Collaborative Privacy Management for Third-Party Applications in Online Social Networks

Pauline Anthonysamy, Awais Rashid, James Walkerdine, Phil Greenwood
School of Computing and Communications, Lancaster University, Infolab21, Lancaster LA1 4WA, United Kingdom.
{anthonys, marash, walkerdi, greenwop}@comp.lancs.ac.uk

Georgios Larkou
Department of Computer Science, Networks Research Laboratory, University of Cyprus, 1678 Nicosia, Cyprus
glarko01@cs.ucy.ac.cy

ABSTRACT

Privacy control mechanisms for online social networks (OSNs) offer little by way of managing access to a user's personal information by third-party applications (TPAs). Most OSNs provide an "accept all or nothing" mechanism for managing permissions from TPAs to access a user's private data. In this paper, we propose an approach that makes all requests for private data from TPAs explicit and enables a user to exert fine-grained access control over what profile data can be accessed by individual applications. Equally importantly, our approach also allows users to share their access control configurations for TPAs with their friends who can reuse and rate such configurations. This is particularly beneficial to novice users or those new to a particular TPA or an OSN. We present an implementation of our approach for managing privacy for third-party Facebook applications and report an initial evaluation (N=50). A significant proportion of our sample (76%) found the collaborative privacy management approach useful in determining the type of applications one might use based on its privacy rankings and noted a raised awareness about data privacy issues arising from use of TPAs.

Categories and Subject Descriptors

H.3.4 [**Information Storage and Retrieval**]: Online Information Services—*Web-based Services*; K.4.1 [**Computers and Society**]: Public Policy Issues—*Privacy*

Keywords

Privacy Management, OSNs, Third-Party Applications, Collaborative Filtering, Recommender Systems.

1. INTRODUCTION

Online social networks (OSNs) are now the norm in modern society. According to [9], in 2008, more than 30% of Internet users were members of at least one OSN. Recent figures show that, at the end of 2011, Facebook alone had 845 million users [1]. As the number and size of OSNs grows, there is increasing concern about the protection of the terabytes of personal data entrusted to OSNs and how the privacy of this data is managed. Studies such as [6, 13] have shown that OSN members generally know little about how OSNs function in regards to transfer (or use) of their personal information. Similarly, Majeski et al. [10] found that privacy controls are inherently difficult for users to understand and configure. However, most of these studies and other privacy management approaches focus on how OSNs themselves utilise the data provided by the user and what control can be exercised by a user on such actions by the OSN. Little or no attention has been paid to how the privacy of users' data as shared with third-party applications (TPAs) is managed. This is not to be confused with recent examples such as the discovery of token leakages from Facebook to TPAs that enabled them to collect users' profile information [2] – though this was a major privacy breach it was an accidental leakage. In contrast, our focus is on empowering users to not only have fine-grained control over their personal data shared by an OSN with TPAs but also enabling users to share such third-party privacy management configurations with other users. Currently, most OSNs offer an *"accept all or nothing"* mechanism for managing permissions from TPAs to access a user's private data. In other words, even if a user is explicitly informed that a TPA would access certain pieces of information, s/he has no control over sharing only a subset of that information – the only alternative being not installing and using the application.

In this paper, we focus on addressing the issue of fine-grained privacy management for TPAs in OSNs and how to engender trust in specific privacy management configurations for novice users or those users who are new to a particular application. The novel characteristics of our approach are as follows:

1. It makes all requests for private data from TPAs explicit and enables a user to exert fine-grained control over what profile data can be accessed by them.
2. Users can share their access control configurations for TPAs with their friends who can reuse and rate such configurations.

This fine-grained control and collaborative management of "best configurations" empowers users to manage what data

PSOSM '12, April 17 2012, Lyon, France

is shared by an OSN with TPAs and which privacy configurations would enable them to utilise the functionality they need from a TPA without sharing all the information requested by the application.

We have implemented our approach as a prototype framework for managing privacy settings for TPAs on Facebook. Our initial evaluation with users (N=50) found that a significant proportion (76%) found the collaborative privacy management approach useful in determining the type of applications one might use based on its privacy rankings and noted a raised awareness about data privacy issues arising from use of TPAs.

The rest of this paper is structured as follows. Section 2 provides an overview of our approach and presents details on our prototype implementation for Facebook. Section 3 reports on results of initial evaluation of our prototype. Section 4 discusses related work in the area while Section 5 concludes the paper and identifies directions for future work.

2. COLLABORATIVE PRIVACY MANAGEMENT (CPM) FRAMEWORK

2.1 Overview of the CPM Framework

An overview of our collaborative privacy management framework (CPM) is shown in Fig. 1. In essence, the framework provides an interceptor mechanism that acts as a *membrane* between the TPAs and an OSN. All information requests pass through this membrane and are intercepted. At the time of installation, the framework makes explicit all the personal data items that an application will access. The user now has the choice to not allow certain permissions or where the permissions are mandatory choose to return dummy data. An example of the latter is the scenario where an application requires access to the user's name to personalise certain elements of the application but also requests access to the user's hobbies or friends list without a clear need for it. The user can now choose to return empty sets or even dummy data so as to safeguard his/her privacy[1]. The user can then also choose to share this new privacy configuration with others in the OSN. The user can also change the privacy configuration for the application in due course (e.g., to share more or less data as desired).

Alternatively, when a user installs a TPA, instead of manually configuring the privacy settings, s/he can choose to search for existing privacy configurations for the application shared by other users in the OSN and how they have been ranked by others utilising them. The user can then load a chosen configuration (either by popularity or based on the kind of information s/he might be willing to allow access) and either use it as is or make modifications before deploying it for that particular application. The user can then also share the new configuration with others in the OSN.

2.2 CPM Prototype for Facebook

Our prototype implementation of the CPM Framework is realised as a Facebook platform application. Facebook applications interact directly with Facebook Servers through their Graph API. The Graph API allows application developers to access all public information about a user i.e. user name and profile picture. Access to a user's private data is permitted (upon acceptance of permissions request by the user) with an access token. However, these permissions request are grouped in a coarse-grained manner. For instance, the application requests access to one's "Basic Information", which includes name, profile picture, networks and friends list. Here, the user's options are either to accept or cancel this access request. S/he is not allowed to selectively grant or deny access to individual profile information i.e. a user is willing to share her/his name but not her/his friend's information. Our CPM implementation for Facebook mitigates this existing coarse-grained access control by allowing one to define fine-grained access control (or permission) on individual profile information and share these configurations with other Facebook users.

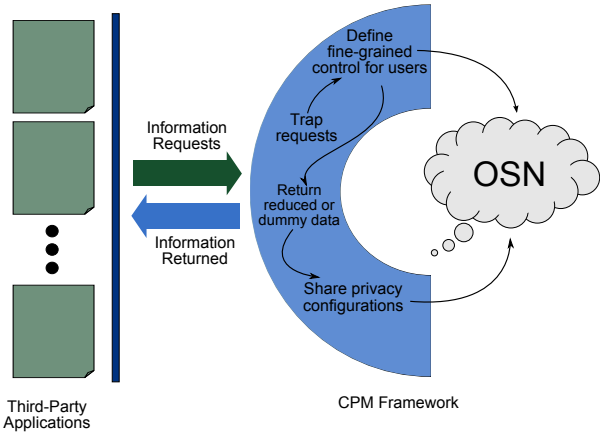

Figure 1: Overview of the Collaborative Privacy Management (CPM) Approach.

2.2.1 *Interceptor Implementation*

The framework is presented exactly as any other application inside an IFrame but in reality it sits between TPAs and Facebook servers. Fig. 2 illustrates the interception mechanism which operates as follows.

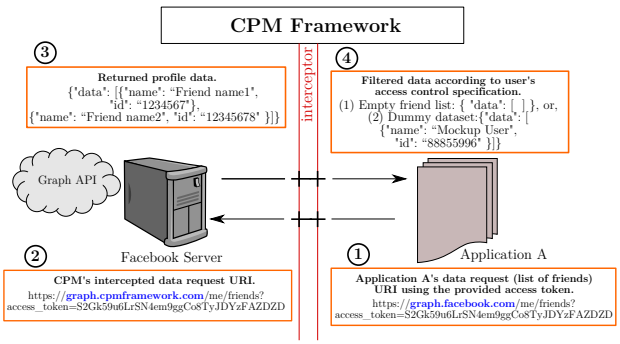

Figure 2: CPM Interceptor

To prevent applications from interacting directly with Facebook's Graph API, the CPM Interceptor exports an identical API through which it intercepts all outgoing application data requests, noting the access token of each one (Step 1 in Fig. 2). Using this access token CPM extracts the application ID and the user information (userID) from whom the request was initiated. The Interceptor then forwards each such request to Facebook's Graph API using the same access token (Step 2) and retrieves the correspond-

[1]We recognise that returning false information has other potential consequences in terms of criminal behaviour – these are issues we are tackling in other projects such as Isis: http://www.comp.lancs.ac.uk/isis

ing data items (Step 3). Note that the mechanics employed by the Interceptor are completely transparent in the sense that existing third-party applications only need to replace the original Graph API URL (http://graph.facebook.com) with CPM's (http://graph.cpmframework.com). Having retrieved the data, CPM Framework evaluates and filters it according to the user's (previously specified) access control rules before returning this filtered data to the original TPA (Step 4).

3. EVALUATION

The evaluation of our Facebook CPM implementation was divided into two phases: *Phase 1* was designed to establish the extent to which privacy leakage is a concern amongst TPA users; *Phase 2* involved evaluating the effectiveness of our CPM Framework in addressing these concerns and discovering whether the participants' attitude towards TPA privacy changed after using our framework. In both phases 50 participants were recruited from a pool of OSN users with the youngest participant being aged 14 and the oldest aged 28. In the following subsections we examine some of the results in detail. Note that the results presented are not exhaustive but rather highlight selected interesting observations.

3.1 Phase 1

In this phase a general questionnaire was developed to understand users' privacy attitudes towards TPAs and their perception on data leakages through those applications. The questionnaire was implemented using Survey Monkey [3] and comprises mostly closed questions, using multiple choices.

Fig. 3 illustrates how participants answered the questions about their awareness of the information collected by third-parties and whether they always accept the default permissions requested by an application. What is immediately obvious from Fig. 3(a) is that a significant number of users (45%) lack awareness of the information that is collected by TPAs. On the other hand 65% of the participants answered 'No' on accepting the permissions requested by these applications. This illustrates that participants are concerned about the exposure of their profile data.

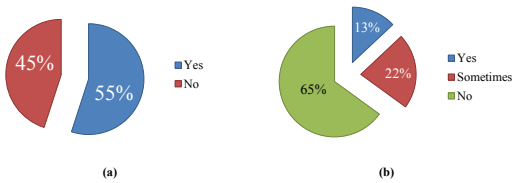

Figure 3: (a) "Are you aware of the information that is collected by third-party applications when you connect to them?". (b) "Do you always accept the permissions requested by third-party applications?."

3.2 Phase 2

In this phase 50 Facebook users utilised our CPM Framework for managing privacy in TPAs. The participants were divided into two groups (N=25 each). We developed and deployed two third-party Facebook applications. Each application had two versions, a standard version and a CPM-enhanced version.

The first group was instructed to install and use the standard application versions first and the CPM-enhanced versions second. The second group, by contrast, was instructed to use the CPM-enhanced versions first and the standard versions second. This methodology was intended to remove any potential bias arising from the order in which the standard and enhanced versions of our applications were used.

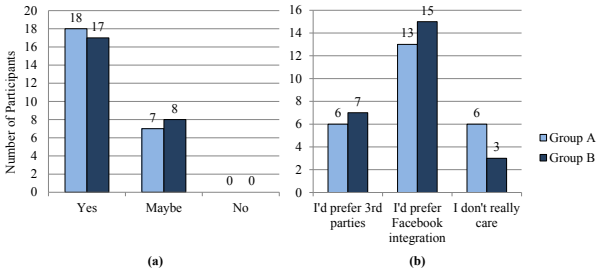

Figure 4: (a) "Do you find CPM Framework's fine-grained access control provides a better mechanism in managing your profile data than the existing functionality on Facebook?." (b) "Do you trust third-party applications to manage your profile data or would you prefer the CPM framework to be integrated into Facebook?".

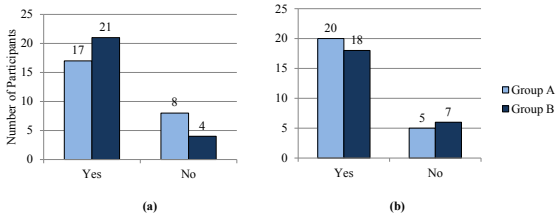

Figure 5: (a) "Do you find CPM Framework's permission sharing mechanism useful?." (b) "Do you find sending dummy data a good approach to continue using third-party applications without compromising your privacy?".

Our results (in Fig. 4(a)) show that the majority of participants from both groups A and B (70%) preferred the fine-grained access control mechanism provided by our framework. This clearly highlights the fact that data owners (i.e. users) prefer to impose specific access control on their profile information shared with TPAs. The results from Fig. 4(b) further support this finding by illustrating the demand for an alternative access control mechanism for TPAs as opposed to Facebook's existing coarse-grained mechanism.

Our approach of allowing users to share their access control configurations for TPAs with their friends who can reuse and rate such configurations also achieved significant positive response (76%)(ref. Fig. 5(a)). This demonstrates the effectiveness of the approach in engendering trust through collaborative creation of privacy configurations for TPAs.

Most participants also found sending dummy data items instead of revealing their profile data to TPAs encouraging (shown in Fig. 5(b)) as this enabled them to safe-guard their personal information and indirectly increases their trust in using these applications.

4. RELATED WORK

A number of studies have highlighted privacy issues in OSNs. Luders et al.'s [6] study of primarily Norwegian users on Facebook has shown that users' knowledge on how social

media functions in regards to use, disclosure and transfer of their personal information is largely inadequate. Majeski et al. [10] found that, in their study, every one of the participants had at least one sharing violation based on their stated sharing intentions. Our recent exploratory study [4] of four OSNs examined has also highlighted the general disconnect between privacy policies and privacy controls. All the above studies highlight the need to both empower users to manage their privacy in OSNs and to provide mechanisms that enable novice users or those new to an application (or even OSN) to navigate the plethora of privacy settings. The CPM framework proposed in this paper responds to both these challenges by offering fine-grained access control over personal data shared by OSNs with TPAs and enabling the OSN users to jointly identify "best fit" privacy configurations for such applications.

PoX [7] is a client-side browser plug-in that acts as a proxy between TPAs and Facebook. In order for users to benefit from PoX, a TPA must use the PoX library to channel all requests to Facebook. Similarly, xBook [12] introduces a secure hosting platform; developers deploy and host their TPAs on this platform which intercepts not only interactions between the application and the OSN but also with other web sites. Both these works contrast our implementation which simply intercepts all the calls to the Facebook Graph API. Most significantly, they do not support sharing and reuse of privacy configurations for TPAs.

MockDroid [5] is a modified version of the Android Operating System, which enables the user to create access templates for each application or just accept/deny access to a resource at run-time. This is similar to the fine-grained privacy control offered by the CPM framework. However, unlike CPM, sharing and collaborative management of privacy configurations is not supported.

Recommender systems, in general, and Collaborative Filtering (CF) techniques, in particular, utilise the preferences of users to derive associations between them and in turn make recommendations based on similar interests or disinterests (see, for example, [8, 11]). The approach closest to ours is that of Walkerdine and Rodden [14] which allows users to create and share generic queries with each other, allowing them to re-use and rate them in a specific query-sharing environment. In contrast, our focus is on sharing and reuse of privacy configurations and, significantly, doing so in the open landscape of third-party applications supported by large OSNs.

5. CONCLUSIONS AND FUTURE WORK

Privacy management in OSNs is an area that is of increasing importance owing to their large user base and the amount of data stored in such OSNs. It is not just sufficient to provide effective privacy control mechanisms for OSNs but, equally or perhaps more importantly, engender trust in such mechanisms. Our proposed CPM framework addresses both these needs by not only enabling users to have fine-grained control over their data shared with TPAs but also utilising the social construct of "friends" to identify "best-of" configurations that can be trusted by other users. Our work in the future will focus on further evaluating the effectiveness of our framework through larger user studies both on Facebook and (through additional prototypes) on other large OSNs. We also aim to develop techniques/mechanisms to expose the "privacy context" to users as they undertake certain actions on OSNs or TPAs to further improve their awareness about the visibility of their personal data as a result of specific actions on their part.

6. ACKNOWLEDGEMENTS

This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) cross-disciplinary account: *UDesignIt: Social Media, Social Good - Ultra-Large Scale Public Engagement Systems to Challenge Anti-Social Behaviour*. Pauline Anthonysamy is supported by a Lancaster University 40th Anniversary Scholarship.

7. REFERENCES

[1] http://techcrunch.com/2012/02/01/facebooks-s-1-845-million-users.

[2] http://www.symantec.com/connect/blogs/facebook-applications-accidentally-leaking-access-third-parties.

[3] Survey monkey: http://www.surveymonkey.com.

[4] P. Anthonysamy, A. Rashid, and P. Greenwood. Do the privacy policies reflect the privacy controls on social networks? *IEEE International Conference on Privacy, Security, Risk and Trust, 2011 IEEE International Conference on*, 2011.

[5] A. R. Beresford, A. Rice, and N. S. Sohan. MockDroid: trading privacy for application functionality on smartphones. In *Proceedings of HotMobile 2011*, Mar. 2011.

[6] P. B. Brandtzaeg and M. Lüders. Privacy 2.0: Personal and consumer protection in new media reality. Tech. Rep. SINTEF A12979, Nov'09.

[7] M. Egele, A. Moser, C. Kruegel, and E. Kirda. Pox: Protecting users from malicious facebook applications. In *Pervasive Computing and Communications Workshops, 2011 IEEE International Conference on*, pages 288 –294, march 2011.

[8] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Commun. ACM*, 35:61–70, December 1992.

[9] S. Gurses, R. Rizk, and O. Gunther. Privacy design in online social networks: Learning from privacy breaches and community feedback. In *ICIS 2008 Proceedings*, New York, USA, 2008. ACM.

[10] M. Majeski, M. Johnson, and S. M. Bellovin. The failure of online social network privacy settings. Technical Report CUCS-010-11, Feb. 2011.

[11] U. Shardanand and P. Maes. Social information filtering: algorithms for automating "word of mouth". In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '95, pages 210–217, New York, NY, USA, 1995.

[12] K. Singh, S. Bhola, and W. Lee. xbook: redesigning privacy control in social networking platforms. In *Proceedings of the 18th conference on USENIX security symposium*, SSYM'09, pages 249–266, Berkeley, CA, USA, 2009. USENIX Association.

[13] R. Singh, M. Sumeeth, and J. Miller. A user-centric evaluation of the readability of privacy policies in popular web sites. *Information Systems Frontiers*, pages 1–14, 2010.

[14] J. Walkerdine and T. Rodden. haring searches: Developing open support for collaborative searching. In *Proceedings of Interact.*, Tokyo, Japan, 2001.