

Collective Privacy Management in Social Networks

Anna C. Squicciarini
College of Information
Sciences & Technology
Pennsylvania State University
University Park, PA, USA
acs20@psu.edu

Mohamed Shehab
Department of Software &
Information Systems
University of North Carolina
Charlotte, NC, USA
mshehab@uncc.edu

Federica Paci
Department of Computer
Science
Purdue University
West Lafayette, IN, USA
paci@cs.purdue.edu

ABSTRACT

Social Networking is one of the major technological phenomena of the Web 2.0, with hundreds of millions of people participating. Social networks enable a form of self expression for users, and help them to socialize and share content with other users. In spite of the fact that content sharing represents one of the prominent features of existing Social Network sites, Social Networks yet do not support any mechanism for collaborative management of privacy settings for shared content. In this paper, we model the problem of collaborative enforcement of privacy policies on shared data by using game theory. In particular, we propose a solution that offers automated ways to share images based on an extended notion of content ownership. Building upon the Clarke-Tax mechanism, we describe a simple mechanism that promotes truthfulness, and that rewards users who promote co-ownership. We integrate our design with inference techniques that free the users from the burden of manually selecting privacy preferences for each picture.

To the best of our knowledge this is the first time such a protection mechanism for Social Networking has been proposed. In the paper, we also show a proof-of-concept application, which we implemented in the context of Facebook, one of today's most popular social networks. We show that supporting these type of solutions is not also feasible, but can be implemented through a minimal increase in overhead to end-users.

Categories and Subject Descriptors

K.4.1 [Computers and Society]: Public Policy Issues;
K.6.4. [Management of Computing and Information Systems]: System Management

General Terms

Security

Keywords

social networks, privacy, collaboration, game theory.

1. INTRODUCTION

Social networks (SNs, for short), including Friendster.com, Tagged.com, Xanga.com, LiveJournal, MySpace, Facebook,

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2009, April 20–24, 2009, Madrid, Spain.
ACM 978-1-60558-487-4/09/04.

and LinkedIn have developed on the Internet over the past several years. SNs have been successful in attracting users. According to ComScore Media Metrix, more users visit MySpace than Yahoo, MSN or Electronic Arts gaming site [24]. SNs provide a form of self expression and help users to socialize and interact with other users. Users can define a personal profile and customize it as they wish. Through SNs, users may engage with each other for various purposes, including business, entertainment, and knowledge sharing. The commercial success of SNs depends on the number of users it attracts, and encouraging users to add more users to their networks and to share data with other users in the SN. End users are however often not aware of the size or nature of the audience accessing their data and the sense of intimacy created by being among digital friends often leads to disclosures that may not be appropriate in a public forum. Such open availability of data exposes SN users to a number of security and privacy risks [22, 34, 19].

A significant privacy threat is raised by an increasing amount of media content posted by users' in their profile. User-provided digital images are an integral and exceedingly popular part of profiles on SNs. For example, Facebook hosts 10 billion user photos (as of 14 October 2008), serving over 15 million photo images per day [2]. Pictures are tied to individual profiles and often either explicitly (through tagged labeled boxes on images) or implicitly (through recurrence) identify the profile holder [1]. Such pictures are made available for other SN users, who can view, add comments and, by using content annotation techniques, can add hyperlinks to indicate the users who appear in the pictures.

In current SNs, when uploading a picture a user is not required to ask for permissions of other users appearing in the photo, even if they are explicitly identified through tags or other metadata. Although most social networking and photo sharing websites provide mechanisms and default configurations for data sharing control, they are usually simplistic and coarse-grained. Pictures, or in the more general case, data, are usually controlled and managed by single users who are not the actual or sole stakeholders, raising serious privacy concerns. Data stakeholders may be unaware of the fact that their data (or data that is related to them) is being managed by others. Even when the stakeholders are aware of the fact that their data is posted and controlled by other users, they have limited control over it and cannot influence the privacy settings applied to this data. The privacy breach due to poor or no access control of shared data in Web 2.0 is well documented in the public news media [34].

Letting one user taking full responsibility over another's

privacy settings is extremely ineffective. Even if two users know each other, their social relationship often does not imply that they have the same privacy preferences. The average number of friends of Myspace users is 115 friends, which indicates that the friend relationship is being stretched to cover a wide range of intimacy level [21]. Consequently, users who share content may have different privacy preferences, and as a consequence their privacy preferences on some data content they share, may be conflicting. Based on such considerations, in this paper we focus on how to enable collaborative privacy management of users' shared content.

We believe this is an important contribution in the realm of Web 2.0, since to date, current SNs support privacy decisions as individual processes, even though collaboration and sharing represent the main building blocks of Web 2.0.

Designing a suitable approach to address this problem raises a number of important issues. First, co-ownership in SN platforms should be supported. Second, the approach should promote fairness among users and be lightweight. Moreover, the approach should be practical and promote co-ownership, since users knowingly do not enjoy spending time in protecting their privacy [35].

We analyze these requirements from a game theoretical perspective [26], and model the process of collaborative privacy management of shared data as a mechanism design problem. We map the user collaborative policy specification to an auction based on the Clarke-Tax [7, 8] mechanism which selects the privacy policy that will maximize the *social utility* by encouraging truthfulness among the co-owners.

The Clarke-Tax mechanism is appealing for several reasons. First, it is well suited to our domain, in that it proposes a simple voting scheme, where users express their opinions about a common good (i.e., the shared data item). Second, the Clarke-Tax has proven to have important desirable properties: it is not manipulable by individuals, it promotes truthfulness among users [11], and finally it is simple. Under the Clarke-Tax, users are required to indicate their privacy preference, along with their perceived importance of the expressed preference. Simplicity is a fundamental requirement in the design of solutions for this type of problems, where users most likely have limited knowledge on how to protect their privacy through more sophisticated approaches. We integrate our design with inference techniques that exploit folksonomies, automating collective decisions, thus freeing the users from the burden of manually selecting privacy preferences for each picture. As part of our assessment, we implement a proof of concept application, in the context of Facebook, one of today's most popular social networks and show that supporting these type of solutions is not also feasible, but can be implemented through a minimal increase in overhead to end-users.

The rest of the paper is organized as follows. In the next section we provide an abstract representation of SNs. Then in Section 3, we discuss data co-ownership in SNs. In Section 4, we highlight the requirements for the design of an effective solution supporting collaborative privacy management. In Section 5, we describe our proposed framework which is based on the Clarke Tax mechanism. We present our applied approach, detailed system implementation and experimental results in the context of Facebook in Section 6. We discuss the related work in Section 7 and conclude the paper in Section 8.

2. REPRESENTATION OF SNs

In this section we provide an abstract representation of a SN. Our intent is not to represent any concrete system, but to identify the key elements of a SN, upon which to build our solution. A SN is characterized by the following core components:

- U . The set of users. The community composing a SN is represented as a collection of users. Each $i \in U$ is uniquely identified.
- RT . The set of relationship types supported by the SN. Users in a SN are possibly connected among each other by relationships of different types.
- Ψ . It denotes the functional assignment of a relationship among a couple of users. Specifically: $\Psi : RT \rightarrow U \times U \cup \emptyset$. Given a pair of users i, j we denote their relationship as $i \text{ } Rt \text{ } j$, where Rt is a relationship name of one of the supported RT . The same pair of users can be related by different type of Rt . We assume all the relationships in general to be binary, non-transitive and not-hierarchically structured. Unary relationships are also enabled, such as for example $i \text{ } is_fan_of \text{ } U2$, although not relevant for us.
- $Profile_i$. The profile of a user i . We represent it as a tuple $Profile_i = (GRelType_1, \dots, GRelType_k, Set)$ where $GRelType_l$ represents the list of users having a relationship Rt_l such that $i' \text{ } Rt_l \text{ } i$ where $Rt_l \in RT$. S represents the data set posted on i 's profile. We denote the profile components of a user i by means of the dot notation. For example, i 's friends are represented as $Profile_i.Friends$ while the data set S as $Profile_i.Set$.
- D . The set of data types supported. Supported content types are multimedia -video and music files - images, documents and hypertext.

Users in SNs are connected among each other by means of direct or indirect relationships. Direct relationships hold when two users $\langle i \text{ } Rt \text{ } j \rangle$ are tied with each other according to a relationship Rt supported by the SN. Two users $1, k$ are indirectly related if there exists a path connecting them of the form: $\langle (1 \text{ } Rt_1 \text{ } 2), (2 \text{ } Rt_2 \text{ } 3), \dots, (k-1 \text{ } Rt_{k-1} \text{ } k) \rangle$, where each tuple $\langle i \text{ } Rt_l \text{ } j \rangle$ denotes an existing relationship of type Rt_l between users i and j . Provided that there may be multiple paths connecting two given users, the *users' distance* between i and j is the path with the minimal number of users between them. In the rest of the paper we always refer to the minimal path, unless stated otherwise.

EXAMPLE 1. Consider users Alice, Bob and John who are part of a social network. Alice and Bob are friends while Bob and John are colleagues. The distance between Alice and John with respect to the relationships *Friend_Of* and *Colleague_Of* is 2 because their minimal connecting path is the social path $\langle (Alice \text{ } Friend_Of \text{ } Bob), (Bob \text{ } Colleague_Of \text{ } John) \rangle$.

2.1 Expressing privacy policies in SNs

In our reference model, each user $i \in U$ enforces locally specified privacy policies over their data posted in $Profile_i$. Such privacy policies are simple statements specifying for each locally owned data item who has access to it, and, in certain cases, which kind of operations can be performed on

the data. In current SN sites, users have little flexibility when specifying such privacy policies (also referred to as access rules or privacy settings), and can choose among a limited set of predefined options, such as: friends, friends of friends etc. Additionally, access rights in a SN are limited to few basic privileges, such as read, write and play for media content.

Here, in order to provide a model that is as general as possible, we assume that users are able to specify *Distance Based* access conditions in their privacy policies. That is, the users allowed to access the data, are identified by means of the notion of users’ distance, discussed in previous section. We omit specifying the type of access privilege, as it is not significant in our case, and assume generic *viewer* rights for users who can access another’s profile. A privacy policy is summarized by the predicate $PrP(i, n)_{RtSet}$, which indicates all the users who are connected with i with a minimum path of length n , by relationships in $RtSet$ ¹. In case i leaves the data public to the whole SN, the predicate will be of the form $PrP(i, \infty)$, while in case accessibility is restricted to owner(s) only, the predicate will be set as $PrP(i, 0)$. We say that a user j **satisfies** a distance based condition $PrP(i, n)_{RtSet}$ if the minimal length of the path between i and j is within n hops according to the relationships listed in $RtSet$.

EXAMPLE 2. *Suppose Alice wants her friends of friends to be able to view her pictures. She will enforce a policy of the type $PrP(i, 2)_{Friend_Of}$. Bob, in Example 1, satisfies the policy, while John does not, since John and Alice are indirectly connected by means of a *Colleague_Of* relationship.*

3. DATA CO-OWNERSHIP IN SNS

In this section we introduce the notion of collaborative data sharing in SNS. We present the notion of co-ownership in SNS and discuss how to detect co-ownership of data in a semi-automated manner².

In SNS, users post data on their profiles, this data is usually considered owned by the profile owner. The profile owner is also expected to take the responsibility of managing the access of the posted data content. However, data posted on a user’s profile often conveys content not belonging only to the profile’s owner. For example, documents can be co-authored and belong to multiple individuals. Several users may appear in a same picture, and the same applies to other media content, such as videos. However, if Alice posts a document in her profile which belongs also to Bob, she is in charge of setting the privacy policy for the document, regardless of whether Bob is happy with her policy or not.

These simple observations naturally lead to the idea of supporting co-ownership (or stakeholders) in SN, to indicate the set of users who are owners of a piece of data, regardless of *where* (i.e., in which users’ profile) this data has been originally posted.

In order to identify co-owners of a given piece of data s , we provide a general classification of users based on their

¹Distance-based access control rules are employed both in real-world SN, where for example, one can indicate the visibility of friends of friends, and in recent access control models proposed for SN sites [5].

²Note that ownership in our discussion is not defined in terms of legislation, but in terms of the information and its relationship with users.

relationship with s . From now on, we focus our presentation on photo images or pictures, although the main idea behind our solution is general enough to be applied for other data types. Users can be classified as *viewers*, *originators* and *owners*. Users who are authorized to access the data s are defined as *viewers*. The *originator* is the user who originally posted data s on a given profile. Finally, the *owners* are the users who share ownership privileges with the originator within the social network and maintain control over s .

The potential owners of a data item posted on a profile are identified using tagging features supported by current SNS. In general, tagging consists of annotating social content by means of set of freely chosen words [38], associated with the data denoted as $TSet$. Their semantic can be analyzed by means of similarity tools [28]. In the case of pictures, we employ a specific type of tags widely used in Facebook [13]. These tags, known as *id-tags*, give the ability for users to add labels over pictures to indicate which users appear in them. Therefore, each id-tag essentially corresponds to the unique user id. By leveraging id-tags, one can easily identify the potential owners in a given picture. We formally define potential owners as follows.

DEFINITION 1 (POTENTIAL OWNERS). *Let s be a shared data item posted on user’s i profile $Profile_i$. Let $TSet$ be the set of tags associated with s . The set of the potential owners of s , $Pot_Own_s^i$ is defined as the set of users whose id-tags are in $TSet$.*

For data types other than pictures, the set of potential owners can be identified by using the meta-data associated with the content, or by the originator’s initiative. A user j belonging to the set of potential owners is qualified as an owner if the originator i agrees to grant ownership for a piece of data s to user j . Ownership privileges are exclusively granted by the originator to ensure that ownership is managed with users who in fact are not complete strangers, but related only by a number of relationships that the originator believes acceptable. This network of admitted owners can be automatically specified by the originator using distance based policy conditions, which indicate the type of relationships and the distance among the users. That is, the originator i can decide to grant the ownership of s to some user j only if j has a certain distance $PrP(i, j)_{RtSet}$ within k hops with respect to a certain set of social relationships $RtSet$. In order to mitigate the risk of originators not sharing ownership with entitled users, in Section 5 we propose an incentive-based mechanism to motivate sharing of ownership rights. The definition of data owners is very intuitive and we thus omit its formalization. In our context a set of owners, denoted as $Own_{U, Set}^s, USet \subseteq U$, do not only decide whether to post/edit/delete s , but more importantly they share the responsibility of managing access of s , by specifying the data *privacy settings* (or privacy policies).

EXAMPLE 3. *Consider Alice, Bob and John who are part of *FactBook* social network. Alice and Bob are friends while Bob and John are colleagues. Alice has participated to a Christmas party organized for the employees of the company where Bob and John are employed. Alice has taken pictures with Bob in which also John appears and posts them on her *FactBook* profile. John requests to Alice to become an owner of the pictures in which he appears. Alice has decided to give the ownership of the pictures contained in the album of the Christmas party to all the users x such*

that $PrP_{\{Friends_Of, Colleague_Of\}}(Alice, 2)$. Since Alice and John have a degree of separation equal at most to two, John is granted the ownership.

In the next section we investigate how collaborative management of data with multiple owners can be achieved.

4. COLLABORATIVE SHARING REQUIREMENTS

In case of single-user ownership, enforcing of privacy policy for a piece of data s is straightforward. The user sets his/her privacy policy according to his/her privacy preference. The privacy policy states who can view the user's data, by indicating the distance and the type of relationships viewers' should have with the owner. On the other hand, a shared data object s has multiple owners where each owner might have a different and possibly contrasting privacy preference. Designing an approach which combines different owners' privacy preferences into a unique privacy policy is a challenging task. In particular, it is unclear how to compose the overall privacy preferences for s without violating individuals' preferences. Furthermore, if multiple owners share more than a single data item, the decisions made in past interactions may be factored.

Several intuitive approaches are not suitable, due to the specific constraints of the SN domain, and the data for which the privacy policy is to be specified. For example, selective disclosure is not desirable, and often not feasible. If the data in question is a picture, cropping or blurring it would result in an altered picture, likely decreasing its intrinsic value to users and owners. Similarly, if a document is co-authored, it is not always possible to separate the different contributions of the authors and disclose portions of it without making it unintelligible. Note that, cryptographic techniques may theoretically solve the problem of selective data disclosure to entitled viewers. However, these approaches will not compose a unique privacy policy that incorporates the preferences to the different co-owners, and will result in a very unpractical approach, with a very large number of encryption keys for users to manage.

A database-like approach, where different owners could enforce their local "views" would not work either, as this approach may result in privacy violations. For example, Alice may require only friends to view a party picture, while Bob may not care and leave the picture public to any SN member. Clearly, as John -who is not a friend of Alice - logs into the social network and accesses the picture through Bob's profile he violates Alice's privacy preference, although the picture is itself not available for Bob to view from Alice's profile.

Based on these considerations, we identified the following core requirements for collaborative privacy management:

- **Content Integrity:** The data s should not be altered, or selectively disclosed. In other words, we cannot assume to blur a picture or crop it to remove certain subjects appearing in it. Nor can we alter a document text or data to satisfy conflicting individuals' preferences.
- **Semi-automated:** The access policy construction process should not solely rely on user's manual input for each data, but should leverage users' past decisions and draw from the existing context.

- **Adaptive:** When a new owner is added for s , his/her input should be taken into account, even if the access policy for s has been already set up.
- **Group-Preference:** The algorithm must leverage the individuals' information to develop a collective policy.

In the next section, we propose an approach that satisfies the above requirements. Building upon mechanism design literature, we suggest a mechanism that collects users' privacy preferences and assigns a unique privacy policy that aggregates the users' individuals' input.

5. ALGORITHMS FOR COLLECTIVE PRIVACY DECISIONS

The most intuitive approach to aggregate users' decisions is to let co-owners iteratively disclose their preferred settings and explicitly agree on the set of possible viewers' each owner proposes to include. Owners could update their preferences as they view other owners' preferred settings, and try to reach a common decision on a single policy after a few rounds of revision of their internal settings. This approach however is hardly applicable in that it requires all the owners to agree on a single set of privacy policies, which may sometimes be an endless task. Since SN users typically access the network independently it is also hard to force synchronization, without introducing unacceptably long decision processes. A more conservative solution is to construct a privacy policy that allows viewers' rights only to the set of users who satisfy each of the owners' preferences, avoiding the need of the owners explicit consent on the final set of viewers'. However, even this approach is pretty simplistic and fails to leverage the individuals' preferences within the co-owners' group. In addition to the identified drawbacks, in general majority and ranking-based approaches such as the ones described above, have proved to be unfair, in that astute individuals may manipulate outcomes at their advantage.

We suggest an approach that is characterized by two main parts:

1. First, we present an algorithm that promotes certain desirable behaviors, such as granting ownership when conditions for co-ownership hold and truthfulness of co-owners when expressing their privacy preferences. Specifically, we suggest an application of the Clarke-Tax [8] mechanism to enforce collective privacy decisions.
2. Second, to avoid users having to input the same privacy settings multiple times for similar data, we suggest a simple inference technique to leverage users' previous privacy decisions, when certain similarity conditions hold true.

5.1 Numeraire and payoffs in privacy contexts

We now describe the basic notions for our incentive-based mechanism for users to share data in the SN and make thoughtful decisions about their privacy. We introduce a credit-based system where the user earns credits proportional to the amount of data (e.g., pictures, documents etc.) the user decides to expose, as a co-owner, and to the number of times he/she grants co-ownership to potential owners.

A user i is assigned an initial virtual numeraire $k_i \in R$ to track the credits upon joining the SN. There are well defined mechanisms to credit and debit the numeraire. For each posted data item s , shared with n co-owners, the originator i gains:

$$c = m_i + (\beta \times m_i) \times n$$

where, $m_i \in R$ are the credits assigned as he/she loads a data item, while $\beta \times m_i$ corresponds to the numeraire assigned for each user accepted as a co-owner, $\beta \in [0, 1]$. Each user accepted as a co-owner for s gains $\alpha \times m_i$, where $\alpha \in [0, 1]$. As shown, the more the user shares ownership, the more he/she gets rewarded. The user's numeraire is credited (taxed) based on how pivotal the user's preferences were in making the group decision.

EXAMPLE 4. Assume that in *FactBook*, each uploaded picture is worth 100 while α and β are set to 0.7 and 0.5, respectively. When Alice posts her picture, she grants ownership to Bob and John, who are id-tagged. Her bid score, initially set to 1000 is raised of $m=100$ for posting her picture, and of 70×2 for both Bob and John. That is, Alice totalizes 140 for posting the picture. Bob and John receive 50 each.

The owners make a collective decision on whether posting a data item and they also agree on the exposure preferences (i.e., distance based conditions) to be imposed to potential viewers. Users associate a *value* with each data preference, represented by $v_i(g)$, this value represents the perceived benefit of the user by exposing a data item with preferences g . For example, a user who is interested in maximizing disclosure of his photos would assign a high value to data settings g that do not limit disclosure and allow more users to view this photo. When multiple users are involved for a single decision, they may select different optimal choices. Therefore, we need to design a collective function $F(\cdot)$ (also known as social welfare function) which outputs a unique outcome, in light of the individuals privacy preference inputs. $F(\cdot)$, known as the *social* function is a function over the individuals' value functions, and outputs a certain collective output X :

$$F(v_1(g), \dots, v_n(g)) = X \quad (1)$$

A fundamental requirement of any decision function is that it should have an "optimal" in some sense. Different kinds of desirable attributes of decision functions that characterize optimality have been suggested in Game Theory, Economics, and Voting Theory. Typically, the attributes are concerned with the influence of an individual user on the outcome, and the impact of the outcome on the individual. Some common criteria include Pareto Optimality, Symmetry, Fairness, and Individual Rationality.

In contexts such as ours, it is not obvious how to measure global utility. Considerations other than pure utility values (such as income and fairness) might need to be taken into account. One simple approach, common in game-theory, (due to Nash [26]) is to choose the outcome that maximizes the collective values (utilities). We take this approach, since, as we will see, it satisfies three important properties [11]: 1) it guarantees a relatively fair distribution of the mutually earned utility, 2) it is simple, and 3) is non-manipulable.

5.2 Privacy as a Tax Problem

Our goal is to formulate a mechanism that "aggregates" all the individuals preferences into single representative *group preference*, which builds upon how each user values the different data exposure preferences. Our approach requires each owner i to associate a value $v_i(g)$ to preference g proportional to how important this preference is for him. The value function $v_i(g)$ corresponds to the estimated numeraire value that the user would benefit from adopting setting g .

Given n co-owners of a data item s for which privacy preferences $g \in G$ need to be setup, each co-owner i can essentially opt for the different possible privacy preferences by assigning their value $v_i(g)$ for each $g \in G$. In this paper, we consider the additive social utility, which for a given preference g is the sum of value $v_i(g)$ for all the co-owners, where $F(v_1(g), \dots, v_n(g)) = \sum_{i=1}^n v_i(g)$. In our case, since we cannot assume synchronization, we let the users express their net values privately (that is, each user does not know the numeraire exposed by others). The outcome that maximizes the social value is the outcome to be selected and represented by:

$$g^* = \arg \max_{g \in G} \sum_{i=1}^n v_i(g) \quad (2)$$

In essence, we wish to maximize the sum of the value for each user's bid over the picture's privacy, where the outcome g^* is the privacy setting that maximizes the social utility. If an outcome g is adopted then each user i is required to pay tax π_i , the utility of the choice $c = (g, \pi_1, \dots, \pi_n)$ is the value of the g minus the tax numeraire, given by: $u_i(c) = v_i(g) - \pi_i$. We utilize the *Clarke Tax* mechanism that maximizes the social utility function by encouraging truthfulness among the individuals, regardless of other individuals choices. This algorithm requires each user to state the net value $v_i(g)$ for their preference simultaneously. Unlike the original Clarke Tax mechanism, our formulation does not require a fixed cost to be paid by the n co-owners. We consider the fixed cost to be equal to 0. The tax levied by user i is computed based on the Clarke Tax formulation as follows:

$$\pi_i(g^*) = \sum_{j \neq i} v_j(\arg \max_{g \in G} \sum_{k \neq i} v_k(g)) - \sum_{j \neq i} v_j(g^*) \quad (3)$$

Note user i 's tax $\pi_i(g^*)$ for selecting outcome g^* is composed of two portions, that are computed over a group of users excluding user i . The first portion computes the new outcome that would have been the societal if user i 's values had been ignored and then computes the social utility for such an outcome. The second part computes the social utility for the outcome g^* over the subgroup of users excluding user i . The tax $\pi_i(g^*)$ defined as the difference between the first and second portions.

Assume each co-owner i can essentially opt for privacy preferences stated in terms of connecting path distance, which take values from $g \in \{0, n_{RSet}, \infty\}$, denoting owners only (0), n -distant viewers of relations in $RSet$ and public (∞), respectively. In case $n_{Friends}$ is the winning option, the set of final viewers is identified as the conjunction of the *pivotal users* friends' set. That is, $Profile_1.Friends \cap \dots \cap Profile_n.Friends$. Each user indicates a value $v_i(g)$ for each of the preferences in $(g \in \{0, n_{RSet}, \infty\})$. Figure 1, shows an example including three users, each user i places their values $v_i(g)$ as indicated in the figure. Note that the outcome $g = \{n\}$ maximizes the social value with a value of 7. The users u_1 and u_3 are the *pivotal users* and get taxed

u_i	$v_i(g)$			$\pi_i(g^*)$
	0	n	∞	
u_1	4	2	0.5	0.5
u_2	0	1	4	0
u_3	0.5	4	1.5	1.5
$\sum_i v_i(g)$	4.5	7*	6	
$\sum_{i \neq 1} v_i(g)$	0.5	5	5.5	
$\sum_{i \neq 2} v_i(g)$	4.5	6	2	
$\sum_{i \neq 3} v_i(g)$	4	3	4.5	

Figure 1: Clark Tax Example.

for their contributions to the social value function. User u_2 only contributed $v_2(n) = 1$ which was not pivotal to the decision made, thus user u_2 is not taxed.

The Clarke-Tax approach ensures that users have no incentive to lie about their true intentions. We can briefly show why the Clarke-Tax approach maximizes the users' truthfulness by an additional, simpler example. Consider two individuals a, b : user a feels that the privacy settings on the picture should be private (option $g = 0$), and $v_a(0) = 20$ is what he is willing to spend in order to keep the picture private among the owners. User b , on the other hand, is willing to spend $v_b(\infty) = 10$ to keep the picture public (option $g = \infty$). We refer to maximum users a and b are willing to spend by \bar{v}_a and \bar{v}_b respectively. Additionally, we refer to the best response for users a and b by \hat{v}_a and \hat{v}_b respectively. The charge mechanism in this case is as follows:

$$\pi_a = \begin{cases} 0 & \hat{v}_a < \hat{v}_b \\ \hat{v}_b & \hat{v}_a \geq \hat{v}_b \end{cases} \quad (4)$$

Essentially, if user a wins he will be charged an amount that is as equal to the loss of the other owner, user b follows a similar formulation. In this case, user a 's best response is:

$$\hat{v}_a = \begin{cases} [0, \hat{v}_b), & v_a < \hat{v}_b \\ [max\{0, \hat{v}_b\}, \bar{v}_a], & v_a \geq \hat{v}_b \end{cases} \quad (5)$$

Notice that $v_a = \hat{v}_a$ is always assured to fall in the range for the best response in both cases. If a and b declare the truth, a option will prevail, and a will have to pay tax to the SN $\pi_a = 10$ in order to see his option enforced. If a aims at spending less and declares, falsely, $\hat{v}_a = 11$, a will still win, but according to equation since $11 > 10$, a still have to pay a tax $\pi_a = 10$. So, underestimating the real value is not going to change the result of the voting process. Similarly, even if b declares less than what he thinks the real value is, since the numeraire is not going to be reimbursed at him, he is not going to get any advantage by lying. That is, truthful revelation is weakly dominant, a more general proof is available in [11]. The simplicity of strategy is highly desirable in the design of solutions for this type of privacy problems, where users most likely are going to make intuitive and simple decisions to address their privacy considerations. Additionally, the Clarke-Tax mechanism satisfies several other desirable criteria, including the "Condorcet winner" (a choice that would have beaten every other choice in pair-wise votes is guaranteed to be chosen by the mechanism [6]), "independence of irrelevant alternatives" (removal of any unchosen preference from the set of alternatives will not change the outcome [33]) and that the identity of a voter has no influence on the outcome.

The Clark-Tax approach is far from perfect. One significant drawback is the assumption of users' should be able to compute the value of the different preferences. We assume users can map the value to the number of users able to access the shared data, and this is possible using several social network indicators, such as the set of friends, set of common friends, and on several small world network metrics such as node degree, centrality, betweenness, trust paths, mixing patterns, and resilience [31, 3].

5.3 Inference of privacy policies

The approach proposed in previous section requires manual input for each of the pictures co-owned. Users may have up to hundreds of pictures, and a significant percentage of them may be co-owned. As such, asking users to bid for each of them may be, in the long run, very cumbersome. An effective idea to overcome this limitation is to utilize inference-based techniques; and leverage previous decisions to free the users from the burden of going through the voting process numerous times. It is easily verifiable, that most users appear in pictures with more or less the same small set of users (typically directly related among each other), and, that the sensitivity of a given picture depends also upon the context in which the picture has been taken. Building upon these observations, we suggest using tags and similarity analysis to infer the best privacy policy to use for pictures shared among owners who have an history of shared pictures.

As discussed in Section 2, users add words, referred to as tags, to associate a context or a topic with their content. In the case of pictures, content tags can be added to each picture, or at the album level³. For simplicity we focus on the case where users add up to one tag each per picture. As such, for a given picture owned by k users, we associate at most k tags, $\{t_1, \dots, t_k\}$. This meta-data is used to conduct similarity analysis with pictures already shared by the same set of users.

For convenience, we represent each picture as a vector of tags. That is, let $T = \{\vec{t}_1, \vec{t}_2, \dots, \vec{t}_n\}$ be a set of pictures shared among the set of owners $Own_{U_{set}}$. Let \vec{t} be the picture whose policy is to be defined. In order to identify the best policy to associate with \vec{t} , we conduct similarity analysis among the pictures in T and \vec{t} . Similarity analysis requires two major steps to be undertaken. First, tags' similarity needs to be conducted. To be able to utilize similarity metrics, we rely on the informal classification system resulting from the practice of collaborative tagging. This user-generated classification system, is referred to as *folksonomy* [27], and is generally defined in terms of a collection of posts, each associated with one or more tags.

DEFINITION 2. A *folksonomy* is a tuple $F := (U; T; R; Y)$ where U , T , and R are finite sets, whose elements are users, tags and resources, respectively. Y is a ternary relation between them, i. e., $Y \subseteq U \times T \times R$. A post is a triple $(u; T_{ur}; r)$ with $u \in U$, $r \in R$, and $T_{ur} := \{t \in T | (u; t; r) \in Y\}$

By relying on a folksonomy, we can compare two pictures and assign them a similarity score, based on the tags associated with each of them. Tags relatedness can be constructed

³Content tags are not to be confused with id-tagging, which we used to identify pictures' potential owners.

according to several metrics [32, 23]. In our case, we employ the following notion, which is based on occurrence of tag pairs.

$$w(t_1; t_2) := \text{card}\{(u; r) \in U \times R | t_1; t_2 \in T_{ur}\} \quad (6)$$

Based on these notions, we define similarity of pictures as the overall relatedness among the tags associated with the pictures. Given two pictures \vec{t}, \vec{t}' their similarity is determined as follows.

$$\text{sim}(\vec{t}, \vec{t}') = \sum_{i=1}^k \sum_{j=1}^n w(t_i, t'_j) \quad (7)$$

Note that similarity is commutative, i.e., $\text{sim}(\vec{t}, \vec{t}') = \text{sim}(\vec{t}', \vec{t})$. The equation 7 returns a similarity value expressed as non-negative number.

Second, once the list of similarity values among all the pictures in T shared by Own_{Uset} is computed, the picture $champ = \max\{\text{sim}(\vec{t}, \vec{t}'), \text{sim}(\vec{t}, \vec{t}_1), \dots, \text{sim}(\vec{t}_n, \vec{t})\}$ with the highest similarity score is selected.

EXAMPLE 5. With reference to Example 4, let us assume Alice tags the shared picture as party, while Bob uses the word fun and John night. Suppose that Alice, Bob and John already share two pictures, say \vec{t}_1 and \vec{t}_2 , tagged using other freely-chosen words. \vec{t}_1 was tagged using gathering, fun, game; while \vec{t}_2 using words friend, beer, home. Let us assume that the $\text{sim}(\vec{t}, \vec{t}_1) = 100$ and that $\text{sim}(\vec{t}, \vec{t}_2) = 92$. Since \vec{t}_1 is the most similar to \vec{t} , its privacy policy will be proposed to the three owners.

The privacy policy associated with $champ$ is prompted to all the users in Own_{Uset} . If the users agree on the inferred privacy policy, the same is used, and the numeraire intakes is the same as the one originally spent for the championed picture. If the users do not agree, or a picture significantly similar to \vec{t} is not found, the auction mechanism is proposed to the end users. A temporary policy, chosen among previously adopted ones is then used, until a final decision is not taken.

6. SYSTEM IMPLEMENTATION AND EXPERIMENTAL RESULTS

We have implemented a proof-of-concept social application of the proposed approach for the collaborative management of shared data, referred to as *Private Box*. *Private Box* is fully integrated with Facebook social network platform [13]. *Private Box* supports the following features: controlled sharing of pictures; automatic detection of pictures' co-owners based on id-tags; collective privacy policies enforcement over shared pictures based on auctions.

Private Box has been implemented in PHP and uses Facebook platform REST-like APIs for PHP and Facebook Markup language (FBML). REST-like APIs are used to retrieve and prompt all the information related to a Facebook user profile such as the Facebook user identifier and its friends identifiers. FBML is an evolved subset of HTML that gives our *Private Box* application the same style of Facebook web site. The information related to a Facebook user profile such as the user identifier, list of friends identifiers, the user photos

and albums identifiers are stored in a MySQL database. The implementation consists of a set of PHP files where each file implements one of the main features of *Private Box*. Figure 2 represents the interaction flow of a user with *Private Box* application. First, **AddPhotos** page allows a user to select those photos from his/her Facebook albums on which he wants to have a fine-grained control. Once photos have been selected, *Private Box* determines the set of potential co-owners of the photos based on the id-tags, as described in Section 3. Each potential co-owner is notified through a standard Facebook notification message about the possible co-ownership. Then, **PrivateBox** page displays the photos stored in the *Private Box*, including the pictures added by him/her, and those have been added into *Private Box* when the user was granted ownership.

Finally, the **Auction** page is the core of the application, and it enables the collaborative enforcement of privacy policies on co-owned data as it is described in Section 5 (see Figure 2). **Auction** page shows the user's updated bid score (i.e., numeraire) each time the user adds pictures, grants ownership or obtains ownership. Moreover, it allows a user to start an auction using the Clarke-Tax for a co-owned photo by specifying a bid value $v_i(g)$ for each possible privacy preference g associated with the photo. $v_i(g)$ represents the perceived benefit of the user by exposing the photo with privacy preference g . The only possible privacy preferences g that are supported by *Private Box* are "share with co-owners" and "share with friends" because in Facebook it is not possible to connect users based on social relationships other than "friends". The user can monitor anytime the progression of an auction that the user has started which is not completed yet. To ensure correctness of the mechanism, however, he/she can only bid once, and cannot view others' bids. During an auction, the photo under auction is visible only to the co-owners that appear in the photo to avoid that any of the co-owners privacy preferences are violated before the privacy setting that maximizes the social utility $F(\cdot)$ is determined. The user can also view the ongoing auctions started by its friends (but not the bids), and choose to join one of them. When the user joins an auction he/she has to specify the bid score for his/her privacy preference g associated with the photo under auction. Finally, the user can also view the results of previous completed actions. Note that only when an auction is completed the user can see the $v_i(g)$ specified for each privacy preference g by the other users (Figure 2, step 5).




Figure 3: Auction execution times

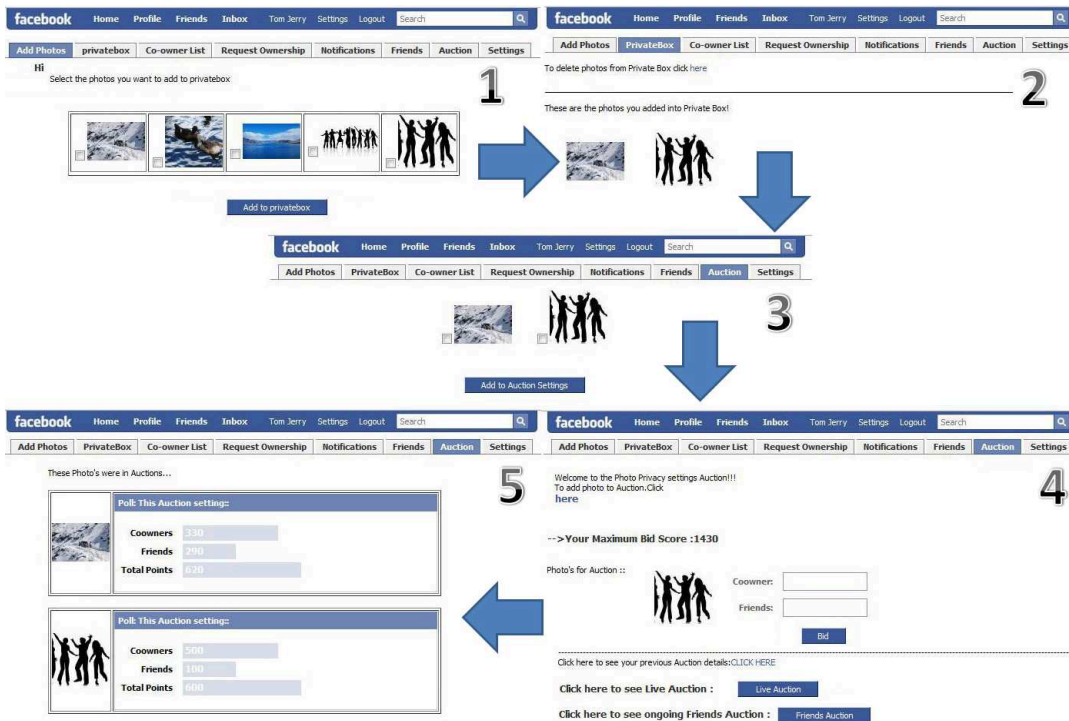


Figure 2: *Private Box's* execution flow.

Private Box has additional functionalities, to visualize friends' and co-owners' pictures. The **Co-owner list** page, for example, displays the list of the co-owners. Once a co-owner is selected, the photos the ownership of which is shared between the co-owner and the current user are visualized. Another supported feature is the ownership request, managed in the **Request Ownership** page. The **Request Ownership** displays a list of pictures where the current user has been tagged, i.e., is a potential-owner. The user can select the pictures of which he wants to obtain co-ownership. A notification is sent to the user who has uploaded the picture and it is displayed in his/her **Notifications** tab. Finally, **Friends** page displays the pictures that the friends of the user have uploaded in *Private Box*. The inference component of the system is not currently implemented, and its deployment is part of our future work.

According to research related to face recognition in online albums there are between 2 to 4 faces per photo [29, 9]. We have evaluated the scalability of the collaborative privacy policies enforcement based on auctions by varying from 2 to 12 the number of co-owners that appear in a photo under auction. Figure 3 reports the execution times to perform Clarke-Tax algorithm once all the co-owners have placed a bid, while varying the number of co-owners. In other words, the graph shows the execution time of finding a privacy setting which satisfies each co-owner privacy preference, and of calculating the bid score to be levied to the pivotal users. The execution time linearly increases with the increase of the number of co-owners because the Clarke-Tax algorithm has to find the maximum for function $F(.)$ over a greater number of co-owners bid scores. However, the increase is negligible with respect to the number of co-owners. The ex-

ecution time is so fast that the collaborative enforcement of privacy policies is transparent to the user.

7. RELATED WORK

Security and privacy in Social Networks, and more generally in Web 2.0 are emerging as important and crucial research topics [15, 4, 14, 19]. SNs have been studied by scholars from different disciplines: sociologists, HCI, computer scientists, economists etc. In this section we overview some of previous work that is most relevant to collaborative privacy management for SNs. Several studies have been conducted to investigate users' privacy attitudes, and possible risks which users face when poorly protecting their personal data [34] in SNs. Gross et al. [1] provided an interesting analysis of users' privacy attitudes across SNs. Interestingly, Ellison et al. [30] have highlighted that on-line friendships can result in a higher level of disclosure due to lack of real world contact. According to Ellison et al. [30] there are benefits in social capital as a result of sharing information in a social network that may limit the desirability of extensive privacy controls on content. Following such considerations, our proposed approach does not simply block users' accessibility to shared data, but it ensures that sharing occurs according to all the stakeholders' privacy interests. The need for solutions addressing the problem of information leakage in this context is also reported in [22] where an extensive analysis of the more relevant threats that SNs users currently face is reported.

To cope with security and privacy problems, SNs sites are currently extending their access control based mechanisms, to improve in flexibility and limit undesired information disclosure. There is a general consensus that in SNs a new paradigm of access control needs to be developed [18, 15, 5].

A first attempt along this direction has been taken by Gollu et al. [18], where a social-networking based access control scheme suitable for online sharing was presented. They proposed an approach that considered identities as key pairs, and social relationship on the basis of social attestations. Access control lists are employed to define the access lists of users.

Carminati et al. [5] have proposed a rule-based access control mechanism for SNs that is based on enforcement of complex policies expressed as constraints on the type, depth, and trust level of existing relationships. Furthermore, Carminati et al. proposed using certificates for granting relationships' authenticity, and the client-side enforcement of access control according to a rule-based approach. In this paper, we employ privacy policies using a simplified version of the access rules used by Carminati et al. More recently, Carminati et al. [4] have extended their previously proposed model to make access control decisions using a completely decentralized and collaborative approach. Their proposed work is orthogonal to the work proposed in this paper. Our analysis of collaborative privacy management does not relate to the privacy of users' relationships. Rather, we focus on collaborative approaches for privacy protection of users' shared content.

Recently, Gates [16] has described relationship based access control as one of the new security paradigms that addresses the requirements of the Web 2.0. Hart et al. [21] proposed a content-based access control model, which makes use of relationship information available in SNs for denoting authorized subjects. However, those works do not address collaborative privacy issues.

Another interesting work related to ours is HomeViews [17], an integrated system for content sharing supporting a light-weight access control mechanism. HomeViews facilitates ad hoc, peer-to-peer sharing of data between unmanaged home computers. Sharing and protection are accomplished without centralized management or coordination of any kind. This contribution, although very interesting, is designed around a very different environment, and it considers sharing of content without taking into account multi users privacy implications.

Mannan et al. [25] proposed an interesting approach for privacy-enabled web content sharing. Mannan et al. leveraged the existing "circle of trust" in popular Instant Messaging (IM) networks, to propose a scheme called IM-based Privacy-Enhanced Content Sharing (IMPECS) for personal web content sharing. This approach is consistent with our ideas of sharing of privacy controls, and presents an interesting implementation design. On the other hand, IMPECS is a single-user centered solution: that is, only one user is involved in the decision of whether to share his/her content within his/her trust circle.

Finally, with regards to game theoretic approaches related to our solution, our work is related to [20, 36]. Varian [36] conducts an analysis of system reliability within a public goods game-theoretical framework. Varian focuses on two-player games with heterogeneous effort costs and benefits from reliability. He also adds an inquiry into the role of taxes and fines, and differences between simultaneous and sequential moves. Grossklags et al. in [20] generalize [36] and build from public goods literature to model security interactions through three well-known games, introducing a novel game (weakest target, with or without mitigation)

for more sophisticated scenarios. Similarly, in our work we model the collective privacy problem as a new game, using the results from game security economics. The adoption of our carefully selected technique ensures the design of a N-player game, in which truthfulness and correctness are the winning strategies.

The Clarke-Tax algorithm [8, 7] has been recognized as an important social decision protocol. The approach has been applied to address problems of different nature [10, 12, 37]. To the best of our knowledge, however, this is the first time a voting protocol of this kind is utilized for collective privacy problems. In [10, 12] the Clarke-Groves mechanism has been introduced into artificial intelligence, using it to explore multiagent planning. At each step, instead of negotiating over the next joint action, each agent votes for the next preferred action in the group plan and individual preferences are aggregating using a voting procedure. Recently, Wang et al. [37], proposed an interesting secure version of the Clarke-Tax voting protocol. Following the security requirements identified by Wang in [37], we implement a system which guarantees full protection of users' privacy and universal verifiability. However, Wang's solution heavily relies on cryptographic primitives, and encryption techniques, implying a level of sophistication of users which may not be appropriate in Web 2.0 settings.

8. CONCLUSION

In this paper we discussed a novel model for privacy management across social networks, where data may belong to many users. We presented a theoretical representation of the collective privacy management problem, and proposed a solution which builds upon well-known game theoretical results. We implemented a tool prototype hosted in Facebook, and carried out performance analysis. Our next step is to conduct extensive user studies, to assess the users' perspective of this type of approach. In a preliminary investigation, we observed high interest from users toward approaches allowing users' control over shared content.

As part of future work, we also would like to extend our mechanism to support more sophisticated and flexible policies. We will investigate policies network-based policies, to include predicates related to the users' geographic locations. Finally, we will investigate further the implications of our approach in case of revocation or leave of some co-owners.

9. ACKNOWLEDGEMENTS

This work was partially funded by the National Science Foundation (NSF-CNS-0831360) and National Security Agency (NSA H98230-07-1-0231). We would like to thank Marco Rossi for his useful advices on the Clarke-Tax algorithm and Shitij Kulshreshtha for his help with the development of the private box application.

10. REFERENCES

- [1] A. Acquisti and R. Gross. Imagined communities: Awareness, Information Sharing, and Privacy on the Facebook. In *Proc. of Privacy Enhancing Technologies*, pages 36–58. Springer, 2006.
- [2] D. Beaver. 10 billion photos. http://www.facebook.com/note.php?note_id=30695603919, October 2008.

- [3] S. P. Borgatti and M. G. Everett. A graph-theoretic perspective on centrality. *Social Networks*, 28(4):466–484, October 2006.
- [4] B. Carminati and E. Ferrari. Privacy-aware collaborative access control in web-based social networks. In *DBSec*, pages 81–96, 2008.
- [5] B. Carminati, E. Ferrari, and A. Perego. Rule-based access control for social networks. In *OTM Workshops (2)*, pages 1734–1744, 2006.
- [6] L. Chen, X. Den, Q. Fang, and F. Tian. Condorcet winners for public goods. In *Annals of Operations Research*, volume 137, pages 229–242, 2005.
- [7] E. H. Clarke. Multipart pricing of public goods. In *Public Choice 11*, pages 17–33, 1971.
- [8] E. H. Clarke. Multipart Pricing of Public Goods: An example. In *Public Price for Public Products, Urban Inst.*, 1972.
- [9] M. Davis, M. Smith, J. Canny, N. Good, S. King, and R. Janakiraman. Towards context-aware face recognition. In *Proceedings of the 13th annual ACM international conference on Multimedia*, pages 483–486, New York, NY, USA, 2005. ACM.
- [10] E. Ephrati and J. S. Rosenschein. The Clarke-tax as a consensus mechanism among automated agents. In *National Conference on Artificial Intelligence*, pages 173–178, 1991.
- [11] E. Ephrati and J. S. Rosenschein. Voting and multi-agent consensus. 1991.
- [12] E. Ephrati and J. S. Rosenschein. Deriving consensus in multi-agent systems. In *Journal of Artificial Intelligence*, volume 87, pages 21–74, November 1996.
- [13] Facebook. Facebook web site. <http://www.facebook.com/>.
- [14] A. Felt. Defacing Facebook: A security case study.
- [15] A. Felt and D. Evans. Privacy protection for social networking platforms. In *Proceedings of Web 2.0 Security and Privacy 2008 (in conjunction with 2008 IEEE Symposium on Security and Privacy)*, 2008.
- [16] C. Gates. Access control requirements for Web 2.0 Security and Privacy. In *IEEE Web 2.0 Privacy and Security Workshop*, 2007.
- [17] R. Geambasu, M. Balazinska, S. D. Gribble, and H. M. Levy. Homeviews: peer-to-peer middleware for personal data sharing applications. In *SIGMOD Conference*, pages 235–246, 2007.
- [18] K. K. Gollu, S. Saroiu, and A. Wolman. A social networking-based access control scheme for personal content. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP '07)-Work-in-Progress Session*, 2007.
- [19] R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *Workshop on Privacy in the Electronic Society*, 2005.
- [20] J. Grossklags, N. Christin, and J. Chuang. Secure or Insure?: a game-theoretic analysis of information security games. In *WWW*, pages 209–218, 2008.
- [21] M. Hart, R. Johnson, and A. Stent. More content - less control: Access control in the Web 2.0. In *IEEE Web 2.0 Privacy and Security Workshop*, 2007.
- [22] G. Hobgen. Security issues and recommendations for online social networks. ENISA Position Paper N. 1, 2007.
- [23] J. Jiang and D. Conrath. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proceedings of ROCLING X*, Sep 1997.
- [24] A. Lenhart and M. Madden. Teens, privacy & online social networks. Pew Internet & American Life Project, 18 April 2007.
- [25] M. Mannan and P. C. van Oorschot. Privacy-enhanced Sharing of Personal Content on the Web. In *WWW*, pages 487–496, 2008.
- [26] A. Mas-Colell and M. D. Whinston. *Micro-Economic Theory. Chapter 23*. Oxford University Press, fourth edition, 1998.
- [27] A. Mathes. Folksonomies: Cooperative classification and communication through shared metadata. <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>, 2004.
- [28] G. A. Miller. Wordnet: a lexical database for English. *Commun. ACM*, 38(11):39–41, 1995.
- [29] M. Naaman, R. B. Yeh, H. Garcia-Molina, and A. Paepcke. Leveraging context to resolve identity in photo albums. In *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pages 178–187, New York, NY, USA, 2005. ACM Press.
- [30] C. L. NB Ellison, C Steinfield. Benefits of Facebook "Friends":social capital and college students' use of online social network. *Journal of Computer Mediated Communication-Electronic*, 2007.
- [31] M. E. J. Newman. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical Review E*, 64(1):016132+, June 2001.
- [32] G. Pirro' and N. Seco. Design, implementation and evaluation of a new semantic similarity metric combining features and intrinsic information content. In *Proceedings of On the Move to Meaningful Internet Systems*, 2008.
- [33] P. Ray. Independence of irrelevant alternatives. In *Econometrica*, volume 41, pages 987–991, 1973.
- [34] D. Rosenblum. What anyone can know: The privacy risks of social networking sites. *IEEE Security and Privacy*, 5(3):40–49, 2007.
- [35] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 38–47, New York, NY, USA, 2001. ACM.
- [36] H. R. Varian. System Reliability and Free Riding. In *Economics of Information Security*, pages 1–15. Kluwer Academic Publishers, 2002.
- [37] C. Wang and H. fung Leung. A secure and private Clarke-Tax voting protocol without trusted authorities. In *Proc. of 6th International conference on Electronic Commerce*, pages 556–565, New York, NY, USA, 2004. ACM.
- [38] X. Wu, L. Zhang, and Y. Yu. Exploring social annotations for the semantic Web. In *WWW*, pages 417–426, 2006.