Sanitizable Signed Privacy Preferences for Social Networks

Henrich C. Pöhls[*], Arne Bilzhause, Kai Samelin, Joachim Posegga
Chair of IT Security, University of Passau, Germany
(hp|ab|ks|jp)@sec.uni-passau.de

Abstract: Privacy preferences are the handling rules and constraints under which a data subject allows a third party to process, store, and use his personal data. We have analysed Facebook and show how the Social Network System fails to collect, manage, and hand-over to third-parties user's consent. Todays technical solutions of collecting the consent on the Internet can be argued to fullfil the regulatory requirements of an informed consent to the service's Privacy Policy and Terms of Service. We found no change in Facebook's processes for collecting and managing user consent from 2009 to 2011. The technical solutions used today neither allow to manage, thus change this consent over time, nor allow to hand-over the consent to a third party. We sketch one technical solution, which lends a lot from public key infrastructures. A social network is already trusted by users to keep or federate their data. Hence, we describe the next step of Social Networks becoming an authority and sign the consent collected from its users to making the available data verifiable for third-parties. Better yet, if you do not trust the Social Network a user himself can run his own certificate authority or a group of users can provide one as a community service.

1 Introduction

If collection and automated processing of personal data occurs, the question of privacy shall always be raised. According to the definition given by the EC Directive 95/46/EC [EU95], we call data items that are personally identifying *personal data*, meaning "any information relating to an identified or identifiable natural person (*data subject*)" [EU95]. The data subject needs to express his consent for processing; if and only if his consent has been acquired, a third party is legally allowed to use the subjects's personal data: "the *data subject's consent* shall mean any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed" [EU95]. In the following the term *privacy preferences* is defined as follows: Privacy preferences are the handling rules and constraints under which the data subject allows a third party to process, store, and use his personal data. We call the third party that handles data *processor*; In our case the social network user is the data subject and the SNS is the processor. Hence, the data subject's privacy preferences are defined by the user's `privacy settings` in conjunction with the social network's `Privacy Policy` (PP) and `Terms of Service` (ToS). Nevertheless, the number of SNS, like MySpace, Facebook, or XING[1] and their number of users, forming their "digital footprint" [MFSV07] by submitting personal data is not a new phenomenon.

[*] Is funded by BMBF (FKZ:13N10966) and ANR as part of the ReSCUeIT project

[1] `www.myspace.com`, `www.facebook.com`, `www.xing.com`

1.1 Our Contribution

This paper's contribution consists of two parts:
(1) We will examine how Facebook collects users' consent and how the data subject can manage it, i.e. how a user can revoke or grant certain processing rights. In Sect. 2 we compare Facebook's current state of consent collection and management with a previous analysis from 2009. We highlight todays existing legal and technical problems of consent collection and management. An important open problem is how a SNS can proof that consent was given by the data subject. We formulate six requirements in Sect. 3.
(2) We propose a technical solution based on sanitizable signature schemes [ACdMT05] and an infrastructure comparable to a public key infrastructure (PKI). The solution sketched in Sect. 4 allows to persistently manage and verify consent. This will be extended to a data distribution scheme. Our proposed solution builds upon previous work that suggested X.509 certificates [Pöh08] and/or sanitizable signature schemes [HJPdM10]. However, this should be considered as just one technically sound implementation. Note: Although, we stick to Facebook as an example of a SNS, the highlighted problems and the requirements we propose remain widely applicable, even outside the scope of SNS.

1.2 Motivation

In the digital world, consent is acquired by saving the decisions the data subject makes during "opt-in" or "opt-out" dialogs [LH06]. Consent is not only given during the registration phase, but each time the data subjects provides the service with new personal data or changes his preferences which affect the consented usage of his personal data. The *privacy policies* state how the service will handle the personal data which data subject submits. The simplest case is a direct relationship between one data subject and one data processor, for example Facebook and a user with a Facebook profile. This direct relationship requires that the processor gathers, stores, and manages the user's consent.

More problems occur if personal data is given or sold to third parties for further usage by this entity. The user can "opt-in", hence express his consent, that the third party uses some or all of his data items. Legally, the third party is bound to process only those data items for which the processing was consented to by the data subject. However, the third party does not have a direct connection to the data subject. Hence, the data subject can only hope that the third party acts honestly, while it might not have the methods resp. possibility nor the obligation to pro-actively verify that the consent exists. If the third party suffers damages due to breaching data privacy law by non-consented processing of a data item it could claim for compensation from the company who sold them the data set, unless they specified that such processing was not consented. However, the third party is liable to the data owner, if consent was not given.

2 Facebook - Two Years of Collecting User's Consent

We analyze the account registration, configuration, ToS, PP, and account deletion process of Facebook. In the following we introduce our evaluation criteria. We evaluated Facebook alongside banking and online shopping sites in 2009. Each criterion is split into sub-criteria, which together yield the overall rating of the particular criterion. The possible ratings are "+" for positive, ”O” for neutral, and ”−” for negative. There are always two ratings for each criterion; the first represents the state of Facebook in September 2009 and the second represents the current state evaluated April 2011.

2.1 Registration

To register at Facebook a user has to provide personal data. The registration form as shown in Fig. 1 asks for first name, last name, e-Mail address, and a self-assigned password; followed by a couple of selection boxes for gender and date of birth. Note, information about Facebook's ToS or PP are available from small links located in the page's footer. At this point of time, when the user provides his personal data for the first time he has to click the *Sign Up* button. Hence, he is not yet asked to consent to any processing. Only in the second step of the registration there is a small printed text below a second *Sign Up* button which can only be used successfully if a CAPTCHA[2] is solved.

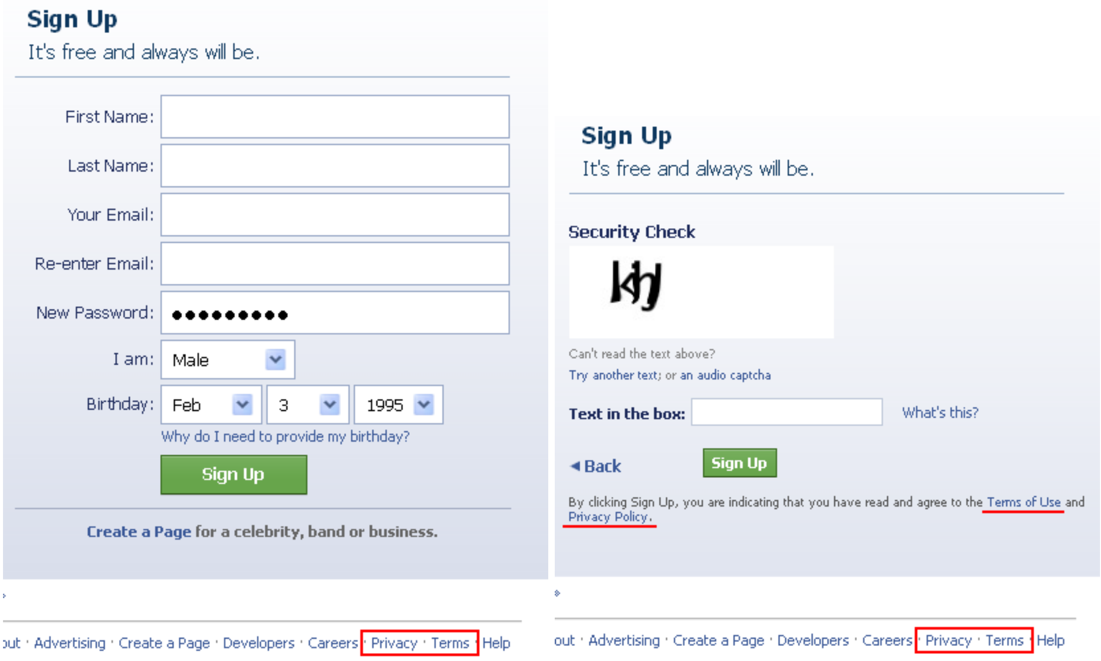

Sign Up

It's free and always will be.

First Name:

Last Name:

Your Email:

Re-enter Email:

New Password: ●●●●●●●●●

I am: Male

Birthday: Feb 3 1995

Why do I need to provide my birthday?

Create a Page for a celebrity, band or business.

Sign Up

It's free and always will be.

Security Check

Can't read the text above?

Try another text; or an audio captcha

Text in the box:

What's this?

By clicking Sign Up, you are indicating that you have read and agree to the Terms of Use and Privacy Policy.

Sign Up
out · Advertising · Create a Page · Developers · Careers · Privacy · Terms · Help
kyj
◂ Back
Sign Up
out · Advertising · Create a Page · Developers · Careers · Privacy · Terms · Help

Figure 1: Screenshot of Facebook's Two Step Registration Dialog

Facebook directly explains that the CAPTCHA has the sole purpose of making it harder for creators of fake accounts[3]. However, one might rethink the functionality of CAPTCHAs: The Federal Court of Justice in Germany issued that the statement that a user is giving his informed consent to must be highlighted. However, courts have ruled that filling in personal details in forms shall raise a user's awareness. CAPTCHAs could be used as a mean to raise a user's

[2] Completely Automated Public Turing test to tell Computers and Humans Apart

[3] Located behind the link labelled *What's this?*.

awareness that he is performing an important action, i.e. giving his consent. Why not build a CAPTCHA as a question and answer game, that not only tells computers and humans apart, but draws its questions from the statements that the user is about to give his consent to?

The second *Sign Up* button is what we call a combo button: It has more than one functionality. Its functionality is described by the small printed text *"By clicking Sign Up, you are indicating that you have read and agree to the Terms of Use and Privacy Policy."*. Both, *Terms of Use* (ToU) and *Privacy Policy* (PP) are accessible through links. After clicking the second *Sign Up* button, and hence agreeing to ToU and PP the user receives an e-Mail. The e-Mail includes no further information regarding ToU or PP, its sole purpose is that the SNS verifies your e-Mail adress. The described registration procedure has not changed from 2009; so 2009 and 2011 lead to following coincided rating shown in Fig. 2.

Year of Investigation	2009	2011
Registration (overall)	O	O
accessibility of terms of use	O	O
accessibility of PP	O	O
consent provisioning action	–	–
registration confirmation	–	–

Figure 2: Account Registration

Year of Investigation	2009	2011
Configuration (overall)	O	O
predefined privacy setting	–	–
privacy setting change	O	O
security relevant change	O	O

Figure 3: Configuration

The accessibility of ToU and the accessibility of the PP would be rated rated positive for a well visible direct link, and negative if there are two or more clicks necessary to reach the information. Facebook's links are below the *Sign in* button and the size of the text is considered to be too small for a sufficient visibility. The action of giving consent would be rated positive for an active component like an opt-in checkbox and neutral for an already filled opt-out checkbox. Here, we have a drive-by consent by clicking the combo button. The registration confirmation would be rated positive when an archivable confirmation information like affected user data and the ToS and PP agreed to would be contained in the confirmation. Just as a contrast, in the banking area this is done physically by sending a letter by postal mail. Using out-of-band confirmations, such as an e-Mail, even without the personal data and the ToS or PP, would still be rated neutral. Not even the latter is the case with Facebook: the user is just greeted by his first name and solely confirms that he is in possession of the e-Mail account.

2.2 Configuration of Settings

After registration, a user's account has a predefined standard privacy configuration. The presentation of the privacy settings has improved in 2011 compared to 2009. Nevertheless, the user had and still has no option to change his personal privacy settings before the data collected during registration is available for others. Thus, a new user's first-, middle-, last-name, and gender are available to everybody on the internet until he changes the defaults, which results in a bad mark. A difference between 2009 and 2011 in the field of configuration is the quantity of user data which is available for others in the standard predefined configuration. Facebook of 2011 is more loquacious [McK11].

After the first login the user can precisely alter his account's settings including his privacy settings. Every successful change is confirmed directly in the web interface. While a password change is followed by an e-Mail confirmation, all other changes are not followed by an archivable affirmation. The overall rating of 2009 and 2011 is still the same as shown in Fig. 3.

Other ratings are based on the criteria described next. The basic configuration would be positive if each data item which is shared later on is selectable by the user. This would be an opt-in approach. An opt-out variant where a users can explicitly deselect data items would be rated neutral. The change of privacy settings would be rated positive if there would not only be an affirmation, but in an archivable form which lists the current changes in detail. This would improve subsequent traceability for the users. In contrast, no affirmation, not even from the web interface, would be rated negative.

2.3 Terms of Use / Terms of Service

The ToU are always directly reachable in the footer of each page. Both in 2009 and 2011, explanations or an annotated version, which would try to help the user to understand them, are missing. Furthermore, Facebook does not keep a track of changes in its ToU nor provides a history. There is also no option to export or print the terms. A user who wants to export the terms has to copy and paste or utilize the browser's print functionality. Fig. 4 shows the results.

Year of Investigation	2009	2011
ToU (overall)	−	−
accessibility	+	+
history	−	−
explanations	−	−
export functionality	−	−

Figure 4: Terms of Use

Year of Investigation	2009	2011
Privacy Policy (overall)	O	O
accessability	+	O
history	−	−
explantions	+	+
export functionality	−	−

Figure 5: Privacy Policy

The accessibility would be rated neutral if there are two clicks necessary, and negative if there are more than two clicks necessary to reach the terms. A partial history of former version would be rated neutral and a complete history would be rated positive. Furthermore, detailed annotations would be positive and superficial annotations would still be neutral. An extra export or print version of the terms would be rated positive, while an extra button or link that invokes the browser's built in print function would be neutral.

2.4 Privacy Policy

The accessibility of the PP has changed between 2009 and 2011. In 2009 there was a direct link in the footer, while in the current version of 2011 the link "Privacy" leads to a page explaining the PP. This contains helpful annotations and further links to videos. However, the link to the PP itself is hidden between other links. History and export functionality for the PP were and are still missing. We use the same criteria as when rating the terms of use. Hence, our comparison shows a difference in PP reachability between 2009 and 2011 as shown in Fig. 5.

2.5 Account Deactivation and Deletion

The menu item "Account Settings" provides the possibility to deactivate the account. The user's data is no longer available online, but Facebook continuously stores the data [fac07]. A complete deletion of the account and corresponding data was, and still is, harder to achieve. The link to the deletion page can be found by searching the help section. However, every hit in the help section at first refers to the account deactivation functionality. In 2011, the link to the deletion page has been added to the PP between plenty other helpful links. The deletion page itself contains a button, which leads to a confirmation page where the user has to enter his password and solve a CAPTCHA. The user gets the affirmation of the account deletion request in two ways: direct feedback in the web interface and via an e-Mail for the "scheduled account deletion". Both affirmations state that the account is now deactivated for two weeks and can be reactivated within this time frame. Only after this two weeks the data is permanently deleted. Our evaluation of the account deletion procedure can be seen in Fig. 6.

Year of Investigation	2009	2011
Account Deletion (overall)	O	O
accessibility	–	O
info about consequences	O	O
procedure of deletion	+	+
affirmation	O	O

Figure 6: Account Deletion

Year of Investigation	2009	2011
Facebook (overall)	O	O
registration	O	O
configuration	O	O
terms of use	–	–
privacy	O	O
deletion	O	O

Figure 7: Facebook Summary

Accessibility would get a positive mark for direct links in the menu. If the deletion is not possible using the online interface, the accessibility would be marked as negative. The information about consequences would be rated positive if the user would be informed about the whereabouts of his data in detail and which data was collected. No information about consequences would be negative. The procedure of deletion would be neutral if it would be more complicated, e.g. if it would be required to sent an e-Mail. A negative rating would be given, if the procedure is complicated and unsecured, e.g. no password request. For positive affirmation an archivable document should be sent to the user. This document shall contain sufficient information about the consequences. No affirmation would be negative. Note: Facebook offers as a separate service to download all the information stored on Facebook. This is new in 2011, however this service is not advertised during the deletion procedure.

2.6 Overall results for 2009 and 2011

As we have shown there are only minor differences between the evaluations of 2009 and 2011. The overall results can be seen in 7. We can only speculate why Facebook is not further improving their privacy policies. A possible reason may be that Facebook classifies the usability of the provided services more important for the user than complicated but perfectly informed privacy settings. There seems to be no or not enough legal pressure on Facebook to induce change.

3 Requirements vs. Current State of Consent Collection and Management

In the previous section we looked in detail at the collection of consent to privacy preferences for users of Facebook. This was part of a comparative study done in 2009 to find differences in consent management between different web application segments. We have analyzed web based banking, an environment with heavy regulated duties especially with respect to ToS, ToU and user identity verification, as well as online-shopping portals. Herein we found, that, except for banking, all segments suffer from the problem of no persistent capture of consent (see 3.3).

All changes to once consented processing and usage of personal data, subtle or big, shall result in a new informed consent. This means that management of previously given consent should be considered as critical as initial registration (see 3.4). How one could technically consent to the aggregation of data was presented in previous work [Pöh08]. In this work we go a step further and collect the necessary requirements for a secure management of consent. We sketch our proposed technical solution later in Sect. 4.

We build upon the observations made during analyzing how one could automate the *right of access* [HJPdM10]. The right of access is is a legal obligation [EU95]. Every company that processes personal data must provide the data subject with means to query the company about what information is stored, processed, and exchanged with other companies. This information is not to be released to anybody else than the data subject. However, in inter-organizational business processes and worldwide supply chains answering requests manually cannot and does not scale. A third party needs to recognize and authenticate a data subject even if the third party does not have a direct relationship with the data subject. Hence, we need means for data processors to authenticate a data subject in later interactions and also for indirect relationships (see 3.5).

Last but not least, in this work we postulate to put an end to unintended data-misuses by clearly marking personal data that is legally safe to use according to consented privacy preferences. We foresee that personal data is obtained with a persistently captured consent. Hence it will provide a third party with a verifiable consent for data usage (see 3.6). Further we assume that data obtained on the black-market or grey-market will lack this verifiable consent. This will make data without the data subject's consent less attractive and decrease the market value of black- or gray-market data.

3.1 Requirement: Privacy by Default

We would like to repeat a long standing request: For a true privacy aware SNS it should offer private-by-default. As shown, Facebook still fails here. We are aware of recent social studies that suggest that the younger generation does no longer expect that data entered into a SNS is private, but it is more cumbersome to first enter bogus data during account registration (i.e. a false name and wrong gender), than adjust the privacy settings for it, only to later enter the correct data.

3.2 Requirement: Informed Consent

First and foremost, we repeat that we need to capture not just consent but "informed consent" [EU95]. Friedmann et al. [FFM00] defined "informed" as the knowledge of and the comprehension of the disclosure in 2000. We reflected only the accessibility of information by Facebook in our rating, not if a user can actually comprehend it. According to Friedman et al. [FFM00] "consent" requires voluntariness, competence, and agreement. They require that "the activities of being informed and giving consent should happen with minimal distraction, without diverting users from their primary task or overwhelming them with intolerable nuisance." [FFM00]. Facebook is already violating Friedman et al.'s concept by placing the button of agreement above the link to acquire the information.

3.3 Requirement: Persistent and Verifiable Capture User's Privacy Preference

Precisely speaking, this requirement postulates two things: (1) We have to capture if and what the user consented to in a persistant and archivable way. (2) It requires the service which collected the data to retain some way of proof for a given consent. As said, the data subject's privacy preferences are defined by the user's privacy settings within the SNS and the SNS's PP and SNS's ToS. During Facebook registration, consent is collected by agreeing to ToS and a PP, implying default PP, by clicking a *Sign In* button. In the simple case the data subject accepts the processor's PP as a suggestion for his privacy preferences when giving his initial consent.In the study conducted in April 2009, we detailed that all data items contained in the user's Facebook profile can be set to be private, for friends only, for friends-of-a-friend, and many can be additionally fine-grained by white- or blacklisting. However, most of the time the service's presets define privacy preferences that share the data subject's data rather than keeping it private. Although, we did not conduct a study on the usability, the fine-grained settings that where possible in 2009 seemed not easy to understand. Facebook adjusted this in 2011 and offered presets. An in-depth comparison if and how these presets or the customizable settings resemble the 2009 privacy settings was out-of-scope and is left for future work. So, once these privacy preferences are determined, they need to be stored and consulted whenever the data is used. A concept better know as "sticky policies" would allow to bind privacy preferences to the data entered.

Next, we will discuss the term of 'notion of verifiability': In case of a dispute, it is not easy for the processor to construct a verifiable proof of the data subject's given consent. Even in a direct relationship, the processor has no verifiable paper-trail, i.e. no paper lottery ticket with a ticked "opt-in" box exists. A processor can assume that a consent was given as the data exists only after the service's sign-up process; In detail consent was given to the processor's PP known to the user at the time of account creation or data submission. To serve as a proof, the processor's sign-up dialogs and the PP need to be timestamped, documented, and their suitability need to be assessed. This can be done by trusted third-parties, even if this was just the simple case. For the moment, we just consider the verifiability between the user and the service which are in a direct relationship. The verifiability between a data item's data subject (user) and the data item's actual holder (no longer the service) is discussed in the requirement of late indirect authentication of the data subject in Sect. 3.5.

In case of a legal dispute the consent collecting actor can find itself in a case where it has to fight off liability. We are not aware of liability cases where the web based process was presented as a proof of given consent. However, there are legal cases that argue against certain technical processes, i.e. cases where a process was not found to allow the collection of an informed consent. Future research with a scope on the legal side of this issue could shed more light on this and maybe bring up missing regulatory frameworks or self-regulation. Especially interesting are cases which deal with changes in ToS or the PP, which leads us to the next requirement.

3.4 Requirement: Managing and Changing Previously Given Consent

As stated above, and also documented with cases by Langheinrich and Karjoth [LK09], collecting the consent is not enough. We will use the image hosting service twitpic[4] as an example to highlight that the current state of the art in managing and changing previously given consent can be flawed. In May 2011 twitpic fell victim to the fact that informing users about a change in the terms of service might be nothing more than a nuisance and minor changes may not lead to a dispute. But if the impact of the changes are heavy, users will be displeased: Twitpic changed its rights and allowed themselves to commercially use the images generated by their users[5]. That change resulted in an upset of the twitpic user community[6]. Assumably, the ToS can be changed with a small note and consented to by still using it, i.e. login after accepting the new terms. It remains unclear if the users indeed expressed their consent. However, this points out that they did not feel “informed”.

We propose that subsequent changes affecting the user’s privacy preferences require a re-collection of consent with at least the same care as the initial consent is collected during the registration. This includes changes to copy- or usage-rights to data in the ToS, as well as PP changes introduced by the service. This may result in more work, confirmation e-Mail to be clicked on, new PDFs with updated ToS to be explained and sent to the user, i.e. via the provided e-Mail address. However, it is not limited to changes that the service does, i.e. change of default PP or ToU/ToS, the same care should be taken if a user changes his settings. Next, all privacy preference changes should be documented and it should be possible for a data subject to acquire the current privacy preferences. This should also be extended to see changes he made over time, introducing a “history of changes”. Legally, the former is codified in the legislation of EU member states, a history is not.

Last, an account deletion requires that consent can be revoked. If personal data is being transferred between services with no direct relationship to the data subject the consent might only exist as long as the data subject has given his consent to the first service. A suggested approach is to time limit the consent apriori, but most often the time of revocation cannot be predetermined. Thus, especially other services that use this data need a way to determine the consent’s freshness automatically. This leads to the next requirement.

3.5 Requirement: Later Indirect Authentication of the Data Subject

We have covered the verifiability between the data subject and the processor that have a direct relationship with each other. In the SNS the data subject gives his personal data to a *primary service*. The primary service then forwards it, either fully or partially, to its successors, and so forth. Each company that is given any part of that data, which is still considered personal,

[4] www.twitpic.com

[5] www.pressgazette.co.uk/story.asp?sectioncode=1&storycode=47089&c=1

[6] www.spiegel.de/netzwelt/web/0,1518,761721,00.html

legally becomes a *controller*. A controller is required by law to grant the data subject certain rights [EU95]. Among them is the *right of access*: The right of access allows the data subject to determine which *categories of data* are being processed and the *purpose of the processing*.

For compound services, like most web services, to work the original processor wants to engage third-parties to provide additional sub-services. For those connected sub-services a direct relationship with the data subject is often not existing or its establishment would be a hindrance to both, the user and the processor. Each sub-service needs to respect the data subject's privacy, hence it would need a proof of consent before processing the data. In SNS we see aggregations of content from other services, or the pushing of content from the SNS to other services, such as Twitter[7]. Due to such exchanges personal data, e.g. pictures, will reach other services. In todays SNS and the current Web these other services do have a direct relationship with the data subject, i.e. Twitter's users also have an account on twitpic used to host photos you want to tweet about on Twitter. Third party services can under potentially different privacy preferences (due to different PP, ToS or privacy settings) combine or re-combine data items from different sources and construct new data sets; a process Pöhls denoted as *aggregation* [Pöh08], while the party doing so will be referred to as the *aggregator*.

We require that an aggregator should, under circumstances set fourth by the data subject, be able to derive the consent for the aggregated data from the consent given to the source data without the need of a direct interaction with the data subject [Pöh08]. We will further refine this by introducing the property of "reconstructability" of data in conjunction with sanitizable signatures in Sect. 4.4.

Based on observations made by Herkenhöhner, Jensen, Pöhls, and De Meer in [HJPdM10] we need, amongst other, a way to authenticate requests as an answer must only be given to a request from an authorized data subject. In consequence we postulate that a solution requires that (1) the consent remains verifiable indirectly for third-parties; (2) that after aggregation of personal data, if consent to aggregation was given, this is still indirectly verifiable.

3.6 Problem: Third Parties Gain No Verifiable Consent for their Data Usage

A third party, receiving personal data, cannot distinguish between black-market personal data that the user has not consented to, and data that is given to a third party with the user's consent. We need to white-flag the positive consented data usage. We need to make gray-market and black-market data detectable, as they would have a missing verifiable consent. Hence, allow data usage only with a verifiable consent statement. Finally, this requires a third party to be able to obtain a confirmation, that the consent has not been revoked by the data subject.

4 Solution: Sanitizable Signed Privacy Preferences

We assume a set of data items D, i.e. the Facebook profile data, gets bound to the consent given, i.e. Facebook's ToU, PP, and the user's privacy settings. Data items $d_i \in D$ are not only held and used internally by Facebook, but handed over to others, i.e. Facebook's Applications, other users, or globally to the Internet. We want to achieve that the data subject who created that *data set* (= a set of data items) can later be authenticated as the legal data subject. Hence, we

[7] www.twitter.com

technically chose a signature scheme which has the desired properties such as origin authentication, integrity protection, and non-repudiation, with respect to identifying the key-pair used for signing. In general, we assume the public key for verifying the digital signature is known. Our signal to consent is a verifiable digital signature over the data set and the accompanied privacy preferences of the data subject. The verification process works as with regular digital signature schemes.

We propose to use a sanitizable digital signature scheme like [ACdMT05; MIea05]. The interview with Miyazaki [Miy08] the author of [MIea05] provides an quick introduction. In a nutshell, a sanitizable signature scheme allows to retain a verifiable valid digital signature on a signed document even if it has undergone changes. The changes are done by a third party called the sanitizer, not the signer, nor the verifier. One of the possible actions a sanitizer might be allowed to carry out is "redaction". Redaction removes original data and leaves only a blinded version $D' \subsetneq D$ of the original data set D. The actual sanitizers and the changes that do not harm the validity are chosen by the signer by choice of parameters or the scheme itself. For example, the signer can either allow everyone to blind sanitizable data items or exercise some form of disclosure control, thus restricting the potential changes in more detail. We assume that once sanitization has happend it is known to the verifier. So, we do not need a sanitizable signature scheme with the property of transparency [ACdMT05]. However, all sanitizable signature schemes allow this sanitization step to take place without an involvement or an interaction with the original signer, i.e. the data subject only signs his data set and privacy preference of it once. Finally, to validate the signature in a sanitizable signature scheme either the original data or a blinded version must be present. For ease of use we use a sanitization scheme which allows everyone to blind data items, so without exercising disclosure control, i.e the Miyazaki Scheme [MIea05].

Additionally, we require the existence of a key pair (sk, pk) as usual for asymmetric infrastructures. We generally assume the use of secure cryptographic primitives. As a result the digital signature verifies under a given public key only if the corresponding private key was used for signature creation and the signed message was not modified in an unauthorized way. We can simply allow users to have more than one identity [Che07] by using a different key-pair for each identity. For better understanding, we will now turn to an example scenario.

4.1 An Example Scenario

Assume the data subject Mr. Luxe has a profile on a social network service like Facebook and consented that Facebook can uses his personal data as stated in Facebook's PP and ToS. However, the SNS is no walled garden and data items like the date of birth and Mr. Luxe's gender are accessible by sub-services. This could be the so called "apps","widgets" or "gadgets" that third parties offer within the social network. We assume that a sub-service $\mathcal{APP}$ needs a date of birth, a name, and a profile identifier (i.e. Facebook's `profile_url`) to connect Facebook users that share the same date of birth. From Mr. Luxe's point of view $\mathcal{APP}$ is a third party. If he enables $\mathcal{APP}$ within his SNS, he grants some processing rights for certain data items. Assuming that Mr. Luxe has consented to this usage, the third party receives Mr. Luxe's `name`, `birthday`, and `profile_url`. This is the actual case today, however this comes with all the consequences.

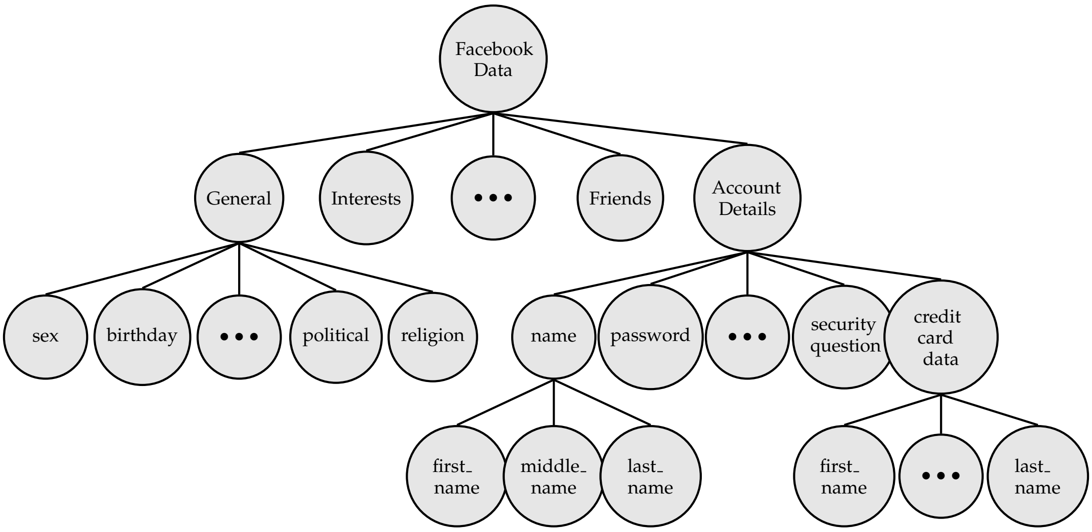

Figure 8: Subset of Facebook Profile Data Items; Names Corresponding with `Users.getInfo`

4.2 Sanitizable and Signed User Data Tree Structure

Our solution assumes a sanitizably signed data set. Hence, we use a suitable sanitizable signature scheme to calculate a signature σ over all the data items $d_i \in D$, while the privacy preferences are attached to them. This signature forms the user's consent to processing with respect to the privacy preferences. Due to the nature of a sanitizable signature scheme, we only have to generate this signature when either a data item d_i is changed, added, or the user modifies his privacy preferences. Fig. 9 visualizes this for profile data items from the general section of Fig. 8. In the above example $\mathcal{APP}$ was about to be given some, but not all data items. Here, the sanitizable signature comes into play; the SNS can remove all the data items d_i which are not needed for $\mathcal{APP}$, while $\mathcal{APP}$ can still verify the signature σ and hence can verify that consent was given to the remaining data items. With the data items the SNS also removes all the privacy settings that are making statements about data items not given to $\mathcal{APP}$. Fig. 10 shows a possible remainder after redaction.

4.3 Social Network Service as the CA

Our assumption requires that there exists a signature which is verifiable using public keys. This implies that we need some sort of Public Key Infrastructure (PKI). How such a solution could piggyback on existing X.509 certificates was already described by Pöhls in [Pöh08]. We have enhanced this with sanitizable signatures. However, the burden of having to create and use private keys on the user side still remains and has been criticized by Karjoth and Lang in [LK09]. Hence, we opted for a solution where another stakeholder for only consented data processing is involved on behalf of the user: the social network service itself. For legal and reputation reasons the social network service tries to enrich and secure that the data items leaving their network

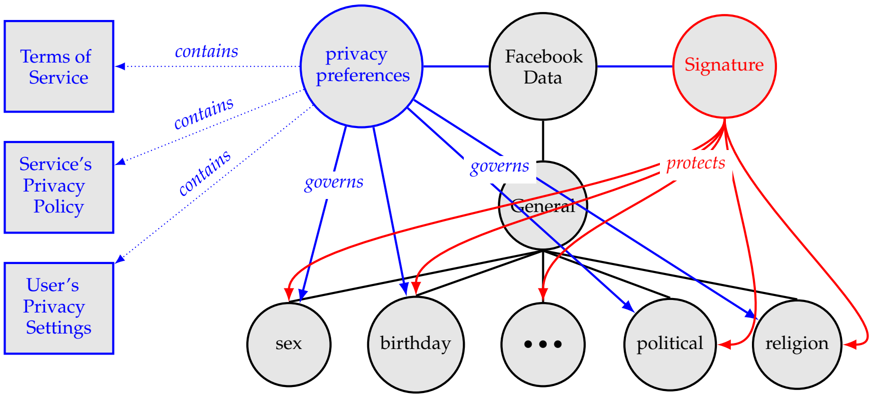

Figure 9: Adding the Signature and Privacy Preferences to the Facebook Data Set

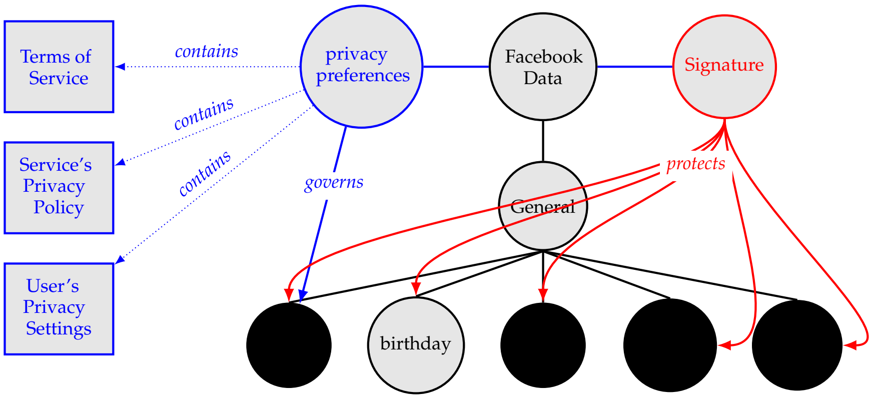

Figure 10: Redaction of all Data Items Except of `birthday` from the Sanitizable Signed Data Set of Fig. 9

are easy to process according to the consent. Hence, third parties receiving this data have all the tools to verify if the data subject still consents to the data's use. The user does not even have to know that this system is in place; it could be done transparently for the user. However, we want to stress that it needs to be an open system and that user should be able to use their own or third party provided authorities, that would sign their data. Hence, storing and using the signed personal data must be separable from the creation of signed data.

4.4 Notion of Reconstructability

The mentioned sanitizable signatures schemes allow that the signed data set can be sanitized more than once. This allows that an provider of applications has access to two data sets - the old and the new one, where different data items/vectors are present, i.e. d_1, d_2 along with their signatures σ_1, σ_2. Consider that the gender data item is only available in data set D_1, while is has been redacted in D_2. Additionally, the date of birth d_2 has not been available in D_1, while it has not been redacted by the user in D_2. Hence, data set D_1 is not a strict subset of D_2 - and vice versa. This could lead to the desire to construct a data set D_3 with its signature σ_3, which

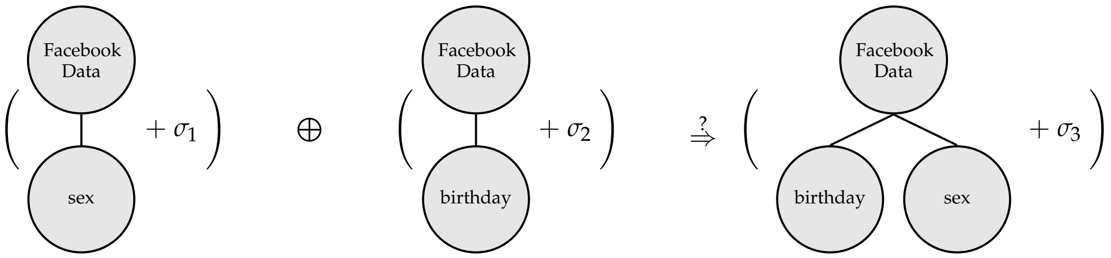

Figure 11: Linkability and Reconstructability of Data Items from the same Data Set

contains the merged information d_1 and d_2, such that $d_1 \in D_3 \wedge d_2 \in D_3$, while the resulting signature σ_3 keeps valid over the data set D_3. Hence, D_3 forms a subset of the original data set D, i.e. $D_3 \subseteq D$. This actually requires two steps; (1) one must be able to link two data sets - in particular one must be able to decide whether data items d_1, d_2 belong to the same original data set D, i.e. if $d_1 \in D \wedge d_2 \in D$ yields. Brzuska et al. called this property "linkability" in [BFLS10]. However, since a data set received is always associated to exactly one identity (respectively signing key), the data sets are always "semantically linkable" in the sense that the verifier knows that the data belongs to exactly one account. Especially, $\sigma_1 = \sigma_2$ if a redaction invariant signature and the same key is used. This leads to the next step: (2) Produce a new data set D_3 and σ_3, where d_1, d_2 are both contained in D_3, i.e. $(d_3 = d_1 \cup d_2) \in D$, where σ_3 is a valid signature on D_3. This has been illustrated in Fig. 11.

In [BFLS10] Brzuska et al. implicitly state that "Linkability" is equivalent ($\equiv$) to "Reconstructability". We assume that this may not always be true, since σ_3 must be generated as well. Reconstructability describes how to (re-)generate a verifiable signature over aggregated data in a way that it remains transparent that separation has taken place before. Even though, a scheme with this property is invading privacy, our new notion allows to identify controlled reconstructability for existing and future schemes.

5 Conclusion

We analyzed Facebook and showed how this social network collects and manages user's consent for processing personal data. While consent collection and out-of-band confirmations (i.e. e-Mails) are unchanged from 2009 to 2011, the only change are the privacy settings; Facebook is now a bit more elaborate. From this current state we postulated some old and many new requirements that would ease proving and forwarding collected consent to third parties. This may look like letting the fox guard the henhouse, as we proposed that the social network service could act as a certificate authority and issue verifiable vouchers that a user consented to third party data processing. Some might even say this is like letting the fox into the henhouse. However, this would be a first step towards white listing consented usage and allowing third parties to verify the user's consent. Social networks are already seen by users as trusted authorities to handle their personal data. For example, adding Facebook's digital signature over a Facebook data item together with Facebook's privacy preferences and hand it to a third party Facebook application does not stop data leaking, neither at Facebook, nor at the application. However, we gain a differentiator between black-market data and data that was consented to for certain processing. Data buyers no longer need to trust sellers to have the data subject's consent, it allows automatic verification. Hence, the next time the selling of un-consented personal data makes headlines and politicians want to "forbid the selling of personal data" completely, our solution could help. Of course, once we have a interoperable framework in place we can outsource the certificate authority to the social networks users. Finally, we showed that the new sanitizable signature schemes can be put to yet another good use.

References

[ACdMT05] G. Ateniese, D. H. Chou, B. de Medeiros, and G. Tsudik. Sanitizable Signatures. In *Proceedings of ESORICS*, pages 159–177, 2005.

[BFLS10] Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. Unlinkability of Sanitizable Signatures. In *Public Key Cryptography*, pages 444–461, 2010.

[Che07] Zhikui Chen. A Scenario for Identity Management in Daidalos. In *CNSR '07: Proceedings of the Fifth Annual Conference on Communication Networks and Services Research*, pages 176–183, Washington, DC, USA, 2007. IEEE Computer Society.

[EU95] EU. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, Official Journal of 23 November 1995, L 281, page 31 - 50, Nov. 1995.

[fac07] facebook. Policy. `www.facebook.com/policy.php`, Dec. 2007.

[FFM00] B. Friedman, E. Felten, and L. I. Millett. Informed Consent Online: A Conceptual Model and Design Principles. Technical report, CSE Technical Report (University of Washington), 2000.

[HJPdM10] Ralph Herkenhöner, Meiko Jensen, Henrich C. Pöhls, and Hermann de Meer. Towards Automated Processing of the Right of Access in Inter-Organizational Web Service Compositions. In *IEEE 2010 International Workshop on WebService and Business Process Security (WSBPS 2010)*. IEEE, Juli 2010.

[LH06] Y.-L. Lai and K. L. Hui. Internet opt-in and opt-out: investigating the roles of frames, defaults and privacy concerns. In C. Shayo, K. Kaiser, and T. Ryan, editors, *CPR*, pages 253–263. ACM, 2006.

[LK09] M. Langheinrich and G. Karjoth. Einwilligung und ihre technische Umsetzung. *digma – Zeitschrift fuer Datenrecht und Informationssicherheit*, 9(4), 2009.

[McK11] Matt McKeon. The Evolution of Privacy on Facebook. `mattmckeon.com/facebook-privacy`, Apr. 2011.

[MFSV07] M. Madden, S. Fox, A. Smith, and J. Vitak. PEW Internet & American Life Project Report: Digital Footprints. `www.pewinternet.org/pdfs/PIP_Digital_Footprints.pdf`, Dec. 2007.

[MIea05] K. Miyazaki, M. Iwamura, and et al. Digitally Signed Document Sanitizing Scheme with Disclosure Condition Control. *IEICE Transactions*, 2005.

[Miy08] Kunihiko Miyazaki. Redactable Digital Signatures for Secure and Easy-to-use Digital Document Systems. `www.hitachi.com/rd/sdl/people/suminuri/index.html`, May 2008.

[Pöh08] Henrich C. Pöhls. Verifiable and Revocable Expression of Consent to Processing of Aggregated Personal Data. In M.D. Ryan L. Chen and G. Wang, editors, *ICICS 2008*, LNCS 5308, pages 279–293. Springer, 2008.