DIW Berlin

German Institute
for Economic Research

Discussion Papers

698

Sören Preibusch
Bettina Hoser
Seda Gürses
Bettina Berendt

# Ubiquitous Social Networks: Opportunities and Challenges for Privacy-Aware User Modelling

Berlin, June 2007

Opinions expressed in this paper are those of the author and do not necessarily reflect views of the institute.

IMPRESSUM

DIW Berlin
German Institute for Economic Research
Königin-Luise-Str. 5
14195 Berlin
Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
http://www.diw.de

ISSN print edition 1433-0210
ISSN electronic edition 1619-4535

Available for free downloading from the DIW Berlin website.

# Ubiquitous social networks – opportunities and challenges for privacy-aware user modelling

Sören Preibusch[1] Bettina Hoser[2]
Seda Gürses[3] Bettina Berendt[4]

**Abstract**. Privacy has been recognized as an important topic in the Internet for a long time, and technological developments in the area of privacy tools are ongoing. However, their focus was mainly on the individual. With the proliferation of social network sites, it has become more evident that the problem of privacy is not bounded by the perimeters of individuals but also by the privacy needs of their social networks. The objective of this paper is to contribute to the discussion about privacy in social network sites, a topic which we consider to be severely under-researched. We propose a framework for analyzing privacy requirements and for analyzing privacy-related data. We outline a combination of requirements analysis, conflict-resolution techniques, and a P3P extension that can contribute to privacy within such sites.
**Keywords**: World Wide Web, Privacy, Social Network Analysis, Requirements Analysis, Privacy Negotiation, Ubiquity, P3P
**JEL classification**: C8, L86

First published in the Proceedings of the Workshop on Data Mining for User Modelling at UM 2007, Corfu, Greece, June 2007. http://vasarely.wiwi.hu-berlin.de/DM.UM07/ http://vasarely.wiwi.hu-berlin.de/DM.UM07/Proceedings/05-Preibusch.pdf

[1] German Institute for Economic Research (DIW Berlin), Mohrenstraße 58, 10117 Berlin, Germany, spreibusch@diw.de
[2] Universität Karlsruhe (TH), Institute of Information Systems and Management, Kaiserstraße 12, 76128 Karlsruhe, Germany, hoser@iism.uni-karlsruhe.de
[3] Humboldt University Berlin, Institute of Information Systems, Spandauer Str. 1, 10178 Berlin, Germany, seda@wiwi.hu-berlin.de
[4] Humboldt University Berlin, Institute of Information Systems, Spandauer Str. 1, 10178 Berlin, Germany, berendt@wiwi.hu-berlin.de

# Contents



# List of Figures

# 1 Introduction

With networked computers becoming more and more ubiquitous around the globe, digital social networks are gaining increasing importance for many people's work and leisure, as they allow for interaction independently of a fixed location. In parallel with their huge and growing acceptance among a wide range of users, social networks (SNs) are becoming a focus of attention for researchers and practitioners (especially in marketing). Also, governments and law enforcement (re-)awaken to the need to analyze the SNs of terrorists and other criminals [11]. What is important to all three groups is the huge amount of knowledge that can be discovered by investigating people's textual/multimedia contributions to SNs and the links they set to their "friends" – in this sense, social network analysis is an important topic for Knowledge Discovery for Ubiquitous User Modelling. (In this paper, we focus on SNs on the Web and thus on the ubiquity of the Web; see [8] for a differentiation between different notions of "ubiquity" that are relevant for user-centric analyses.[1])

Surprisingly, a topic that has received a lot of attention over the last years in all other areas of computer and Internet use, is scarcely attended to in current discussions on SNs: privacy. In the privacy statements of social network sites (SNSs), it appears that SNs are just another application on the Web (where "of course your privacy is very important to us"); the implication being that privacy challenges and problems are comparable to other Web applications, such as eCommerce, and therefore can be solved with the same privacy preservation methods.

In this position paper, we argue that while SNs share many privacy problems (and therefore solution possibilities) with other Web applications, there are also important new challenges. Using some simple examples, we highlight the extent of the current commercial interest in SN, point to the interest in SNs in ubiquitous computing environments, and discuss the resulting new challenges. Finally, we outline new research directions for currently existing methods for privacy-preserving data mining / data analysis.

[1] Many SN platforms are currently moving to mobile environments, e.g. [16]; in a separate paper, we will investigate how the issues raised here resurface or change in mobile SNs.

# 2 Background: The importance of social networks

In this section, we want to give an impression of the currently perceived importance of SNs. To this end, we focus on those examples that have recently received the most attention, in particular in terms of the monetary magnitude of takeover deals.

MySpace has grown to be the largest SNS in the world.[2] News Corp. invested 580 mio. USD in MySpace in mid 2005 [6], and one year later, Google signed a 900 mio. USD deal with News Corp. for the search feature [7].

Flickr[3] and del.icio.us[4] are both Yahoo!-owned; LinkedIn and Facebook are other prominent examples of SNs.

In Germany, two deals happened in the last year. First the SNS known as OpenBC went public and changed its name to Xing[5]. It has 1.5 million users; their market capitalization reached 164 mio. Euro [27]. The second deal was even more interesting. The publishing company Holtzbrinck has recently bought StudiVZ[6], a student community with more than 1 million accounts, for around 100 mio. Euro [33].

The next step appears to be Second Life[7]. This Web site is evolving into a parallel world, as more and more companies, universities, and users join. What differentiates this site from all other SNSs is its own flourishing economy. People earn real money within this virtual world. Second Life is likely to generate an enormous und unprecedented amount of social-network data.

The first interesting question behind all these deals is the economic rationale. Marketing and advertisement appear to be the major trigger behind

[2]Estimates of the real number of users vary widely. While a much-cited blog of August 2006 stated that the threshold of 100 million accounts had been surpassed [12] – a number which was changed in all-too-many subsequent articles into more than 100 million *users*, an analysis of 303 random accounts showed that only between 30 and 40% of accounts are likely to belong to real users [4].

[3]close to 7 million accounts as of 10 Feb 2007, see `http://www.flickr.com/search/people/?q=+`

[4]1 million accounts as of 25 Sep 2006, see `http://blog.del.icio.us/blog/2006/09/million.html`

[5]`http://www.xing.com`

[6]`http://www.studivz.de`

[7]`http://www.second-life.com`

these deals. It seems that companies want to use three characteristics of those sites to their advantage. First, all users voluntarily give information about themselves. This is more information than any company could collect without great expenses. Second, especially in sites for professional SNs like Xing, the company can rely on the correctness of the data, as only a true profile enables successful networking. Finally, networks are made visible through the analysis of simple interactions in the network, and thus provide supporting data sets for validating the classification of potential customers.

# 3 Marketing in social networks

Social Network Analysis has emerged from sociology in the 1970s. But the ground work has been laid in the 1930s when Moreno introduced graph-theoretic approaches to sociology. Since then the analysis of network structures based on mathematical indices has been of growing interest. With the Internet and thus the availability of ever-growing data sets in conjunction with the evolution of computer technology and algorithm design, social network analysts are now capable of analysing structures of large networks of small as well as large networks. This field of research has become highly multi-disciplinary, with research from mathematics, physics, sociology, information sciences and economics, e.g., [3, 19, 22, 23, 38].

The most common use of user data is in marketing, for which profiles, as collected in traditional eCommerce, are supported by data-mining the explicit self-descriptions, the behaviour, and the ratings of users (e.g., Amazon, Yahoo!, Google, and Google Mail). This use is explicitly mentioned, for example, in the MySpace privacy statement: "MySpace.com also collects other profile data including but not limited to: personal interests, gender, age, education and occupation in order to assist users in finding and communicating with each other. [...] MySpace.com also logs non-personally-identifiable information including IP address, profile information, aggregate user data, and browser type, from users and visitors to the site. [...] This non-personally-identifiable information may be shared with third-parties to provide more relevant services and advertisements to members."[8]

Marketing initiatives also actively utilize the relational information in user profiles. (We believe that the under-specification of "profile" in the above privacy statement – 'profile information is information including, but not

[8] `http://www.myspace.com/Modules/Common/Pages/Privacy.aspx` [10 Feb 2007]

restricted to, ...' – legally allows MySpace to subsume network information under the profile that may be handed over to third parties. To the best of our knowledge, no legal investigation or lawsuit on this question has been published.)

Developing a functioning marketing strategy for an SNS requires at least two things: First, to find out how to address people in an environment geared towards "friends", who also tend to be highly Internet-savvy and hence may not respond to traditional forms of marketing. Second, to utilize the information inherent in linkage patterns to discover and target high-value customers.

The first strategy can be subsumed under "Guerrilla marketing": unconventional ways of performing promotional activities, often on a very low budget, with high entertainment value and leaving people unaware that they have been marketed to ("undercover marketing"), see [21]. This is one of the currently most-hyped marketing strategies (see the study by the German Society for Consumption Research, [15]), and recommendations specifically tailored to the SNS MySpace exist [14].

However, these recommendations rely more on the creativity and motivation of marketing employees to engage in an SN, than on the utilization of formal models. The question arises what kind of information *is* contained in the network structure. This is a typical question of social network analysis [37]. Combining social network analysis and data mining, [29] proposed to "mine the network value of customers" and to use this knowledge for "viral marketing". Viral marketing denotes "marketing techniques that use existing SNs to produce increases in brand awareness, through self-replicating viral processes, analogous to the spread of pathological and computer viruses. It can often be word-of-mouth delivered and enhanced online; it can harness the network effect of the Internet and can be very useful in reaching a large number of people rapidly" [40]. The core idea of [29] is to exploit measures of "opinion leadership" inherent in SNs and to translate them into measures of customer value. Thus, they distinguish between a customer's intrinsic value (based on the products s/he is likely to purchase) and the network value (the expectation that s/he has a positive influence on others' probabilities of purchasing).

"Customer network value" is but one example of measures of node importance. In the social network literature, many other measures are currently being discussed; it is beyond the scope of this paper to enter the discussion of their relative merits.

In ubiquitous environments, marketing companies are hoping for even more detailed information. Ubiquitous information is expected to return higher granularity data with strong identifiers like location and time, which not only allow persons to be easily identified, but also their interactions in the social realm to become overt, including their belonging to groups of which they are not even aware of. An example is a specific group of commuters who pass by strategically-placed digital billboards. The collection and dissemination of ubiquitous information will allow advertising and marketing companies to optimally make use of the time and places at which persons may best succumb to advertisement, as well as to identify those groups or individuals best suited for various viral marketing strategies.

# 4 Privacy challenges in social networks

In what sense is all this a privacy problem? First, because being an SNS user implies being a Web (platform) user, all the problems arise that are already well-known and documented in the Internet at large, e.g., [35]. Summarized briefly, personal data accrue and can be utilized not only for the primary purposes for which they were collected (finding and communicating with other users, cf. the MySpace privacy statement).[9] They can be utilized also for secondary (from the perspective of the user) purposes that are covered in the SNS's terms of use and in that sense accepted by users. Such purposes are usually targeted marketing. However, they can also be utilized for other purposes – illegally or legally for commercial purposes, as many examples, for example from eCommerce, show [26], and by law enforcement, secret services, etc. (the explicit targeting also of information marked as 'private' in law-enforcement analyses of data has been confirmed by leading politicians [28]).

Technically, the use of SNS data for novel purposes is even simpler than in traditional eCommerce. The very essence of social media is that user-profile information is public (as opposed to, for example, Amazon's usage data which

[9]Even accesses that at first sight look like a legitimate usage in this sense are not without problems, and people are beginning to be wary of this. The following is a good example of the new intricacies of the shifting notions of "private" and "public": boyd [10] pointed out that many US teenagers (due to their heavy usage of MySpace both the most sophisticated and the most vulnerable users of SNSs today) feel strongly about preserving a certain form of privacy: they want to be visible and searchable for their friends but not their parents.

are an important and secret business asset of the company). Moreover, the data often carries semantic markup and/or is presented in a uniform (hence easily minable) manner, for example as RSS feeds. Thus, while the legal issues at this level are the same in SNSs as in other sites, technical (ab)uses become simpler.

So at first sight, social-network data describe a person in the same way as other data. For example, a "person" record in a database may contain the attributes "health status", "favourite book", and a (probably set-valued) attribute "friend". The values of these attributes are properties of the data subject of this record (say, person A).

For the subsequent analysis, we propose to extend the common classification of confidentiality levels into "private data" and "public data" agreed upon between the customer (user) and the site operator. We propose to use two further levels that we call "community data" and "group data", specific to SNSs:

**Private data** is disclosed to the SNS operator for its internal purposes only. This data must not be disclosed unless explicit consent is given. An example is the user's email address provided upon registration.

**Group data** is disclosed to the SNS operator and can be accessed by other users of the same SNS that are also in the same group as the user: data disclosure is limited to the group. Here we imagine messages shared among a certain group, almost like a closed mailing-list.

**Community data** has been disclosed to the SNS operator and is available to all registered and logged-in users of the SNS. The data is not accessible for anonymous SNS visitors. Examples are the user's online status, her contacts, her member page details, photos, etc.

**Public data** has been disclosed to the SNS operator and is made accessible for all SNS visitors, including anonymous visitors: this may include the fact that the user is registered in the SNS, her user name, or her guestbook.

The concrete details and the application of these confidentiality levels to data depends on the SNSs implementation. One may not always find disclosure examples of all levels.

A priori, the site operator has diverging privacy goals. On the one hand, he needs enough personal user information to be disclosed in order to attract new users. On the other hand, some information must be kept at the community level to create sufficient benefit from community membership. At the time of signing up, the perceived benefits, including access to secured personal information, must exceed registration costs. A typical situation is that one searches for someone's email address by entering her name in a search engine. The contact is found in an SNS like Xing, but the email address is secured to registered users.

The privacy challenges in the Web portrayed so far arise from the operator-user interaction. In the context of SNSs, new problems arise because of the semantics of social-network relations, i.e. the user-user interaction. As an example, consider friendship relations which are – at least in real life – symmetric. Thus, the record of person A that states that person B is a friend also contains information that is part of B's record. Another example is groups of users. Group attributes may be changed by any member of the group. A user whose group membership is public thereby discloses interests, preferences, or other personal information (for a worked-out example, see Section 6). This means that if A discloses information about himself or groups including himself, he (whether willingly or inadvertently) also discloses information about someone else. Expressed differently, A's treatment of his privacy has a direct effect on B's privacy.

Such social-network data usually concern people who also have an ID in the same system, i.e. this privacy dependency is a problem that affects different users of the same system.

In addition, problems arise when systems support the interaction with the world outside the system. For example, Google Mail (Gmail) users consent to their emails' data being analyzed by Google; however, all incoming mails of a Google Mail account (whether sent by another Gmail user or by somebody else) are also analysed. Thus, A's treatment of his privacy also has direct external effects on the privacy of C, who is a non-user of the system.

The distinction between "in the system" and "outside the system" vanishes in case of loosely coupled networks where members may engage in relationships spontaneously and without a central authority. An example are the "friend-of-a-fried" nets built by publishing FOAF files [31]. A FOAF file describes a persons contact information, as well as his/her relationships to other people and details about them in an RDF-based standard format. As users publish their friendship details autonomously, symmetry of relation-

ships is not enforced. However, revelation of private information is likely to occur for instance by combining real names and email addresses, and legal requirements apply [13].

Because SNs are (by definition) built on interaction, they are typically open systems, and have certain semantic characteristics. Each privacy-related declaration has effects beyond the interaction between one individual data subject and one data collector, effects that may concern a number of stakeholders who may or may not be users of the same system.

In a quest for solutions, we identify two essential steps: First, the potential privacy conflicts that arise by social-network interaction must be identified. To do this in a systematic way, methods from requirements analysis are needed. This includes methods for conflict resolution a priori. Second, privacy preferences and requirements must be formalized sufficiently such that software can automatically detect problems, alert the user, and assist her. In data analysis routines, mechanisms need to be implemented to enforce privacy requirements. We believe that this should be based on existing standards or de facto standards for privacy-enhancing technologies (PETs), in order to make a large-scale adoption of such technological solution approaches realistic. In the following two sections, we investigate the two parts of our solution proposal in turn.

This method of analysis draws attention to an important question: what "privacy" actually means. In the following, we emphasize that privacy is not just about data protection, or about restricting the access to, or the processing of, personal data. It is also about who can edit which data (e.g., information about individuals or groups), how people want to and can interact with a site and other users (e.g., identified, pseudonymized, or anonymized), i.e. what different private, public, and shared spaces they can create for their lives, how they can separate and share identities between these spaces, etc. For an extended discussion of our notion of privacy, see [17].

# 5 Identifying privacy conflicts in the interaction of requirements for social network sites

As mentioned in the examples above, identifying privacy conflicts in SNSs is not trivial. In order to do this in a systematic manner, we make use of the Multilateral Security Requirements Analysis (MSRA) method [17]. The main idea of the MSRA method is to consider the security and privacy interests or needs of all stakeholders related to the system. An important aspect of the method is to identify interest conflicts among these stakeholders and develop mechanisms for negotiating these conflicts. Here we introduce aspects of conflict identification and negotiation mechanisms in multilateral security requirements analysis.

## 5.1 Stakeholders and their privacy interests

In MSRA, *stakeholders* of a system are all persons who have some functional, knowledge, security or privacy interest in the system. This encompasses all persons involved in the conception, production, use and maintenance of the system. Stakeholders encompass more than *users* (those who will use the functionality of the system).

Stakeholders, for example, include all persons who have a privacy interest in the system. This could be stakeholders representing legal requirements as well as non-users whose data is processed by the system – i.e. patients in a Hospital Information System or customers in a Customer Relationship Management System. As mentioned in Section 4, the sender of an email to a Gmail account may count as a stakeholder of the Gmail platform, although she is not a user of that platform. This stakeholder is likely to have different privacy interests towards the Gmail platform than a user or provider of the platform.

The inclusion of an external sender of an email as a stakeholder of an email platform also points to the fact that further stakeholders may be acquired once the system is running. Subsuming the privacy interests of all prospective stakeholders is not possible during the development of the system. Nevertheless, the potential of discovering new stakeholders requires the conception of negotiation mechanisms during the development process that anticipate potential divergences in privacy interests during run-time.

Moreover, the introduction of new stakeholders and their requirements often demands a review of all security and other requirements and hence an iterative approach.

The analysis of the stakeholders security and privacy requirements can be compared to viewpoints-oriented requirements analysis [32]. The collection of different privacy interests from the viewpoints of the stakeholders results in a complex list of requirements that are likely to include inconsistencies, repetitions and conflicts. To identify these, requirements interaction management is necessary.

## 5.2 Identification of privacy conflicts through requirements interaction management

Requirements interaction [30] can be understood through direct comparisons of requirements descriptions, or through the analysis of the underlying components that can satisfy these requirements. According to the definition in [30], a requirement $R$ is satisfied by a component $C$ if the component exhibits all the properties specified in the requirement. There may be degrees of satisfaction of a requirements and this can be mapped to a range:

**Definition.** $Sat_R : C \rightarrow [0, 1]$

As a result, requirements interaction can be defined as follows:

**Perceived interaction** : Two requirements, labeled $R_1$ and $R_2$ *interact* if and only if the satisfaction of one requirement affects the satisfaction of the other.

**Operational interaction** : If component $C_1$ satisfies $R_1$ and component $C_2$ satisfies $R_2$, and the run-time behaviour of $C_1$ affects the run-time behaviour of $C_2$, then $C_1$ interacts with $C_2$, and indirectly $R_1$ interacts with $R_2$.

The definition of operational interaction points to the dependency between the requirements and design phases of systems development. Interactions may have different degrees of intensity, and run-time interactions may have varying probabilities of appearing. Interactions between requirements

may be positively correlated (they strengthen each other), negatively correlated (they are in conflict), the correlation may be unspecified (the effect is unclear but exists) or non-existent (no effect).

Privacy requirements can be articulated in terms of security goals [17]. Security goals also interact, and may be correlated positively or negatively. For example, the anonymous use of a resource and the accountability for that use – the possibility to prove to a third party the use of the same resource – are conflicting requirements. Design solutions that partially satisfy both are possible; we will refer to these later.

*Example 1.* In an SNS, the stakeholders may have conflicting interests concerning the authoring and editing of entries:

**R1** The members of the SN may edit parts of entries of other authors that contain information about themselves.

**R2** The authors of entries want to be the sole editors of their own entries.

**R3** All members of the SN want the accountability of authors for all their entries towards all other members.

**R4** All members of the SN want to be able to use the SNS services anonymously.

| | R1 edit others entries | R2 only authors editors | R3 authors accountable | R4 anonymous authors |
|---|---|---|---|---|
| R1 edit others entries | 0 | - | ? | - |
| R2 only authors editors | - | 0 | + | - |
| R3 authors accountable | ? | + | 0 | - |
| R4 anonymous authors | - | - | - | 0 |

+ positive correlation
- negative correlation
? unspecified correlation
0 no correlation

Figure 1: Requirements interaction for a social network

Figure 1 gives an overview of the interactions between these initial requirements for various sets of stakeholders. For example, the anonymity requirement R4 is

obviously in conflict with all the other requirements. If a user uses the services of the SNS anonymously, it is not possible to prove that information in an entry is about oneself (requirement R1), it is not possible to authenticate the users who edited entries through their identities (requirement R2), and accountability for requirements is not possible through user identities (requirement R3). Hence, some negotiation is necessary to resolve the negative and unspecified correlations between the different requirements. Resolutions of conflicts may also introduce new conflicts. Thus, an iterative requirements interaction management approach is needed.

## 5.3 Activities and negotiation techniques

In [30], Robinson et al. suggest six activities:

**Requirements partitioning** : a subset of the requirements are analysed depending on scenarios, stakeholder views etc ("episodes" in MSRA).

**Interaction identification** : the different kinds of correlations between the requirements are identified.

**Interaction focus** : the requirements are prioritized, since not all interactions can be resolved.

**Resolution generation** : different approaches are used to generate resolutions. A value-oriented approach considers alternative goals, whereas a structure-oriented approach considers new operators and resources.

**Resolution selection** : different methods are used to prioritize generated resolutions, for example, utility theory or decision theory.

**Requirements update** : further stakeholders and/or requirements may become apparent through the requirements interaction management process; these are considered in this activity.

Based on their study of approximately conflict resolution 30 methods, the authors suggest the following methods for resolution generation:

**Relaxation** : the conflicting requirements are relaxed or generalized to avoid conflict.

**Refinement** : the conflicting requirements are partially satisfied.

**Compromise** : a compromise is found between the requirements.

**Restructuring** : a set of methods are used to modify the conflict context, which includes assumptions and related requirements.

**Other** : conflict resolution is postponed, either to later stages of the development, or the attempt is abandoned entirely.

*Example 1 (contd.)* In the example, the conflict between the requirements R1 through R3 with the anonymity requirement R4 can be solved with one of these resolution methods. A relaxation of the anonymity requirements can be reached by replacing anonymity by pseudonyms of different strength. Refinement could be reached by allowing certain services to be used anonymously, i.e. authoring reserved entries anonymously but using services for which accountability is important with registered pseudonyms. This could also be seen as a compromise. In restructuring, one could divide the services of the SNS into those which include anonymous interactions, and others which exclude anonymous interactions. Further restructuring could be done through keeping the community so small and protected that anonymity ceases to be a requirement.

Recognizing interactions in privacy and security requirements written in natural language is not a trivial activity. We need an adequate modelling language that makes the identification of interactions easier [24]. Further, the interaction between the high-level security requirements of the stakeholders and the data that are related to these requirements needs to be analyzed, which inevitably requires inference analysis to be undertaken.

# 6 Enhancing privacy in social network sites using P3P

What happens after requirements have been analyzed and conflicts identified? How can technology help to resolve conflicts during run time? In this section, we focus on restructuring: a modification of the SNS application logic (and hence the interaction/conflict context) that can help avoid the occurrence of conflicts. We concentrate on privacy in the sense of data protection, i.e. as a restriction on data access and data processing.

Appropriate measures need to be taken to satisfy privacy requirements in an operational SNS. This includes the conception and adaptation of technologies and processes, mainly privacy languages and tools to interpret and enforce these languages. The design goal is twofold:

First, we need mechanisms to ensure that data/information of one privacy level must not be made accessible via data/information of a lower privacy level. For example one should not be able to perform *(data) inferences* [34] towards personal information that is private on a "community level", from personal information that is private on a "public level". The AOL privacy breach [5] gives evidence that trivial anonymization is insufficient for preventing data inferences that may even lead to the identification of individuals.

Second, we need mechanisms that prevent users from disclosing personal information about other users inside an SNS.

Both objectives should be addressed within the existing technological and legal infrastructure of Privacy Enhancing Technologies (PETs), Privacy Protocols (especially P3P and APPEL / XPref), and mandatory legislation.

P3P, the Platform for Privacy Preferences, is a protocol designed to inform Web users about the data-collection practices of Web sites. It provides a way for a Web site to encode its data-collection and data-use practices in a machine-readable XML format known as a P3P policy [39]. Moreover, P3P enables Web users to understand what data will be collected by sites they visit, how that data will be used, and what data/uses they may "opt-out" of or "opt-in" to [39]. An SNS operator will post a P3P policy on its Web site to communicate its data handling practices. Visitors and users can receive this policy in a textually presented format. Their decision whether to send data to the site or not can be supported by APPEL rules: APPEL, A P3P Preference Exchange Language, allows a user to express her preferences in a set of preference rules, interpreted by her user agent to make automated or semi-automated decisions regarding the acceptability of P3P Privacy Policies [20]. XPref [1] is a newer privacy preference language, more expressive than APPEL yet easier to use.

In a P3P policy, one or several statements describe data practices that are applied to particular types of data. A statement indicates recipients, usage purposes, and a retention time for data elements. Every potential data usage must be indicated by an appropriate statement; hence statements span a superset over the actually implemented data usage. P3P hereby translates the privacy concepts of, e.g., European privacy legislation and the OECD Fair Information Practices into a machine-readable policy.

*Example 2.* Consider the P3P fragment below, which expresses the data collection and usage scenario outlined in section 4. A professional SN collects the username, publicly accessible, and the details about the user's job, the latter being secured. Users may join special interest groups based on their industrial and departmental focus, e.g. "Helpdesk Professionals Group", "Data Protection Officers Group", or "CEO VIP Club". Group membership, expressed by the data categories `<political/><preference/>` is public.

Listing 1: P3P Policy fragment

```
<STATEMENT>
  <PURPOSE>     <current/>  </PURPOSE>
  <RECIPIENT>   <ours/>
                <public/>  </RECIPIENT>
  <RETENTION>   <indefinitely/>  </RETENTION>
  <DATA-GROUP>  <DATA ref="#user.login.id"/>
                <DATA ref="#dynamic.miscdata">  <CATEGORIES>
                <political/> <preference/>  </CATEGORIES>  </DATA>  </DATA-GROUP>
</STATEMENT>
<STATEMENT>
  <PURPOSE>     <current/>  </PURPOSE>
  <RECIPIENT>   <ours/>  </RECIPIENT>
  <RETENTION>   <indefinitely/>  </RETENTION>
  <DATA-GROUP>  <DATA ref="#user.login.password"/>
                <DATA ref="#user.jobtitle"/>
                <DATA ref="#user.business.employer"/>
                <DATA ref="#user.business.department"/>  </DATA-GROUP>
</STATEMENT>
```

However, group details must be public so that users can decide whether they want to join a given group. Even if these details are hidden, the group name is often explicit enough ("Data Protection Officers Group").

Thus, we can formulate the following inference rule with infix relation notation of is_member_in, focusses_on, and works_in:

$$\forall\, g{:}Group, u{:}User, d{:}Department :$$
$$u \text{ is\_member\_in } g \wedge g \text{ focusses\_on } d \Rightarrow u \text{ works\_in } d$$

Using this data inference rule, one can infer a user's department from her group membership details; the group details are public and can be accessed freely. The confidentiality of the user's department is not guaranteed any more, and the P3P Policy does not accurately reflect that the recipient of the department information is effectively broadened to `<public/>`.

To avoid such degradation of privacy levels, we have proposed an extension to P3P, the INFERENCE element, together with a logic that blocks the use of data described in the INFERENCE [36]. The new INFERENCE

element, realized by P3P's built-in extension mechanism and thus backward-compatible, codes a data inference inside the P3P policy. A user-agent may parse the inference rule and alert the user to possible privacy breaches. Inside an analysis framework, inference rules can be used to lift privacy levels. For instance, access to the group membership information should be restricted to `<ours/>`.

*Example 2 (contd.)* In P3P, the inference rule is coded as follows:

Listing 2: P3P Policy Extension for coding inferences

```
<EXTENSION optional="no">
<INFERENCES xmlns="http://preibusch.de/namespaces/SIMT/inferences">
  <INFERENCE>
    <CONSEQUENCE> If group membership is known, group details let
                  conclude on the user's details. </CONSEQUENCE>
    <GIVEN>  <AND>
      <DATA-GROUP>
        <DATA ref="#dynamic.miscdata">
          <CATEGORIES>  <political/>  <preference/>  </CATEGORIES>  </DATA>
      </DATA-GROUP>  </AND>
    </GIVEN>
    <INDUCED>
      <DATA-GROUP>  <DATA ref="#user.business.department"/>  </DATA-GROUP>
    </INDUCED>
  </INFERENCE>
</INFERENCES>
</EXTENSION>
```

*Example 3.* We now consider the second problem of personal information about oneself to be disclosed by other users. Again, we observe `<public/>` as a new recipient where a higher privacy level was intended. As a remedy, the users A and B have to agree on a privacy policy that B will not disclose their friendship. Note that a privacy policy between A and the SNS operator does not cover B's privacy obligations. Nevertheless, the operator may provide privacy policy templates and implement measures to ensure that B does not make public his friendship to A unless A has given her consent.

The choice between an open (public) or hidden (private) friendship can be offered via the mechanisms provided in [25]. Similar to the coding of inferences in P3P, different usage options for the SNS are coded in a single valid P3P Privacy Policy. Therefore, a user agent can seamlessly parse those alternative scenarios of friendship making and select the most appropriate option for the user (see Listing 3 below). The policy negotiation and the choice of the right option is automated so that the "overhead" is transparent

to the SNS user. The necessary XML schemas and namespaces are available, see [25].

*Example 3 (contd.)* The scenarios of friendship making are described in P3P as follows:

Listing 3: Different friendship alternatives (public/hidden) are coded in a single P3P Privacy Policy

```
   <POLICY xmlns:PRINT="http://preibusch.de/namespaces/PRINT/PRINT.xsd">
  <EXTENSION optional="no">
    <PRINT:NEGOTIATION-GROUP-DEF  id="friendship"
     standard="public_friend"  fallback="public_friend"  selected="public_friend"
     description="Choosing public (open) or private (hidden) friendship" />
  </EXTENSION>
  <STATEMENT>  <EXTENSION optional="no">
      <PRINT:NEGOTIATION-GROUP id="public_friend" groupid="friendship"
       serviceuri="/make-friend/public"
       description="Make this user a public friend of yours" />
    </EXTENSION>
    <CONSEQUENCE>Other visitors will see that you are friends</CONSEQUENCE>
    <RECIPIENT>  <ours/>
                 <public/>  </RECIPIENT>
    <PURPOSE>    <contact/>
                 <other-purpose> friendship </other-purpose>  </PURPOSE>
    <RETENTION>  <indefinitely/>  </RETENTION>
    <DATA-GROUP> <DATA ref="#user.login.id"/>  </DATA-GROUP>
  </STATEMENT>
  <STATEMENT>  <EXTENSION optional="no">
      <PRINT:NEGOTIATION-GROUP id="hidden_friend" groupid="friendship"
       serviceuri="/make-friend/hidden"
       description="Make this user a hidden friend of yours" />
    </EXTENSION>
    <CONSEQUENCE>Other visitors will not see that you are friends</CONSEQUENCE>
    <RECIPIENT>  <ours/>  </RECIPIENT>
    <PURPOSE>    <contact/>
                 <other-purpose> friendship </other-purpose>  </PURPOSE>
    <RETENTION>  <indefinitely/>  </RETENTION>
    <DATA-GROUP> <DATA ref="#user.login.id"/>  </DATA-GROUP>
  </STATEMENT>
</POLICY>
```

As the friendship making process is realized through the SNS, the SNS operator can record the chosen option and integrate enforcement mechanisms into the site [2]. When displaying a user's friends list, only public friends will be listed. The scenario demonstrates that privacy enhancements can be implemented without disturbing the user. The standard-compliant coding in machine-readable privacy policies allows for computer-supported decision-making. Moreover, the content presentation becomes semantics-driven as it is governed by semantic policies; policies will provide for privacy even if the friendship may no longer exist.

# 7 Conclusion

In this paper, we have shown that privacy in SNSs is of growing interest as these sites gain economic relevance. As companies buy SNSs for the inherent marketing potential and sites like Second Life create parallel economic worlds, it should be of interest to the user and even more to researchers and software developers how to implement techniques that provide users with "digital privacy". If this is not achieved, a backlash could result as we observed for eCommerce in the late 1990s.

While SNSs already use some privacy functions and have their own privacy policies, these are still centered around the individual, although SNSs clearly take into account network effects. If for example one user reveals data about himself, as well as a list of his friends, this "network" information could lead to revelations that had not been intended by his friends. Such leaks can prove bothersome or disastrous for individual users. In addition, these users may lose trust in the SNS and leave, which in turn creates problems for the operators of the site and the marketing initiatives financing them (this happened in one of the sites mentioned in Section 2, StudiVZ). This shows that both sides have a vital interest in effective privacy measures. In this paper, we aimed at contributing to the discussion about privacy in SNSs, a topic which we consider to be severely under-researched. We proposed a framework for analyzing privacy requirements and for analyzing privacy-related data.

To build on a comprehensive notion of "privacy", we investigated desired properties of (inter)actions on the one hand and issues of data confidentiality on the other hand. We developed a data confidentiality taxonomy to capture the privacy specificities in SNs: The (intended) interaction with other users, especially with "friends" inside the network, can result in personal data being disclosed by third parties and in other data being inferred from users' communication patterns. We outlined methods for multilateral requirements analysis for identifying, negotiating, and – if possible – resolving conflicts already during system design. The dichotomous distinction between "public data" and "private data" was refined to a set of tiered confidentiality levels. We provided an extension to the Privacy Policy language P3P to code data inferences that may result in confidentiality level breaches. The machine-readable coding of inferences allows for a better-informed consent, as the user becomes aware of side-effects. In particular, symmetric relations like "friendships" are potential privacy pitfalls as one user's disclosure makes it

possible to draw conclusions about other users' data. We provided mechanisms how privacy policies can be integrated seamlessly into the interaction among users. These policies give semantics to confidentiality and can be enforced by SNS operators.

Many challenges lie ahead. They include further investigations of the formal characteristics of the proposed inference (avoidance) schemes, practical applications and the development of best practices in requirements analysis and conflict resolution, and last but not least extensive user studies on the usability of concrete, implemented privacy options (these studies could build on the methods of, e.g., [9, 16, 18].

# References

[1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. XPref: a preference language for P3P. *Computer Networks*, 48, 2005.

[2] Anne Anderson. The Relationship Between XACML and P3P Privacy Policies, 2004. `http://research.sun.com/projects/xacml/XACML_P3P_Relationship.html`.

[3] A. Barabasi and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999.

[4] Barbarian. Debunking the MySpace Myth of 100 Million Users, 2006. `http://www.netscape.com/viewstory/2006/09/27/the-real-number-of-myspace-users`, 27 Sep 2006.

[5] M. Barbaro and T. Zeller. A face is exposed for AOL Searcher No. 4417749. *New York Times*, 9 August 2006.

[6] BBC. News Corp in USD 580m internet buy. *BBC News*, 2005. `http://news.bbc.co.uk/2/hi/business/4695495.stm`, 15 Feb 2007.

[7] BBC. Google signs USD 900m News Corp deal. *BBC News*, 2006. `http://news.bbc.co.uk/2/hi/business/5254642.stm`, 15 Feb 2007.

[8] B. Berendt and E. Menasalvas. Introduction. In *Proceedings of the Workshop on Ubiquitous Knowledge Discovery for Users at ECML/PKDD 2006*, pages 1–2, 2006. `http://vasarely.wiwi.hu-berlin.de/UKDU06/Proceedings/UKDU06-proceedings.pdf`.

[9] Shlomo Berkovsky, Nikita Borisov, Yaniv Eytani, Tsvi Kuflik, and Francesco Ricci. Examining users' attitude towards privacy preserving collaborative filtering. In *Proceedings of DM.UM'07*, 2007. `http://vasarely.wiwi.hu-berlin.de/DM.UM07/Proceedings/berkovsky.pdf`.

[10] danah boyd. Identity production in a networked culture: Why youth heart MySpace. In *Annual Meeting of the American Association for the Advancement of Science, St. Louis, MO. February 19*, 2006. `http://www.danah.org/papers/AAAS2006.html`.

[11] Bundesregierung, 16. Wahlperiode. Antwort der Bundesregierung auf die Kleine Anfrage der Abgeordneten Jan Korte, Petra Pau, Kersten Naumann und der Fraktion DIE LINKE. Drucksache 16/3787. Rechtmäßigkeit und Anwendung von Online-Durchsuchungen, 2006. http://dip.bundestag.de/btd/16/039/1603973.pdf.

[12] Pete Cashmore. MySpace Hits 100 Million Accounts, 9 Aug 2006. http://mashable.com/2006/08/09/myspace-hits-100-million-accounts/.

[13] European Court. Judgment of the Court of 6 November 2003. Criminal proceedings against Bodil Lindqvist, 2003. http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:62001J0101:EN:HTML.

[14] Janie Gafford. Guerilla Marketing on MySpace-Smart Do-It-Yourself Online Marketing, 2007. http://ezinearticles.com/?Guerilla-Marketing-on-MySpace-Smart-Do-It-Yourself-Online-Marketing&id=433759.

[15] GFK Marktforschung GmbH Bereich Online Research. Marktforschungsstudie zur Nutzung alternativer Werbeformen [market research study on the use of alternative forms of marketing.

[16] G. Groh. Groups and group-instantiations in mobile communities – detection, modeling and applications. In *Proceedings of the International Conference on Weblogs and Social Media 2007*, 2007. http://www.icwsm.org/papers/paper7.html.

[17] S.F. Gürses, B. Berendt, and Th. Santen. Multilateral security requirements analysis for preserving privacy in ubiquitous environments. In *Proceedings of the Workshop on Ubiquitous Knowledge Discovery for Users at ECML/PKDD 2006*, pages 51–64, Berlin, September 2006. http://vasarely.wiwi.hu-berlin.de/UKDU06/Proceedings/UKDU06-proceedings.pdf.

[18] Indratmo and Julita Vassileva. A usability study of an access control system for group blogs. In *Proceedings of the International Conference on Weblogs and Social Media 2007*, 2007. http://www.icwsm.org/papers/paper33.html.

[19] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, 1999.

[20] M. Langheinrich. A P3P Preference Exchange Language (APPEL), 26 February 2001. W3C Working Draft., http://www.w3.org/TR/P3P-preferences.

[21] Jay Conrad Levinson. *Guerrilla marketing*. Houghton Mifflin, 1984.

[22] S. Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.

[23] S. Milgram and J. Travers. An experimental study of the small world problem. *Sociometry*, 32(4):425–443, 1969.

[24] Adeniyi Onabajo and Jens H. Jahnke. Modelling and Reasoning for Confidentiality Requirements in Software Development. In *ECBS*, 2006.

[25] Sören Preibusch. Privacy Negotiations with P3P. In *W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement*, 2006. http://www.w3.org/2006/07/privacy-ws/papers/24-preibusch-negotiation-p3p/.

[26] Privacy Rights Clearinghouse. A Chronology of Data Breaches, 2005–2007. http://www.privacyrights.org/ar/ChronDataBreaches.htm.

[27] Inken Prodinger. Ins Netz gegangen. *Börsen-Zeitung*, 15 Feb 2007. http://corporate.xing.com/fileadmin/image_archive/review_Boersen-Zeitung_050107_eng.pdf.

[28] Christian Rath. "Terroristen sind auch klug" [Terrorists are clever too] – Interview with Wolfgang Schäuble. *die tageszeitung*, 2007. http://www.taz.de/pt/2007/02/08/a0169.1/textdruck [10 Feb 2007].

[29] Matthew Richardson and Pedro Domingos. Mining Knowledge-Sharing Sites for Viral Marketing. In *Proc. of the Eighth Intl. Conf. on Knowledge Discovery and Data Mining (SIGKDD'02)*, 2002.

[30] William Robinson, Suzanne Pawlowski, and Vecheslav Volkov. Requirements Interaction Management. *ACM Computing Surveys*, 35(2), 2003.

[31] Joseph Smarr. Technical and privacy challenges for integrating FOAF into existing applications, 2001. http://www.w3.org/2001/sw/Europe/events/foaf-galway/papers/fp/technical_and_privacy_challenges/.

[32] Ian Sommerville and Peter Sawyer. Viewpoints: Principles, Problems and a Practical Approach to Requirements Engineering. *Annals of Software Engineering*, 3:101–130, 1997.

[33] Christian Stöcker. Community-Millionendeal: Holtzbrinck schnappt sich StudiVZ [Community Million-Euro Deal: Holtzbrinck snatches StudiVZ]. *Spiegel Online Netzwelt*, 2007. `http://www.spiegel.de/netzwelt/web/0,1518,457536,00.html`, 3 Jan 2007.

[34] Latanya Sweeney. Achieving k-anonymity privacy protection using generalization and suppression. *International Journal of Uncertainty*, 10(5):571–588, 2003.

[35] Maximilian Teltzrow and Alfred Kobsa. Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study. `http://www.ics.uci.edu/~kobsa/papers/2004-PersUXinECom-kobsa.pdf`.

[36] Maximilian Teltzrow, Sören Preibusch, and Bettina Berendt. SIMT – A Privacy Preserving Web Metrics Tool. In *Proceedings of CEC*, pages 263–270. IEEE Computer Society, 2004.

[37] Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications*, volume 8 of *Structural Analysis in the Social Sciences*. Cambridge University Press, Cambridge, 1 edition, 1999.

[38] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.

[39] Rigo Wenning and Matthias Schunter. The Platform for Privacy Preferences 1.1 (P3P1.1) Specification, 2006. W3C Working Group Note 13 November 2006, `http://www.w3.org/TR/P3P11/`.

[40] Wikipedia. Viral Marketing, 2007. `http://en.wikipedia.org/wiki/Viral_marketing` [10 Feb 2007].