

Spiny CACTOS: OSN Users Attitudes and Perceptions Towards Cryptographic Access Control Tools

Ero Balsa*, Laura Brandimarte†, Alessandro Acquisti†, Claudia Diaz* and Seda Gürses*‡

* KU Leuven, Dept. of Electrical Engineering (ESAT), COSIC, iMinds, Leuven, Belgium

Email: {ero.balsa, claudia.diaz}@esat.kuleuven.be

† Heinz College, Carnegie Mellon University, Pittsburgh, PA, USA

Email: {lbrandim, acquisti}@andrew.cmu.edu

‡New York University, Dept. of Media, Culture, and Communication, NY, USA

Email: seda@nyu.edu

Abstract—Cryptographic access control tools for online social networks (CACTOS) allow users to enforce their privacy settings online without relying on the social network provider or any other third party. Many such tools have been proposed in the literature, some of them implemented and currently publicly available, and yet they have seen poor or no adoption at all. In this paper we investigate which obstacles may be hindering the adoption of these tools. To this end, we perform a user study to inquire users about key issues related to the desirability and general perception of CACTOS. Our results suggest that, even if social network users would be potentially interested in these tools, several issues would effectively obstruct their adoption. Participants in our study perceived that CACTOS are a disproportionate means to protect their privacy online. This in turn may have been motivated by the explicit use of cryptography or the fact that users do not actually share on social networks the type of information they would feel the need to encrypt. Moreover, in this paper we point out to several key elements that are to be considered for the improvement and better usability of CACTOS.

I. INTRODUCTION

In the last 10 years, the success of online social networks (OSN) has raised significant privacy concerns. The wealth of data that can be collected by OSN providers has raised questions regarding the balance of power between data subjects and data holders, their respective trade-offs, and how to balance their respective interests. Computer scientists and security engineers have developed cryptography-based access control tools for OSNs (CACTOS) [6], [19], [25], [33], aimed at providing users greater control over privacy in OSNs. These tools are grounded on robust security foundations that try to provide the average Internet user with greater autonomy and freedom from online surveillance. The designers of these tools often

present them as especially needed in OSNs given the sensitive nature of the information shared on these platforms. Users upload information related to private realms of their lives, stored and processed by the service providers of these platforms, in some cases without the users' awareness or informed consent [3]. CACTOS can also solve more mundane privacy problems, such as preventing employers from finding out about blatant criticism from their employees or avoiding that parents intrude the social space of their children. However, existing tools—even those which are free and publicly available—have seen little adoption in practice [2], [34]. Security and privacy experts have hinted at possible explanations why OSN users lack interest in privacy technologies. Some argue that there is scarce interest for privacy enhancing technologies—or, more broadly, for protecting privacy—in OSNs—which are meant for sharing and communicating [21]. Others point at the poor understanding of privacy problems that many users have [26]. Another plausible explanation anchors the problems in the (lacking) usability of these tools, that prevents users from being able to use them. Despite the wealth of user studies related to privacy on OSNs, little research has been done so far on the perceptions and attitudes of users towards CACTOS [34].

In this paper, we report on the results of a user study aimed at discovering the hitches and obstacles to the adoption of CACTOS. We have investigated participants' attitudes towards the “privacy properties” CACTOS aim to provide, such as the protection of communication content from the service provider itself and additional controls over users' information flows without the burden of relying on a third party for enforcement. In the execution of our study, we provide participants with an access control tool for Facebook called *Scramble!* [6] in order to discuss their attitudes with respect to a concrete CACTOS. Hence, our goal and focus are less to perform a usability study of a specific tool, and more to investigate users attitudes towards the privacy properties and benefits of cryptographic privacy preserving tools for OSNs. Still, usability is a key factor on users' attitudes towards a given tool. Hence, in our work we try to separate the impact of usability from users' attitudes towards privacy properties, and still unveil usability issues that result in recommendations for better, more

usable, cryptographic tools. The study results suggest that OSN users perceive CACTOS as disproportionate and inefficient for protecting online privacy—either because of high costs “average end-users” associate with learning, adopting, and using cryptographic tools, or because of the low benefit expected in comparison to alternative privacy enhancing strategies [22], e.g., self-censorship. Despite these results, it is possible to draw from the study suggestions towards designing more usable and useful CACTOS.

This paper is structured as follows. In Sect. II we review the literature relevant to our research. In Sect. III we provide the set of definitions and terminology that we use throughout the paper. In Sect. IV we describe the goals of the user study and the experimental setting. We describe our analysis methodology in Sect. V and present our results in Sect. VI. Lastly, we discuss the main findings of our study and provide recommendations for the design of better CACTOS in Sect. VII.

II. RELATED WORK

A. Privacy Preserving Tools for Online Social Networks

Computer scientists and security researchers have proposed a wide range of privacy preserving tools for OSNs. The mechanisms these tools rely on vary widely. Some of them piggy back on existing OSN infrastructures—either relying solely on the existing infrastructure [6], [14] or complimenting them with external infrastructures [33]. Other tools are based on “alternative” or decentralized architectures, and require users to migrate from their current OSNs to new platforms [5], [11], [15]. Despite the differences, these tools share a common goal: to allow users to control who has access to their content on an OSN (e.g., personal information, posts, messages) without the need to depend on the service provider’s privacy controls. This also implies that the OSN provider is barred from gaining access to the users’ content. In sum, these tools undo the delegation of access control enforcement to OSN providers by taking advantage of cryptographic properties. Despite the mathematically and technically undergirded “privacy properties” that these tools guarantee, to the best of our knowledge, none of them have experienced significant adoption.

B. User Attitudes and Perception towards Privacy

Many researchers have studied users’ attitudes towards privacy. Acquisti and Grossklags [4] found supporting evidence for the dichotomy between privacy attitudes and behaviour. Lampinen et al. [23] performed a study on the attitudes of OSN users towards *context collision*, and the strategies of users to manage this problem. Besmer et al. [7] examined the privacy issues that users encounter when sharing photos in OSNs, finding out that these issues were motivated by identity and impression management. Raynes-Goldie [27] studied how Facebook users understand and tackle privacy problems, finding out that users are more concerned about controlling access to personal information than to how Facebook uses and processes that information. Xu et al. [36] revealed that users’ privacy concerns are not solely determined by their individual perceptions and attitudes, but also informed by organizational factors such as privacy policies. Our study differs, in that we study user attitudes towards using CACTOS to protect their privacy in the context of OSNs.

C. Usability Studies of Cryptographic Tools

Usability, defined as *ease of use*, is a key factor to ensure that a certain tool is adopted. Previous research has shown that this property is especially critical in the case of cryptographic tools. The complexity and obscurity of cryptography makes most crypto tools difficult to use. Whitten and Tygar’s seminal study on the usability of PGP [35] pointed to several design flaws that would prevent the general Internet user from successfully using PGP. The authors blamed the inadequacy of general user interface design principles for security tools, which result in dangerous mistakes, such as users sending secrets without encrypting them first. Fahl et al. [17] performed an analysis of privacy preserving tools for OSNs and a usability study of existing approaches. They found that users tend to prefer tools that minimally disrupt their user experience. Our results support this finding. We provide evidence that users desire seamless, nearly automatic integration between the access control tool and the social network.

Lastly, based on previous user studies, Vemou and Karyda provide a classification of factors that may affect the low adoption of privacy preserving tools for OSNs [34]; such as unawareness of privacy enhancing technologies (PETs) or the fact that users lack the technical skills required to use PETs.

III. PRELIMINARIES

A. Definitions and Terminology

We informally refer to *content* to denote the different bits of data a user keeps on an OSN, such as messages, posts, photos and any other types of media. We refer to *privacy settings* to denote any set of mechanisms that define access control rules over the content of a user in an OSN—that is, they allow users to express who should be able to access their content. We define *unintended recipients* as any person or entity that, according to a set of privacy settings, is not authorized by the user to access a certain piece of content (e.g., a post) on the OSN. Lastly, we refer to *Cryptographic Access Control Tools for Online Social Networks* (CACTOS) to denote any tool that tries to address privacy problems that arise due to the following design decisions:

- *Delegation of privacy settings enforcement*: users expect or need to rely on the OSN to enforce their privacy settings.
- *Disclosure of all content to service provider*: all user content is by default disclosed to the OSN provider and it is not possible to express access control rules towards the OSN provider.

In this paper, we refer to these as *SNP-problems*.

The objective of designers in developing CACTOS is to mitigate SNP-problems by (1) allowing users to enforce their privacy settings independently of the mechanisms already available on the OSN, (2) allowing users to do so without relying on the provider or any other third parties, (3) allowing users to prevent the service provider or other third parties from accessing their content. Throughout this paper we refer to these features as *CACTOS-properties*.

Further, we consider that users may adopt CACTOS for a variety of *reasons*. Users may want (a) to keep their content

confidential from the OSN provider, i.e., have the option to set access control rules that apply to the OSN provider. This may also be used (b) to guarantee that the user’s content cannot be indexed by the OSN or any other parties. Users may want (c) to delegate the enforcement of their privacy settings to the OSN provider no more. This may be due to past experiences with the OSN provider, i.e., the latter may have changed the semantics of the privacy settings or the defaults in undesirable ways or applied unreasonable rules to (new) features leading to privacy violations. Last, users may want (d) to have greater granularity in setting new access control rules, being CACTOS’ access control rules more expressive than those of the OSN provider. We will compare these reasons against the privacy problems currently perceived by OSN users to assess whether CACTOS could effectively help OSN users mitigating those problems.

Lastly, we note that the definitions above are *informal*, i.e., we do not attempt to formalise or establish strict limits for any of the concepts above. Rather, we provide these definitions to loosely summarize the objectives behind CACTOS design, to define the scope of this study (as our findings may have implications beyond the tools considered here), and to improve the readability of this paper.

B. Scramble!

Scramble! [6] is a publicly available CACTOS tool developed by security engineers as a browser plug-in (concretely, a Firefox extension).¹ *Scramble!* features a hybrid cryptosystem (based on OpenPGP [12]) to enforce access control lists [28]. This allows *Scramble!* to enforce privacy settings independently of the OSN provider. Even if *Scramble!* features a standard interface design, usability testing, as is typical to many CACTOS, was mostly limited to the immediate environment of the designers of the tool and was not systematically part of the development. *Scramble!* can be installed as any other Firefox extension. After installation, users need to initialise the tool, which involves generating (or importing) a pair of public-private keys. Users then need to add their friends’ public keys to *Scramble!*’s key ring; either by asking *Scramble!* to perform a lookup in a public key server or by loading them from a local file. After the initialisation, *Scramble!* is ready to use. To encrypt a message before they post it to the OSN, users select the text of the message and choose *Scramble!* on the right-click drop down menu. This makes a window appear, where users can choose the people they wish to encrypt the message for. Afterwards, *Scramble!* automatically replaces the plaintext with its corresponding cyphertext, which the user can then post to the OSN. That way, all the user’s communications are stored encrypted on the server of the OSN. This provides several benefits: users do not need to trust the OSN provider to guarantee the confidentiality of their content, they do not need an alternative trusted server to store their data, and they can independently manage who is able to access their information. Furthermore, *Scramble!* provides automatic decryption (whenever users have the right cryptographic keys to decrypt the content), releasing users from the burden of dealing with the decryption of messages.

¹Both the Firefox extension and the source code can be downloaded at <http://www.cosic.esat.kuleuven.be/scramble/>

IV. THE STUDY

A. Goals

The ultimate goal of the study presented in this paper is to gain insight into the reasons that can foster or hamper the adoption of CACTOS. To this end, we set out to assess whether OSN users share designers’ attitudes towards the importance of and the best ways to mitigate SNP-problems. Specifically, we wanted to find out whether OSN users (Q1) share the concerns of CACTOS designers with respect to SNP-problems; (Q2) think they should be responsible for taking measures to mitigate the SNP-problems; (Q3) think they should use a (technical) tool to address SNP-problems, and (Q4) comprehend and agree with the way CACTOS address SNP-problems.

To address these questions, we performed a user study consisting of two questionnaires and a guided tour to a CACTOS (*Scramble!*). Through the questionnaires, we (S1) asked participants about their privacy problems on OSNs, and assessed whether they relate to SNP-problems; (S2) asked participants about what they feel personally responsible towards protecting their privacy, and assessed whether their answers relate to SNP-problems; (S3) asked participants about the strategies they use to tackle SNP-problems, and assessed if these strategies overlap with CACTOS-properties; (S4) asked participants about their experience using one CACTOS (*Scramble!*), and assessed whether they find these tools useful and appropriate.

B. Study Design

Our study consisted in three phases: an entry questionnaire, a guided tour to *Scramble!*, and an exit questionnaire.

1. The goal of the entry questionnaire was to address questions Q1, Q2 and Q3. This questionnaire consisted in 36 questions² of which, given space limitations, we only analyse a subset in this paper.

2. The goal of the guided tour to *Scramble!* was to provide participants with a hands-on experience of a concrete implementation of a CACTOS, thus providing them with an example to evaluate the benefits and drawbacks of such tools. We expected this evaluation would hint to possible obstacles that hinder the adoption of CACTOS by revealing discrepancies between users’ communication needs and the communication models inherent to CACTOS. We chose *Scramble!* for this tour because it features the basic CACTOS-properties, its source code is freely available and we could rely on the assistance of his main developer, Filipe Beato.

Because we wanted participants to evaluate *Scramble!* in the context of the SNP-problems it tries to solve (S4), we provided them, right before the guided tour, with a short list of potential privacy issues arising from SNP-problems on Facebook. This list included issues such as: the fact that Facebook has access to all the information they upload to the site; that privacy settings afford little protection against system vulnerabilities or someone breaking into Facebook; that

² Due to space limitations, rather than providing the full list of questions of the questionnaires in the annex, we transcribe each of the questions inline when we present the results.

users can change the settings but need to rely on Facebook to properly enforce them, or that Facebook could reveal their information to third parties. Afterwards, we gave participants a *manual* to use *Scramble!*, including both an introduction to what *Scramble!* is and how it works and the instructions for the guided tour.³

The tour involved all the steps that a new *Scramble!* user needs to follow to be able to use the tool. Therefore, users were guided through the download and installation process and then provided with instructions on how to encrypt messages. After the installation, participants were encouraged to send encrypted messages to one or more participants in the same session, and hold a brief conversation.

One may wonder why we designed a guided tour instead of a traditional usability environment in which participants are required to perform a set of tasks on their own. Previous research has shown that users often have a hard time using encryption technologies, specifically at first contact (i.e., learning from scratch) [17], [31], [35]. In the case of *Scramble!* we feared a similar outcome; and also usability testing was not our main objective. Hence, we guided users step-by-step on the usage of the tool. Usability experts may fear that this decision may have affected the ability of the participants to judge *Scramble!* based on its ease of use. However, usability is just one of the factors that may affect the adoption of CACTOS [34]. Moreover, we show in Sect. VI that, in spite of the guidance, participants raised numerous usability issues. Hence, incidentally, we also provide suggestions towards improving the usability of *Scramble!*.

3. The goal of the exit questionnaire was to address question Q4. It consisted of a total of 11 questions⁴ plus the system usability scale (SUS), based on a 10-item attitude Likert scale, a widely used metric in usability studies [10]. SUS yields a score between 0 and 100, allowing a quick, coarse evaluation of usability. This allowed us to perform a sanity check on the degree of usability of *Scramble!*, although we insist that this was only a secondary objective of our study.

C. Study Setup

The phases mentioned above occurred in a laboratory environment during the first week of September 2013. At arrival, participants were instructed to log into one of the 8 available computers in the laboratory and received an email with a link to the entry questionnaire. They were given 15 minutes to complete it, at which point they received the documentation for the guided tour to *Scramble!*. They had 30 minutes for the tour. Then, they received a second email with the link to the exit questionnaire, to be completed within 10 minutes.

We invited 52 students (42% female, average age = 21.5, SD = 2.6) from the Center for Behavioral Decision Research Pool at Carnegie Mellon University to participate in a study to “Test Scramble! - A Facebook app.” Participants were paid \$10. All participants had been using Facebook for at least 2 years and identified themselves as active Facebook users.

³If needed, all the documentation is available from the first author.

⁴See footnote 2

V. SURVEY ANALYSIS METHODOLOGY

The main purpose of the analysis was to find common features or patterns in the participants’ responses. For instance, after asking participants if they found *Scramble!* to be a useful tool and why, we categorized their responses based on patterns in their justifications. In summary, our methodology was largely inspired by the method of *emergent coding* [24]. Emergent coding is especially useful to work on a topic for which there is limited prior research, as it is the case of user studies for CACTOS.

Specifically, during the first round of analysis, we developed codes to capture the essence of each of the responses individually. Each code, which consisted in one or two keywords, summarized the key points in a given response. For this first round of coding, we heavily relied on *in-vivo* codes, namely, terms provided by the participants themselves that summarise accurately the concept they are referring to. During the second round of coding we used thematic, hierarchical coding. This involved grouping the *in-vivo* codes in themes that emerged during the first round. Themes were in turn grouped into more general, broad themes, until we found that further generalisation was not possible or made no sense.

There are some caveats that are typical to such a study [32]. Throughout the presentation of the results we refer to the concepts that emerged from the analysis. However, these concepts were not drawn from previous theory (e.g., [34]) but rather resulted from our own interpretation. In order to contextualize our analysis in the responses, for each category we quote our study subjects. On the other hand, in the discussion of our results we heavily rely on research-denoted concepts. There we also only provide our own interpretation and opinion of the implications of this study.

VI. RESULTS

In this section, we present the results of our study. We use “quotation marks” to refer to questions in the questionnaire and both “*italics and quotation marks*” to refer to participants’ responses. We place the participants’ quotes between parentheses when we give several examples of a certain attitude, perception or position.

Further, note that even though the objective of the study is not to provide quantitative results, at times we mention numbers to indicate whether articulated positions were voiced by a majority or a small minority of the participants.

A. Responsibility vs Control (Q1, Q2)

Following strategies *S1* and *S2*, we explored whether users would be interested in having control over their privacy as enabled by CACTOS-properties (*S1*), and compared it to whether they feel responsible for taking measures to mitigate SNP-problems (*S2*). To test the former, we asked participants “Who should decide...?” for a set of decisions related to the visibility of their data on Facebook such as “[who should decide] who is able to see what you post on the site?” or “who is able to see your personal details (age, phone number, hometown, etc)?”. For each decision, users were able to choose among the following options: *You, Facebook, or Others - Please specify*. Multiple choice selections were

allowed. For all decisions but one, most participants thought that they should have control over those decisions themselves, being the average percentage of *You* and *Facebook* across all decisions $M = 81.3\%$ ($SD = 12.9\%$) and $M = 23.9\%$ ($SD = 14.1\%$), respectively. This is consistent with previous research suggesting that users desire control over personal information, even though they may not actually make use of that control [8]. This desire for control seemingly clashes with the unwillingness of participants to be the only responsible for many privacy-related issues we proposed to them. Following *S2*, we asked participants “Who should be responsible for the following [privacy related] decisions?” such as “setting the proper privacy settings on your profile” or “making sure private companies do not have access to the data you post to the site without your permission.” Most participants considered that Facebook should be responsible for issues such as “making sure your privacy settings work” or “making sure strangers cannot see your photos/posts online.” Since these are matters that are outside of the authority of the users, it is not surprising that most of the study participants saw these matters as part of the responsibility of service providers. However, a few participants declared that Facebook should be responsible for “setting the proper privacy settings on your profile” or even “making sure your friends do not post photos of you that you do not like.” In general, most users attributed to Facebook greater responsibility than control. Across all decisions, the percentage of *You* and *Facebook* was $M = 67.4\%$ ($SD = 21.1\%$) and $M = 52.2\%$ ($SD = 27.8\%$), respectively. Furthermore, we noticed certain trends in how users assigned responsibility. Users attributed to Facebook the responsibility for issues such as “preventing strangers from logging in to your account” and “preventing people other than your friends from reading your messages and seeing your photos,” (and more generally, issues that are out of their control by default on Facebook), whereas they attributed to themselves the responsibility for “what your friends can see in your profile.” (and more generally, matters for which Facebook provides privacy controls).

Regardless of the amount of control users desire over their privacy, they may not consider themselves responsible for the SNP-problems tackled by CACTOS. Following *S2*, we asked participants “On Facebook, what do you feel responsible for with respect to your own privacy?” Participants’ answers varied widely but can be classified in two main categories: *what they post*, and *how they use the functionalities provided by Facebook*. Answers belonging in the first category were the majority (47%). Participants referred to the nature of the content they post on Facebook and, specifically, to *self-censorship* practices, e.g., “*I am responsible for not posting very intimate and personal photos [...]*”, “*I [...] feel responsible for the kind of posts and personal information I put up online*”. Answers in the second category referred to practices such as “*getting my privacy settings right*”, “*the access I set for each of my photos/posts etc.*” as well as other options provided by Facebook such as “*making sure that people who I don’t wish to have contact with are blocked*” or “*making sure that my profile is not [...] able to be searched on google*”. Moreover, some participants suggested that they are not responsible for making sure that the privacy settings are actually enforced according to their preference —rather, they answered, that this is Facebook’s duty. For instance, one participant wrote: “*[Block people. It is then facebook’s job to make sure that*

they cannot message me from that point on”, or “*I am not the one who can guarantee the execution of [the privacy settings], Facebook does it. At this level what choice do I have? To trust Facebook*”. Interestingly, respondents often referred to feeling responsible for managing photos, links, comments they post but not the private communications they take part in. In fact, some participants explicitly mentioned that Facebook’s *private domain* falls beyond the scope of their responsibility, charging Facebook instead, e.g., “*I feel responsible for the content of my public posts/comments and photos posted to the public. I feel I should not have to further manage private messages [...] which I want to remain private and have selected as such*”. It is also interesting to note that users feel responsible for social privacy problems and barely mention surveillance problems [20]. Participants acknowledged being responsible for the configuration of their privacy settings, but limited their self-censorship practices to the public domain. They consistently disregarded the fact that privacy settings do not prevent the service provider (Facebook) from being able to access all their content, both public and private. The response of one participant nicely summarises this trend: “*I would rather friends send me private messages if they want to share something fun with me*”.

All in all, we caution that different people may have had different interpretations of what “responsibility” means. More studies are needed to surface further evidence of the control-responsibility dichotomy.

B. Limitations of Facebook’s Privacy Settings (*Q1*, *Q2*)

We asked participants “Which privacy problems, if any, do you encounter using Facebook?”. The goal of this question was, on the one hand, to find out whether they would mention or refer to any SNP-problems (*S1*) and, on the other hand, to get a glimpse of what other problems they might be dealing with. We posit that all responses but “None” can be effectively used as input for improving CACTOS. For each response, we also reflect on the potentials and shortcomings of CACTOS in addressing these problems. The responses of the participants can be classified in two main categories: *lack of control over information flows* and *lack of control over how their information is used*. Within the former category, we found several groups of problems: Participants referred to the lack of control over the privacy settings (“*privacy settings always seem to be reset to a lower level after each update*”) and the lack of granularity of the privacy settings (“*[...] would be very helpful to be able to decide exactly who can see each post/friend’s post/photo/piece of info [...]*”). Note that, on Facebook, individual privacy controls are available for items like posts, but not for others such as specific photos in an album. Also, the ability of a user to manage the audience of a certain post is constrained by the predefined categories available on Facebook, e.g., one can share a post with specific *friends* but it is impossible to select specific *friends of friends*: one must share with all *friends of friends* or none of them (same for the *Public* category, which does not allow to select specific people outside Facebook). Other participants referred to the difficulty of effectively deleting information (“*even you delete something still it can be seen by a search tool*”). We note that these are the type of issues expected to motivate the adoption of CACTOS, in line with the reasons mentioned in Sect. III (reasons *c*, *d* and *b*, respectively). However, we also

found problems in this category that one cannot solve with CACTOS, namely, the lack of control over somebody else’s activities, e.g., “*I am able to stop people from seeing my posts, but not posts I’m tagged in*”. Similarly, the category *lack of control over how information about them is used* comprises both problems that can and cannot be solved with CACTOS. A couple of participants mentioned privacy problems related to their information being processed to target advertising to them (“*my data being used to target ads at me*”). One could use CACTOS to mitigate the processing of posts for advertisement by concealing them from the service provider (reason *a*). However, CACTOS are often not intended to prevent “*Facebook tracking you across the web*”. Finally, other participants pointed out to strangers sending them messages or friend requests (e.g., “*random people message me even though I don’t know them*”, or “*Some random strangers sending me friend requests*”). As yet, no CACTOS we are aware of provides a solution to this problem.

We asked participants “What, if anything, would you add to, modify or delete from the Facebook privacy settings?” to assess whether they would be interested in any of the CACTOS-properties (S1). Many participants pointed to the lack of granularity of the settings and the fact that these settings do not allow to control the visibility of certain items, e.g., “*want to put limit to the photos one by one not the whole album*”, “*cover photos are all public. I would change that*”, “*I would make it possible to hide specific things from specific people*”. This fine-grained control over the visibility of any specific item is one of the features of CACTOS (reason *d*). Some participants pointed out to enhanced protection from search engines, e.g., “*[...] a function that makes your profile unsearchable for a certain amount of time [...]*”, “*[...] that no one can find [my information] even if googled*”. Encryption can definitely disrupt the ability of search engines to index content. However, CACTOS are weak in the flexibility they offer in managing disclosure decisions over time (e.g., flexible revocation). Some participants suggested the ability “*to see who has viewed my page*”, “*if any random user [...] viewed my pictures*”, while another participant requested the opposite “*would never allow people to see the profiles I’ve looked at*”. In any case, CACTOS do not provide control over these features. Some participants referred to Facebook’s terms of service and, more specifically, to “what Facebook does with your data”. One can use CACTOS to conceal data from the service provider (reason *a*), thus preventing the collection in the first place.

Moreover, we asked participants “What privacy issues you have, if any, that you are not able to solve with Facebook’s privacy settings?” (S1) followed by “Which strategies do you use to solve those privacy issues?” (S2). Most of the participants reported to have “*none*”, i.e., no issues unsolvable with Facebook’s privacy settings. On the other hand, responses from those participants who did have some issues can be classified in three categories: Some participants referred to (1) the lack of control over other people’s activities (“*Friend requests from strangers*”, “*Individuals comment inappropriately on [my] status*”). Coping strategies included “*reporting [to Facebook]*”, “*deleting the post*” but also “*nothing*”. CACTOS do not provide solutions to these problems. Other participants referred to (2) the lack of control or knowledge on Facebook’s uses of data (“*to not sell my data to companies*”, “*Tracking*

you across the web”, “*Seeing who my top friends are on chat or on my profile*” or “*The adds [sic] I see in facebook are related to even my google searches, they interfere in to every space of mine.*”). Coping strategies mentioned were “*ask explicitly if they are okay*”, “*Firefox add-ons*”, “*nothing*” and “*never login in to facebook*”, respectively. One could use CACTOS to mitigate some of these problems by concealing data from the service provider (reason *a*). Lastly, a couple of participants mentioned (3) being unable to hide their cover photos and to use self-censorship as a solution (“*Only post [...] appropriate cover photos [...]*”). CACTOS are tailored to address this last type of problem (reason *d*).

C. Reliance on the OSN Provider (Q1)

We asked participants “How concerned would you be if Facebook changed the privacy settings?” (S1). Most participants (83%) answered “moderately concerned”. Other responses were evenly distributed between “slightly”, “very” and “hugely concerned”. As for changes on FB’s privacy policy (“How concerned would you be if Facebook changed the privacy policy?”), participants seemed to be generally more concerned, with a shift of opinion from “moderately concerned” to “very” and “hugely concerned.” A desire to shield themselves from these threats could also function as a motivation to adopt CACTOS (reason *c*).

D. Awareness of and Attitude towards Alternative Privacy Controls (Q2, Q3)

We asked participants “Are you aware of any strategies or mechanisms, currently not provided by Facebook, that can help you better protect your privacy?” (S2, S3). All users but one reported to be unaware of any such strategies or mechanisms. Still, that one participant mentioned “DoNotTrack”, which does not fall under the category of what we consider CACTOS—rather, it is a technique to express a preference towards online trackers [30]. We asked participants “Which strategies or mechanisms do you know, even if you do not use them, to prevent unintended recipients from having access to your messages and information you send or post on Facebook?”. Most participants (81%) responded “None”. Other participants pointed out to strategies such as “*limiting the amount of people [added] as friends*”, tightening their privacy settings (including blocking people) and “*deleting facebook*”. Only one participant mentioned “*encrypting messages/posts*”. Further, we asked them “Why would you, or would you not, use such a tool?” Most of the participants that would install such a tool (30% of the total) provided a simple motivation for doing so: to increase their privacy (“*I like to increase my privacy*”), even if some provided more elaborate answers (“*it would add an extra layer of protection against marketing [sic] companies and online hackers*”). At this point in the questionnaire it was obvious that privacy played a central role in the study, thus we suspect these answers were motivated by a strong social desirability bias [18] and we deem them irrelevant for our study. A few other participants however added a shade of scepticism, e.g., “*All ready, I’m concerned with one such thing I’m using. I don’t want to involve something else and provide my data to more sources*” or “*if it could actually protect me from something I needed to be protected from*”. CACTOS are typically open-source, thus under scrutiny by

anybody. Computer scientists often rely on this property to justify that one does not need to trust the developers, as anybody can examine the code (and change it) to make sure it effectively does what it is supposed to. This in turn distributes the trust users need to place on a single developer to the whole community. However, the general Internet user may not be aware or willing to rely on this property. Current CACTOS designs do not address this issue. Those participants who responded that they would *not* install such tools (20%) also declared that the tool itself could be unreliable or leak their data (“*I would not like to broadcast my privacy settings. I would use the app only if it remains anonymous.*”, “*Not sure if it is safe to install*”) even suggesting that “*they can be unreliable unless facebook certifies the tool themselves*”. Other participants dismissed the usefulness of such tools as “*I control what I share and I trust facebook to a certain extent*” or “*too much effort for a trivial thing*”. Lastly, one participant justified that unintended recipients collect users data because “*sometimes those recipients have to use the data to improve facebook itself, or the community.*” In short, the fact that users take advantage of alternative, non-technical strategies to manage their privacy, and the mistrust in the effectiveness of alternative technical tools seem to offer little support for adoption of technical tools to address SNP-problems (Q3).

E. Attitudes and Perceptions towards Scramble! (Q4)

After the guided tour to *Scramble!*, participants were prompted to comment on their experiences with the tool. We asked participants “Can you describe, in a few words, your experience using *Scramble!*?”. They expressed a wide range of opinions that can be classified in three groups: *negative*, *positive* and *mixed* responses. This categorisation is not particularly informative, but it does reflect a heterogeneous spectrum of perceptions. Overall, responses we deemed as negative alluded to the lack of *usability*: a steep learning curve (“*the learning curve took too much time*”), how cumbersome *Scramble!* was to use —*cumbersome* was indeed a very popular word or that *Scramble!* requires “*too many steps [...] to send a message*”. Both positive and negative opinions similarly hinted at the lack of usability, framing it as the main obstacle to what actually is “*a proper tool*”, “*a valuable tool*”, “*a solution to facebook’s problem*”. The perception that *Scramble!* was useful but only for *very private information* was also recurrent, e.g., “*It’s a very great idea, but only useful for messages that really needed to be protected*”, “*awesome for the people who want to send private [i]mportant [...] text messages*”. However, this notion was also linked to poor usability, e.g., “*It was effective, but too complex to be integrated into my everyday routine. I am not THAT concerned about my private messages to go through the hassle*”. Some participants called into question the benefits of *Scramble!*, e.g., “*Easy, interesting, not sure about the benefits, though*”. Lastly, positive responses described the experience as both *easy* and *fun* (“*it was fun using it and its easy [...] too*”). In summary, we cannot conclude there was a unanimous response that reveals whether participants perceived *Scramble!* as a suitable tool or not. However, *ease of use* seems to be the central feature around which users judged their experience.

Further, we asked participants “What would be the advantages, if any, of using a tool like *Scramble!* over, or in combination with, other privacy controls?” (S4). Many

participants (70%) gave succinct answers (“*more private communication*”, “*better privacy*”) or simply repeated what they found in the material they had been given earlier (privacy issues arising from SNP-problems). However, some provided more insightful answers. A couple of participants mentioned that the advantages of using *Scramble!* were “*None*”, or that “*I wouldn’t worry so much about my Facebook messages being in the open.*” One participant wrote that “*There seems to be no advantage, as again facebook has all our public names and email ids associated with that*”. This is a fair point which reveals the *scepticism* of some OSN users (despite the fact that this particular critique reveals a flawed understanding of the protection offered by *Scramble!*). On the other hand, however, such a comment may also be interpreted as reaffirming the fact that participants did perceive protection against Facebook as a benefit (Q1).

Participants’ scepticism became patently notable when we asked them “*Scramble!* encrypts messages before you send or post them on Facebook. Do you think this is a secure way to prevent unintended recipients from having access to them?”. Even if most of them (77%) simply replied “*yes*”, other participants had some reservations, e.g., “*having a simple private and public key mechanism may not be robust enough*”, “*yes, kind of, I am sure they will find another way to decode it*”, “*To a point... unless someone can figure it out and decrypt it*”, “*that would require a second, and third level of encryption, which I think is illogical*”. Other participants did not see cryptography as the source of mistrust, but rather its particular implementation on *Scramble!* and the people behind it, e.g., “*Probably yes, but remember we don’t know whether scramble is a government controlled plug-in or actually developed by Facebook itself*”, “*no, till proper and full information about, what scramble is, how and why it encrypts our data*”. Note however that our sample consisted of CMU students —potentially with a solid background in computer science—, so these answers may not be representative of most users. Further, the study was done as revelations about NSA’s surveillance programs were hitting headlines, which may have led to greater scepticism towards cryptographic tools. Lastly, some participants’ concerns derived from a misunderstanding of how *Scramble!* works, e.g., “*what if you accidentally send a message to someone who has scramble but they were an unintended recipient can they still read your message? or do you have to add them to your contact list first?*”. In fact, it does not matter who installed *Scramble!* or who is on the contact list as long as the public keys chosen to encrypt a given message correspond to the intended recipients.

We also asked participants “What do you think are the differences, if any, between what *Scramble!* does and the privacy settings of Facebook?”. Once again, most participants (70%) referred to the CACTOS-properties mentioned by us, e.g., the fact that *Scramble!* is independent from Facebook or that it provides greater control over their privacy settings. Interestingly, some participants pointed out to mistaken technical or implementation details to frame those differences, e.g., (“*it doesnt save the information on a server rather keeps it local*”, “*Scramble data is saved on my computer so it protects me from facebook itself*”). In reality, only the private key is stored locally, all (encrypted) messages are uploaded to Facebook. The users may have also conflated OSN access to encrypted data to being equivalent to having no access to

the same data. Some participants also pointed to encryption, e.g., “Encryption at the user end is a win for this tool”. In particular, one participant mentioned that “there is no encryption in facebook”, however, note that this is not strictly true, as Facebook does implement SSL⁵, thus it is unclear to what extent this participant understood or was aware of the differences between using encryption as in SSL and as in PGP. In fact, some participants admitted not being able to understand the differences between what *Scramble!* does and what the privacy settings of Facebook do, e.g., “I don’t know if Facebook uses encryption between two individual users like Scramble does”, “im not exactly sure of how facebooks privacy stuff operates”.

In short, the inability to assess whether the protection mechanism a CACTOS relies on delivers the protection it promises may discourage users to entrust a CACTOS with the protection of their privacy.

We further asked participants to (“Overall, explain in a few words why you find a tool like *Scramble!* to be useful or not useful”). Overall, there was a slightly greater inclination to deem *Scramble!* a useful tool. Participants responded *Scramble!* was useful because it provided *better privacy*. A couple of participants conditioned its usefulness to their own purpose definition, namely, to protect *very private messages* (“it’s useful if you are really sending private messages, probably not as useful if your messages aren’t super private”, “maybe sometime i might have to talk about some secret info”). Others pointed out to encryption being fun (“I like a modicum of privacy, it is a good thing. Also encryption is fun.”), and the feeling of security afforded by *Scramble!* (“i feel more comfortable online”). Lastly, one participant provided an innovative benefit of using *Scramble!*, i.e., one we had not considered: “if you accidentally send a message to someone it was not intended for, they won’t be able to read it”. On the other hand, participants that found *Scramble!* not useful referred to three main themes: lack of usability (“it is too difficult to use”, “it isn’t easy enough to use”), excessive protection (“I don’t really seem to send any such messages which needs encryption and protection”, “I am not working in any cove[r]t operations. So there is no “very private” stuffs that I actually share in fb”, “i don’t send to[o] many messages that need to be kept too private”), and scepticism about *Scramble!* being really effective (“the internet itself is not safe, either public or private key could be captured by service providers”, “It is not efficient; people can easily hack into accounts and read messages” or that “they are likely to be cracked eventually”). We note that some responses combined elements from either of the three categories, e.g., “it’s too complicated for facebook users who just want to talk about trivial things”, or “it is just extra steps for people if they care enough”. Both responses suggest that a combination of lack of usability and an excessive protection render *Scramble!* hardly useful to them.

Addressing primarily those users who deemed *Scramble!* “not useful”, we asked participants “Which cases, purposes or people do you think a tool like *Scramble!* could be useful for?” (S4). The most popular theme across all responses was

“to send confidential information”, or “messages containing sensitive personal information” or “[...] really private information [...]”. Some participants even went further and saw *Scramble!* as a tool “for high security messages”, “[...] something top secret”, “[...] really really secret confidential information”. Incidentally, many participants mentioned the word *confidential* although it was not present in our documents or explanations. A couple of participants thought *Scramble!* would be a tool for “[i]llegal activities”, for “criminals”, but also for “Secret agents” and “intelligence”. Other participants pointed out to conscious or paranoid people (“Most conscious and concerned facebook users”, “paranoid people”). Hence, there seems to be a mismatch between the purpose and audience that developers have in mind when they develop CACTOS (i.e., access control on OSNs for the general OSN user) and those that participants in our study perceive. In fact, very few participants thought *Scramble!* would be useful for “[...] the everyday person”, “students, business people”, “[m]yself” or “everyone!”.

Lastly, we asked participants about the elements they liked, those which they disliked and those they missed. To the question “What, if anything, did you like about *Scramble!*?” participants replied *the idea* behind it, some of them being rather vague in their responses (“[t]he main idea behind it!”, “[t]he idea behind the product”) and others pointing to specific benefits such as its independence from Facebook (“[I] liked that it hides data from Facebook”, “[...] it’s nice to not be under the thumb of organizations like Facebook”) or implementation details (“[t]hat it keeps [...] local [the] key”, “the encryption”). Others valued the increased security, some participants being more convinced than others, e.g., “[g]uaranteed security” and “[t]he messages seemed to be more secure; the feeling of security was there”, respectively. Interestingly, and despite the generalised complaints about the lack of usability, some participants reported to like the *simplicity* of *Scramble!* (“It’s easy to just right click and encrypt a message”, “Simplicity and explanation of its workings”). In fact, to the question “What, if anything, did you dislike about *Scramble!*?” almost all participants responded with a complaint about the procedure required to use the tool or its interface. Participants said *Scramble!* was “cumbersome” (“[a] bit cumbersome to use”, “too cumbersome”), complex (“A bit too complex”), slow (“It is a lil [sic] bit slow”, “it was more time-consuming to send a message”), or that it required “many steps”. Regarding its interface, participants considered it was “clunky, unattractive”, or “GUI and ergonomics are real bad”. They also provided suggestions as how to solve these issues, but we discuss these in the next section. One participant mentioned that “[b]ecause messages sent to me were automatically scrambled, I wasn’t sure if the person sent a scrambled message or a normal one. I’m also not sure if anyone with scramble would be able to read a scrambled message, or I would have to have them on my contacts list first”. One participant asked “[w]hat if someone hacked into my facebook?” while another wondered “[whether it would] be possible for third parties to figure out the encryption mechanism”. Both responses suggest a lack of transparency and feedback that would enable a user to better understand how *Scramble!* works. This contrasted with the fact that many participants demanded a higher *automation* of the whole process. We elaborate on this tension between feedback and

⁵Facebook announced in July 31, 2013, that they use https by default for all Facebook users. [1]

transparency against automation in the next section. Finally, we asked participants “What features did you miss in *Scramble!* or you think that such a tool should have?”. In short, most of their demands correlated with what they did not like about the tool. However, participants who were positive about *Scramble!* pointed out to additional features such as being able to encrypt photos and also provided specific suggestions on how to improve its usability, which we discuss in the next section.

SUS: In addition to the questions above, participants were asked to fill in the standard questionnaire *SUS*. The average score for *Scramble!* fell barely above the middle score ($M = 52.9$, $SD = 18.35$, $MAX = 95$, $MIN = 15$). This supports the previously reported variety of attitudes of the participants towards the tool, its workings and its interface.

VII. DISCUSSION

Based on the study, we can be confident that our participants’ privacy preferences and attitudes are not completely divergent from the objectives and design principles that inform CACTOS. Participants reported several privacy concerns related to SNP-problems (*Q1*) that would provide them with plausible *reasons* to adopt a CACTOS. Prior to the study, we speculated that users may not adopt CACTOS because of the responsibility burden that comes with greater control over privacy settings (*Q2*). This is a complex hypothesis that needs further investigation. We observe that the complexity and obscurity of the cryptographic mechanisms CACTOS feature may in fact leave users worse off. By trying to protect their privacy with tools they do not understand, users may end up actually losing control over their information flows.

In fact, participants found that CACTOS like *Scramble!* provide too great a degree of protection, at too high a usability cost, to match their actual needs (*Q4*). We conjecture that these attitudes may be motivated by the perception people have of cryptography. The guided tour to *Scramble!* included an explanation of the cryptographic mechanisms *Scramble!* relies on, because we considered that participants needed some information in order to minimally understand the initialisation process (i.e., why they needed to generate cryptographic keys). However, had the participants not heard about cryptography and just about the CACTOS-properties provided by *Scramble!*, would their perception still be the same? To what extent their perception of cryptography biased their perception of *Scramble!*? It is not trivial to study whether the formulation of how the tool works or the participants’ perceptions of cryptography inform their reactions. This as an interesting future research question.

Moreover, another factor that influenced users’ attitudes towards *Scramble!* is the perception that what they share on OSNs is not sensitive enough to require encryption. Participants reported to control what they share; even to refrain from sharing very sensitive information on Facebook. These responses are in accordance with widespread use of self-censorship to mitigate the risk of privacy breaches reported in previous studies [16], [22], [29], and may help explaining why participants do not opt for technical tools to overcome the Facebook’s privacy settings limitations they pointed to (*Q3*). Hence, it would be a challenge to explore whether the use of CACTOS enables or motivates less repressed communication practices. Such a study may be especially difficult, as

participants may not be truthful about their self-reported self-censorship practices, and may in fact disclose the very kind of sensitive information they thought CACTOS would be useful for (*S4*).

Even if people agree with the fundamental reasons why one would use a CACTOS, would it be sensible for them to use them? Our results suggest *no*. Participants showed a great deal of scepticism towards the effectiveness of CACTOS. Some participants believed that by using CACTOS they would have to trust the tool developers instead (or on top) of the current OSN provider. For other participants, the complexity and obscurity of cryptography failed to reassure them that CACTOS deliver the protection they promise. Our participants signaled that they are unlikely to use such a tool if they doubt whether it delivers the protection promised. This is further aggravated if such a tool is seen as *a hassle*. Usability for security [35] is thus key in fostering the adoption of CACTOS.

A. Potential Recommendations to CACTOS Designers

Users need to be reassured that CACTOS are effective and deliver the protection they promise. Transparency is therefore a key element to be considered when designing CACTOS. Designers should aim at communicating better the workings of cryptography, e.g., through the use of mental models and metaphors, as was previously suggested in the literature [9], [13]. Beyond the utmost importance of a usable interface, attractive manuals could also help. In fact, participants missed in *Scramble!* “a manual? [...] the printed instructions helped the most”, and that “[the manuals we gave them] are not available [in *Scramble!*] by default.” Moreover, it may be more attractive to communicate the benefits of cryptography in terms other than those informed by security, such as *attack*, *secret* or *confidential*, that keep reinforcing the widespread idea that cryptography is only for *top-secret* use.

Moreover, all trust issues that users may have with an OSN provider, are also relevant for CACTOS developers. The study participants raised concerns with respect to having to trust yet another entity with their data. While *Scramble!* does not collect users’ posts or communication patterns, users have to rely on the *Scramble!* developer community to provide a tool that works and is well maintained. This is a challenge for CACTOS like *Scramble!* that are developed by a small community of researchers and depend on the availability of incoming funds. Further, OSN providers may enhance their trustworthiness through quality of service and branding activities. The independence and transparency of CACTOS is imminent in building trust in them, yet it may be interesting to explore how OSN providers can contribute to supporting the availability and use of CACTOS .

Participants complained about how cumbersome *Scramble!* was, the *extra steps* they needed to perform “just to send a single message” and how slow it was (“make it fast!”). As a solution to all these problems, participants expressed a desire for more automation, e.g., “there should be a function (like a switch) that turns [encryption] on throughout a conversation, so that the user doesn’t need to keep scrambling each line”, “[a tool like *Scramble!* should have] automatic encryption” or “automatic [encryption] as you type”. Automation would indeed release users from the burden of dealing with cryptography. However, more studies should further explore to what

extent better interfaces can be designed without resorting to automation, as the latter is likely to diminish users' oversight of the security mechanisms. In fact, we must be cautious about the amount of automation in a security tool, as this usually comes at the expense of better control for the user, e.g., deciding which keys are to be used to encrypt a message and securely storing private keys. Developing a tool that automatically infers the recipients of a message and manages keys is a challenging problem, if not an additional vulnerability from the point of view of privacy protection.

However, if users prefer not to get familiar with cryptographic protocols, and designers do not use proper communication strategies (e.g., mental models) that make tools intuitive and easy to manage by the general Internet user, automation is unavoidable. Future research should explore, on the one hand, how to securely automate some tasks and, on the other hand, how to better engage users in the decision-making required by CACTOS by leveraging transparency and feedback mechanisms. Users should understand the principles and basic workings of the tools they are using or they will most surely fail to use them properly. In the case of privacy and security tools, this issue is especially critical.

We finally conclude that there is still a lot of work to do in the development of CACTOS to avoid putting extra burdens on the users while mitigating their privacy concerns.

Acknowledgements

The authors would like to thank Ralf De Wolf and Laurence Claeys for their valuable comments on methodology, Filipe Beato for his technical assistance with *Scramble!* and Janel Sutkus for her valuable comments towards improving the questionnaires used in this study. This work was supported by the project IWT SBO SPION.

REFERENCES

- [1] "Secure browsing by default," <https://www.facebook.com/notes/facebook-engineering/secure-browsing-by-default/10151590414803920>, accessed: 2013-12-08.
- [2] A. Acquisti, "Privacy and security of personal information," in *Economics of Information Security*. Springer, 2004, pp. 179–186.
- [3] A. Acquisti and R. Gross, "Imagined communities: Awareness, information sharing, and privacy on the Facebook," in *Privacy Enhancing Technologies*. Springer, 2006, pp. 36–58.
- [4] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *Security & Privacy, IEEE*, vol. 3, no. 1, pp. 26–33, 2005.
- [5] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin, "Persona: an online social network with user-defined privacy," in *SIGCOMM*, vol. 39, no. 4. ACM, 2009, pp. 135–146.
- [6] F. Beato, M. Kohlweiss, and K. Wouters, "Scramble! your social network data," in *Privacy Enhancing Technologies*. Springer, 2011, pp. 211–225.
- [7] A. Besmer and H. Lipford, "Tagged Photos: Concerns, Perceptions, and Protections," in *CHI '09*, ser. Extended Abstracts on Human Factors in Computing Systems. New York, NY, USA: ACM, 2009, pp. 4585–4590. [Online]. Available: <http://doi.acm.org/10.1145/1520340.1520704>
- [8] L. Brandimarte, A. Acquisti, and G. Loewenstein, "Misplaced Confidences Privacy and the Control Paradox," *Social Psychological and Personality Science*, vol. 4, no. 3, pp. 340–347, 2013.
- [9] C. Bravo-Lillo, L. F. Cranor, J. S. Downs, and S. Komanduri, "Bridging the gap in computer security warnings: A mental model approach," *Security & Privacy, IEEE*, vol. 9, no. 2, pp. 18–26, 2011.
- [10] J. Brooke, "SUS-A quick and dirty usability scale," *Usability evaluation in industry*, vol. 189, p. 194, 1996.
- [11] S. Buchegger, D. Schiöberg, L.-H. Vu, and A. Datta, "PeerSoN: P2P social networking: early experiences and insights," in *EuroSys Workshop on SNS*. ACM, 2009, pp. 46–52.
- [12] J. Callas, L. Donnerhacke, H. Finney, and R. Thayer, "OpenPGP message format," RFC 2440, November, Tech. Rep., 1998.
- [13] L. J. Camp, "Mental models of privacy and security," *Technology and Society Magazine, IEEE*, vol. 28, no. 3, pp. 37–46, 2009.
- [14] M. Conti, A. Hasani, and B. Crispo, "Virtual private social networks," in *CODASPY*. ACM, 2011, pp. 39–50.
- [15] L. A. Cuttillo, R. Molva, and T. Strufe, "Safebook: A privacy-preserving online social network leveraging on real-life trust," *Communications Magazine, IEEE*, vol. 47, no. 12, pp. 94–101, 2009.
- [16] S. Das and A. Kramer, "Self-Censorship on Facebook," in *CSCW*, 2013, pp. 120–127.
- [17] S. Fahl, M. Harbach, T. Muders, M. Smith, and U. Sander, "Helping Johnny 2.0 to encrypt his Facebook conversations," in *SOUPS*. ACM, 2012, p. 11.
- [18] R. J. Fisher, "Social desirability bias and the validity of indirect questioning," *Journal of Consumer Research*, pp. 303–315, 1993.
- [19] S. Guha, K. Tang, and P. Francis, "NOYB: Privacy in online social networks," in *Proceedings of the first workshop on Online social networks*. ACM, 2008, pp. 49–54.
- [20] S. Gurses and C. Diaz, "Two tales of privacy in online social networks," *IEEE Security & Privacy*, vol. 11, no. 3, pp. 29–37, 2013.
- [21] W. Hartzog and F. Stutzman, "Obscurity by Design," *WASHINGTON LAW REVIEW*, vol. 88, p. 385, 2013.
- [22] M. L. Johnson, S. Egelman, and S. M. Bellovin, "Facebook and privacy: it's complicated," in *SOUPS*, 2012, p. 9.
- [23] A. Lampinen, S. Tamminen, and A. Oulasvirta, "All my people right here, right now: management of group co-presence on a social networking site," in *SIGGROUP*. ACM, 2009, pp. 281–290.
- [24] J. Lazar, J. H. Feng, and H. Hochheiser, *Research methods in human-computer interaction*. Wiley, 2010.
- [25] W. Luo, Q. Xie, and U. Hengartner, "Facecloak: An architecture for user privacy on social networking sites," in *CSE*, vol. 3. IEEE, 2009, pp. 26–33.
- [26] A. Mazzia, K. LeFevre, and E. Adar, "The pviz comprehension tool for social network privacy settings," in *SOUPS*. ACM, 2012, p. 13.
- [27] K. Raynes-Goldie, "Aliases, creeping, and wall cleaning: Understanding privacy in the age of Facebook," *First Monday*, vol. 15, no. 1, 2010. [Online]. Available: <http://firstmonday.org/ojs/index.php/fm/article/view/2775>
- [28] R. S. Sandhu and P. Samarati, "Access control: principle and practice," *Communications Magazine, IEEE*, vol. 32, no. 9, pp. 40–48, 1994.
- [29] M. Sleeper, R. Balebako, S. Das, A. L. McConahy, J. Wiese, and L. F. Cranor, "The post that wasn't: exploring self-censorship on facebook," in *CSCW*, 2013, pp. 793–802.
- [30] C. Soghoian, "Secure browsing by default," <https://www.facebook.com/notes/facebook-engineering/secure-browsing-by-default/10151590414803920>, accessed: 2013-12-08.
- [31] R. Stedman, K. Yoshida, and I. Goldberg, "A user study of off-the-record messaging," in *SOUPS*. ACM, 2008, pp. 95–104.
- [32] R. Suddaby, "From the editors: What grounded theory is not," *Academy of management journal*, vol. 49, no. 4, pp. 633–642, 2006.
- [33] A. Tootoonchian, S. Saroiu, Y. Ganjali, and A. Wolman, "Locker: better privacy for social networks," in *Co-NEXT*. ACM, 2009, pp. 169–180.
- [34] K. Vemou and M. Karyda, "A Classification of Factors Influencing Low Adoption of PETs Among SNS Users," in *Trust, Privacy, and Security in Digital Business*, ser. LNCS. Springer, 2013, vol. 8058, pp. 74–84.
- [35] A. Whitten and J. D. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0," in *USENIX*, vol. 99. McGraw-Hill, 1999.
- [36] H. Xu, T. Dinev, J. Smith, and P. Hart, "Information privacy concerns: linking individual perceptions with institutional privacy assurances," *JAIS*, vol. 12, no. 12, p. 1, 2011.