

Respect My Privacy: Establishing Social Norms for Privacy in Social Networks

Ted Kang and Lalana Kagal

MIT CSAIL, Cambridge
Massachusetts, USA

`tkang@mit.edu`, `lkagal@csail.mit.edu`

Abstract. Most social networks have implemented extensive and complex privacy controls in order to battle the host of privacy concerns that initially plagued their online communities. These privacy controls have taken the form of access restriction, which allow users to construct barriers preventing unwanted users from viewing their personal information. However, this system leaves users unprotected in cases in which the access restriction mechanisms are bypassed or when the access restrictions are met but the data is later misused. Our framework, Respect My Privacy, offers a different approach to privacy protection and is founded on the philosophies of Information Accountability. The strategy taken by Respect My Privacy is similar to how legal and social rules work in our societies. The vast majority of these rules that form the fabric of our societies are not enforced perfectly or automatically, yet most of us follow the majority of the rules because social systems built up over thousands of years encourage us to do so and often make compliance easier than violation. Our project aims to support similar functionality in social networks by allowing users to clearly declare the policies that govern the use of their private data, implementing mechanisms that make people who use this data aware of the policies, and notifying the user of misuse of this data. Instead of focusing on preventing data from being accessed, we focus on helping users conform to existing policies by making them aware of the privacy policies associated with the data. We have defined a simple and easily extensible vocabulary for defining restrictions on the use or misuse of data. In its current state, the framework has been implemented across three platforms: Facebook, OpenSocial, and the Tabulator Semantic Web browser. We have Facebook and OpenSocial applications that enable users to specify their privacy preferences. These preferences then appear on their profiles alerting viewers to the fact that a privacy policy exists. We also have FOAF applications that export privacy annotated profiles as Friend Of A Friend profiles in RDF. In Tabulator we've developed a social pane that allows social data as well as RMP settings to be modified and a sidebar that uses colors as subtle warnings that the explored data can or cannot be used for certain purposes.

1 Introduction

From their inception, social networks have suffered from a host of privacy issues. When the first social networks were gaining in popularity, privacy mechanisms

were sparse with most profiles being publicly available to all members. As social networks like Facebook and MySpace exploded in popularity, however, many users were shocked to find that the information that they had posted on their profiles were coming back to have unintended consequences in their real life: employers were reported to be using Facebook as a way to vet possible employees; universities were using Facebook pictures to identify people that had attended illegal functions; and children were being preyed on by registered sexual offenders on MySpace.

In response to these highly publicized privacy concerns, social networks responded by implementing complex privacy controls that allowed users to construct barriers preventing unwanted users from looking at private information. This method of privacy protection, called access control, seeks to close off and hide information from those that are not explicitly given access to it. It is a binary system in which those that obtain access to the data, legitimately or not, have full reign over the use of that data while those without access cannot view anything.

These access restriction systems, while useful in blocking out unwanted viewers, are ineffective for a large, decentralized system like the World Wide Web. It is simply too easy to copy or aggregate information, and it is often possible to infer "private" information without actually having explicit access to the information itself. In addition, there are always human factors that a technical access restriction system will have trouble handling. For example, Facebook's access restriction systems did not prevent users from unwillingly publicizing their purchases when Facebook introduced Beacon, a controversial advertising program. Only a mass protest from users caused Facebook to readjust their privacy controls [3]. Given all these ways for data to escape from behind access restriction systems, it is troubling that access restriction mechanisms are powerless in cases where information is compromised. An apt analogy from *Information Accountability* [1] compares sole reliance on access restriction to "focusing all one's attention on closing the barn door and ignoring what might happen to the horses after they've escaped."

The Respect My Privacy (RMP) framework hopes to offer an alternative approach to protecting privacy in social networks. The Respect My Privacy framework is grounded in Information Accountability, which argues that in addition to access control, there need to be ways of ensuring that people know exactly what they can and cannot do with personal or sensitive information. In addition, Information Accountability also argues for mechanisms that notify users of misuse of their data after the fact.

This approach is similar to the system in place for legal and social rules in society. In society, a set of legal or social norms govern what we can or cannot do, and they have been ingrained into our way of thinking such that most people go through life without any problems. When problems occur, there are mechanisms that ensure that those who broke the set of legal or social norms are reprimanded. Likewise, social networks would benefit from having mechanisms in place that allow users to declare the 'social norms' that they expect to be respected by

those that use their private information. Once these 'social norms' are adequately declared, Information Accountability calls for further mechanisms that will assist users in detecting misuse of their information.

Currently, the mechanisms that assist users in detecting misuse of their information are based on societal pressures. In the future, there will need to be technical mechanisms to assist users in detecting when their private information is misused.

The current implementation of the RMP framework consists of four main parts.

The first part of the framework is the RMP ontology. This ontology represents the five restrictions that are currently offered in RMP: *no-commercial*, *no-depiction*, *no-employment*, *no-financial*, and *no-medical*. Users can apply these restrictions over their social network profiles to declare that they do not want their personal information used in certain ways. The ontology allows users to apply the restrictions over their Friend of a Friend (FOAF) documents.

The second part of the framework is the RMP applications in Facebook and OpenSocial that allow users to create RMP restrictions and display them on their profile pages. The aim of the RMP applications on the mainstream social networks is to gain a wider acceptance for the RMP restrictions with which users can declare their personal information.

The third part of the framework is the FOAF converter. Initially developed by Matthew Rowe, the FOAF converter takes Facebook profiles and converts all the information to RDF in the format of FOAF files [9]. This FOAF converter was altered to include the RMP restrictions that a user has declared for his Facebook/OpenSocial profiles. This will hopefully increase the number of FOAF users and offer initial privacy protections for members of a decentralized FOAF network.

The final part of the framework is on the Tabulator Semantic Web Browser[11]. The Tabulator extension acts as a framework upon which a decentralized social network can be built[12]. The RMP pane in the Tabulator extends the previously developed social pane, which allows users to view FOAF files in a format similar to a social network profile page. The RMP adds to the Tabulator by allowing users to declare or modify restrictions using the RMP ontology over their FOAF files. In addition, a privacy aware pane highlights restricted data with different colors to notify users of those restrictions and makes it easier to responsibly browse Semantic social data.

2 Design

The RMP framework hopes to enhance privacy protections in social networks by establishing a set of 'social norms' that the users will hopefully adopt and embrace. This requires offering users a simple but clear way to declare how they want their private information handled. In addition, it requires that enough members are knowledgeable about this vocabulary such that they can easily realize the limitations on using a certain piece of data.

To achieve this goal of widespread adoption, we aimed to create a simple and recognizable brand. To create this brand, we looked to the Creative Commons project that has been successful in creating a brand that is now fairly well known across the Internet. Creative Commons is a nonprofit corporation founded by Lawrence Lessig that is "dedicated to making it easier for people to share and build upon the work of other, consistent with the rules of copyright." [2] One of the main reasons that Creative Commons has reached its level of widespread acceptance is its ease of use and easily recognizable icons.

To receive a free Creative Commons license, users merely have to log onto the Creative Commons website and choose one of their six licenses. The user chooses whether they want to allow commercial uses of their work and if they want to allow modification of their work. Once the user makes two choices, the user is given a small, simple icon that notifies users about the type of license that is placed on piece of work. The icon is a simple rectangle with the common "CC" logon on the left and smaller intuitive icons that represent the type of restrictions on the license. By clicking on the icon, users are redirected to a page that gives additional details about the license.

Following the Creative Commons model, the RMP application aims to offer an easy user experience that allows users quickly declare the restrictions they wish to place on their data. In addition, the RMP application needs to establish a brand by offering a simple, recognizable icon that links to a page that offers additional information. As with the Creative Commons licenses, we hope that as more and more people view the RMP icons attached to profile pages, it will gain widespread traction and acceptance. This also means that we need to have the first implementations of the RMP application in well established social networks.

In the future, we would also like to provide technical mechanisms that could reason about how a piece of data is used and automatically notify users when it finds that personal information has been misused. This implies that a certain amount of control is needed over the actual personal information itself as it would need to be machine readable.

This goal is more difficult to accomplish as all large social networks are currently centralized. Facebook, MySpace, and the other major social networks "own" the personal information that is placed on their sites and are very protective of it [6]. Without more control over one's personal information, it is difficult to apply technical mechanisms that could detect misuse. For instance, Information Accountability argues for attaching semantic data with provenance, or a trace, such that reasoners would be able to trace a data's history of usage and ensure that data has been appropriately used at each step. This attaching of provenance data is impossible in a centralized social network unless social networks like Facebook and MySpace become actively involved in the project. Until that occurs, another aim of the RMP framework is to use Semantic Web technologies to represent the restrictions.

Thus, our main design goals are two fold. The first concentrates on creating a simple user experience in already established social networks in order to gain acceptance of the RMP restrictions. The second concentrates on also applying the

RMP framework to Semantic Web technologies such that as the RMP framework begins to gain traction, we can offer more extensive accountability mechanisms that take advantage of Semantic Web technologies.

3 Architecture

The current architecture of the RMP framework consists of two distinct parts. The first is the Respect My Privacy applications on the mainstream social networks: Facebook and OpenSocial. These applications are there to introduce the RMP restrictions and attempt to spread some familiarity with the restrictions. The second is in the Tabulator extension. The RMP framework on the Tabulator extension attempts to plant the foundations for a decentralized social networks consisting of FOAF files that can be navigated through the Tabulator extension. The Tabulator extension also offers the door for future work on accountability systems as users have greater control and flexibility over their personal information.

Both of these parts depend on the set of restrictions that we have developed, which allow users to declare certain uses that are forbidden with their data. They are represented in RDF in the ontology and can currently be attached over FOAF documents.

Also connecting these two parts is the FOAF converter, originally developed by Matthew Rowe, that is a part of the RMP applications on Facebook. The FOAF converter takes the personal information stored in Facebook and creates a FOAF file that is stored on a Web server. This, in effect, becomes a profile in a decentralized social network if navigated through the Tabulator Extension. The hope is that this pane will encourage the adoption of decentralized social networks and show the benefits of privacy protections based on Information Accountability offer.

3.1 RMP Restrictions and Ontology

In keeping with the Creative Commons model, RMP offers predetermined restrictions for RMP users. Offering predetermined restrictions will ensure a better user experience and allow the creation of simple and recognizable icons that will assist in the widespread adoption of RMP. There are currently five restrictions that are implemented on RMP: *no-commercial*, *no-depiction*, *no-employment*, *no-financial*, and *no-medical*. Users may choose to apply none or any combination of the five restrictions on their social network profiles and related pages. The lack of any of these restrictions on a profile page implies that the use is allowed. For example, not including the no-financial restriction would imply that you are willing to allow your personal information to be used for financial purposes.

This set of restrictions is currently aimed primarily at protecting individuals from organizations as organizations have a vested interest in protecting privacy interests. This can be seen by the current trend of appointing CPOs, or Chief Privacy Officers, who ensure that companies are aware of and protect privacy



Fig. 1. The pictures for the five restrictions: no-commercial, no-depiction, no-employment, no-financial, and no-medical.

interests. There is a market place for privacy that is being created and surveys report that consumers prefer organizations that respect privacy interests to those that do not [8]. The hope is that the competition between organizations to protect privacy will make them respect privacy declarations, fearing bad public relations, more so than individuals would.

no-commercial: The no-commercial restriction is similar to its counterpart in the Creative Commons. At the time when the RMP restrictions were being developed, there was no way to apply Creative Commons restrictions on the content that one posted to a social network. This restriction states that the user does not want anything on his profile or related pages to be used for a commercial purpose.

no-depiction: The no-depiction restriction allows a user to declare that he does not want his pictures taken and used for any reason and he they does not want his private information used to identify him in an image. This restriction was meant to specifically protect the pictures that users often post on social network sites. These posted pictures have been the most troubling with universities using student photos as evidence for infractions and employers using Facebook pictures to prove that employees were not doing what they claimed to be doing.

no-employment: The no-employment restriction declares that the user does not want any personal information used for the purposes of any kind of employment decision. For example, this would make companies aware that the user does not want them using their social network page as a way to vet them for a job. In addition, this would imply that an employer could not use personal information from a social network as justification for a firing. This restriction was meant again as a response to common incidents of users not being hired or being fired from a job owing to something they posted on a social network.

no-financial: The no-financial restriction declares that the user does not want any personal information used for any financial purposes. For example, the user would not want banks using personal information from the social network to influence a loan or credit decision, or have any influence in divorce proceedings.

no-medical: The no-medical restriction declares that the user does not want any personal information used for any medical purposes. For example, the user

would not want hospitals or insurance providers using personal information from the social network to research into her lifestyle habits or more.

Each of these restrictions have a corresponding picture as seen in Figure 1. These pictures are combined to create a simple icon that can be placed on social network pages and link to additional information. A sample of these icons are shown in Figure 2. When clicked, these icons link to a page that offers additional information about each of the restrictions that has been applied.



Fig. 2. RMP icons that can placed on a social network profile page to indicate a certain privacy policy. They link to a page containing additional information about the policy.

These restrictions are represented in RDF in the RMP ontology. The ontology can currently be used to apply the RMP on FOAF documents. In Figure 3, we can see a sample FOAF file where the user has selected to apply the no-commercial and no-medical restrictions on his FOAF file.

▼ http://dig.csail.mit.edu/2008/02/rmp/tkang-foaf.n3			
Ted Kang	Error Reports To	mailto:leigh@ldodds.com	
	Generator Agent	http://www.ldodds.com/foaf/foaf-a-matic	
	type	PersonalProfileDocument	
	maker	Ted Kang	
Ted Kang	primary topic	Ted Kang	
	restricts	No-Commercial	
	type	No-Medical	
	family_name	Person	
Ted Kang	Given name	Kang	
	homepage	Ted	
	knows	http://web.mit.edu/tkang/www	
		http://csail.mit.edu/~lkagal/foaf.rdf#me	
Ted Kang	seeAlso	Person	
	sha1sum of a personal mailbox URI name	http://web.mit.edu/shauni/www/foaf.rdf	
	name	832846fe2f1f5b578a0b1353b39a7a17dc8de328	
	type	I. Shauni Deshmukh	
Ted Kang	seeAlso	Person	
	sha1sum of a personal mailbox URI name	http://web.mit.edu/shirleyf/Public/foaf.rdf	
	name	9d53ea39f02565a81a9e53dd9b51fb1046342e74	
	type	Shirley S. Fung	
Ted Kang	knows	Person	
	type	Person	
	name	Shirley S. Fung	
	type	Person	

Fig. 3. A sample FOAF files with a user that has applied the no-commercial and no-medical restrictions on his FOAF file.

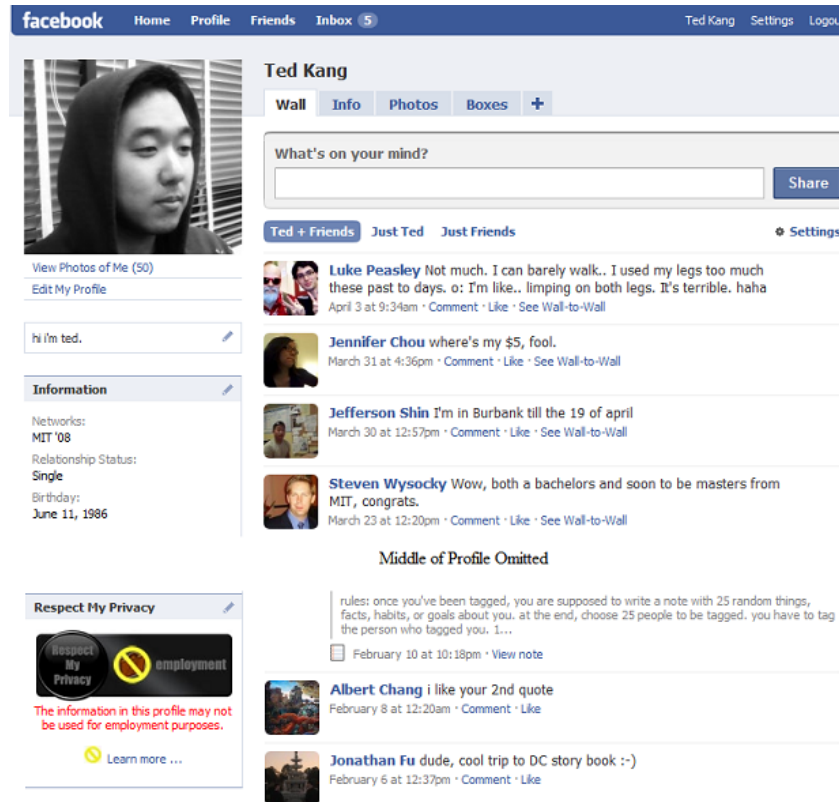


Fig. 4. A user's Facebook profile with the RMP icon in the lower left. Any Facebook user that clicks on that icon is directed to a page that offers additional information on the applied restrictions.

3.2 RMP on Facebook and OpenSocial

The RMP application on Facebook is a MySQL/PHP driven web application that uses the Facebook Application API and is hosted on the Decentralized Information Group's server. In order to mimic the ease of use for Creative Commons, the creation of a RMP setting is simple, taking mere minutes. When a user decides to add the RMP application, they are directed to a page that explains the philosophy behind RMP. This page is very important as RMP on Facebook is currently a project entirely dependent on its members. As more users create the RMP restrictions and expect their restrictions to be respected, organizations will feel more pressure to actually respect those restrictions. Thus, the introductory text attempts to instill the idea that the user is part of a movement that will improve everyone's social network experience the more the user respects others restrictions. The user is then directed to a page that lists the five restrictions with descriptions of each. Each restriction has an accompanying checkbox,

which allows the user to decide whether they want to apply that restriction or not. Once they have chosen the restrictions, they are done. The restrictions are saved into the MySQL database and the appropriate icon is pushed to the profile page so that everyone that visits a user's profile page can clearly see the restrictions that the user has placed on his personal information.

Once a user has created a set of RMP restrictions, there are several features that become available to them. First, the user's RMP icon is pushed onto their profile page along with some informative text in the following context: "The information on this profile may not be used for ... purposes." Now anyone that visits the user's profile page will be able to view the RMP icon. A sample Facebook page with the RMP icon in the lower-left is shown in Figure 4.

If any visitor clicks on the RMP icon they will be directed to a page that lists the restrictions that the user has decided to apply and a paragraph giving more information on the restrictions chosen. The users are then invited to join the RMP movement on Facebook by creating their own set of licenses.

The RMP application on OpenSocial is driven primarily by Javascript. OpenSocial's API and data storage is fairly different from that of Facebook, relying heavily on Javascript and not an actual database, but the OpenSocial application was designed to be exactly like its counterpart on Facebook.

3.3 The FOAF Converter

The RMP application on Facebook allows users to port their Facebook profiles to Friend of a Friend (FOAF) files. This acts as a bridge between the RMP applications in Facebook and the corresponding extensions on the Tabulator. Users of Facebook automatically become members of decentralized social networks by choosing to create a FOAF file. They can choose to download the FOAF file and host it on their own or have it hosted on our group's Web server.

This will hopefully introduce the members of mainstream social networks to the idea of decentralized social networks. With further work on the Tabulator, users might be able to see the advantages of having complete control of their data especially as methods of attaching provenance and more sophisticated accountability mechanisms are developed.

3.4 RMP pane and sidebar for Tabulator

The Tabulator is a generic Semantic Web data browser and editor for RDF data, similar to how a web browser is used to navigate HTML pages. The Tabulator is currently implemented as a Firefox extension. When a user installs the extension and uses Firefox to go to a URI that contains RDF triples, the Tabulator offers an easy interface with which to view the RDF data and allows users to easily explore triple relationships to obtain more data. The Tabulator recently implemented a social pane that becomes available when users browse upon FOAF data. The social pane displays the FOAF data in a format similar to social network profiles, allowing the typical social network experience in a decentralized setting.

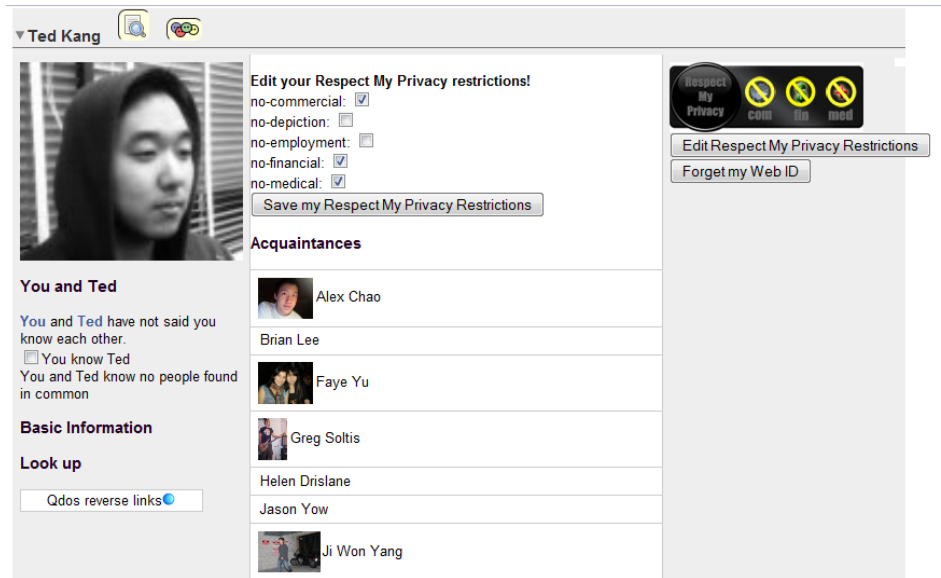


Fig. 5. The social pane on the Tabulator offers a social network profile like view of FOAF information and also allows users to create and edit RMP restrictions.

The RMP pane attempts to make it easier for users to attach RMP restrictions, using the aforementioned ontology, to their FOAF files. Users with editable FOAF files can use the Tabulator to identify a FOAF files as their identity. Once an identity has been established, users can host editable FOAF files on a Web-DAV servers and use SPARQL updates to create or edit the restrictions they place over their FOAF profiles. The social pane also displays the RMP icons if restrictions are detected, similar to the RMP applications on the mainstream networks. An example social pane that would occur when browsing upon a FOAF file with the Tabulator is shown in Figure 5.

In addition, a policy aware sidebar makes it easier to recognized protected data while browsing semantic information on the Tabulator. The sidebar detects instances of RMP restrictions or Creative Commons licenses as data is browsed and allows users to choose a color for each restriction or license. When users do so, any data that is protected by the restriction or license will be highlighted in that color, allowing users to instantly recognize the information that is protected under certain policies. An example of the the policy aware sidebar with the highlighting functionality is shown in Figure 6.

4 Related Work

The Creative Commons released a Facebook application on May 18th, 2009 [7]. Users are able to choose from the six Creative Commons licenses and apply them over their entire profiles. Users are given the recognizable Creative Commons

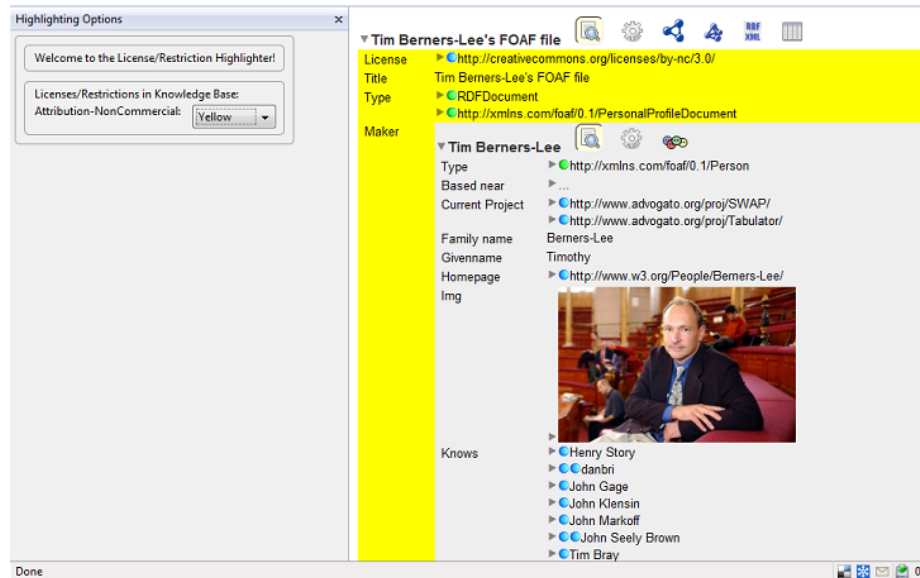


Fig. 6. The privacy aware sidebar detects protected data, like a Creative Commons license, and highlights data that is protected by different policies

icon and are able to publish it on their profiles linking to a page with more information. This application highlights the need for users to be able to declare how they want their data used. This application is definitely a step in the right direction as it does not solely rely on access control to prevent unwanted users from viewing data but declares to all viewing users certain restrictions on how they want their licensed information used. The RMP will hopefully offer more protections as the restrictions are aimed primarily at social network users while the Creative Commons licenses are focused on creative works.

Another example of users being more proactive of how their personal information is handled is in the discussion of Google AdSense ads. Google AdSense includes a notion of policy-awareness by putting a hyperlink "Ads by Google" on all its advertisement [5]. When clicked, the user gets general information about why these ads were displayed and is able to slightly modify how further targeting is performed. Turow proposes an approach in which each ad will have an icon that when clicked displays exactly what information was used in order to choose that ad ¹. These approaches are related to our framework in that they attempt to make policy explicit but they focus on the use of search/clickstream data for targeted advertisements versus the use of personal data available in social networks.

¹ <http://www.asc.upenn.edu/ascfaculty/FacultyBio.aspx?id=128>

5 Future Work

The RMP project still has some areas that require significant work. First, the applications in the framework expect privacy annotations to be associated with a foaf:Person, and are unable to handle more finely grained annotations. We would like to use N3Logic [13] to allow any RDF sub-graphs to be annotated with RMP privacy restrictions and support N3Logic in our Facebook, Open Social applications as well as in the Tabulator pane.

As social networks move from being centralized hosted applications to more decentralized applications [12], the role of RMP becomes more important. In these decentralized networks, accessing, copying, and reusing data inappropriately becomes even easier. Using RMP will enable users to specify their privacy requirements and encourage third party application developers to develop tools such as the RMP Tabulator sidebar that clearly identifies the restrictions of social data.

Another area of future work is in building accountable systems that will be able to use the provenance attached to data to pinpoint cases of misuse. Systems that are able to model and implement policies have been created and demonstrated in the case of data mining, but there has been no clear conclusion on how to detect misuse in the context of social networks. One problem in this area is the difficulty in pinpointing "use" in a social network. In many contexts, such as an employer doing a background check on a potential employee, merely looking at unflattering data on the prospective employee's social network profile might be enough to eliminate him from getting a job. A possible solution is to allow users to create multiple FOAF files on the decentralized social network, and have servers hosting the FOAF files to query visitors for their intent in viewing the profile. Based on the visitor's intent, the server can display one of the user's more appropriate profiles. This strategy is similar to those taken by niche social networks, like LinkedIn, that are made specifically for a certain purpose, such as networking. Nonetheless, determining what defines a "use" in the context of social network information is a problem that requires future work.

6 Conclusion

The efforts of social networks in protecting privacy could be greatly improved with the adoption of information accountability techniques. One of the tenets of accountable systems is transparency in policy, and the Respect My Privacy project currently offers users a clear and simple way to define the restrictions they want to place on their data. In addition, the current RMP implementations employ the model used by the Creative Commons to try to gain widespread acceptance of the defined restrictions. This will hopefully lead to legal mechanisms to protect users from malicious misuse of their personal information and encourage adoption of a set of social norms online, which depend on people responsibly helping each other protect their privacy. Finally the Respect My Privacy project's new direction in decentralized social networks offers a foundation upon which to build a fully accountable system for social networks.

Acknowledgements

The authors wish to thank their group members for their contributions to this project. This work was carried out with funding from NSF Cybertrust Grant award number 04281 and IARPA award number FA8750-07-2-0031.

References

1. D. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, and G. Sussman *Information Accountability*. MIT CSAIL Technical Report.
2. "About - Creative Commons" <http://creativecommons.org/about/>
3. Louise Story and Brad Stone. "Facebook Retreats on Online Tracking" <http://www.nytimes.com/2007/11/30/technology/30face.html>
4. "Reciprocal Privacy (ReP) for the Social Web" <http://dig.csail.mit.edu/2007/12/rep.html> 12 Dec 2007
5. Saul Hansell. "An Icon That Says They're Watching You" <http://bits.blogs.nytimes.com/2009/03/19/an-icon-that-says-theyre-watching-you/>
6. Erick Schonfeld "Zuckerberg on Who Owns User Data on Facebook: It's Complicated" <http://www.techcrunch.com/2009/02/16/zuckerberg-on-who-owns-user-data-on-facebook-its-complicated>
7. Frederic Lardinois. "Creative Commons Releases Facebook App: Choose a License for Your Photos, Videos, and Status Updates - ReadWriteWeb" http://www.readriteweb.com/archives/creativecommons_releases_facebook_app.php
8. Harris Interactive Inc. *Privacy Notices Research Final Results. Privacy Leadership Initiative (PLI)*. 2001
9. Rowe, Matthew "FOAF Generator" <http://www.dcs.shef.ac.uk/~mrowe/foafgenerator.html>
10. Berners-Lee, Tim. "Notation 3 (N3) A readable RDF Syntax." <http://www.w3.org/DesignIssues/Notation3.html>
11. "The Tabulator Extension". <http://dig.csail.mit.edu/2007/tab/>
12. Ching-man Au Yeung, Ilaria Liccardi, Kanghao Lu, Oshani Senevirante, Tim Berners-Lee. *Decentralization: The Future of Online Social Networking* MSNWN Position Paper. <http://dig.csail.mit.edu/2008/Papers/MSNWS/>
13. Tim Berners-Lee, Dan Connolly, Lalana Kagal, Yosi Scharf, and Jim Hendler". *N3Logic: A Logical Framework For the World Wide Web*. Journal of Theory and Practice of Logic Programming, 2007.