D4.1 Context Aware Adaptive Privacy Requirements

Project	SWELL
Project leader	Wessel Kraaij (TNO)
Work package	4
Deliverable number	D4.1
Authors	Wouter Bokhove (Novay), Bob Hulsebosch (Novay), Maya Sappeli (TNO), Kees Wouters (Philips)
Reviewers	Miriam Vollenbroek (RRD), Sanne Huveneers (TNO), Linda Kool (TNO)
Date	December 23rd, 2011
Version	V1.5
Access Rights	SWELL Project Members
Status	Final

Summary

Context-aware applications raise serious privacy issues that must be addressed both to appease public concern and to comply with current legislation such as the 95/46/EC directive which requires the explicit consensus for the use of privacy-sensitive context information.

CAAP (Context-Aware Adaptive Privacy) is to satisfy users' need to control their context (e.g. location, time or heart rate) privacy and at the same time the need to minimise the demands made of user interaction for privacy authorisation. This deliverable describes the requirements for CAAP.

In this deliverable, we discuss the different requirements of privacy control in context-aware services architectures. Further, we present the different functionalities needed to facilitate this control. The main objective of this control is to help end-users make consent decisions regarding their private information collection for the purpose of well-being and well-working applications. Although the actual privacy settings might be different for these different applications, the requirements for the controls are the same.

Contents

1 Introduction

1.1 Motivation

Well-being applications at work and at home are expected to help people to continue contributing to society, the marketplace and the economy (e.g. by allowing elderly people to live at home, by working from different locations improving the combination of work and private life or by improving lifestyle – and thus health – increasing the productivity of employees). Furthermore, these applications may help suppressing the rising costs of chronic disease and ill-health.

User-centric sensing and reasoning techniques can help to improve the efficiency and acceptability of physical and mental well-being (mostly in a private context) and well-working (in a work context) applications. 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. Moreover, it requires advanced distributed reasoning algorithms that can deal with heterogeneity of sensed context information, partial context information and resource limitations regarding the processing of the collected context information such as location, behavior, mood, activity and weather conditions.

The use of multi heterogeneous sensory devices 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. Data collection and processing with respect for privacy [5] and data protection [6, 7], 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 is in control of data collection, processing and distribution. 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 these aspects.

Without such control functionality, privacy concerns may become an inhibitor for the success of context aware services for well-being and –working. Such privacy concerns will likely depend on a combination of the following factors [1, 28]:

- Relevance and usefulness of the shared context information;
- Type and level of detail of the context information that is shared, i.e. the sensitivity of the data;
- Inquirer of the context information, e.g. a friend, an unknown service provider, a family member, the boss, or one's physician;
- The trustworthiness of the inquirer;
- The time and place of the request for information;
- Contractual and legislative restrictions or obligations;
- Cultural and social expectations;
- Gender and age;
- The security deployed in the infrastructure;
- The user personal privacy sensitivity, i.e., is he a privacy pragmatist, fundamentalist or unconcerned.

For instance the user's boss is allowed to know where the user is and what he is doing during working hours but not in the weekend. Or the patient's physician has access to his ECG recording when in the hospital or in the presence of the patient.

These user-related contextual factors for a given context affect the necessary degree of privacy protection on a continuous scale. However, the relationship between these factors and user controlled privacy measures that should be deployed can be complex. The balance should be found between ease of use, amount of control and intrusiveness towards the user as shown in Figure 1 [1]. In many cases this results in conflicts. For example, a more fine-grained control will make it more difficult to understand for the end-user and will most likely require more time and focus from the user and will thus be regarded as more intrusive.
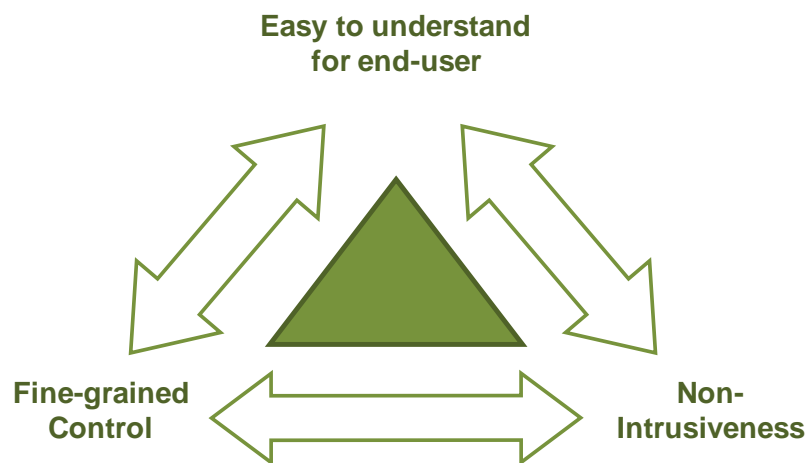

Figure 1: Conflict of interests in user controlled privacy.

Somehow the user should be empowered to control his privacy, e.g. by delegation, consent, policy management, or via existing user centric identity management solutions. User privacy control however is not easy to achieve in context aware service environments.

1.2 Goals

In emerging context aware or context-based reasoning systems it is for the user almost impossible to manage a controlled release of privacy sensitive information. Sometimes, the user may not even know all contextual information that has been sensed or inferred about her. Key to a technical framework that can support personalized service provisioning based on distributed context information will be 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 which parties for what purposes and under

which provisions. From the usability perspective, it is required that the system is transparent enough so the user understands what they have specified, and to encourage the user to actively use the system instead of relying on default settings. Obviously, the trade-off triangle as shown in Figure 1 will play an important role in these considerations.

Privacy control in context aware environments may require solutions different from those of today's systems. Today's most advanced privacy control systems are seen in popular (social) networking sites. However, these systems are based on relatively stable, well-defined, consistent configurations and static contexts that barely take the participants' personal privacy preferences into account. What is needed can be characterised by the term 'context aware adaptive privacy control', in which the degree and nature of privacy control associated with any particular type of action will change over time, with changing circumstances, and with variable available information so as to suit the context. Privacy control has to be non-intrusive, intelligent, and able to adapt to the rapidly changing contexts of the environment.

1.3 Approach

This deliverable provides an overview of the requirements for making privacy control context aware and adaptive such that it becomes more intuitive and user friendly. The requirements are inspired by current examples of privacy control as well as general privacy principles. The examples are taken from (social) networking sites as well as existing software and web applications from both the mobile and non-mobile domain. Each of the requirements will be described and – where possible – illustrated with examples and classified in terms of nice to have or essential (must haves) according to the terminology from RFC 2119 [2]. Classification is based on user perception taking into account the privacy vulnerabilities associated with the described functionality and the perceived feasibility. Even though security is a very important aspect in managing privacy and privacy control settings, this document will not describe specific security requirements; the focus is on privacy control requirements.

1.4 Structure of the report

The rest of this report is structured as follows:

Chapter 2 describes the vulnerabilities of privacy and presents a number of accepted privacy principles frameworks from different organisations before coming up, in chapter 3, with user requirements for context-aware adaptive privacy control. Chapter 4 summarizes the major findings of this report and hints towards future work in SWELL.

2 Privacy in context

2.1 Privacy vulnerability

In order to understand privacy vulnerabilities, it is good to imagine how far privacy invasion can go. An example of a more or less complete privacy invasion with consent of the persons themselves is the monitoring of the astronauts who flew to the moon. Many of us have seen the motion picture Apollo 13 about the failed trip to the moon. Every word being said, every heartbeat is monitored by the mission control people. Next to that the whole environment is monitored as well, ranging from spaceship temperature, to the people being talked to. Most if not all of this information including information regarding the activities of mission control is stored for later use. Analysis can be conducted on this information which may even be published. Even today people can still study the stored material.

In short, a complete privacy invasion means that only your thoughts are your own, and even those may be not so private when filling in a questionnaire.

In the above example, all the different aspects of privacy invasion can be found. The aspects are [3]:

- monitoring (i.e. monitoring the person's context);
- storing (i.e. storing the monitored information);
- aggregating (i.e. combining, possibly stored, monitored information with other sources);
- transferring (i.e. transferring aggregated or raw monitored information to a third party).

All aspects aim to collect and/or personalize the gathered information. The collected and/or personalized information is of value in many ways, amongst others to marketing and direct sales organizations. Moreover, for all aspects the person, whose context it concerns, is not in control of the dissemination of his context information.

Specific examples of the different privacy invading aspects are:

- Monitoring:
  - Checking which websites are being visited and in case these are not within policy raise a flag.
  - Checking the body temperature and when this exceeds 38 ^{o}C, warn a nurse.
- Storing:
  - Storing the addresses of all webpages visited.
  - Storing the measured temperature with a five minute interval.
- Aggregating:
  - Collecting the names of people the user chats with about cars.
  - Determine areas of interest of the user by inspecting the stored web pages.
- Transferring:
  - Forward the list of chatters interested in cars to make them a deal for visiting an auto show.
  - Forward the areas of interest of the user to a call enter.

Possible results of privacy invasion can be benign or malicious. As the examples show, some of them are done with the intention of helping the user, some are for the benefit of both the user and the collector of the information, and others are for the sole benefit of the collector of the information. The latter can go as far as identity theft, e.g. a malicious party has determined a user with the same car (brand, colour, identity plate) and uses that identity plate for a hit and run robbery.

The impact of the potential loss of privacy as a consequence of an exploited vulnerability can be used to evaluate the risks. If the risks are high, countermeasures should be taken to mitigate the risks. Countermeasures could be defined in the security realm or in the definition of specific privacy control requirements.

2.2 Privacy protection principles

Privacy architectures try to meet the fair information practices principles developed since the 1970s. 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. The most recent US privacy guidelines have been issued by the FTC in the Fair Information Practice Principles (FIP) cover five basic areas [4]:

- Notice and Awareness: The user should have clear notice or be aware of the type of information collected, its use, and an indication of third parties other than the original collector who will have access to the data. The following aspects of notice and awareness are recognized as essential:
    - o identification of the entity collecting the data;
    - o identification of the uses to which the data will be put;
    - o identification of any potential recipients of the data;
    - o nature of the data collected;
    - o means by which data is collected (if not obvious);
    - o whether the provision of the requested data is voluntary; and
    - o steps taken by the data collector to ensure the confidentiality, integrity and quality of the data.

This list is suggestive of the kinds of privacy requirements that context aware applications should satisfy. The mechanism by which users are typically made aware of such practices is the service provider/context consumer's privacy policy.

- Choice & consent: The ability to choose not to have data collected. The principle of choice and consent ensures that users are given the option to decide what personal information collected about them is to be used and whether it may be used for secondary purposes. The collection of personal information in itself can be an invasion of privacy, one over which users should have some control.

- Access & Participation: The ability for the user to see what personal information is held about him/her, to correct errors, and to delete the information if desired. Second use refers to use and sharing of data after initial access has been made. Second use also includes passing data from one party, who might have authorized access to the data, to another party, who might not. Consequently, data owners often have very little control over second use. Important decisions made at this phase include who else should be able to access the data, what they can do with it, and whether they should be allowed to share it even further with others.

- Integrity & Security: Reasonable measures taken to secure (both technically and operational) the data from unauthorized access. The following aspects of integrity/security are recognized as essential:

  - providing consumer access to data;
  - destroying untimely data or converting it to anonymous form;
  - managerial measures to protect against loss and the unauthorized access, destruction, use, or disclosure of the data; and
  - technical measures to protect against loss and the unauthorized access, destruction, use, or disclosure of the data.

- Enforcement & Redress: There must be a mechanism in place to enforce privacy, otherwise a policy is merely suggestive, rather than prescriptive. The FIP principles list three specific types of enforcement and redress: self-regulation, private remedies, and government enforcement.

Many European countries have implemented data protection legislation to protect civil liberties since the OECD publication in 1980 of guidelines on the protection of privacy and trans-border flows of personal data [5]. Most of these laws, for example the European Data Protection Directive[1] [6] and the ePrivacy Directive [7], are based on the OECD's privacy principles and form the basis for consumer privacy protection:

1. Collection limitation: Data collectors should only collect information that is necessary, and should do so by lawful and fair means, i.e., with the knowledge or consent of the user.
2. Data quality: The collected data should be kept up-to-date and stored only as long as it is relevant.
3. Purpose specification: The purpose for which data is collected should be specified (and announced) ahead of the data collection.
4. Use limitation: Personal data should only be used for the stated purpose, except with the user's consent or as required by law.
5. Security safeguards: Reasonable security safeguards should protect collected data from unauthorized access, use, modification, or disclosure.
6. Openness: It should be possible for users to learn about the data controller's identity, and how to get in touch with him.

---

[1] The EU is currently working on an update of this directive (refer to http://tweakers.net/nieuws/78621/europa-komt-met-scherpe-privacyregels.html). This update will most likely give more rights to the user regarding his privacy and more obligations for organizations consuming personal information.

7. Individual participation: users should be able to query data controllers whether or not their personal information has been stored, and, if possible, challenge (i.e., erase, rectify, or amend) this data.
8. Accountability: Data controllers should be accountable for complying with these principles.

The following (non-exhaustive) list contains organisations which have done important work on these issues:

- Kim Cameron has developed the seven "Laws of Identity" [8] as a framework on how to manage and use digital identities.
- NIST recently has issued a draft specification for privacy control of personally identifiable information [9].
- AICPA and its Canadian counterpart CICA have developed the Generally Accepted Privacy Principles (GAPP) from a business perspective [10].
- The APEC has developed a Privacy Framework [11] representing a minimum standard containing nine high level Privacy Principles.
- The US Department of Commerce in consultation with the EU has developed the Safe Harbor Privacy Principles [12] for US companies to comply with EU legislation. This framework is based on seven principles.

Appendix **Error! Reference source not found.** contains an overview of these principles.

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 [13] has tried to develop a comprehensive set of guidelines for designing privacy-aware ubiquitous systems based on a number of the aforementioned guidelines. He describes the following set of five principles:

- Notice: no single data collection can go unnoticed of the subject that is being monitored.
- Choice and Consent: any collection and usage of personal information is prohibited, except for certain legal procedures or when explicitly consented by the individual.
- Anonymity and Pseudonymity: levels of anonymity are an important option when offering users a choice, but they also allow legal collection of certain types of data without requiring user consent.
- Proximity and Locality: proximity could allow the collection of data from certain context information providers as long as the user is present while locality could prevent spreading context information outside the area where it has been collected.
- Adequate Security: the complexity of security can be reduced by employing robust security only in situations with highly sensitive data transfer according to the principle of proportionality.
- Access and Recourse: only collect data for a well-defined purpose, only collect data relevant for this purpose and keep this data only as long as it is necessary for this purpose.

Any solution for context-related privacy control should try to adhere to these principles.

3 Privacy control requirements

A privacy control system for context enhanced well-being and well-working application will address several needs. The privacy policy representation must be flexible and scalable, able to operate over loosely structured and dynamic context data rather than rigid data types and categories. It must also be rich enough to capture the many relationships between content owners and viewers, and to express a wide range of privacy preferences. The policy authoring and control tools must adapt as users restructure their social networks, and be efficient and simple enough for users with varying levels of technical knowledge.

This chapter describes privacy control requirements. Each requirement is boxed and includes a level of necessity. While most requirements are considered to be 'must haves', some requirements are marked as 'nice to have', following the terminology of RFC 2119 [2]. Each requirement is followed by a description and, if available, an example. Many examples are taken from the social network companies, which – although their behaviour with respect to changing privacy policies and settings is often criticized [14, 15] – have developed quite an impressive range of tools to control privacy settings. The requirements are grouped and ordered in a logical manner.

3.1 Control

> ***User must have control over the privacy of context information***
>
> ***Level: must have***

Where privacy used to be about secrecy, nowadays control is one of the factors that can reduce privacy concerns. Many users don't mind sharing personal information as long as they control how, where, when and with whom information is shared [16]. 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 Facebook [17] or the contents of emails, but is also applicable to a user's context information. This results in the prime requirement that users must have control over the privacy of their context information. This will translate to a user, application or service asking for consent before privacy sensitive context information is accessed. Additionally, when taking data ownership into account, it will translate into a user having not only control over its privacy settings, but also over the actual context information the user has provided to the service provider. The user may ask to service provider to remove/delete the provided context information.

In many situations user control may result in a conflict of interest between the user and others requiring access to the information of the user. For example, a patient may want to refuse access to his or her medical data which is needed by a treating physician. Also, employers may require access to the location of professional drivers (e.g. trucks, money transport, and packet delivery and pick-up). In situations like these, user control may not be feasible. However, users should still be made aware of the context information which is acquired, who has access to this data, etc. (see section 3.3).

Examples for asking for consent can be found with many applications installed on a smartphone. As shown in following figure, the application indicates it requires access to 'Your location' in order to function properly. Even though consent is asked, there are several limitations to this method. For example, in many cases the only option for the user to deny access to his location information is to not install the application on his mobile phone (see section 3.8, 3.9 and 3.10). Also, this request for consent does not include enough information for the user to determine whether or not consent should be given (see section 3.3).
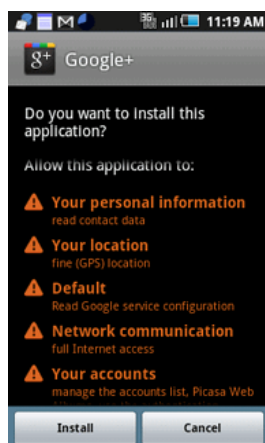

Figure 2: Android - Requested permissions on installation of an app.

Figure 3 shows a consent question from a large scale pilot for SURFnet [18] in which the user has been given more control.

Figure 3: SURFfederatie - Requested permissions for accessing an online service.

3.2 User friendliness – unobtrusiveness

> **Users must not perceive privacy control as annoying or interruptive**
>
> **Level: must have**

When using applications for well-being of well-working these applications may need access to several different types of context and at different times during the day (or night) and with different frequencies. The user should not need to grant or refuse access each and every time access is requested as this will make controlling your 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 and therefore it should be done using a friendly user interface which enables an unobtrusive control of the privacy settings. For example, in case a single services requires multiple aspects of the available context information or multiple services require the same aspect of the context at the same time or a service requires a piece of context multiple times during the day or for a continuous period of time, these questions should be aggregated and thus presented to the user or answered automatically based on the known privacy preferences of the user.

Timed consents [18] as shown in the following figure could be a solution for reducing the obtrusiveness of privacy control questions.

Figure 4: SURFfederatie – Timed consent.

3.3 Easy to understand – informed consent

> **Users must be able to understand the provided privacy controls**
>
> **Levels: must have**

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 [5].

Accordingly, a user should be asked to give his or her informed consent before any context collection. When this consent is asked the user should be informed about the following aspects:

- which context information is requested;
- which parties will have access to this information;
- what the requested information will be used for;
- for how long access will be granted;
- at what interval this information will be acquired;
- which level of detail of the context will be used;
- for how long the requested information will be retained at the requestor (i.e. service provider).

The user should be able to translate this information into short-term benefits and long-term implications to make a well-balanced decision.

From a usability point of view, it will be difficult to present all this information in a form which is acceptable for most users although privacy icons similar to the well-known Creative Commons and other tools to visualize privacy settings are being researched [19, 20]. Different users may need more or less information to make a consent decision (as described in section 3.7) and will accept different default values.

Registering for a new account with for example Google can serve to show how informed consent is difficult to achieve. The user needs to accept the "Terms of Service" and the "Privacy Policy" when registering. The privacy policy can be accessed with an additional two mouse clicks which lead to a website with a statement of more than 1500 words. Most users won't read this policy or won't understand the ramifications of this policy in case they do decide to read it [21, 22].
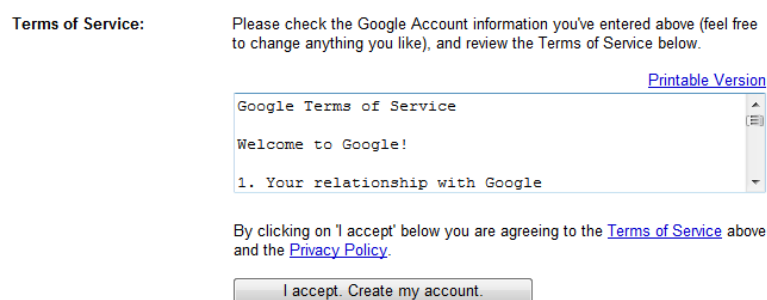

Figure 5: Google – Accepting the Terms of Service.

3.4   Just in time (JIT)

Users must be asked for consent at the time the context is requested (JIT)

Level: nice to have

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

Of course, this requirement is in direct contradiction with the requirement of unobtrusiveness (see section 3.2). In order to achieve both the system will need to learn from responses and thus increasingly develop its privacy policies.

For example, many software applications include an option to remember the response or apply the response to other situations when presenting a question to the user as shown in the following figure.
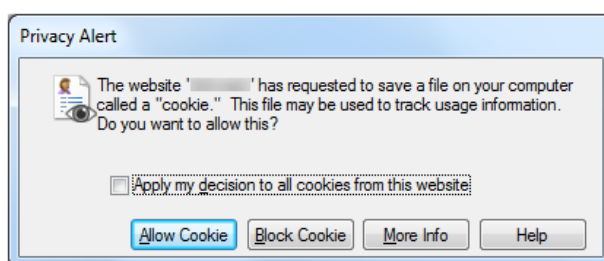

Figure 6: Internet Explorer – JIT consent.

3.5 Overview

> **Users must be able to get an overview of all their context-related privacy control settings**
>
> **Level: must have**

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 these settings and consents, preferably in a single overview. Such an overview will provide insight into the different users, applications or services having access to a user's (aggregated) context information and preferably also into the times and frequency this context information is accessed. The user must also be able to modify given consents at any time.

For example, Facebook provides an overview of applications which have been authorized to interact with a user's account and allows the user to review the specific consents which have been given to each application as shown in the following figures.
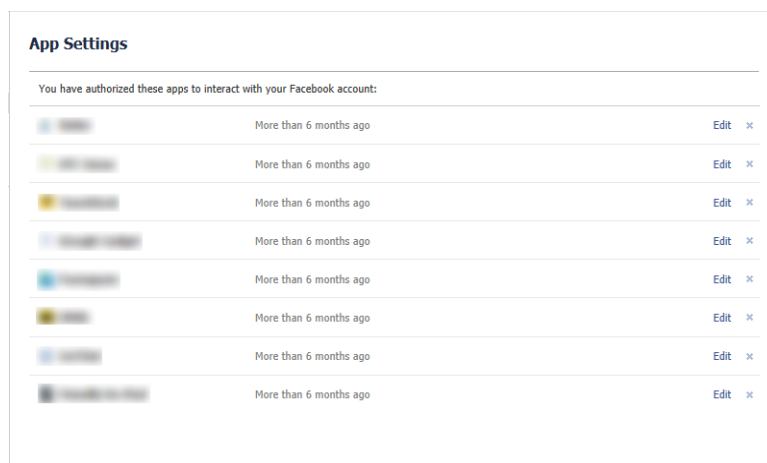
Figure 7a: Facebook – Overview of apps having access to personal information.
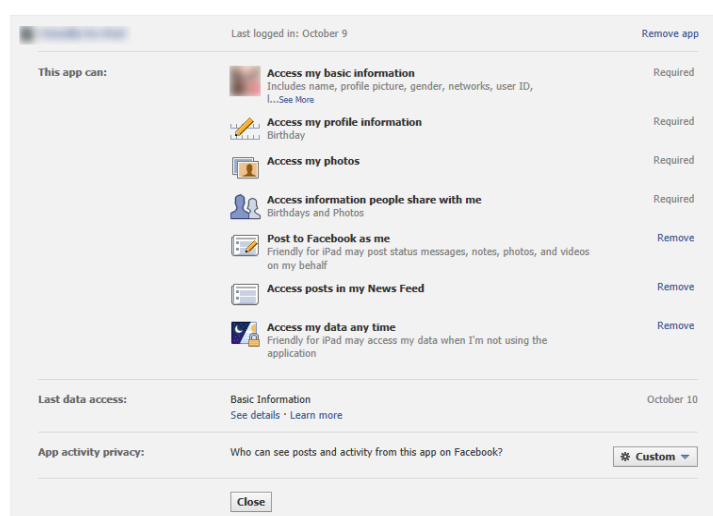
Figure 7b: Facebook – Overview of (general) consents given to a specific app.

> **Users must be able to get an overview of context-related privacy data provided to or accessed by a specific service**
>
> **Level: nice to have**

When users give consent to a service to access some of their context-related privacy data they may not be aware on the frequency this data will be used or the quality of context (see section 3.9) of the information. Therefore, it would be nice to have the possibility to have access to an overview which data is used by a specific service. A step further would be to also have insight in what is derived from the collected data by the context information consuming and/or aggregating parties.

Facebook has implemented a feature which allows the user to show the last time (date) which type of data has been accessed. It does however not show how often it's been accessed, which data has been accessed or how long the retrieved data will be retained.
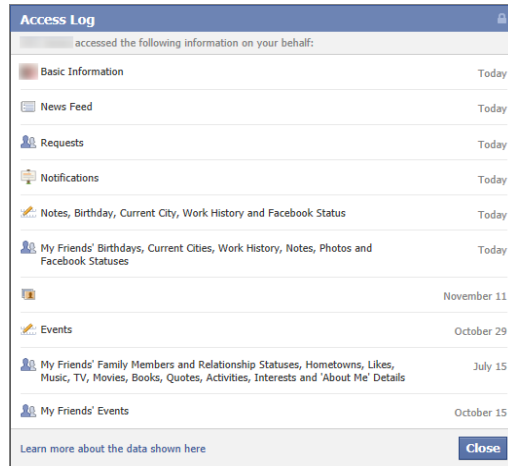
Figure 8: Facebook - Access log for an app with access to personal information.

3.6 Recovery

> **Users must be able to modify and revoke their consents**
>
> **Level: must have**

When a user gives consent to access context-related information the user should be able to revoke or modify this consent at any time. If consent is revocable, research shows that this can reduce risk perception [23]. 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.

As shown in section 3.5 Facebook not only shows the given consents, but also allows the user to revoke consents. LinkedIn [24] also allows prior consents to be revoked as shown in the following figure. While Facebook's control is a bit more fine grained (see section 3.8), LinkedIn's control is limited to take-all or leave-all.
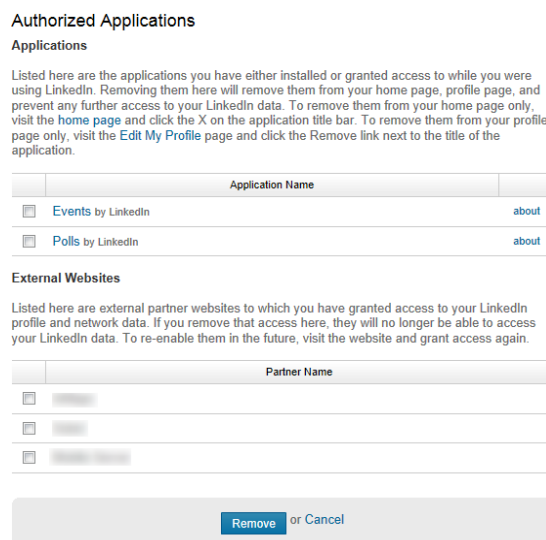
Figure 9: LinkedIn – Revocation of authorized applications.

Another example is Yahoo! Fire Eagle [25]. When subscribing to this service the user is offered the possibility to get a recurring question (by default once a month) to check whether the user is still comfortable sharing location information. This option can be changed at any time within the settings section of this service. Fire Eagle also offers the possibility to purge the current location information and to hide the current location (until the user chooses to "unhide").

Figure 10: Yahoo! Fire Eagle – Purging location information

None of the above examples offer the possibility to purge all the actual context information accessed and stored by service providers nor is any information given on how long the information will be retained.

3.7 Personalised

> ***Privacy policies must be personalised***
>
> ***Level: must have***

Many studies have investigated the attitude of users towards privacy issues. It's generally accepted to classify a person as being a privacy unconcerned (approximately 25%), pragmatist (approx. 50%) or fundamentalist (approx. 25%) [26].The first group of persons is not concerned with privacy at all and is willing to share most information. The second group is aware of both the costs and the benefits of sharing information and considers both before making the decision to share information. The final group is very conservative with privacy and wants to share as few information as possible. For internet users a similar classification has been developed [27] labelling users as unconcerned, circumspect, wary and alarmed. Although this classification shouldn't be used as a predictor for disclosing location information [28], it is found that different users have a different attitude towards privacy.

Additionally, the willingness of a user to give up some of his or her privacy can depend on:

- How much the user trusts the receiver of the privacy sensitive information;
- The privacy sensitivity of the context information;
- The context of the user (i.e. what's the time of the request, who's in the company of the user, what's the user doing at the time of the request, etc.);
- The experienced level of control over the provided context information;
- The benefit or usefulness for the user of sharing this information.

Service Providers may try to influence the last aspect by tempting the user to share information using (financial) awards. This may result in a system in which a wealthy user has more privacy than a poor user as the poor user might sooner be tempted to sell his privacy-sensitive context information.

All of the above factors are very personal – what is useful to one person is not useful for another. Even the interpretation of privacy policies can be personal [29]. It is therefore recommended to start with tight default settings and allow the user to personalize these settings.

Most applications do not personalize their policies themselves, but do offer the possibility of personalization by allowing control over privacy settings.

For example, Facebook allows you to set a wide variety of privacy options such as what information is shared with people and apps (see Figure 11). However, these privacy settings are sometimes hard to find and difficult to interpret. Additionally, Facebook changes its privacy policy and default privacy settings from time to time [30], making it difficult for people to keep track of their personal privacy settings.

For people to gain trust in an application, it is important that their personal information is handled with care. Personalised privacy policies may increase the confidence potential users have in an application.

Figure 11: Privacy settings in Facebook.

3.8 Fine-grained control

Users must have fine grained privacy control

Level: nice to have

The different attitude towards privacy (see section 3.7) can also be translated to a requirement with respect to the level of control a user wants to exert with respect to privacy settings.

Following the common practice of many applications different levels could be defined for controlling ones privacy: normal, detailed and expert. Each higher level could add more detail into the control mechanisms. Expert users may want to control access to their context information based on more aspects than normal users. The default privacy control settings should be set to minimal disclosure and thus maximal protection of the user's privacy. It's then up to the user to adapt these settings to personal preferences.

Although using multiple levels of detail is quite common for (configuring) software applications, it is not (yet) seen implemented for privacy control settings. The following figures show an example of using multiple levels from Microsoft Internet Explorer. The user can modify how the browsers handles cookies using a high-level setting (Allow all cookies, Low, Medium, Medium High, High, Block all cookies), but also allows more advanced users to override the automatic settings and specify the rules for specific types of cookies. It is also possible to manually allow or block cookies from certain websites.
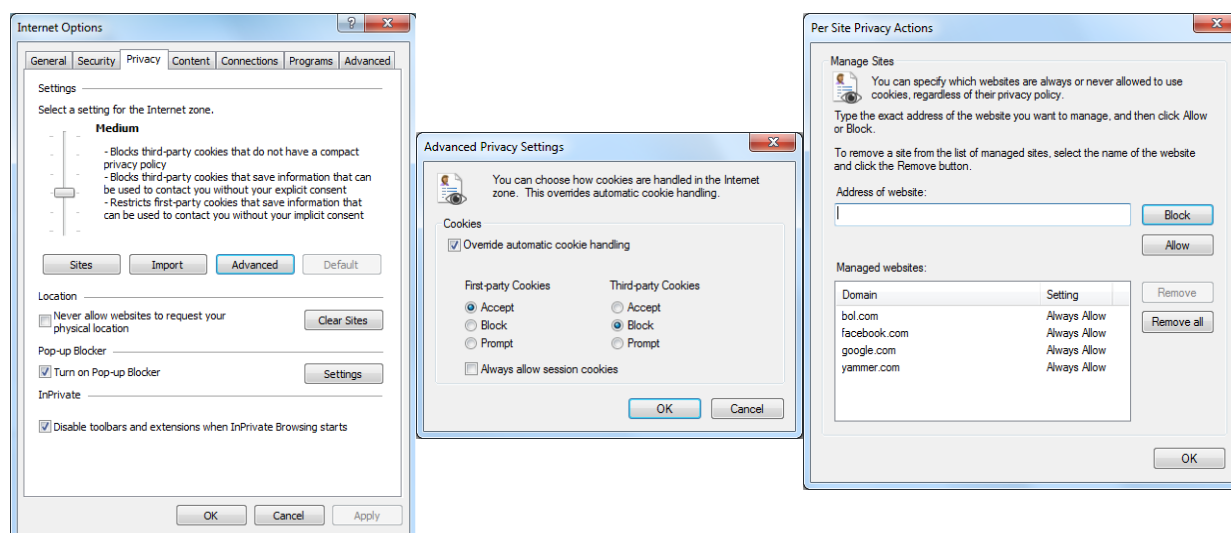
Figure 12: Internet Explorer – Fine grained privacy control.

3.9   Quality of context (QoC) control

> **Users must be able to define the granularity of the context information**
>
> **Level: nice to have**

Besides the level of detail of the privacy control settings (see section 3.8) also the context information itself can be more or less detailed. Users should have control of this granularity. With respect to location users may want to provide their exact GPS location or maybe just a (descriptive) derivative: home or work, neighbourhood or city or region or country. Similar granular levels (Quality of Context) can also be defined for other context information. Each context provider should specify all the levels of granularity of context information it is able to provide, and each context consumer should specify all the levels of context information it is able to consume.

However, presenting a list of all the potential levels for a specific context type may result in an unusable and useless user interface. Moreover, for some services the specific quality of context may not be relevant as long as other users will recognize it.

Currently, Android apps can be allowed access to two levels of location information where the 'fine' information is based on GPS coordinates and the 'coarse' information is based on network information. Even though this different level of detail exists, users do not have the option to select one or the other as shown in the following figure (and applications often require access to both). Yahoo! Fire Eagle on the other hand does allow the user to define the QoC for the provided location information.

Figure 13a: Android – QoC for location information.
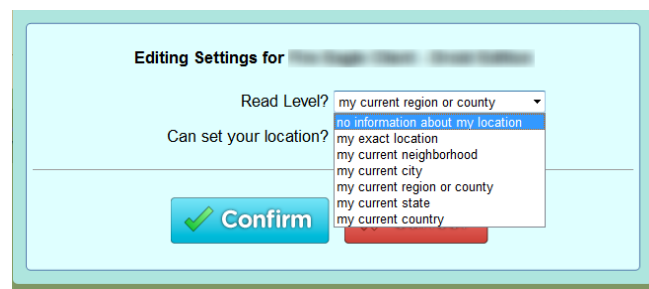
Figure 13b: Yahoo! Fire Eagle – QoC for location information.

3.10 Required and optional attributes

> *Any context-related privacy control solution must support required and optional attributes*
>
> *Level: nice to have*

There is a growing consensus [31, 32, 33] around an interpretation of EU privacy regulations that each service provider might have required attributes as well as optional attributes. The regulations seem to require that the user must be able to individually approve for the release of each of the optional attributes and must inform the user of the required attributes.

Figure 3 shows an example of how it's possible to make a distinction between required and optional attributes, although it would even be better when the benefit for providing the additional, optional attributes would be made clear.

3.11 Multiple context providers for the same context information

> *The user must be able to determine which context provider should provide the requested context*
>
> *Level: nice to have*

When multiple context providers can provide the same type of context the user should have the option to decide which provider to use.

For example, location information could be provided by the telecom operator (based on triangulation calculations), a GPS sensor in the user's smartphone, an ISP (based on the IP-address) or a location broker like Google Latitude, Foursquare, Yahoo! Fire Eagle or one of the many others. Different providers could have different features, performance, reliability and different policies with respect to storage and data ownership of location information. This could influence the user in deciding which provider to use for which service.

3.12 Combined / aggregated data

> ***Aggregation of context data must be made explicit in privacy policies***
>
> ***Level: nice to have***

Objectively measured data may not be privacy sensitive, but if this data is interpreted as "stressed" or "not stressed" this gives a new dimension to privacy sensitivity. These interpretations may or not be correct resulting in an uncertainty of the correctness of the data that is send. Users need to be aware of which data is automatically interpreted and how certain the application is about the correctness of this interpretation.

In the same way, a single context attribute may not be privacy sensitive. However if it is combined with other context attributes it may result in privacy issues [34].

For example, knowing that stress increases may or not may not be privacy sensitive in itself. However if this information is combined with the person with whom the user is speaking, this will be much more privacy sensitive, since this may lead to interpretations. Users need to be aware of possible consequences of sharing combinations of information. Another example of a combination of anonymous sources resulting in privacy issues can be found in the "AOL search data leak" in which anonymized search queries were linked to specific users [35].

3.13 Anonymisation

> ***Users must be able to provide privacy sensitive context information anonymously***
>
> ***Level: nice to have***

Even when the user gets personalised services from a service provider by providing privacy sensitive context information, it should be possible to provide this information without disclosing the user's identity to the service provider using anonymisation or pseudonymisation.

For example, when using a weather-app on a mobile phone, the application just requires information regarding the location of the user to provide the service of showing the local weather information. The user can thus remain anonymous. It should be noted that anonymity may limit the added value of personalised services [36].

3.14 Technical compatibility

> *Any context-related privacy control solution must be compatible to common standards*
>
> *Level: nice to have*

Although the technologies currently used for access control may not be equipped to handle the above user requirements with respect to privacy control, any privacy control solution should be compatible with industry standards and major ad hoc implementations. There are a number of technologies which may be used for privacy control, although each of these has its own limitations [37]. Possible technologies may be OAuth 2.0, OpenID 2.0, OpenID Connect, OStatus, OpenSocial and others.

3.15 Decentralised privacy control

> *Any context-related privacy control solution must allow for decentralised control and enforcement of privacy policies*
>
> *Level: nice to have*

As context providers will typically provide only one type of context information, multiple context providers from multiple domains will be involved in a single service. This leads to the requirement that the exchange of context information should be decentralized to avoid compromising the privacy of the user.

Also, in order to scale, distributed privacy policy management is required. Privacy policies in all domains must be adhered and enforced [38].

4   Summary

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. Every solution for privacy control will need to adhere to numerous requirements indifferent of whether it's for well-being or well-working applications. Even though the actual privacy control settings will be different for well-being and well-working applications both will require users to be in control of his or her personal context information. It seems inevitable that both applications will require a trade-off between the control of the user and others requiring access to the context data of this user.

The willingness of users to share context information is not only personal, but also depends on the context of the user. For example, normally a person might not want to share their complete medical history with anyone. However, in a life-threatening situation, the same person is likely to be more than willing to share that information if it helps to save his/her life. Similarly, some people might not want to share their eating behaviour during work with their employer or colleagues, but they may wish to share this information with their friends and family when at home.

Therefore, it would be interesting 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. 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.

Facebook and Google+ [39] already experiment with limited forms of context aware privacy policies. By adding friends to so called circles or lists the user can determine what this particular friend is allowed to see. This way the user can distinguish between work-related posts that are meant for colleagues or pictures of parties that the user only wants to share with friends. The context awareness is manual, i.e. the system is not aware of the context automatically, but the user is allowed to add context (friends, family, colleagues) to pictures and messages to adapt the privacy policy as shown in the following figure.

Figure 14: Facebook – Context managing.

Obviously, these manual privacy control systems are not very flexible, intuitive and user friendly. Making them more adaptive to changing contexts is challenging. An important aspect which should

not be overlooked when using context-aware adaptive privacy policies is the possibility to fake context information and as a consequence to manipulate the release of personal context information. For example, if it's possible to manipulate a user's heart rate and thus fake a heart attack, the user's medical records may become accessible by all users in the direct surroundings of this user according to his or her context aware privacy policies. By manipulating the location information of this user, the attacker need not even be close to the user to access the user's medical records.

The next step is to identify possible approaches and current best-practices for context-aware adaptive privacy control. Although social networking sites like Facebook, Google Plus, LinkedIn are often criticized for their lack of respect for the privacy of their users, these sites provide inspiring examples on how privacy can be controlled. The challenge is to extrapolate the lessons learnt from these examples into the domain of context aware adaptive privacy control. Furthermore, criteria for using context information for privacy policy management should be established. For instance the trustworthiness, integrity, confidentiality, quality, and availability of the context information should be taken into account when used for privacy control.

5   References

[1] Wegdam, M. et al, "Empowering users to control their privacy in context-aware systems through interactive consent", Freeband AWARENESS deliverable Dn3.21, 2008 (https://doc.novay.nl/dsweb/Get/Document-107335/Dn3.21%20-%20Interactive%20consent.pdf)

[2] "Request for Comments 2119", IETF, March 1997 (http://www.ietf.org/rfc/rfc2119.txt)

[3] Antón, A.I. et al, "Analyzing Website Privacy Requirements Using a Privacy Goal Taxonomy", Proceedings of the 10th Anniversary IEEE Joint International Conference on Requirements Engineering, p.23-31, September 09-13, 2002

[4] "Privacy Online: Fair Information Practices in the Electronic Marketplace, a report to Congress", Federal Trade Commission, 2000 (http://en.wikipedia.org/wiki/FTC_Fair_Information_Practice)

[5] "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data", OECD, September 23, 1980 (http://www.oecd.org/document/18/0,3746,en_2649_34255_1815186_1_1_1_1,00.html)

[6] Data Protection Directive (officially "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, October 24, 1995 (http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:NOT)

[7] Directive on privacy and electronic communications (officially "Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector"), EU, July 12, 2002 (http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0058:EN:NOT)

[8] Cameron, K. "The Laws of Identity", November 5, 2005 (http://www.identityblog.com/?p=354)

[9] "Security and Privacy Controls for Federal Information Systems and Organizations", NIST Special Publication 800-53, Draft Appendix J (http://csrc.nist.gov/publications/drafts/800-53-Appdendix-J/IPDraft_800-53-privacy-appendix-J.pdf)

[10] "Generally Accepted Privacy Principles (GAAP)", AICPA/CICA, May 2009 (http://www.aicpa.org/interestareas/informationtechnology/resources/privacy/generallyacceptedprivacyprinciples)

[11] "APEC Privacy Framework", Asia-Pacific Economic Cooperation (APEC), 2004 (http://www.apec.org/Groups/Committee-on-Trade-and-Investment/~/media/Files/Groups/ECSG/05_ecsg_privacyframewk.ashx)

[12] "Safe Harbor Privacy Principles", U.S. Department of Commerce, July 21, 2000 (http://export.gov/safeharbor/)

13[13] Langheinrich, H., "Privacy by Design – Principles of Privacy-Aware Ubiquitous Systems", Proceedings of the 3rd international conference on Ubiquitous Computing, p.273-291, September 30-October 2, 2001, Atlanta, Georgia, USA

[14] Liz Gannes in "The Apologies of Zuckerberg: A Retrospective" on AllThingsD.com, November 29, 2011

[15] Tim Bradshaw in "The first Google+ privacy flaw" on Financial Times Tech Hub, June 29, 2011 (http://blogs.ft.com/fttechhub/2011/06/google-plus-privacy-flaw/)

[16] Bruce Schneier in "Google and Facebook's Privacy Illusion" on Forbes.com, April 6, 2010 (http://www.forbes.com/2010/04/05/google-facebook-twitter-technology-security-10-privacy.html)

[17] Facebook (http://www.facebook.com)

[18] Wegdam, M. et al, "User controlled privacy for the SURFfederatie: the user perspective", SURFnet, 2011 (http://www.surfnet.nl/nl/Innovatieprogramma's/gigaport3/Documents/20110125%20-%20GP3-UCP-2010-1.pdf)

[19] Hansen, M., "Putting Privacy Pictograms into Practice – a European Perspective", Proceedings of GI Jahrestagung 2009, pp.1703-1716

[20] Kelley, P.G., et al, "Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach", In Proceedings of the International Conference on Human Factors in Computing Systems (CHI), ACM, p1573–1582 (http://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab09014.pdf)

[21] O'Dell, J., "The Real Reason No One Reads Privacy Policies [Infographic]", Mashable.com, January 28, 2010 (http://mashable.com/2011/01/27/the-real-reason-no-one-reads-privacy-policies-infographic/)

[22] Morphy, E., "Consumers Trust Brands, Not Policies", CIO Today, January 29, 2004 (citing research at Michigan State University)

[23] PrimeLife EU project (http://www.primelife.eu/)

[24] LinkedIn (http://www.linkedin.com)

[25] Yahoo! Fire Eagle (http://fireeagle.yahoo.net)

[26] Kumaraguru, P. et al., "Privacy Indexes: A Survey of Westin's Studies", Institute for Software Research International, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, Technical Report CMU-ISRI-5-138, December 2005 (2005) (http://reports-archive.adm.cs.cmu.edu/anon/isri2005/CMU-ISRI-05-138.pdf)

[27] Sheehan, K., "Toward a Typology of Internet Users and Online Privacy Concerns", The Information Society, Vol. 18, No. 1, 2002, pp.21-32

[28] Consolvo, S. et al., "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

[29] Gnanasambandam, N. & Staddon, J., "Personalized Privacy Policies: Challenges for Data Loss Prevention", AAAI Spring Symposium Series, 2010

[30] Kurt Opsahl in "Facebook's Eroding Privacy Policy: A Timeline" on Electronic Frontier Foundation, April 28, 2010 (https://www.eff.org/deeplinks/2010/04/facebook-timeline)

[31] "Attribute Release Recommendations", Terena, accessed on November 28, 2011 (https://refeds.terena.org/index.php/Attribute_Release_Recommendations)

[32] "Technical Specifications and Attribute Specification", HEAL-Link Federation Authentication and Authorization Infrastructure (AAI) (http://www.heal-link.gr/journals/aai/docs/HEAL-LinkTechnicalSpecs.pdf)

[33]"AAI - Authentication and Authorization Infrastructure - Attribute Specification", version 1.3, SWITCH, June 2010 (https://www.switch.ch/aai/docs/AAI_Attr_Specs.pdf)

[34] Beresford, A. et al., "Location Privacy in Pervasive Computing", IEEE Pervasive Computing, vol. 2(1): pp. 46-55, 2003

[35] "AOL Search Data Leak" on Wikipedia, accessed on December 7, 2011 (http://en.wikipedia.org/wiki/AOL_search_data_leak)

[36] Dritsas, S. et al., "Protecting privacy and anonymity in pervasive computing: trends and perspectives", Telematics and Informatics 23, 196-210, 2006

[37] Grossfeld D. et al, "Privacy and Protection of Personal Data using Open Social Web Technologies", Proceedings of Federated Social Web 2011 (http://d-cent.org/fsw2011/agenda/papers/)

[38] Hesselman, C, Eertink, H., and Wibbels, M., "Privacy-aware Context Discovery for Next Generation Mobile Services", 3rd SAINT2007 Workshop on Next Generation Service Platforms for Future Mobile Systems (SPMS 2007), Hiroshima, Japan, January 2007 (http://www.novay.nl/publicaties/privacy-aware-context-discovery-for-next-generation-mobile-services/11444)

[39] Google Plus (http://plus.google.com)

6   Appendix

6.1   Laws of Identity

Kim Cameron described the following seven laws of identity [8] which explain the successes and failures of digital identity systems:

1. User control and consent: information identifying a user can only be revealed with the user's consent.
2. Minimal disclosure for a constrained use: only the minimally required identifying information must be released.
3. Justifiable parties: identifying information should only be released to those parties which actually need to use it.
4. Directed identity: discovery of identities must be facilitated without unnecessary release of correlation handles.
5. Pluralism of operators and technologies: multiple technologies from multiple providers must be allowed to operate.
6. Human integration: the human user is integral part of the system and should be protected against identity attacks.
7. Consistent experience across contexts: users must have a consistent experience through multiple operators and technologies.

6.2   Privacy controls (NIST)

The privacy controls outlined in the publication of NIST [9] are primarily for use by organizational privacy officials when working with program managers, information system developers, information technology project staff, and information security personnel to determine how best to incorporate effective privacy protections and practices (i.e., privacy controls) within those programs and/or systems. These controls facilitate the organization's efforts to comply with privacy requirements affecting those programs and/or systems that collect, use, maintain, share, or dispose of personally identifiable information (PII).

Cntl No.	Privacy Controls
TR	Transparency
TR-1	Privacy Notice
TR-2	Dissemination of Privacy Program Information
IP	Individual Participation and Redress
IP-1	Consent
IP-2	Access
IP-3	Redress
IP-4	Complaint Management
AP	Authority and Purpose
AP-1	Authority to Collect
AP-2	Purpose Specification
DM	Data Minimization and Retention
DM-1	Minimization of Personally Identifiable Information
DM-2	Data Retention and Disposal
UL	Use Limitation

Cntl No.	Privacy Controls
UL-1	Internal Use
UL-2	Information Sharing
UL-3	System Design and Development
DI	Data Quality and Integrity
DI-1	Data Quality
DI-2	Data Integrity
SE	Security
SE-1	Inventory of Personally Identifiable Information
SE-2	Privacy Incident Response
AR	Accountability, Audit, and Risk Management
AR-1	Governance and Privacy Program
AR-2	Privacy Impact and Risk Assessment
AR-3	Privacy Requirements for Contractors and Service Providers
AR-4	Privacy Monitoring and Auditing
AR-5	Privacy Awareness and Training
AR-6	Privacy Reporting

Table 1: NIST – Privacy control principles

6.3 GAPP

The following are the (ten) Generally Accepted Privacy Principles [10] from AICPA/CICA:

1. Management. The entity defines, documents, communicates, and assigns accountability for its privacy policies and procedures.
2. Notice. The entity provides notice about its privacy policies and procedures and identifies the purposes for which personal information is collected, used, retained, and disclosed.
3. Choice and consent. The entity describes the choices available to the individual and obtains implicit or explicit consent with respect to the collection, use, and disclosure of personal information.
4. Collection. The entity collects personal information only for the purposes identified in the notice.
5. Use, retention, and disposal. The entity limits the use of personal information to the purposes identified in the notice and for which the individual has provided implicit or explicit consent. The entity retains personal information for only as long as necessary to fulfil the stated purposes or as required by law or regulations and thereafter appropriately disposes of such information.
6. Access. The entity provides individuals with access to their personal information for review and update.
7. Disclosure to third parties. The entity discloses personal information to third parties only for the purposes identified in the notice and with the implicit or explicit consent of the individual.
8. Security for privacy. The entity protects personal information against unauthorized access (both physical and logical).
9. Quality. The entity maintains accurate, complete, and relevant personal information for the purposes identified in the notice.
10. Monitoring and enforcement. The entity monitors compliance with its privacy policies and procedures and has procedures to address privacy related complaints and disputes.

6.4   APEC

The Data Privacy Sub-Group developed the APEC Privacy Framework [11]. In developing the Framework the Sub-Group took as its starting point the 1980 OECD Guidelines on the Protection of Privacy and Trans-Border Flows of Personal Data, which also underpins the *Privacy Act 1988*. The Framework contains nine high level Privacy Principles, as follows:

1. Preventing Harm
2. Notice
3. Collection Limitations
4. Uses of Personal Information
5. Choice
6. Integrity of Personal Information
7. Security Safeguards
8. Access and Correction
9. Accountability

6.5   Safe Harbor

The US Department of Commerce drew up the Safe Harbor Privacy Principles [12] in conjunction with the EU in order to prevent US companies from accidental information disclosure or loss. The seven Principles are:

1. Notice - Individuals must be informed that their data is being collected and about how it will be used.
2. Choice - Individuals must have the ability to opt out of the collection and forward transfer of the data to third parties.
3. Onward Transfer - Transfers of data to third parties may only occur to other organizations that follow adequate data protection principles.
4. Security - Reasonable efforts must be made to prevent loss of collected information.
5. Data Integrity - Data must be relevant and reliable for the purpose it was collected for.
6. Access - Individuals must be able to access information held about them, and correct or delete it if it is inaccurate.
7. Enforcement - There must be effective means of enforcing these rules.

7 Abbreviations

Abbreviation	Description
AICPA	American Institute of Certified Public Accountants
APEC	Asia-Pacific Economic Cooperation
CAAP	Context Aware Adaptive Privacy
CICA	Canadian Institute of Chartered Accountants
ECG	Electrocardiogram
EU	European Union
FIP	Fair Information Practices
FTC	Federal Trade Commission
GAPP	Generally Accepted Privacy Principles
GPS	Global Positioning System
IETF	Internet Engineering Task Force
JIT	Just In Time
NIST	National Institute of Standards and Technology
OECD	Organisation for Economic Co-operation and Development
PII	Personal Identifying Information
QoC	Quality of Context
US	United States of America