Collaborative Data Privacy for the Web

Clemens Heidinger
Karlsruhe Institute of Technology (KIT), Germany
heidinger@kit.edu

Erik Buchmann
Karlsruhe Institute of Technology (KIT), Germany
erik.buchmann@kit.edu

Klemens Böhm
Karlsruhe Institute of Technology (KIT), Germany
klemens.boehm@kit.edu

ABSTRACT

While data privacy is a human right, it is challenging to enforce it. For example, if multiple retailers execute a single order at Amazon Marketplace, each retailer can use different agencies for shipment, payment etc., resulting in unmanageable flows of personal data. In this work, we present the Privacy 2.0 system, which enables people to share experiences, observations, and recommendations regarding the privacy practices of data collectors. The basis of Privacy 2.0 is a folksonomy where a user community tags web sites on the Internet with privacy-related labels, e.g., "no privacy policy" or "collects too much personal data". Privacy 2.0 evaluates this folksonomy, and issues a warning if a user is about to enter a web site that has been marked with alarming tags by the majority of users. We have evaluated an operative implementation of our approach by means of a user study. The study indicates that the Privacy 2.0 system helps to assess the privacy practices of service providers and adapts well to a wide range of privacy threats.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Design

Keywords

Privacy, Folksonomies

1. INTRODUCTION

Privacy is one of the biggest challenges for today's society. Enterprises that offer services on the Internet often have a multitude of affiliates, and delegate tasks like billing, shipping, delivery, customer-relationship management and advertising to other companies. For example, eBay or Amazon Marketplace provide the technical and organizational infrastructure for a huge number of independent retailers. Each retailer might use different providers for all kinds of after-sales services. As each member of this complex network of enterprises might have different privacy practices and it is often unknown in advance which company will execute which task, it is very difficult for the individuals concerned to overlook the whereabouts of their personal data. Furthermore, many Web 2.0 services like social network sites, photo portals or online communities let their users disclose personal information, e.g., texts or photos. As these services often allow to create mash-ups with other services, it is frequently unclear who is able to access which personal data. For example, Facebook supports third-party applications and provides programming interfaces to internal functions and data structures. Thus, information disclosed to Facebook can flow to other parties. Many legal privacy regulations exist, e.g., based on the OECD principles [8] or the EU directive 95/46/EC [6]. However, the complexity of the regulations and the sheer amount of transactions and data collectors result in a serious lack of enforcement [3]. This calls for privacy enhancing technologies (PETs) [16]. Such PETs must support the users in managing their privacy in intuitive ways [2].

In this work, we present the Privacy 2.0 system. It lets people share experiences, observations and recommendations about the privacy practices of service providers on the Internet. The Privacy 2.0 system incorporates a folksonomy [13] that allows users to tag URLs on the Internet with privacy-related tags, e.g., "no privacy policy" or "collects too much personal data". Since a folksonomy does not impose any restrictions on the tags generated by the users, we expect Privacy 2.0 to cover a wide range of different privacy threats in an intuitive way. By evaluating the privacy tags provided by a large community of users, Privacy 2.0 promises to identify data collectors whose privacy practices violate the privacy standards of society, and it indicates if a data collector handles personal information with care. The generic concept of Privacy 2.0 is part of our previous work, cf. [1]. In this work, we describe a realization of Privacy 2.0 that is fully operational. Furthermore, we show the results of an early user study. In particular, we make the following contributions:

- We analyze the requirements for the Privacy 2.0 system that is tailored for the use case web.
- We describe the design and system architecture of our Privacy 2.0 system according to these requirements.
- We evaluate the Privacy 2.0 approach by the means of a user study.

PAIS'10, March 22, 2010, Lausanne, Switzerland.

The remainder of this paper is organized as follows: Section 2 presents our previous work, the Privacy 2.0 framework and its components. In Section 3, we describe the approach of this paper to implement the framework and apply it on the use case web. Section 4 presents evaluation results. Section 5 presents related work. Section 6 concludes the paper.

2. THE PRIVACY 2.0 FRAMEWORK

Our Privacy 2.0 system borrows from previous work [1], which describes the ideas and principles behind Privacy 2.0 as a generic framework. In this section, we briefly sketch the Privacy 2.0 framework, its motivation and its components.

2.1 Motivation

The following observations motivate our approach:

1. Data privacy is endangered in many areas, e.g., the Internet, Sensor Networks or Ubiquitous Computing installations. A wide range of privacy threats exist for each area. Thus, the whereabouts of personal data are intransparent to the individuals concerned.
2. Many countries have established data protection as a fundamental human right [6]. However, the regulatory approach often results in a daunting number of norms and laws that are hard for individuals to overlook and for the data collectors to manage. Thus, the regulatory approach suffers from a lack of enforcement [3]. Furthermore, new technologies often result in legal uncertainty, since no regulations exist at roll-out time.
3. Almost all current PETs require a thorough understanding of the technology. Achieving comprehensive privacy goals requires skilled individuals which install and configure a large number of PETs for different protocols and applications, e.g., a cookie blocker or a JAP proxy.
4. The web features many sites, blogs, and other projects where privacy activists share information on privacy threats. However, as these projects are scarcely connected, they do not provide comprehensive information and cannot generate social pressure on privacy violators.

The generic Privacy 2.0 framework is a holistic approach, designed to cover a wide range of privacy threats. It allows a community of users to intuitively share knowledge about privacy threats in all areas where resources exist that can be tagged electronically. In particular, Privacy 2.0 considers (but is not limited to) URLs on the Internet, geographical positions obtained from a GPS device or RFID-tagged items read via near field communication. Privacy 2.0 users tag these resources with intuitive and schema-free labels. The framework proposes to store these tags in a folksonomy. As with any community-intelligence approach, Privacy 2.0 cannot provide absolute guarantees regarding the completeness or correctness of the information provided. However, we think that a large folksonomy with tags from many individuals would correctly represent privacy issues that are important for large shares of society.

2.2 Components

Privacy 2.0 consists of five components, as shown in Figure 1. The components manage folksonomy records (O, U, T) where O are the objects, U the users, and T the tags: $(O, U, T) = \{(o_1, u_1, t_1), \ldots, (o_n, u_n, t_n)\}$ with $o \in O$, $u \in U$, $t \in T$. We speak of a *tag application* if a user tags an object. Based on the information collected, the framework computes if a tagged resource threatens the privacy of a certain user.

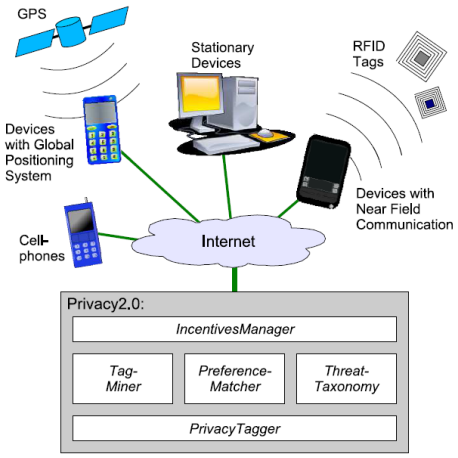

Figure 1: Privacy 2.0 architecture.

PrivacyTagger This component is responsible for collecting data. For every tag application, PrivacyTagger creates a new (o, u, t) record. For example, if a user u_l *applies* the tags t_m="video surveillance" and t_n="no privacy policy" to the geographical position o_k, PrivacyTagger stores the records (o_k, u_l, t_m) and (o_k, u_l, t_n).

ThreatTaxonomy In order to decide about objects that are scarcely tagged, the Privacy 2.0 framework infers privacy threats from one object from the threats from similar ones. Therefore, the ThreatTaxonomy component calculates the similarity of two objects $o_a, o_b \in O$ by implementing the function $tt(o_a, o_b) = x$, where $x \in [0, 1]$. $x = 0$ means that the objects are completely different, $x = 1$ indicates identical objects.

TagMiner This component extracts meaning from tags, e.g., by applying opinion extraction [4, 11], and evaluates whether a tag is threatening regarding privacy or not. Thus, TagMiner implements the function $tm(t) = x$ with $x \in \{positive, negative\}$. To provide an example, tm("no privacy policy") $= negative$.

PreferenceMatcher To find out if the tags of one user are meaningful for another one, PreferenceMatcher computes the similarity of the privacy attitudes of two users a and b. It uses the folksonomy records provided by a and b as preference sets π_a and π_b, and computes the function $pm(\pi_a, \pi_b) = x$ with $x \in [0, 1]$. $x = 0$ stands for complementary preferences and $x = 1$ for an identical attitude towards privacy.

IncentivesManager Finally, this component is responsible for motivating users, to tag relevant objects and to provide tags of high quality.

2.3 Threat Computation

We now specify how the framework evaluates the request of a user u whether object o might be a privacy threat or not. In the next section we will describe how we have refined this procedure for our Privacy 2.0 system tailored for the use case web.

Recall that an object in Privacy 2.0 means a URL, geographical position, or RFID chip that can be assigned to a data collector or a service provider.
Using the folksonomy (O, U, T) and the user preferences $\{\pi_1, \ldots, \pi_n\}$, the framework computes $isThreat(o, u) = x$ with $x \in \{true, false, unknown\}$. The value *unknown* is returned if the folksonomy does not contain information about object o. A return value *true* means that the privacy practices related to o might pose a threat according to the privacy preferences of user u. Otherwise, $false$ is returned.

First, TagMiner partitions the tag application tuples stored in the PrivacyTagger component into two sets of (*object*, *tag*)-pairs with positive (T^{pos}) and negative (T^{neg}) tags:

$$T^{pos} = \left\{(\hat{o}, \hat{u}) \mid (\hat{o}, \hat{u}, \hat{t}) \in (O, U, T) \ \wedge \ tm(\hat{t}) = positive\right\}$$

$$T^{neg} = \left\{(\hat{o}, \hat{u}) \mid (\hat{o}, \hat{u}, \hat{t}) \in (O, U, T) \ \wedge \ tm(\hat{t}) = negative\right\}$$

Second, Privacy 2.0 computes a score of each pair $(\hat{o}, \hat{u})$ in T^{pos} and T^{neg}. ThreatTaxonomy quantifies the similarity between the object in question o and an object $\hat{o}$ that has been tagged before. The PreferenceMatcher provides a measure for the similarity of the current user u and the user $\hat{u}$ who provided the tag. The score is the sum of the products of these values:

$$score = \sum_{(\hat{o},\hat{u}) \in T^{pos}} tt(\hat{o}, o) \cdot pm(\pi_u, \pi_{\hat{u}}) - \sum_{(\hat{o},\hat{u}) \in T^{neg}} tt(\hat{o}, o) \cdot pm(\pi_u, \pi_{\hat{u}})$$

Finally, the result is computed as follows:

$$isThreat(o, u) = \begin{cases} unknown & \text{if } \forall \hat{o} \in O : tt(o, \hat{o}) = 0 \\ true & \text{if } score \leq 0 \\ false & \text{otherwise} \end{cases}$$

More details about the generic Privacy 2.0 framework can be found in [1].

3. PRIVACY 2.0 FOR THE WEB

Privacy 2.0 as described so far is a sketch for a Web 2.0 approach that lets the individuals concerned share experiences, observations and recommendations about privacy practices of service providers.

In this section, we describe a Privacy 2.0 system specifically designed for the use case web. We identify requirements for the application domain web and technologies best suited for this use case. After that, we describe our design decisions and our system architecture.

3.1 Requirements

We have derived four requirements for our Privacy 2.0 system, which stem from the use case web, and the motivation behind the Privacy 2.0 concept:

R1: Convergence of Knowledge Information stored in a folksonomy can be ambiguous, incomplete or misleading. Yet in popular folksonomies the most important issues are well represented [13]. Our implementation has to support this phenomenon. Thus, the knowledge of a large community of users must converge to a state that represents privacy threats important for society.

R2: Effectiveness The approach must be effective in supporting individuals to decide if a data collector handles personal information with care, and in identifying privacy threats that are important for large shares of society.

R3: Usability and Transparency Since the target group of our approach are individuals without in-depth knowledge on PETs, usability and transparency are important (cf. [16]).

R4: Extensible Design Our goal is to develop a system that serves as a baseline for future development and can be used to confirm the applicability of Privacy 2.0 by means of user experiments. Thus, we require a minimalistic yet extensible design that provides interfaces to measure system parameters and allows to switch between different implementations of the components.

3.2 Design Decisions

Our design decisions consider three aspects: How to realize (1) the framework components, (2) the threat computation, and (3) the system architecture.

3.2.1 Components

We now describe our realization of the Privacy 2.0 components according to our requirements.

PrivacyTagger This component stores the tags provided in a folksonomy. According to Requirement **R4**, we need a folksonomy platform that is extensible for future developments, and provides interfaces to measure various system parameters during user experiments, e.g., if a user visits a web site even if our approach says that this site might pose privacy threats. Since no existing folksonomy platform supports such interfaces, we have decided to implement our own component.

ThreatTaxonomy The ThreatTaxonomy provides a measure for the similarity of objects. In the web scenario, the objects are web pages addressed by URIs. We assume that all web pages with the same host address are under control of the same service provider and pose the same privacy threats. Thus, our similarity measure assumes `http://www.google.com/` and `http://www.google.com/search?q=privacy` to be the same object, while `http://www.google.de` is different to `http://www.google.com`. Obviously, this similarity measure does not consider portals like Amazon Marketplace or eBay that provide the technical infrastructure for a large number of independent retailers. We leave aside such cases for the time being.

TagMiner This component finds out if a tag has a positive or negative meaning. For example, the labels "worse policy" and "spammer" carry a negative assessment. A

baseline implementation of this component is to let the users provide information on the semantics of tags. In particular, the folksonomy stores for each $(tag, user)$ pair a value $c \in \{-1, 1\}$ to indicate if the meaning of the tag is negative or positive.

PreferenceMatcher To determine if the tags provided by a user have a meaning for another one, the PreferenceMatcher computes a similarity measure for two users. Since the tags stored in the folksonomy represent the opinions of the users regarding certain objects, it should be possible to extract user preferences from the data set. However, a minimal baseline approach is to consider all users to be equal, i.e., $pm(\pi_a, \pi_b) = 1$ for any two users a and b. This assumption is reasonable, if we perform user experiments with a homogeneous study group.

IncentivesManager Since monetary incentives (cf. Section 4.1) are known to be sufficient for a small study group, we have omitted this component for the time being.

3.2.2 Threat Computation

The Privacy 2.0 framework computes if an object o might be a privacy threat for a user u. That is, it implements the function $isThreat(o, u) = x$ with $x \in \{true, false, unknown\}$. Our baseline implementation of this function is a majority vote, i.e, the number of positive tags for an object minus the number of negative tags for the same object. The function returns *unknown* if the folksonomy does not contain tags for the host address in o.

As described in Section 3.2, we let the user assign a value to each tag that indicates if the tag has a negative or positive meaning. Thus, for each tuple $(\hat{o}, \hat{u}, \hat{t}) \in (O, U, T)$ managed by PrivacyTagger, TagMiner stores a tuple $(\hat{u}, \hat{t}, \hat{c})$ with $\hat{c} \in \{-1, 1\}$. A value of $\hat{c} = -1$ means that the tag $\hat{t}$ is used in a negative sense, $\hat{c} = 1$ refers to positive tags. In order to compute the score from these tuples, we extend the function computed by TagMiner so that $tm(u, t) = c$. Remember that our baseline realization does not consider individual preferences. Thus, given that o is the object in question, the score is:

$$score = \sum_{(\hat{o},\hat{u},\hat{t}) \in \{(O,U,T) | \hat{o}=o\}} tm(\hat{u}, \hat{t})$$

Now, $isThreat(o, u)$ is computed as described in Subsection 2.3.

Note that we assume that users truthfully apply tags and truthfully indicate a positive or negative meaning for tags. This is reasonable for user experiments with a small study group. Furthermore, there exist measures to encounter tag spam, e.g., [10].

3.2.3 System Architecture

We have decided to realize our approach as a server application (TaXor) that is accessed through a web browser add-on. TaXor collects and manages all information required for the components described in the Privacy 2.0 framework (cf. Figure 1). TaXor has a modular design, i.e., it integrates the Privacy 2.0 components as plug-ins.

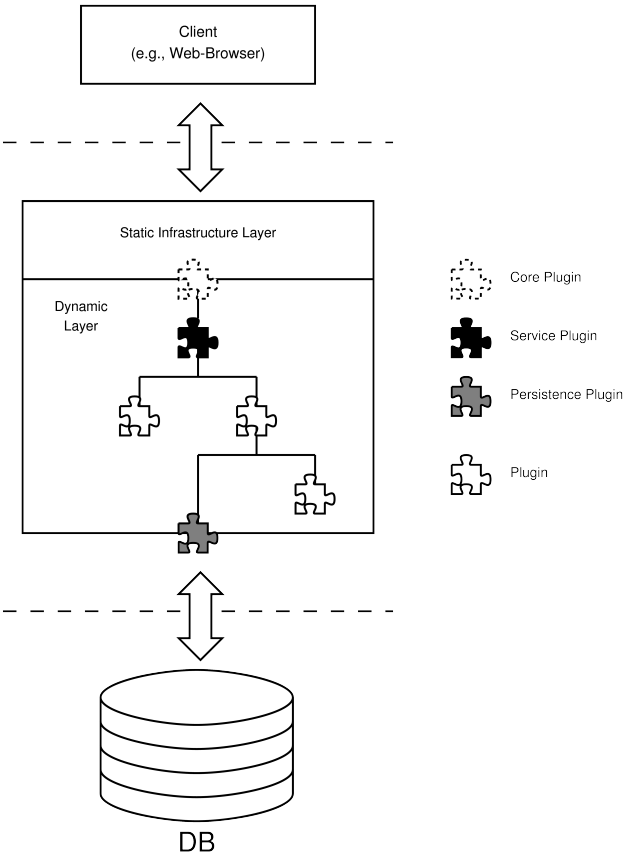

Figure 2: System architecture.

Figure 2 shows the system architecture of TaXor. The architecture follows the classical client-server model. TaXor consists of a static infrastructure layer and a dynamic layer. While the interfaces of the dynamic layer depend on the plug-ins used, the static layer provides a defined interface for the communication between server and clients, e.g., our browser add-on or other external applications. Using those interfaces, the client can send (e.g., a tag application) or receive data (e.g., a tag cloud for an object). Available communication formats include XML, XHTML, JSON or RSS feeds. All server functionality is encapsulated in plug-ins, using the Java Plug-In Framework JPF. For example, there is one plug-in for tag-cloud generation. The plug-ins can be exchanged easily, thus they are part of the dynamic layer of the architecture. The plug-in for persistent storage of the folksonomy in a database is based on the Hibernate library.

We have implemented the client as a Firefox add-on. Browser add-ons are a popular choice to provide extra functionality independent of the web page currently displayed. Figure 3 shows the google.de web site. Our add-on blends in a toolbar at the bottom of the browser window. We will describe its user-interface elements in the next subsection.

Note that our flexible system architecture allows to implement different kinds of clients, e.g., web-portals, standalone applications or add-ons for other browsers. However, for the time being we focus on a single client to gather experience about the design of the user interface.

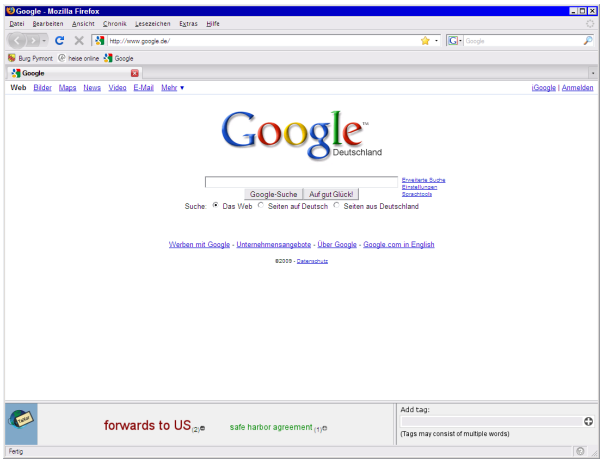

Figure 3: The browser add-on.

3.3 User Interface

The users of our Privacy 2.0 system must be able to (1) point out new privacy issues, (2) substantiate findings of others, and (3) identify web sites that might pose a threat for their privacy. Therefore, the browser add-on of our Privacy 2.0 system offers a set of GUI elements.

The input field in the bottom right corner of the toolbar (see Figure 3) allows to apply privacy tags to the web site that is displayed at the moment.

To cope with Requirement **R2**, we have implemented a visual indicator for $isThreat(o, u)$ (cf. Section 3.2.2). It shows how our approach assesses the privacy aspects of the web site: The color of the logo in the left of the toolbar (see Figure 3) changes to red for $isThreat(o, u) = true$, green for $isThreat(o, u) = false$, and blue for $isThreat(o, u) = unknown$, thus indicating that o is probably threatening for the privacy of u, possibly unobjectionable, or there is insufficient data to make a decision.

Requirement **R1** means that privacy issues important for society must be emphasized, i.e., the knowledge gathered by individuals should converge to knowledge that represents the privacy issues that are most important for society. In the context of folksonomies, tag clouds [14] are the standard way to represent tags important for a given object: Tags are displayed in varying size, position or color according to their importance. This concept is familiar to users due to the popularity of tag clouds on the web (**R3**). Figure 4 shows the tag cloud from Figure 3. The tag *forwards to US* has been applied more often than the tag *safe harbor agreement*. Thus, it is displayed more prominently.

Figure 4: A tag cloud in the browser add-on.

Next to the tags in the tag cloud are small buttons. A minus-button allows a user to remove a tag she has applied before. A plus-button lets a user apply a tag to a web site which has been applied before, i.e., to substantiate the findings of others (**R1**). The colors of the tags indicate their meaning. In the figure, the tag *forwards to US* has been marked by the user the tag cloud is shown to as threatening for privacy (red), and *safe harbor agreement* has been marked unobjectionable (green). Figure 5 shows the interface elements where a user can mark a tag as privacy-threatening or unobjectionable.

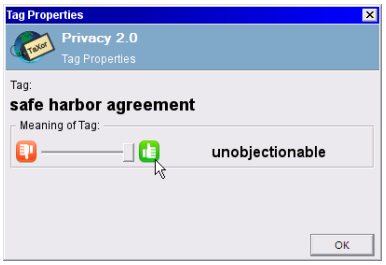

Figure 5: Tag semantics in the browser add-on.

4. EVALUATION

In this section, we describe an evaluation of our approach by means of a user study. Intuitively, we want to know: (1) Is the Privacy 2.0 approach helpful in practice, i.e., does Privacy 2.0 help people to decide if they want to use a service that requires personal information? (2) Does our Privacy 2.0 system fulfill the requirements specified in Section 3.1? Therefore, we have devised three research questions:

Q1: Convergence of Knowledge Does the data set obtained from the users converge to a consistent state for popular privacy issues? In particular, do the users find appropriate tags, and do they tag popular web sites with tags with similar semantics? Can we observe patterns of user activity similar to results published for folksonomies where the knowledge converges?

Q2: Effectiveness Does our approach help a community of users to assess the impact of web sites on their privacy? That is, do the participants appreciate the tags provided for a web site, and do the study participants think that our approach is useful in practice?

Q3: Usability and Transparency Is the usability of our Privacy 2.0 system sufficient to make it practical? Is all functionality transparent to the user, i.e., do the participants understand the controls and results displayed by the interface of our browser add-on? How do the participants estimate the effort of using our system?

To evaluate these questions, we have devised a user experiment with 9 participants. Since we did not test the Privacy 2.0 approach before, we have decided to start with a relatively small study group. We will conduct user studies with a large group of participants in the future.

4.1 Experiment Setup

In the following, we describe the participants, the methodology and the incentives we have used for our experiment.

Participants. One of our design decisions has been to leave individual preferences aside for the time being (cf. Sec-

tion 3.2). Thus, our study group had to consist of people who are likely to have the same preferences regarding data privacy. Additionally, we were looking for participants which have a good understanding of the privacy issues on the Internet. In this respect, we have contacted the audience of a university class on data privacy (as part of the computer-science curriculum) and over the web[1]. Nine persons have participated in all phases of the experiment, i.e., they have actively used our browser add-on and have filled out questionnaires. We have ensured that our participants had the same educational background and were of similar age (between 21 and 28).

Methodology. We have conducted the user experiment in three phases:

1. **Registration & Information**: In this phase, we have contacted our participants by e-mail. We have introduced our approach by referring to a web site which explains the idea of Privacy 2.0 and the interface elements of our Privacy 2.0 system.
2. **Experiment**: Our study participants used the browser add-on for 20 days in their everyday life.
3. **Questionnaire**: At the end of the experiment phase, we handed out a questionnaire. It contained questions regarding utility, effort required, and the understanding of our approach in general.

The introduction of our approach and the questionnaire can be seen at our web site[2].

Incentives. User participation is essential for the success of any collaborative approach. According to [12], participation is motivated by (i) competition, (ii) social factors such as community acceptance, and (iii) the benefit from using the system. Although we expect that those factors suffice to attract a large user community when we deploy our approach on the Internet, such incentives are not applicable for a small study group. Thus, we have decided for monetary incentives. Specifically, we have paid 3 EUR for registration, 5 EUR for the first 20 tag applications, and 0.2 EUR for any further tag application. Furthermore, we have stimulated competition by raffling 30 EUR among the most active participants.

4.2 Evaluation Results

In the following, we describe the results of our user experiment, structured according to the research questions Q1 to Q3.

4.2.1 Q1: Convergence of Knowledge

If we observe that individuals find appropriate tags, we have provided an indication that our approach allows to collect privacy-related knowledge from a user community. Tags are appropriate for our approach if they describe privacy-related aspects of web sites, e.g., if the site displays a privacy policy or uses a service like google analytics. We have asked our participants how often they had problems finding descriptive tags for privacy-relevant aspects of web sites. Therefore, we have provided a 5 point Likert scale. Table 1 shows that most participants rarely had problems finding tags.

[1] `http://experiments.ipd.kit.edu/`
[2] `http://privacy20.ipd.kit.edu/Exp1/`

Table 1: Problems to find appropriate tags.

Answer	# Participants
Very rarely	1
Sometimes	5
Often	1
Very often	1
All the time	0
No answer	1

We can provide an indication for the convergence of knowledge in our Privacy 2.0 system if we can show that it is used like popular folksonomies. It is known in literature that many parameters of popular folksonomies, e.g., the number of tag applications per object or per user, follow a power-law distribution [13].

Figure 6 shows all tags and objects sorted by their rank (in tag applications) on the X-axis, and the corresponding number of tag applications on the Y-axis. The most popular object had 33 tag applications (cf. Table 2), while most objects (nearly 200) had only one tag application. Thus, we have observed a power-law distribution. Furthermore, the figure shows a strong correlation between the number of tags applied and objects tagged. The average count of tag applications per object was 2.07, the one of applications per tag 2.23.

Figure 6: Tag applications per object and tag.

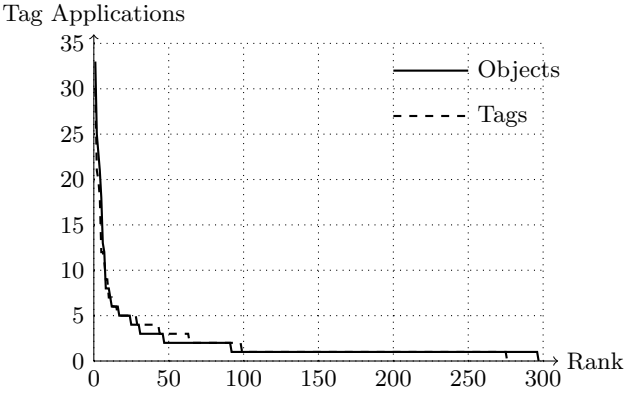

Table 2 and Table 3 show the five most popular objects and tags, respectively. In total, participants have tagged 300 objects, and they have used 277 different tags (in 613 tag applications).

Table 2: Popular objects.

Objects	# of tag applications
www.studivz.net	33
www.youtube.com	25
www.google.de	23
de.wikipedia.org	21
www.google.com	18

Table 3: Popular tags.

Objects	# of tag applications
no google analytics	30
No Google Analytics	21
google analytics	20
privacy threat through data aggregation	18
search terms may possibly be linked to persons	12

An important indicator to measure convergence is the growth of the set of objects indexed and of the set of tags used. The growth rate of both numbers is almost linear over time in systems like CiteULike [7]. Figure 7 confirms this for our experiment. The X-axis of the diagram is the time axis (in days), and the Y-axis is the number of tags and objects in the system at the point in time.

Figure 7: Object and tag growth.

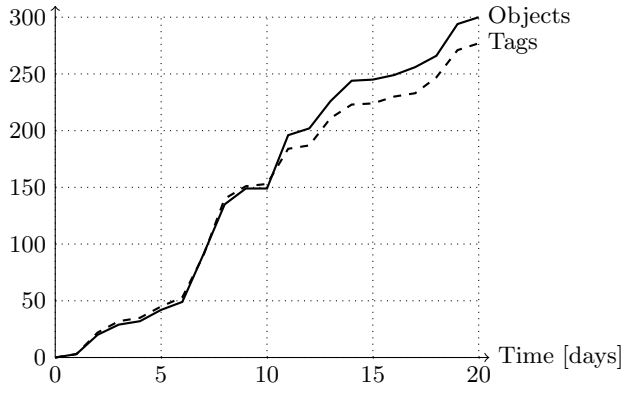

A further measure that usually follows a power-law distribution [7, 13] in popular folksonomies is the user activity, i.e., the number of tag applications per user. However, as we have provided monetary incentives to motivate all of our users to provide a large number of tags, we do not expect to find such distribution for our experiment. Figure 8 ranks the users by the number of tags provided. The two most active users had 155 and 98 tag applications while most users (7 out of 9) had more than 20 tag applications. We assume that our monetary incentives prevent the "long tail" that is typical for power-law distributions.

We have indicated that the distribution of tags and objects follow distributions known from popular folksonomies where knowledge converges. Thus, we can assume that our approach does represent privacy threats well that are important for large shares of society.

4.2.2 Q2: Effectiveness

One of our concerns is to find out if Privacy 2.0 helps users to assess the impact of services like web shops, search engines or other web applications on privacy. To confirm that this is challenging without a system supporting the users, we have asked the participants if they agree to the statement "Frequently I have problems to find out how a service provider on the web handles personal data.". Table 4 shows the answers. Most participants stated to have problems assessing the privacy impact of web sites. We value this as an indicator that people might appreciate an approach that eases this issue.

Figure 8: Tag applications per user.

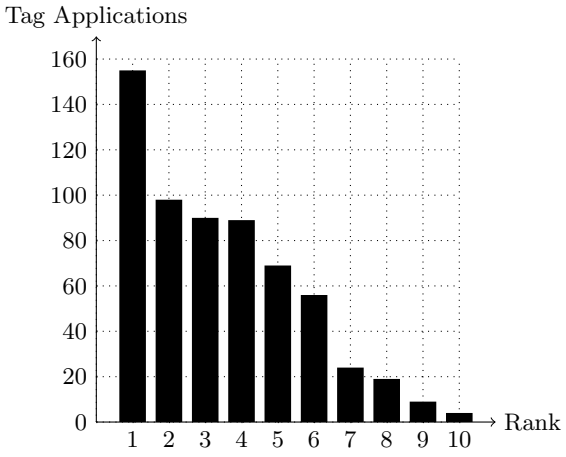

Table 4: Problems in assessing the privacy impact of web sites.

Answer	# Participants
Yes	5
No	3
No opinion	1

Next, we have asked if our Privacy 2.0 system helps to evaluate if a web site complies with the attitude of the user towards privacy or not. As shown in Table 5, all of our participants have agreed with this statement. To complement this finding, we have asked our participants if they agree that the more tags are applied to a web site, the better they can evaluate the privacy impact of the site. Table 6 confirms that most of our participants appreciate a large number of tags.

Table 5: Does Privacy 2.0 help to evaluate the privacy impact of a web site?

Answer	# Participants
Yes, very much	1
Yes	8
No opinion	0
No	0
Not at all	0

We had decided to test our approach with a homogeneous study group. Our approach uses a straightforward majority vote to find out if a web site is threatening privacy. Nevertheless, we are interested to find out how much this simplistic baseline approach complies with the individual preferences of our study group. Table 7 shows the answers to a question if participants agree with $isThreat(o, u)$ as computed by our Privacy 2.0 system (cf. Section 3.2.2) for most web sites.

Table 6: Is a large number of tags better to assess the privacy impact of a web site?

Answer	# Participants
Yes, very much	1
Yes	6
No opinion	0
No	2
Not at all	0

Table 7: Correctness of $isThreat(o, u)$.

Answer	# Participants
Very rarely	0
Sometimes	0
Often	4
Very often	3
All the time	1
No answer	1

We have also conducted an analysis on the data set managed by our Privacy 2.0 system. In particular, we have analyzed each tag application $(\hat{u}, \hat{o}, \hat{t}) \in (O, U, T)$, and we have computed two values: (1) the majority vote returned by Privacy 2.0 (cf. $isThreat()$, Section 3.2.2), and (2) the user opinion. We compute the user opinion as the number of tags with a positive meaning user $\hat{u}$ has applied to object $\hat{o}$ minus the number of tags with a negative meaning $\hat{u}$ has applied to $\hat{o}$. Table 8 shows the aggregated results. We can observe a high degree of conformity between the Privacy 2.0 system and the opinions of the users. For example, in 260 cases, a user found a web site threatening, and our approach decided likewise. In 264 cases, our approach computed that a web site is unobjectionable, and the individual users had the same opinion. Thus, we can observe a strong overlap in user opinion and community opinion, at least for this group of users.

Table 8: User opinion versus Privacy 2.0.

User \ Privacy 2.0	threatening	unknown	unobjectionable
threatening	260	2	0
unknown	13	28	17
unobjectionable	27	2	264

Finally, we were interested to find out if our approach is applicable in everyday life. Therefore, we have asked three questions: (i) How stressful do the users deem the application of tags to web sites, (ii) how do the users value the utility of our approach when it comes to assessing the privacy impact of a web site, and (iii) given the effort, do the users think our approach is useful in practice? The answers are shown in Table 9.

The table indicate that our approach is useful. However, that four users rated the effort as medium or high motivates us to ease the creation of tags, e.g., by implementing methods that suggest or auto-complete tags while the user is typing.

Table 9: Comparison effort and utility.

Answer	Effort	Utility	Utility vs. Effort
Very low	0	0	0
Low	5	0	2
Medium	2	4	2
High	2	3	4
Very high	0	2	1

4.2.3 Q3: Usability and Transparency

Our questionnaire contained several questions to find out if the user interface is intuitive and allows to share information easily. Furthermore, we have asked control questions to confirm that our participants correctly understand the user interface. More specifically, we were interested in the tagging functionality, the functionality to specify whether a tag is threatening or unobjectionable, and the results of the treat computation displayed.

Tagging. The participants rated the usability of the tagging functionality on a 5 point Likert scale. Table 10 shows the answers. Additionally, all participants correctly answered a separate control question how to add a new tag to a web site.

Table 10: Usability of tagging functionality.

Answer	# Participants
Very easy to use	5
Easy to use	4
No opinion	0
Complicated	0
Very complicated	0

To see if the user-interface elements in the tag cloud (buttons for functions "add tag" and "delete tag") were understood correctly, we have asked another control question. Five participants identified the functionality provided by these buttons correctly as "add tag" and "delete tag", while four participants thought that these buttons exist to make statements whether a tag is threatening or harmless. Thus, our user interface needs some improvement to be more intuitive. Nevertheless, we have implemented these buttons for convenience only, and the users were free to enter new tags manually over the input field (cf. Section 3.3). Thus, a misunderstanding of these GUI elements does not impact our study results.

Tag semantics. Since we have left aside methods to extract the semantics from the tags provided, we let our participants assign each tag with a value indicating if the tag has a positive or negative meaning.

It is important that our participants understand this mechanism correctly. Only two of our participants gave a wrong answer to a control question about the tag-semantics control: One participant was thinking to rate the web site in this way, not the tag. The other one stated that this interface element performs a tag application. Yet, a data analysis has shown that these participants used these interface elements

like the other users. Thus, we assume these participants did not understand the question correctly. Furthermore, we have asked if our participants find this interface element difficult to use. As Table 11 indicates, some people found it complicated to rate a tag. This finding encourages us to develop mechanisms that extract the semantics from the tags, e.g., by adapting technologies from knowledge discovery or opinion extraction [4, 11].

Table 11: Usability of tag-semantics control.

Answer	# Participants
Very easy to use	3
Easy to use	2
No opinion	0
Complicated	1
Very complicated	3

Threat computation. It is of utmost importance that the threat computation by our Privacy 2.0 system is perceived correctly. We have asked two control questions about this issue: (1) What are possible values Privacy 2.0 computes, and (2) what is the meaning of the values.

Six participants answered correctly that Privacy 2.0 returns three distinct values (cf. Section 2.3). Two participants gave no answer, and the answer of one participant was wrong. Three participants described the meaning of each value correctly, e.g., they told us that the visual indicator indicates that the web site might pose a privacy threat. Two participants wrongly assumed the values "describe a tag". The other four participants did not answer this control question. We value those answers as a further indicator that parts of our user interface must be reworked in order to be more intuitive.

To sum up the evaluation, we have indicated that the Privacy 2.0 folksonomy data converges, similarly to popular folksonomies. Our users were able to introduce and exchange knowledge about possible privacy threats, and they found the approach effective in general. However, our evaluation has shown that we need to improve the usability of some interface elements. Furthermore, we have to ensure that interface elements cannot be misinterpreted. After having tackled those issues, our approach is ready for another user study with more participants and for a deployment on the Internet.

5. RELATED WORK

Folksonomies have gained popularity over the last years. Popular examples include folksonomies for bookmarks (delicious), music (Last.fm) or pictures (flickr). Marlow et al. [13] provides an introduction and classification into existing folksonomy systems.

We apply Privacy 2.0 on the use case web by introducing a browser add-on. PETs involving many individuals into identifying privacy threats are scarce. Most similar is the WOT Firefox add-on[3] that allows users to rate web sites in four fixed categories: trustworthiness, vendor reliability, privacy, and child safety. Our approach to use tags allows to be more specific in classifying web sites. Other PETs such as P3P-enabled web servers require a thorough understanding of users and cover only some privacy threats in the Internet (cf. [5]).

A few studies have been published with metrics about tag usage in popular folksonomies. [9] analyzed data from delicious.com and published results about user activity, tag frequencies, etc. [15] did the same for vocabulary evolution, tag utility, and tag adoption in the MovieLens system. [7] analyzed over two years of data from CiteULike, proposed tag metrics, and ran the CiteULike data against those metrics. They also looked at other publications so that their metrics can be applied to many different data sets. We thus have compared ourselves mostly against this work. Such studies are an important indicator how well our approach works. This is because in those systems, the folksonomies perform their task well and allow users to manage web sites, movies, or academic publications with ease. We want our system to be equally successful in making users aware of important privacy threats.

6. CONCLUSIONS AND FUTURE WORK

The Internet exposes individuals to many different privacy threats. In particular, a daunting number of interwoven services and web portals that are operated by a large number of independent companies collect, manage and transfer personal data on the Internet.

We have proposed our Privacy 2.0 system to tag web sites collaboratively with privacy-related information, in order to identify privacy threats. Our approach is fully operative, and it helps people to assess if a web site might threat their privacy. Our simplistic yet extensible system serves as a baseline for future developments of the Privacy 2.0 approach.

We have tested our Privacy 2.0 system by means of a user study. Our results are encouraging: The study indicates that Privacy 2.0 indeed helps to identify possible privacy threats. However, our study has also revealed that parts of the user interface can be misinterpreted, and that the effort of using our approach impacts its usefulness. After having dealt with these issues, we plan to have another user study with a large number of participants, and to deploy our approach on the Internet.

Acknowledgments

We thank Andreas Mähler for the implementation of the framework and the browser add-on, and his assistance in data analysis.

7. REFERENCES

[1] E. Buchmann, K. Böhm, and O. Raabe. Privacy2.0: Towards collaborative data-privacy protection. In *Trust Management II*, volume 263 of *IFIP International Federation for Information Processing*, pages 247–262. Springer-Verlag, Boston, MA, USA, 2008.

[2] T. Burghardt, E. Buchmann, and K. Böhm. Why do privacy-enhancement mechanisms fail, after all? a

[3] `https://addons.mozilla.org/en-US/firefox/addon/3456`

survey of both, the user and the provider perspective. In *Workshop W2Trust, in conjunction with IFIPTM'08*, 2008.

[3] T. Burghardt et al. A study on the lack of enforcement of data protection acts. In *Proceedings of the 3rd International Conference on e-Democracy*, 2009.

[4] K. Dave, S. Lawrence, and D. M. Pennock. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In *WWW '03: Proceedings of the twelfth international conference on World Wide Web*, pages 519–528. ACM Press, 2003.

[5] Electronic Privacy Information Center. Pretty Poor Privacy: An Assessment of P3P and Internet Privacy. Available at http://www.epic.org/reports/prettypoorprivacy.html, 2000.

[6] European Parliament and the Council of the European Union. 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. Official Journal L 281, 11/23/1995, p.31., 1995.

[7] U. Farooq, T. G. Kannampallil, Y. Song, C. H. Ganoe, J. M. Carroll, and L. Giles. Evaluating tagging behavior in social bookmarking systems: metrics and design heuristics. In *GROUP '07: Proceedings of the 2007 international ACM conference on Supporting group work*, pages 351–360, New York, NY, USA, 2007. ACM.

[8] O. for Economic Cooperation and D. (OECD). Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, 1980.

[9] S. Golder and B. A. Huberman. The structure of collaborative tagging systems, Aug 2005.

[10] P. Heymann, G. Koutrika, and H. Garcia-Molina. Fighting spam on social web sites: A survey of approaches and future challenges. *IEEE Internet Computing*, 11:36–45, 2007.

[11] M. Hu and B. Liu. Mining and summarizing customer reviews. In *KDD '04: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177, New York, NY, USA, 2004. ACM.

[12] K. Lerman. User participation in social media: Digg study. *CoRR*, abs/0708.2414, 2007.

[13] C. Marlow, M. Naaman, D. Boyd, and M. Davis. Ht06, tagging paper, taxonomy, flickr, academic article, to read. In *HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia*, pages 31–40, New York, NY, USA, 2006. ACM.

[14] A. W. Rivadeneira, D. M. Gruen, M. J. Muller, and D. R. Millen. Getting our head in the clouds: toward evaluation studies of tagclouds. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 995–998, New York, NY, USA, 2007. ACM.

[15] S. Sen, S. K. Lam, A. M. Rashid, D. Cosley, D. Frankowski, J. Osterhouse, F. M. Harper, and J. Riedl. tagging, communities, vocabulary, evolution. In *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, pages 181–190, New York, NY, USA, 2006. ACM.

[16] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 38–47, New York, NY, USA, 2001. ACM.