

Visible Flows: Contextual Integrity and the Design of Privacy Mechanisms on Social Network Sites

Heather Richter Lipford^{*}, Gordon Hull⁺, Celine Latulipe^{*}, Andrew Besmer^{*}, Jason Watson^{*}

^{*}Department of Software and Information Systems

⁺Department of Philosophy

University of North Carolina at Charlotte

Charlotte, NC 28223

{Heather.Lipford, clatulip, ghull, arbesmer, jwatso8}@uncc.edu

Abstract—Social Network Sites have a number of well publicized privacy issues stemming from the overdisclosure of personal information. On one hand, users seem oblivious to their privacy, doing little to protect their personal data. On the other hand, there have been a number of privacy uproars and backlashes due to certain site features or behaviors. In this paper, we explore the privacy issues in social network sites using contextual integrity, a recently proposed privacy framework. We use the framework to highlight a number of privacy issues on social network sites, and to propose a set of design guidelines and potential solutions.

Keywords—contextual integrity, social network site, privacy, Facebook.

I. INTRODUCTION

Online social network sites such as Facebook, MySpace, LinkedIn and others have experienced tremendous user growth in the past several years. Facebook alone reports over 200 million users worldwide, with 100 million of those signing on every day [2], and many smaller sites are catering to specific user populations. Additionally, a number of social network sites are creating platforms for a variety of applications that build off of users' social networks [6]. Thus, online social networking communities are becoming an increasingly important aspect of our digital lives, mediated by an increasingly complex underlying architecture.

Users of social network sites are sharing an amazing amount of personal information including contact information, political and sexual preferences, class schedules, photographs, personal associations and more [9][15]. The benefits of these sites come through building stronger personal connections with offline friends as well as through meeting and building relationships that are purely virtual. However, this proliferation of personal information leads to risks of embarrassment, stalking, spear phishing, and even identity theft [1].

In this paper we explore the privacy issues that arise on social network sites through the lens of contextual integrity, a recent philosophical framework. We explain several privacy issues that arise on social networks sites based on contextual integrity [12], and the design guidelines and solutions that this framework suggests.

II. PRIVACY AS CONTEXTUAL INTEGRITY

Helen Nissenbaum recently introduced contextual integrity as a framework for analyzing privacy with information technology [12]. Nissenbaum's initial exploration of the contextual integrity framework confines the analysis to issues of public surveillance, such as public WebCams. We believe that applying the lens of contextual integrity to social network sites also illustrates a number of privacy issues and challenges, and provides suggestions for mechanisms to potentially alleviate or address privacy problems.

Nissenbaum's account of privacy and information technology is based on what she takes to be two non-controversial facts. First, there are no areas of life not governed by norms of information flow, and these norms are in turn highly context-specific. For example, it is appropriate to tell one's doctor all about one's mother's medical history, but it is most likely inappropriate to share that information with casual passers-by. Second, people move into and out of a plurality of distinct contexts every day. Thus, as we travel from family to business to leisure, we will be traveling into and out of different norms for information sharing. As we move between spheres, we have to alter our behaviors to correspond with the norms of those spheres, and there will always be risks that information appropriately shared in one context becomes inappropriately shared into a context with different norms.

On the basis of these facts, Nissenbaum suggests that the various norms are of two fundamental types. The first, which she calls norms of "appropriateness," deal with "the type or nature of information about various individuals that, within a given context, is allowable, expected, or even demanded to be revealed". So it is appropriate to share medical information with doctors, but generally not appropriate to share religious affiliation with employers. The second set are norms of "distribution," and cover the "movement, or transfer of information from one party to another or others". It may be appropriate to share very personal details with a friend, but it is not appropriate for that friend to broadcast the same information on her radio show. In medical situations, on the other hand, a certain sharing of information – from a radiologist to a primary care physician, for example – is both normal and expected.

Nissenbaum emphasizes two general points that emerge. On the one hand, information is always tagged, as it were, with the context in which it is revealed: there is no such thing as context-free information. On the other hand, the scope of privacy norms is always internal to a context. There is no such thing as a universal privacy norm. To apply Nissenbaum's framework to social network sites, we will examine the social norms surrounding information sharing between people, how that information flows across the sites, and where the sites break or go against expected norms. We will do this for four particular aspects of social network sites – overall profiles, the introduction of the Newsfeed, photo sharing, and social applications.

III. EXAMINING PRIVACY ISSUES

A. Profile Privacy

Several researchers have identified the widespread disclosures of personal information on social network site profiles. For example, Gross and Acquisti examined the use of Facebook by undergraduates at Carnegie Mellon University. By data mining available Facebook pages, they found that 90% contained an image, presumably a photo of the user, 88% revealed their birth date, and 51% a current residence [9]. As they state: "It would appear that the population of Facebook users we have studied is, by and large, quite oblivious, unconcerned, or just pragmatic about their personal privacy. Personal data is generously provided and limiting privacy preferences are sparingly used." Results such as these have led many to speculate that users of these sites do not care about posting their information online, and do not desire any real privacy.

Yet, further research demonstrates that this is not the case [14]. Users do not wish to share everything with everyone. And while users may still not have an adequate picture of all the risks, they do not wish to share information that is too personal or that could lead others to locating or stalking them. So users do expect their profiles to be viewed by a large and public audience, and are attempting to tailor their profile information with that expectation. However, these considerations often only occur when initially filling in the profile information or after a privacy intrusion [14]. Since users are often connecting with offline friends and expect that others are doing the same, few expect that anyone other than their friends will view their profile [10]. So over time, users consider their closest friends as their main audience and overlook the wider privacy implications of their continued activities on the site. Thus, users have a shrinking "perceived audience" based on their regular and ongoing activities [14].

Relating back to contextual integrity, users find themselves acting in one context – sharing information with close friends, when they are actually in another, more public context. This can result in accidental disclosures and privacy violations. Additionally, users rarely go back and clean up their profiles or friend lists. Thus, users are likely to forget about information posted a long time ago, even though it is still available.

B. Newsfeed

The introduction of the Newsfeed on Facebook resulted in a media uproar over users' privacy. Yet, the Newsfeed did not actually change the accessibility of any information, so no information was actually being shared with anyone who could not already see it. So why the uproar? Prior to the introduction of the Newsfeed, a user's friend had to explicitly visit her profile page to see any updates, such as who she was now friends with, what she was doing, and who she was dating. The Newsfeed instead pushes stories of updates to all of a user's friends. While the actual access to information did not change, the introduction of the Newsfeed altered the visibility of information, effectively reducing privacy. Because this change was completely unexpected, this violated the norms of distribution and caused many complaints about the "stalker-ticker" [14]. Yet as users became accustomed to the new flows of information, the controversy died down. People adjusted their behavior, and now modify their profiles and status to take advantage of the fact that such activities will be broadcast to all of their friends. Thus, the Newsfeed per se is not necessarily a privacy intrusion, and is now an integral part of users' activities on Facebook.

C. Photos

Sharing photos is one of the most popular activities on many social network sites. For example, Facebook reports that 14 million photos are uploaded daily [2]! And the privacy problems have been well publicized. Photos on profiles have been used by law enforcement and employers to investigate the behavior of individuals. Students are commonly warned about the consequences of posting photos of partying or activities involving alcohol. Despite these warnings, a recent study of photo use on Facebook revealed that users expect privacy problems to occur and feel almost helpless to stop them [4].

Photo sharing on social network sites adds the complication over traditional online photo sharing in that users commonly post photos of each other, annotating and linking the images to the identities of the people in them. While this "tagging" is an efficient and convenient method of sharing a photo with the people in it, this goes against traditional offline social norms. It would generally be considered acceptable for someone to take a photo at a party, and show that photo to all of the others who were also at that party. Yet, the person would then not send that photo to all of the families and work colleagues of everyone at that party. So photos appropriately shared in one context, may not be appropriate in the context of the entire social network. And the person who uploads the photo is making that determination, not each individual user. The uploader is not likely aware or even considering whether another person's mom or boss is on the social network and appropriate to share with. Thus, the "tagging" of photos reduces the control that people have over the sharing and distribution of pictures of themselves.

Users are unable to completely prevent the disclosures of a photo posted by someone else. They are able to un-tag a photo to remove a reference to themselves, yet the photo still remains close to them as it is cross-linked to others in the photo and still accessible from other profiles. And users have little knowledge or control over the size and scope of the social network of the

other users in the photo, making it difficult to judge the extent of the photo's disclosure. People instead have to ask the photo uploader to remove any undesired photos, and hope that such requests are honored.

D. The Application Platform

A number of social network sites now have an application platform, where services written by third party developers can be added to a user's profile. These applications enhance the social experience on social network sites by allowing users to add additional content to their profiles, play games with their friends, share photos and other media, and much more. In order to complement a user's profile, most platforms allow applications to access most of the user's personal information, as well as that same information from a user's friends. While this allows applications to customize the experience for the user and incorporate information about a user's social spheres into their functionality, few need access to such a wide variety of information to do so [6].

Applications are extremely popular and widespread. Yet, research has shown that users have very little understanding of what information they are sharing, and with whom, as a result of using an application [2]. Users are notified of the information sharing when first accessing applications, and in explanations about the application platform. However, similar to other warnings and notices such as EULAs, few users are likely to thoroughly read and pay attention to those messages [7]. Instead, users build their mental models through their regular interactions with applications. Applications run within the boundary of Facebook, giving users the impression that they are interacting with Facebook and others on Facebook. This effectively obscures the fact that they are also interacting with some third party server and the unknown developers of the application. Additionally, users see that some of their information is used by an application, often their name and photo, and shared with friends who also use that application. But since few applications actually make use of the majority of the profile information, users are not aware that it is actually accessible.

We suspect that if users did fully understand this information sharing, they would find this a violation of their privacy. Indeed, we have found that users do not actually want to share so much profile information with applications [2]. There has yet been little controversy, however, because these flows of information are so obscured. With the overall profile and Newsfeed, the issues are with users sharing information with other users, and what is and is not appropriate to share. With applications, however, users are sharing information with unknown developers, making this risk difficult to explain and understand.

IV. DESIGNING FOR PRIVACY

In each of the examples we discussed, contextual integrity helps to explain the variety of privacy issues that exist due to the invisibility or difficulty in fully understanding the flows of information across the social network site. This implies that one way to improve privacy management, and help users achieve their desired privacy levels, is to make these flows of

information more visible. In this way, users can behave more closely with their desired norms of appropriateness and distribution. Thus, a number of design guidelines arises from these considerations:

- Information flows should be transparent. Users should always be able to determine what information is shared, and with whom.
- Increase the awareness of information flows during regular activities, so that the ongoing decisions users make are informed by the context of their information. This is needed to combat the "shrinking audience" phenomenon.
- Increase awareness of how much information is archived, and still available. This may influence users' current decisions about what to post, and may also influence users to remove old or outdated information.
- Make information and context concrete. Provide examples of the specific pieces of information when revealing information flows, and examples of specific people or organizations with whom it will be shared.
- Provide more control over the information flows. While many sites have some privacy settings, users are still not able to fully control the sharing of all of their information.
- Do not abruptly modify the flow of information. Give users a chance to modify their behavior before changes that could result in privacy problems.

We now highlight several specific solutions and design ideas that illustrate these guidelines.

A. General Profile

One of the privacy issues with the overall profile involves the "shrinking audience," the difficulty in maintaining awareness of the audience of one's information. This suggests that sites should find ways to increase this awareness so users can make appropriate sharing decisions based on understanding the true social context. Users rarely have feedback about who is looking at what aspects of their profiles. While providing full transparency about every page view may be overwhelming, sites could still increase awareness with various interface mechanisms.

For example, we have created AudienceView (Figure 1) a prototype privacy settings interface for social network sites [11]. Rather than modify settings with a set of menus and checkboxes, users view pages of their profiles from the point of view of different audiences, such as their different groups of friends, networks, public, etc. This interface provides a more visual and accurate mental model of what different people can view of them, allowing users to more explicitly and concretely consider the context of their information and adjust that information flow as desired.

While this interface is a step in the right direction, there are still other issues to be addressed. First, while AudienceView helps users with privacy settings, users are still not reminded of their information flows during regular activities such as viewing friends' profiles and posting information. Sites need additional mechanisms, such as visualizations of information



Figure 1. Audience View prototype interface for modifying profile privacy settings.

flows, to provide that awareness so that the context is more salient while users are making decisions about their actions. For example, perhaps a message box on the user’s home page could show the most recent information access, or could summarize the numbers of accesses in a certain time period.

Another challenge in social network sites is the flattening of a user’s social contexts into relatively few categories, such as “friend” and “not friend.” These small sets of contexts do not accurately reflect our offline and overlapping social spheres. This further impacts privacy as users may not be able to behave as desired because of this unnatural context. To allow for more of these social contexts, sites have introduced the notion of different groups of friends. Yet, creating these groupings and modifying and regulating the settings for each is still difficult and time consuming. Privacy management could also be improved by better reflecting the more nuanced and varied social contexts of offline relationships. For example, social spheres could be determined automatically from the social network graph. If such techniques are accurate, information sharing could also be determined automatically based upon aspects of the social network or the ongoing activities of individuals.

B. Photos

As stated earlier, the problem with photo sharing on social network sites is the inability of users to understand and control the disclosure of photos uploaded by others. Thus, users need greater control over those photos, so they can judge for themselves whether a particular image is appropriate to share within the context of their social network. Yet, this is challenging because a photo with multiple people in it is essentially a shared artifact. Each person may have differing opinions on the content and disclosure of that image, based upon their individual social contexts.

Research in security and privacy has explored mechanisms for such shared media to allow groups of users to determine the effective policies on such objects [1]. To be effective, these mechanisms must balance the rights of both the owners of photos and those in them, as well as adequately inform users of the resulting sharing of the photo. We have explored a simpler mechanism and interface that allows individuals to request that certain members of their social network not be allowed to view particular photos [4]. This still allows the photo uploader to share the photo with a variety of people, yet gives the others in the photo the ability to protect their image.

Rather than simply allow all users to restrict access to photos posted of them, we could also explore more graphical techniques that could blur or remove individuals or objects in an image. This would allow the images to remain available, yet still potentially protect the privacy of those in them.

C. Applications

With social applications, users are not only completely unaware of the information flows, they have very little ability to control and restrict such flows of information and still take advantage of applications. Current platforms offer an “all or nothing” approach, where if a user wants to access an application they must agree to share whatever information the site allows, whether the application needs it or not. Thus, a first step in improving privacy is to provide greater controls for users to restrict these information flows.

We have proposed giving users fine-grained control over which pieces of information they would like to share with an application [3]. The interface prototype for the controls is shown in Figure 2. The interface seeks to increase awareness of what is shared by providing concrete examples of the information the application requests, and choosing a random friend to show that their information will also be accessed. Our prototype then allows users to restrict information that does not



By proceeding, you are allowing Circle of Friends to access your information and you are agreeing to the Facebook Platform Terms of Service in your use of Circle of Friends. By using Circle of Friends, you also agree to the iLike Terms of Service.

Figure 2. Prototype interface for controlling information disclosures to applications.

make sense to share within the context of the application. In studying a simulation of this interface, our results indicated that approximately half of users would protect their information and only share what was needed by an application. Yet, the other half of users still ignored the message and controls and proceeded as before. Thus, such an interface does increase awareness and control for some users, but does not do so for others.

There are a number of other possibilities for giving users greater awareness that have not yet been investigated. One challenge is how to communicate that the application is not only sharing information with other users, but with the developers of that application, outside of Facebook. One possibility is to represent this developer more explicitly in the installation message for the application. Another is to alert the users whenever an application queries for a piece of their information. Another is to send messages to all of a user's friends when that user installs an application, letting them know that their information is now accessible and how to protect themselves. However, any of these potential alerts need to be done so that the user is not overwhelmed with too many messages. Users also need mechanisms to remind them of these information flows during their everyday activities, not just at application installation, so that again the context can regularly inform the user's behaviors.

V. CONCLUSION

Users could (and perhaps should) operate under the assumption that any information they put online is public and potentially available to anyone forever. Yet, this would clearly impact the social benefits of social network sites, as users would be forced to share less information than they would like, or simply accept all the privacy problems that result. Sites do not imply that all information is public, and users are not behaving in such a way. Thus, privacy management on social network sites fails to reflect the nuanced and contextual nature of privacy in the offline social world.

Contextual integrity provides a framework for understanding many of the privacy issues on social network sites today. Users must be able to understand the context within

which they are behaving in order to make reasonable decision about what information is appropriate to share. This means that the flows of information that the site allows and supports must be visible, and users must be aware of them during their regular activities. Without this visibility and awareness, users are likely to make decisions that go against their expected social norms, resulting in privacy problems. Contextual integrity also implies that users should be given more control over these flows of information so they can interact and share as they desire, with fewer risks of privacy problems. Our goal is to design improved privacy interfaces and mechanisms that can help users understand the privacy implications of their online behaviors and provide greater control over their information sharing. We seek to reduce the privacy risks and problems that occur on social network sites, while still allowing users to fully enjoy the social benefits that come from interacting with others and sharing personal information.

REFERENCES

- [1] Acquisti, A. and Gross, R. (2006). Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In the Proceedings of the 6th Workshop on Privacy Enhancing Technologies (Cambridge, UK, June, 2006), PET2006, 1-22.
- [2] Besmer, A., Lipford, H.R. (2009). Users' (Mis)Conceptions of Social Applications. Unpublished manuscript.
- [3] Besmer, A., Lipford, H.R., Shehab, M., and Cheek, G. (2009). Social Applications: Exploring a More Secure Framework." In Proceedings of the Symposium on Usable Privacy and Security, July 2009.
- [4] Besmer, A. and Lipford, H.R. (2009b) Tagged Photos: Concerns, Perceptions, and Protections. Extended Abstracts of CHI 2009, Work-In-Progress, April 2009.
- [5] Facebook.com (2009). Press Room, <http://www.facebook.com/press/info.php?statistics>. Retrieved June 11, 2009.
- [6] Felt, A., and Evans, D. (2008). Privacy Protection for Social Networking APIs. In the Proceedings of the Workshop on Web 2.0 Security & Privacy, May 2008.
- [7] Good, N.S., J. Grossklags, D. K. Mulligan, and J. A. Konstan. (2007). Noticing notice: a large-scale experiment on the timing of software license agreements. In Proceedings of the SIGCHI conference on Human factors in computing systems, pages 607-616, April 2007.
- [8] Google.com (2009). OpenSocial. <http://www.opensocial.org/>. Retrieved June 11, 2009.
- [9] Gross, R., and Acquisti, A. (2005). Information revelation and privacy in online social networks. In Proceedings of the 2005 ACM Workshop on Privacy in the Electronic, 71-80, November 2005.
- [10] Lampe, C., Ellison, N., and Steinfield, C. (2006). A face(book) in the crowd: social Searching vs. social browsing. In Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative, 167-170, November 2006.
- [11] Lipford, H.R., Besmer, A., and Watson, J. (2008). Understanding Privacy Settings in Facebook with an Audience View. In the Proceedings of the USENIX Conference on Usability, Psychology, and Security, April 2008.
- [12] Nissenbaum, H. (2004). Privacy as Contextual Integrity. *Washington Law Review* 79 (2004), 101-39.
- [13] Squicciarini, A.C., Shehab, M., and Paci, F., Collective Privacy Management in Social Networks, WWW 2009: 18th International World Wide Web Conference, April, 2009, Madrid, Spain.
- [14] Strater, K. and Lipford H.R. (2008). Strategies and Struggles with Privacy in an Online Social Networking Community. In the Proceedings of BCS HCI 2008, September 2008.
- [15] Stutzman, F. (2005). An evaluation of identity-sharing behavior in social network communities. In the Proceedings of iDMAa and IMS Code Conference, 2005.