privacy mirror is a method that makes the user aware of what information she is sharing and with whom she is sharing it with. It is a method for checking whether your privacy controls are working the way you expect them to work.

M. Privacy Quiz

A privacy quiz can be used to make the user aware of her privacy settings. It ensures that the user understands what happened to his or her data. The privacy quiz can be implemented by asking the user to answer a privacy-related question. Depending on the complexity of the privacy policies, these questions can be very simple or more advanced. Answering the questions should be optional. With many policies there are a lot of questions possible that can automatically be generated. This is particularly the case if context is taken into account in the policy rules. When a user answers several questions wrong, she is expected to update here privacy settings.

N. Notifications

Notifications are part of many privacy regulations and fair information policies and play an important role in raising and maintaining awareness with the user with respect to his or her privacy. The user can be informed of his personal data being accessed and used by a service provider in many different ways. For example, the user can be notified of each of the times a service provider accesses a certain piece of privacy information. In some cases, this will probably lead to the undesired situation in which the user is constantly being notified which will reduce the power of notifications in itself and the user intrusiveness is too large. The number of notifications can be reduced by notifying a user only when a service provider is accessing information in an unusual frequency or after a fixed number of times. It could also be envisioned that a user is notified only when a service provider becomes active after being dormant for some time. The opposite is also possible: a user might be notified when a service provider has been granted access to personal information, but has not actually accessed this information for some time.

O. Making suggestions

The information gathered from previous behavior regarding sharing of information of the user and choices of the user made regarding consent can be used to suggest privacy settings and specific privacy rules for future situations. These suggestions could include a number of previously mentioned privacy controls, such as clustering or time-based consent.

IV. DISCUSSION

The many different approaches that were described earlier may all help in one way or another in increasing the level of awareness of, or control of users over, their personal information. However, it is impossible to simply implement all of these approaches, as this would result in inconsistencies, and may not be necessary to ensure adequate

privacy protection to start with. To determine which approaches are preferred, they are first matched to the requirements and then their application for well-being and well-working applications is discussed.

Discussing how the individual approaches relate to each of the requirements is not feasible considering the large number of combinations that would need to be analyzed and discussed. Therefore, as a starting point of the analysis we mapped the requirements to the approaches. The result is shown in Table I. A ‘+’ indicates that the approach can be used to implement a requirement, a ‘-’ means that it conflicts with a requirement, and a ‘0’ means that there is a dependency on the actual implementation. Empty cells indicate there is no relation between the approach and the requirement.

TABLE I. MAPPING OF PRIVACY CONTROL REQUIREMENTS TO THE DESCRIBED APPROACHES

Approaches	Privacy Control Requirements									
	A1	A2	A3	A4	A5	B1	B2	C1	C2	C3
A. Quality of Context	+	+		+	+			0	+	+
B. Symmetry	+							0	+	+
C. Lying about yourself	+	+			-			-	+	0
D. K-anonymisation	+				+			+	-	+
E. Anonymisation and pseudonymisation	+							0	0	0
F. Consent	+	+	+	0			+	0	+	
G. Privacy control layers	+	-	+	-			+	+	0	+
H. Fine-grained control	+	-	+	+			+	-	-	0
I. Grouping	+	-		+		0		+	+	+
J. Removing Policies	+	-	+				+	0	+	
K. Overview		-				+		+	+	
L. Privacy Mirror		-				+		+	+	+
M. Privacy Quiz		-				+		-	+	
N. Notifications		+				+		-	+	+
O. Making suggestions	+					+		-	+	+

Based on the mapping presented in this table, several observations can be made. A number of approaches appear especially suitable for giving the user control. Asking consent to the user before sharing or processing his personal data implies that the user has control, and can control the handling of his personal data so that it matches his personal preferences. Most approaches that give the user control also allow a certain level of personalized control. Good examples that also provide the user with fine-grained control and fill-in a number of other requirements are: Quality of Context, Privacy Control Layers, or grouping of attributes, people, and service providers. More direct forms of control can for example be implemented using a Kill switch.

Providing the user with a good overview is essential, because consent can only be meaningfully given when it is informed consent.

User friendliness is an essential issue here, as some awareness approaches do not fit this requirement. This leaves approaches such as giving an Overview, or using a Privacy Mirror or Privacy Quiz.

The well-being and well-working applications aim at increasing physical and mental well-being of users. The unobtrusive nature of these applications that consume detailed privacy sensitive information to enhance service experience and effectiveness motivates the need for intuitive approaches able to cope with the high dynamic nature of

situational changes. The approaches are therefore divided into three categories depending on how they are suitable for this application domain.

A. Preferred approaches

Several approaches are very suitable for this application domain.

Quality of context is a very relevant tool, as information needs only to be as precise as required by an application, and no more. For example, if a person's heart is being monitored using ECG (ElectroCardioGram), it may not be necessary to transmit and process the detailed (and more revealing) ECG data. Instead, a derived current heart rate may be sufficient.

Consent is legally and ethically a strong requirement and an essential precondition for the user to be in control. The way in which the user can give consent is important, however. For this, the use of other (more specific) privacy approaches is necessary.

Grouping attributes, people and service providers is especially relevant for well-being and well-working applications. Using (configurable) clusters of attributes, and service providers with which to share personal information, provides a level of control that may offer a suitable implementation of the "control" approach. A balance will have to be found, however, in the level of detail in which grouping takes place. Also, grouping has to be done before the actual processing is done, putting some limits on its use.

Privacy control layers are a way to allow users with different privacy attitudes to translate their personal privacy concerns in a convenient way. As users of well-being and well-working applications will have a diverse attitude and including context in privacy preferences may lead to complex, fine-grained control requirements, privacy control layers will be needed.

Overview is important for users to get an awareness on how their personal information is being processed. Moreover, for consent to be meaningful it needs to be informed consent, so the user must have an understanding of what information is shared with whom.

Notifications may be used to maintain awareness of what is happening with the user's personal information and is a suitable tool for use in well-being and well-working applications. It may also provide a way to give the user just-in-time control. Of course care must be taken not to "spam" the user with notifications that are not relevant.

B. Conditional approaches

Some approaches are suitable only in specific situations or when specific conditions have been met.

Symmetry is a principle which is mainly relevant in sharing information with one's peers. So this may only be useful for some specific well-being and well-working applications, even though it is largely compatible with the requirements.

Anonymisation and pseudonymisation are powerful tools but may be difficult to successfully implement in some applications because of the kind of data that may be monitored. Essential data items in a well-being and well-

working application include many potentially identifying features, such as age, gender, weight, health status, etcetera.

Fine-grained control can be useful for well-being and well-working applications as the number of options, especially when a large number of context sources is being used, will be quite high. Fine-grained control allows for the users of these applications to accurately control their privacy. Fine-grained control requires Privacy Control Layers for usability reasons.

Removing policies and a kill switch may be useful for providing the user with a high level of control in well-being and well-working applications, but depends largely on how this is implemented. As many well-being monitoring applications may depend on the long-term availability of data to discover trends, these may not be suitable control tools for some applications.

A privacy mirror is a potentially very usable tool for increasing user awareness, but this depends strongly on how it is implemented. Also, some personal information related to psychological or mental well-being may not be in a format that gives much insight in what is actually being shared (e.g. detailed sensor information).

Making suggestions is an approach that can be used to support users in making decisions on their privacy "policies". Although this fits with the well-being and well-working applications that aim at supporting the user in similar ways for other ends, it requires a large amount of privacy control settings (particular consents) before becoming useful. Making suggestions can be then used to optimize and to make privacy control more user-friendly.

Ask, but don't tell can be useful for users whose well-being is constantly being monitored as this approach allows them to get off the grid temporarily and thus allows the user to protect his or her privacy in a simple way. As medical or life-style advice is based on the observed information applying this tool frequently could have an adverse effect.

C. Unsuitable approaches

Some approaches are generally not useful or difficult to implement in the well-being and well-working domain. We discuss these below.

Lying about yourself as a privacy approach in a well-being of well-working application environment may have very undesirable effects, as the quality of guidance provided by well-being and well-working applications depend on accurate information. Acting after medical or life-style advice based on incorrect information may have detrimental effects.

k-anonymisation depends on hiding the user's personal information in a large number of other user's personal information. However, as the well-being and well-working applications depend on providing specific users with feedback based on their specific personal information, this is not a useful technique for such applications. This is also true for hiding in the crowd, for the same reason.

Confronting the user with a privacy quiz is intrusive, and therefore useful only in specific circumstances, for example for privacy settings that are very important.

V. CONCLUSIONS AND FUTURE WORK

This paper presented a number of privacy control approaches which are mapped to privacy control requirements. Also, the suitability for well-being and well-working applications is discussed. Based on this analysis it has become clear that controlling one's privacy with respect to context information, while finding the proper balance between being easy to understand for the end-user, being fine-grained and being unobtrusive, is not an easy task.

The logical next step would be to find out if it is possible to use the context of the user to automatically make decisions about sharing his or her context information, i.e. determine the 'context-awareability' of the various privacy control approaches. If this can be done, this will result in more user-friendly and adaptive solutions (e.g., the user will not be asked for consent while in an important meeting or while sleeping). Context-aware adaptive privacy might exploit the ability to sense and use contextual information to augment or replace traditional user privacy control mechanisms by making them more flexible, intuitive and less intrusive. Moreover, we also intend to determine which sensors will be best suitable for this purpose thereby taking into account the quality of the provided sensor information, reliability and its dynamicity. This is currently work in progress and will lead to the development and user-evaluation of well-being and well-working applications that takes into account several privacy control approaches that are context aware.

ACKNOWLEDGMENT

This publication was supported by the Dutch national program COMMIT (project P7 SWELL). We thank the SWELL members Linda Kool, Saskia Koldijk, Maarten Wegdam, and Wessel Kraaij for the inspirational discussions.

REFERENCES

[1] "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data", 2002.
 [2] "Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data", EU, 1995.
 [3] "Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector", EU, 2002.
 [4] M. Wegdam and D.-J. Plas, "Empowering users to control their privacy in context-aware systems through interactive consent", Technical Report TR-CTIT-08-66, Centre for

Telematics and Information Technology University of Twente, Enschede, 2008.
 [5] R. Gellman, Fair Information Practices: a Basic History, April 25, 2012.
 [6] Federal Trade Commission, "Privacy Online: Fair Information Practices in the Electronic Marketplace, a report to Congress", 2000.
 [7] K. Cameron's Identity Weblog (www.identityblog.com), "The Laws of Identity", 2005.
 [8] NIST, "Security and Privacy Controls for Federal Information Systems and Organizations", NIST Special Publication 800-53, Draft Appendix J, 2011.
 [9] AICPA/CICA, "Generally Accepted Privacy Principles (GAAP)", 2009.
 [10] "APEC Privacy Framework", 2004.
 [11] U.S. Department of Commerce, "Safe Harbor Privacy Principles", 2000.
 [12] H. Langheinrich, "Privacy by Design – Principles of Privacy-Aware Ubiquitous Systems", Proceedings of the 3rd international conference on Ubiquitous Computing, pp.273-291, September 30-October 2, 2001.
 [13] B. Schneier, "Google and Facebook's Privacy Illusion", 2010.
 [14] EU, PrimeLife project, see: <http://www.primelife.eu>
 [15] P. Kumaraguru and L.F. Cranor., "Privacy Indexes: A Survey of Westin's Studies", Institute for Software Research International, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, TR CMU-ISRI-5-138, 2005.
 [16] S. Consolvo, I. Smith, T. Mathews and A. LaMarca, "Location Disclosure to Social Relations: Why, When, & What People Want to Share", in CHI '05: Proceedings of the SIGCHI Conference on Human factors in Computing Systems, New York, NY, 2005.
 [17] K. Sheikh, M. Wegdam and M. van Sinderen, "Quality-of-Context and its use for Protecting Privacy in Context Aware Systems", Journal of Software, Volume 3:83-93, 2008.
 [18] A. Kofod-Petersen, E. Klæboe, J. Jervidal, K. Aaltvedt, M. Romnes and T.M. Nyhus, "Implementing privacy as symmetry in location-aware systems", Proceedings of the International Workshop on Combining Context with Trust, Privacy and Security (CAT 2008), volume 371, pp. 1-10, Trondheim, Norway, 2008.
 [19] F. Bustamente and L. Amaral, "Software Improves P2P Privacy By Hiding In The Crowd", 2009.
 [20] F. Li, J. Sun, S. Papadimitriou, G.A. Mihaila and I. Stanoi, "Hiding in the Crowd: Privacy Preservation on Evolving Streams through Correlation Tracking", ICDE, 2007.
 [21] M. Benisch, P.G. Kelley, N. Sadeh, T. Sandholm, L.F. Cranor, P.H. Drielsma and J. Tsai "The Impact of Expressiveness on the Effectiveness of Privacy Mechanisms for Location Sharing", CMU-ISR-08-141, 2008.
 [22] E. Goffman, "The Presentation of Self in Everyday Life", New York: Anchor, Doubleday, 1959.