The VLDB Journal
DOI 10.1007/s00778-010-0193-7

SPECIAL ISSUE PAPER

Privacy policies for shared content in social network sites

Anna C. Squicciarini · Mohamed Shehab · Joshua Wede

Received: 15 August 2009 / Revised: 10 February 2010 / Accepted: 21 April 2010

Abstract Social networking is one of the major technological phenomena of the Web 2.0, with hundreds of millions of subscribed users. 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, they do not provide any mechanisms for collective management of privacy settings for shared content. In this paper, using game theory, we model the problem of collective enforcement of privacy policies on shared data. 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. Our approach enables social network users to compose friendship based policies based on distances from an agreed upon central user selected using several social networks metrics. 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 privacy protection mechanism for social networking has been proposed. We also extend our mechanism so as to support collective enforcement across multiple social network sites. 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. Through our implementation, we show the feasibility of such approach and show that it can be implemented with a minimal increase in overhead to end-users. We complete our analysis by conducting a user study to investigate users' understanding of co-ownership, usefulness and understanding of our approach. Users responded favorably to the approach, indicating a general understanding of co-ownership and the auction, and found the approach to be both useful and fair.

Keywords Social networks · Privacy · Game theory · Clarke-Tax

A. C. Squicciarini (✉)
College of Information Sciences and Technology,
Pennsylvania State University, University Park, PA, USA
e-mail: acs20@psu.edu

M. Shehab
Department of Software and Information Systems,
University of North Carolina, Charlotte, NC, USA
e-mail: mshehab@uncc.edu

J. Wede
Department of Psychology, Pennsylvania State University,
University Park, PA, USA
e-mail: jlw63@psu.edu

1 Introduction

Social networks (SNs, for short), including Friendster.com, Tagged.com, Xanga.com, LiveJournal, MySpace, Facebook, and LinkedIn, have developed on the Internet over the past several years and these SNs have been successful in attracting users. According to ComScore Media Metrix, more users visit MySpace than Yahoo, MSN, or Electronic Arts gaming site [39]. Through SNs, users 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 by 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 through interactions among digital friends often leads to disclosures that may not be appropriate

Published online: 30 June 2010

in a public forum. Such open availability of data exposes SN users to a number of security and privacy risks [28,33,59].

In order to help users protect their personal content, current SN architectures adopt a simple user-centric policy management approach [36,41,15], where a privacy aware user is able to specify a policy that manages access to their posted profile objects. There have been numerous studies concerning privacy in the online world [5,28,70]. A number of conclusions can be drawn from these studies. First, there are varying levels of privacy controls, depending on the online site. For example, some sites make available user profile data to the Internet with no ability to restrict access, while other sites limit user profile viewing to just a set of selected trusted friends. Other studies introduce the notion of the privacy paradox — the relationship between individual privacy intentions to disclose their personal information and their actual behavior [52]. Individuals voice concerns over the lack of adequate controls around their privacy information while freely providing their personal data. Other research concludes that individuals lack appropriate information to make informed privacy decisions [1], and, when there is adequate information, short-term benefits are, more often then not, opted over long-term privacy. People are concerned about privacy but most are not doing anything about it. This can be attributed to many things, e.g., the lack of privacy controls available to the user, the complexity of using the controls, and the burden associated with managing these controls for large sets of users.

A significant privacy threat of SN sites is given 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 [4]. Pictures are tied to individual profiles and often either explicitly (through tagged labeled boxes on images) or implicitly (through recurrence) identify the profile holder [2]. Such pictures are made available for other SN users, who can view, add comments and, 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 even when they are not the actual or sole stakeholders. 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.

Letting one user take full responsibility over another's privacy settings is extremely ineffective. 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 levels [32]. 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 collective 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. Moreover, the approach should be practical and promote co-ownership, since users knowingly do not enjoy spending time in protecting their privacy [61].

We analyze these requirements from a game theoretical perspective [43,68] and model the process of collective privacy management of shared data as a mechanism design problem. We map the user collective policy specification to an auction based on the Clarke-Tax [10,11] 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 [17], 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 problem, 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 and automate collective decisions, thus freeing the users from the burden of manually selecting privacy preferences for each picture.

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. We also discuss how to export collective privacy policies across domains, using

OpenSocial[26]. As we discuss, this involves addressing a number of non-trivial issues. As part of our assessment, we present the results of a user study, where we investigated users' understanding of co-ownership, usefulness and understanding of our approach, and attitudes toward the fairness of the approach. Users responded favorably to the approach, indicating a general understanding of co-ownership and the auction, and found the approach to be both useful and fair.

The rest of the paper is organized as follows. We begin with a discussion of the related work in the next section, followed by an abstract representation of SNs. In Sect. 4, we discuss data co-ownership in SNs. In Sect. 5, we highlight the requirements for the design of an effective solution supporting collective privacy management. In Sect. 6, 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 Sect. 8. In Sect. 9, we report the result of our user study, while we discuss the limitations of our approach in Sect. 10. We conclude the paper in Sect. 11.

2 Related work

Security and privacy in Social networks and more generally in Web 2.0 are emerging as important and crucial research topics [7,20,21,28]. 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 collective 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 [59] in SNs. Gross et al. [2] provided an interesting analysis of users' privacy attitudes across SNs. Interestingly, Ellison et al. [47] 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. [47], there are benefits in social capital as a result of sharing information in a SN that may limit the desirability of extensive privacy controls on content. Following such considerations, the approach we present in this work 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 [33], 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 [8,21,25]. A first attempt along this direction has been taken by Gollu et al. [25], 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. [8] 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. [7] have extended their previously proposed model to make access control decisions using a completely decentralized and collective approach. Their proposed work is orthogonal to the work proposed in this paper. Our analysis of collective privacy management does not relate to the privacy of users' relationships. Rather, we focus on collective approaches for privacy protection of users' shared content.

Recently, Gates [22] 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. [32] 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 collective privacy issues.

Another interesting work related to ours is HomeViews [23], 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 co-ordination of any kind. This contribution, although very interesting, is designed around a very different domain, and it considers sharing of content without taking into account multiusers privacy implications.

Mannan et al. [42] 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 respect to game theoretic approaches related to our solution, our work is related to [29,64]. Varian [64] conducted an analysis of system reliability within a public goods game theoretical framework. Varian focused on

two-player games with heterogeneous effort costs and benefits from reliability. He also added an inquiry into the role of taxes and fines, and differences between simultaneous and sequential moves. Grossklags et al. in [29] built 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 [10,11], which is at the core of our solution, has been recognized as an important social decision protocol. The approach has been applied to address problems of different nature [16,18,67]. 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 [16,18], 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. [67], proposed an interesting secure version of the Clarke-Tax voting protocol. Following the security requirements identified by Wang in [67], we implement a system that 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. As we discuss in Sect. 10, these countermeasures are required for unstable domains, where strong guarantees are required to ensure the users' correct behavior. In the context of SN, we believe these techniques do not apply well, and go against the spirit and philosophy of SN sites, which are primarily entertainment sites.

3 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 $i\ Rt\ j$, 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, for example $i\ is\_fan\_of\ U2$, although not relevant for us.
- $Profile_i$. The profile of a user i. We represent it as a tuple $Profile_i = (GRelType_1, \ldots, GRelType_k, Set)$ where $GRelType_l$ represents the list of users having a relationship Rt_l such that $i'\ Rt_l\ i$ 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 $\langle i\ Rt\ j \rangle$ 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: $(\langle 1\ Rt_1\ 2 \rangle, \langle 2\ Rt_2\ 3 \rangle, \ldots, \langle k-1\ Rt_{k-1}\ k \rangle)$, where each tuple $\langle i\ Rt_l\ j \rangle$ 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
$(\langle Alice\ Friend\_Of\ Bob \rangle, \langle Bob\ Colleague\_Of\ John \rangle)$.

3.1 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. 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.

4 Data co-ownership in SNs

In this section, we introduce the notion of collective 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 manner.[2]

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.

We primarily 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. We discuss technical challenges-related strategies required to support co-ownership for textual content in 4.2.

4.1 Users' classification and image co-ownership

In order to identify co-owners of a given piece of data s, we provide a general classification of users based on their relationship with s. 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 [69], associated with the data denoted as $TSet$. Their semantic can be analyzed by means of similarity tools [45]. In the case of pictures, we employ a specific type of tags widely used in Facebook [19]. 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 notice that although not error-free (one could add the wrong tags), using id-tags has several advantages: 1. anyone can tag a picture, so even a potential stakeholder can tag him/herself on a picture that is not hosted in his/her own profile, to claim the ownership. It is also equally easy to remove wrongly assigned tags. 2. images can be annotated automatically by employing facial recognition techniques [62]. 3. It is relatively easy to leverage the id-tags to detect shared content. Social annotations can be generalized as a method to identify the stakeholders of content of any type, not limited to pictures, as long as a unique tag for each user in the set U of users in a SN is provided. Using images as a case study, annotated with id-tags, and documents, annotated with metadata, we now derive a general notion of users' relationship to data.

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

[1] Distance-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 [8].

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

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 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 Sect. 6, 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_{USet}^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.

4.2 Documents co-ownership

The major challenge related to co-ownership in case of documents and/or text files deals with the proper identification of the stakeholders, who should be entitled to ownership. While in case of images, it is inarguable that who ever appears in a picture may have interest in knowing—and possibly controlling—how the image will be handled on the social network sites, this is not always the case for textual documents. Even if text mining approaches are used, and names—or equivalent users' identifiers—are extracted from the content, this does not always imply that the cited names are in fact of potential stakeholders. A typical example where this approach would fail would be an academic journal, where several authors of related works are cited. Being cited does not relate with ownership rights. An alternative approach, which is more suitable, is to add annotations, so that documents can be annotated with meta-data listing the possible co-owners who participated in the creation of the document or who should be involved in its management. Annotations are themselves not ideal, in that they require for the most part manual input of one or more individuals to input meta-data so as to provide the necessary authorship information. Additionally, annotations should be unforgeable, so that co-owners cannot freely change them to deny privileges to legitimate co-owners. To this extent, we propose a simple approach that incorporates watermarking techniques into documents upload, so as to substantially mitigate the risks of altered meta-data, and allow only the actual stakeholders to do proper editing. Intuitively, in order to support this feature, an integrated component supporting a watermarking technique such as [13] in the annotation system needs to be supported. The annotation process consists of the following simple steps:

- The originator adds a list of users' ids $\{id_1, \ldots, id_n\}$, as part of the uploading process of the document;
- The list of users' ids is embedded as watermark in the uploaded document. The access policy enforced at this time is the one of the originator.
- The document is posted and automatically made available to the annotated stakeholders.
- Stakeholders receive an automated invitation to accept/deny co-ownership. If they accept, the document is "unlocked", and they are also able to add new stakeholders, if additional ones are identified. To avoid uncontrolled diffusion of the document, new stakeholders are to be approved by the co-owners. Co-owners can simply opt in or out for additional stakeholders at the time they accept the co-ownership invitation.

The other important aspect related to data co-ownership deals with the access type that can be allowed. While images cannot be edited, it is important to enable multi-owners to editing access rights. Hence, the document may be open to read-only or editing access to viewers, or it could be accessible for editing among the shared co-owners only. Although documents' collective editing is being addressed in several existing works [51,58], and a number of collective tools are nowadays available for these purposes, the best approach for collective editing across SN sites has yet to be identified. A control mechanism that well suits the need of collective editing within these sites for maintaining consistency of shared data in collective editing is currently represented by opera-

tional transformation. In this approach, local operations are executed immediately after their generation, while remote operations must be transformed regarding concurrently executed operations. Known algorithms based on this approach are SOCT4 [66] and GO/GOTO [63]. These algorithms use a central server for exchanging and timing operations, which well suits the centralized architecture underlying SNs sites. To time operations, a vector clocks is associated to each operation performed against the document.

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

5 Collective sharing requirements

In case of single-user ownership, enforcing of a 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.

Finally, the 'least common denominator' approach is not a satisfying one either, since it does not represent the co-owners' group as a whole, nor does it help in maximizing the group's overall benefit. There are potentially many cases in which this approach does not produce desirable results. For instance, consider a case with 10 stakeholders and all of them indicate that all of their friends should be given viewer privileges. Suppose that all of the stakeholders have many friends but only two or three in common among all the stakeholders. The result is that only two or three users other than the stakeholders will be given viewer access. This is not reasonable. Likewise, consider the scenario where a group of close friends take a group picture at a party and post it. Since they are close friends in real life, it is likely that they have a lot of friends in common on the social network as well. When the intersection is done with each of the party members as stakeholders, the set of users granted viewer access will probably be reasonable; however, assume Bob brings his coworker, Alice, who knows none of Bob's friends, to the party. It is reasonable to assume that Alice's friends list on the SN does not overlap at all with the other members of the party. Now that Alice is also a stakeholder, the result of the intersection is empty; thus only the stakeholders have access to the data. This is a perfectly plausible scenario, but the result is unreasonable. Hence, simply doing intersections to determine the users that can view a data item allows the potential for a single stakeholder's wishes to override those of all others in what seems to be an unreasonable way.

Based on these considerations, we identified the following core requirements for collective 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.
- **Incentive-Compatible** The mechanism deployed should provide incentives to users to be honest and discourage any malicious behavior trying to subvert the outcome of the collective decisions.

In the next section, we propose an approach that satisfies the aforementioned 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.

6 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, this approach is pretty simplistic and fails to leverage the individuals' preferences within the co-owners' group. In addition to the identified drawbacks, majority and ranking-based approaches, such as the ones described above, have proved to be unfair, in that astute individuals may manipulate outcomes to their advantage [30].

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.
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.

6.1 Shared content's privacy as a game theoretical problem

In order to ground our solution, we approach the problem of collective privacy using Game Theory. We can model our problem of collective privacy as a Nash equilibrium problem[43] because each player is assumed to know the strategies available to the other players, and no player has anything to gain by changing only his or her own strategy unilaterally. If each player has chosen a strategy and no player can benefit by changing his or her strategy while the other players keep their strategy unchanged, then the current set of strategy choices and the corresponding payoffs constitute Nash equilibrium. This is most definitely true for shared content, if each co-owner picks a certain privacy setting, none of the co-owners can benefit by unilaterally changing his setting. The games that fit the described scenario usually fall under co-ordination or co-operation games in Game theory. Well-known games for these classes are the Prisoners Dilemma, the Chicken Game (or its Hawk-Dove variant), and the tax Games[43,68]. However, the problem with these types of games is that while the dominant strategies are the ones where parties reach a collective decision, they actually have no gain in co-operating, and while pursuing their individual goals, they may not reach an equilibrium. Since these games provide no gain for collective decision making, they do not serve as a good model for our approach.

In our context, users may have different and conflicting goals (privacy preferences). Our goal is to provide them with an approach that addresses such conflicts, returning an outcome that satisfies individuals interests as much as possible. Hence, our problem is essentially a mechanism design issue [68]. Mechanism design is concerned with the design of mechanisms that favor particular outcomes despite agents pursuing their own interests. It is also known as reverse game theory: while game theory analyzes the strategic behavior of rational agents, mechanism design uses these insights into design games inducing certain strategies and hence outcomes. The main difficulty with mechanism design problems is that agents (users, in our domain) may choose to misreport their valuations in an attempt to affect the outcome to their liking. The tool that the mechanism uses to motivate the agents to reveal the truth is monetary payments. These payments need to be designed in a way that ensures that rational players always reveal their true valuations [3]. To date, only one general method, called VCG [11,31,65], is known for designing such a payment structure. The VCG mechanism is usually used to treat a public choice problem in which a public project's cost is borne by all players. VCG is well suited to model content sharing since it is *the only known model that guarantees not only that truthfulness is a dominant strategy, but also that allocations are efficient* in an economic sense, that is, awards are given to the highest bidders [37].

Since VCG applies for cases where multiple allocations need to be made at the same time, such as combinatorial auctions, we need to deploy a slight variant of VCG. Variations of the VCG that could apply in our context are given by the Groves–Ledyard [27] mechanism and Clarke-Tax [11]. While the Groves–Ledyard mechanism provides a useful balanced incentive-compatible mechanism that solves the free rider problem, it does not guarantee the stability — in an economic sense — required for the problem tackled in this paper [27]. The Clarke-Tax mechanism [11], on the other

hand, is one of the few incentive-compatible mechanisms that maximizes the social utility function by encouraging truthfulness among the individuals, regardless of other individuals' choices [3]. Further, it has the advantage that for each user, equilibrium is a stable dominant strategy, as for the general VCG mechanism. Thus, there are no complications related to stability or multiple equilibria. 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 [9]), "independence of irrelevant alternatives" (removal of any unchosen preference from the set of alternatives will not change the outcome [56]) and that the identity of a voter has no influence on the outcome. As we demonstrate in the following sections, it allows modeling the issue of privacy for shared content in a simple yet effective manner.

6.2 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. In order to create incentives for users to reveal their truth preferences (valuations), some form of monetary payment must be introduced. Therefore, we design a credit-based system where the user earns credits proportional to the amount of data (e.g., pictures, documents) 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:

$$c = m_i + (\beta \times m_i) \times n \tag{1}$$

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 Facebook, 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 1,000 is raised to 100 for posting her picture and of 70×2 for both Bob and John. That is, Alice totals 240 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 $F(.)$ (also known as social welfare function) which outputs a unique outcome, in light of the individuals privacy preference inputs. $F(.)$, known as the *social* function, is a function over the individuals' value functions and outputs a certain collective output X:

$$F(v_1(g), \ldots, v_n(g)) = X \tag{2}$$

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. One simple approach, common in game theory, (due to Nash [43]), is to choose the outcome that maximizes the collective values (utilities). We take this approach, since it satisfies three important properties [17]: 1) it guarantees a relatively fair distribution of the mutually earned utility, 2) it is simple, and 3) it is non-manipulable.

6.3 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.

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 $F(v_1(g), \ldots, v_n(g)) = \sum_{i=1}^{n} v_i(g)$. 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:

$$g^* = \arg\max_{g \in G} \sum_{i=1}^{n} v_i(g) \tag{3}$$

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

u_i	$v_i(g)$ 0	n	∞	$\pi_i(g^*)$
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**	

Fig. 1 Clark tax example

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 $c = (g, \pi_1, \ldots, \pi_n)$ is the value of the g minus the tax numeraire, given by: $u_i(c) = v_i(g) - \pi_i$.

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:

$$\pi_i(g^*) = \sum_{j\neq i} v_j \left(\arg\max_{g\in G} \sum_{k\neq i} v_k(g) \right) - \sum_{j\neq i} v_j(g^*) \tag{4}$$

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 $\pi_i(g^*)$ is 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, $Profile_1.Friends \cup \cdots \cup Profile_n.Friends$. 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 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 $\overline{v}_a$ and $\overline{v}_b$, 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:

$$\pi_a = \begin{cases} 0 & \hat{v}_a < \hat{v}_b \\ \hat{v}_b & \hat{v}_a \geq \hat{v}_b \end{cases} \tag{5}$$

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 as follows:

$$\hat{v}_a = \begin{cases} [0, \hat{v}_b), & v_a < \hat{v}_b \\ \left[max\{0, \hat{v}_b\}, \overline{v}_a\right), & v_a \geq \hat{v}_b \end{cases} \tag{6}$$

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 5 since $11 > 10$, *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 to 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 [17]. 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.

One important assumption of the Clarke-Tax algorithm is that 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 SN 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 [6,54].

Note that the auction is based on the friendship distance, where the auction's outcome is the friendship distance $d = g^*$ from co-owners in involved in the auction. Building on the friendship relationships, the SN can be modeled as a directed graph $G(V, E)$, where the set of vertices V is the set of users and the set of edges E is the set of friendship relationships

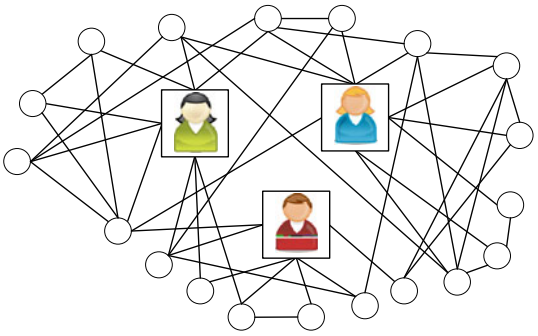

Fig. 2 Example of combined co-owner graph

between users. The edge $(u_i, u_j) \in E$ implies that users u_i and u_j are friends. Given the social graph G and the auction outcome distance d, we are able to extract a subgraph $G_i^d(V_i^d, E_i^d)$ where V_i^d and E_i^d are, respectively, the vertices and edges that are included in the paths that originate from u_i and are of length d or less. This subgraph is easily computed using the breadth first search algorithm [12]. The set of users given access based on the auction's result is given by $G^d(\cup_{i=1}^d V_i^d, \cup_{i=1}^d E_i^d)$. Figure 2 shows an example subgraph centered around three co-owners with auction distance $d = 1$. This approach combines the subgraph computed from each of the co-owners, the issue here is that the there is no unifying graph as the distance is computed from different references. We propose a solution in which the co-owners select a *central common user* from which the distance is computed, this user is to be selected based on the social network metrics such as centrality, betweenness and node degree. This mechanism provides a central point from which the distance is computed and provides a straightforward approach for users to see the effect of their distance selections. Actually, the subgraph G^d approach is a special case of the central common user approach, which assumes that an imaginary common user between only the co-owners and computes breadth first search of distance $d + 1$ from that user.

Note the distance-based approach for policy authoring enables the co-owners to bid on policies for which they can easily compute social utility. The use of the distance-based approach is another form of classifying users into directly connect friends, friend-of-friend, and public. This approach can be further extended to accommodate other forms of policy authoring, such as specifying policies based on common user group. For example, co-owners could specify different groups of friends, such as work, family, school, sports, public, and others. The co-owner preferences g could be formulated as combinations of the specified groups, for example $\{family\}$, $\{family, work\}$, $\{work\}$, and $\{public\}$, and similar to the distance-based approach, the co-owners place their bids on the different group preferences based on the computed social benefit.

7 Automating privacy policies settings

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 shared with others. As such, asking users to bid for each of them may be, in the long run, a cumbersome task. In this section, we discuss a few alternative approaches to help users automate privacy policy settings and free them from the burden of going through the voting process numerous times.

7.1 Inference of privacy policies

An effective idea to avoid users' input for each co-owned image is to utilize inference-based techniques, so as to leverage previous decisions. It is easily verifiable that most users appear in pictures with more or less the same small set of users (typically directly related with each other) and that the sensitivity of a given picture also depends 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 Sect. 3, 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 at each picture or at the album level.[3] 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, $\{t_1, \ldots, t_k\}$. 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 $T = \{\overrightarrow{t}_1, \overrightarrow{t}_2, \ldots, \overrightarrow{t}_n\}$ be a set of pictures shared among the set of owners Own_{Uset}. Let $\overrightarrow{t}$ be the picture whose policy is to be defined. In order to identify the best policy to associate with $\overrightarrow{t}$, we conduct similarity analysis among the pictures in T and $\overrightarrow{t}$.

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 collective tagging. This user-generated classification system is referred to as *folksonomy* [44] 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 $T_{ur} := \{t \in T | (u; t; r) \in Y\}$

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

In our domain, U is the set of users in a social network, while R corresponds to the objects (documents and images) uploaded in the social network site. 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 according to several metrics [35,55]. In our case, we employ the following modified notion of co-occurrence of tags. We consider the tags associated by the co-owners of a given image $\overrightarrow{t}$, $Own^{\overrightarrow{t}}_{USet} \in U$. A couple of tags' weight, say t_1, t_2, is given by its co-occurrence.

$$w(t_1; t_2) := card\{(u; r) \in Own^{\overrightarrow{t}}_{USet} \times R | t_1; t_2 \in T_{ur}\} \quad (7)$$

For a given tag $t \in T$, the tags that are most related to it are thus all the tags $t' \in T$ with $t' \neq t$ such that $w(t, t')$ is maximal. Notice that since we restrict the set of weights to pairs generated by co-owners, we concentrate on the similarities of tags as perceived by co-owners only.

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

$$sim(\overrightarrow{t}, \overrightarrow{t'}) = \sum_{i=1}^{k} \sum_{j=1}^{n} w(t_i, t'_j) \quad (8)$$

Note that similarity is commutative, i.e., $sim(\overrightarrow{t}, \overrightarrow{t'}) = sim(\overrightarrow{t'}, \overrightarrow{t})$. The equation 8 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 $champ = max\{sim(\overrightarrow{t}, \overrightarrow{t'}), sim(\overrightarrow{t}, \overrightarrow{t}_1), \ldots, sim(\overrightarrow{t}, \overrightarrow{t}_n)\}$ 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 $\overrightarrow{t_1}$ and $\overrightarrow{t_2}$, tagged using other freely chosen words. $\overrightarrow{t_1}$ was tagged using *gathering*, *fun*, *game*, while $\overrightarrow{t_2}$ using words *friend*, *beer*, *home*. Let us assume that the $sim(\overrightarrow{t}, \overrightarrow{t_1}) = 100$ and that $sim(\overrightarrow{t}, \overrightarrow{t_2}) = 92$. Since $\overrightarrow{t_1}$ is the most similar to $\overrightarrow{t}$, 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 intake is the same as the one originally spent for the championed picture. If the users do not agree or a picture significantly similar to $\overrightarrow{t}$ is not found, the auction mechanism is proposed to the end-users. Notice that explicit agreement is necessary to avoid the unlikely case of users wishing to override the decision of previous actions. If one of the co-owners wants to control the image, he can theoretically keep adding the same occurrence of tag pairs to multiple images, in the hope to drive process of similarity toward a certain champion image that is associated with the policy of his choosing. Therefore, by requesting an explicit agreement, the other co-owners have a final say on the proposed policy and can decide if the champion policy fits the picture to be protected. A temporary policy chosen among previously adopted ones is then used, until the auction is not taken.

7.2 Collaborative filtering for privacy policies settings

An alternative approach that could be used to help the users' policies setting tasks is to integrate the proposed scheme with collaborative filtering techniques. Collaborative filtering (CF) [38,48,53,57] systems collect ratings from members in a community in order to provide recommendations based on the closest object to the individual's tastes. The input for CF algorithm is a *ratings matrix* containing user profiles represented by ratings vectors, i.e., lists of user's ratings on a set of items. Such ratings can be true (explicit) or can be assumed (implicit). To generate a user's prediction for an item, CF initially computes the degree of similarity neighborhood of K users having the highest degree of similarity with the active user and generates a prediction for a specific item by computing a weighted average of the ratings of the other users in the neighborhood on this item. There exists a wide variety of algorithms for generating recommendations. For example, recommendations can be calculated based on criteria such as others' (or self) ratings, the content being rated, the users' tastes, and history of actions.

In our context, CF can be deployed, either to recommend policies to single stakeholders, or groups of stakeholders (appearing in the same photo).In the case of single user's recommendations, stakeholders' privacy policies can be suggested based on their previous auctions' results. There could be several ways to leverage CF for such purposes. One simple approach can be that of deploying a rating-based system, where policies and bid values are suggested to the stakeholder based on his previous auction history. Aspects such as the common set of co-owners and the auction result (e.g did the user win or did he not win the auction) can be factorized to recommend the best action bid. To derive the user's ratings, this information can be either used as a form of assumed ratings or it can be combined with explicit ratings of the users collected at the end of each auction. While in this context, it is quite intuitive that any user would rate higher a successful auction rather than a lost one, the user may rate auctions with a same outcome differently. For example, users may feel that certain won auctions resulted in a too high numeraire being levied, or they may not care as much about other auctions in which outcome was negative. Hence,

collecting explicit individual ratings may result in a better, more accurate recommendation. In a more sophisticated approach, the filtering criteria can be based on the actual image content and its similarity with the one currently being auctioned.

Regardless of the specific criteria used, each user would be suggested a possible bid amount for a given privacy option, and it would be up to him whether to implement the suggestion or not. That is, even if accurate, the recommendation should require the user's approval, especially in light of the fact that the suggested bidding auction puts the user at risk to loose part of the numeraire. Hence, the auction still needs to be carried out.

In order for the users to completely avoid the auctions and/or any other policy setting task, group recommender systems should be considered. This approach would be similar to the inference technique discussed in the previous section, in that it would leverage users' past decisions and completely eliminate the need of additional input. Group recommender systems are an emerging topic in the recommender system field [34]. However, building a collective recommender that uses the feedback provided by multiple users and which must generate suggestions that satisfy the group introduces a number of interesting challenges, that arise mostly from the need to support multiuser interaction and model both individual and group preferences accurately. In this context, the issue of defining a social value function that describes how the tastes and opinions of individuals affect the group's recommendation is crucial, due to the important nature of the recommendation. One wrong or somewhat unfair recommendation could result not only in unjustified numeraire loss but also cause privacy losses for the involved users. Furthermore, the well-known cold start problem is even more salient here, since if no direct recommendation can be given, the system should be able to supply recommendations for potentially a large number of groups. This issue is similar to the case of no available image for inference (the *champ* image of previous section) that can occur in the case where inference techniques are employed. While some proposals exist [48–50,53] for group-based recommendations, these solutions are quite distant in nature and scope from the problem discussed in this paper. Hence, the applicability of group recommender systems in the context of collective privacy is yet to be investigated.

8 System implementation and experimental results

We have implemented a proof-of-concept social application of the proposed approach for the collective management of shared data, referred to as *PrivateBox*. *PrivateBox* is fully integrated with Facebook social network platform [19]. *PrivateBox* 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. We discuss the system's implementation and our performance evaluation in the remaining of the section, followed by a discussion on cross-site enforcement of collective policies.

8.1 Private box implementation and evaluation

PrivateBox 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 *PrivateBox* 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 *PrivateBox*. Figure 3 represents the interaction flow of a user with *PrivateBox* 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, *PrivateBox* determines the set of potential co-owners of the photos based on the id-tags, as described in Sect. 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 *PrivateBox*, including the pictures added by him/her, and those have been added into *PrivateBox* when the user was granted ownership.

Finally, the `Auction` page is the core of the application, and it enables the collective enforcement of privacy policies on co-owned data as it is described in Sect. 5 (see Fig. 3). `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 *PrivateBox* 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

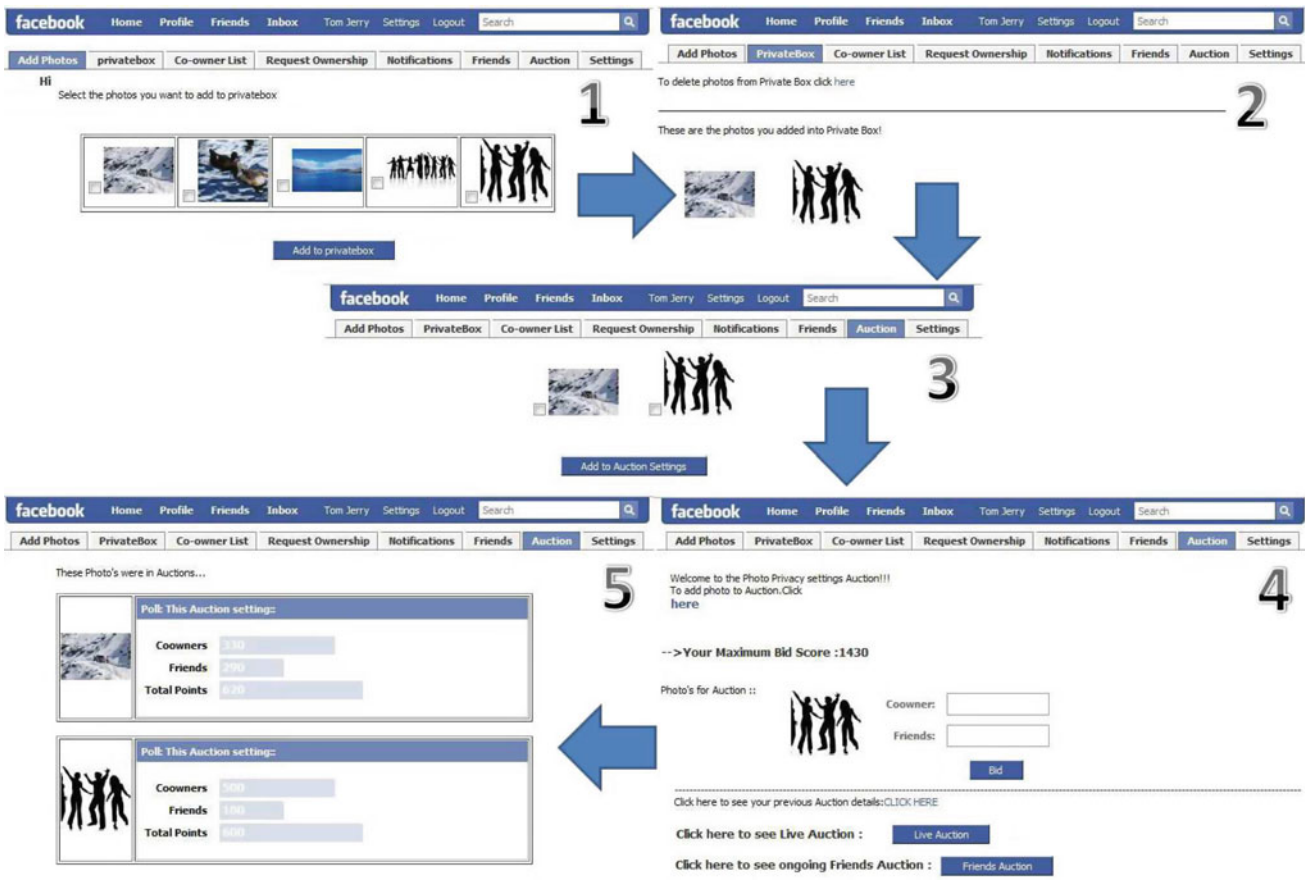

Fig. 3 *PrivateBox*'s execution flow

the photo to avoid that any of the co-owners privacy preferences are violated before the privacy setting that maximizes the social utility $F(.)$ 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 (Fig. 3, step 5).

PrivateBox 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 *PrivateBox*. 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 and 4 faces per photo [14,46]. We have evaluated the scalability of the collective 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 4 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 in 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 execution time is so fast that the

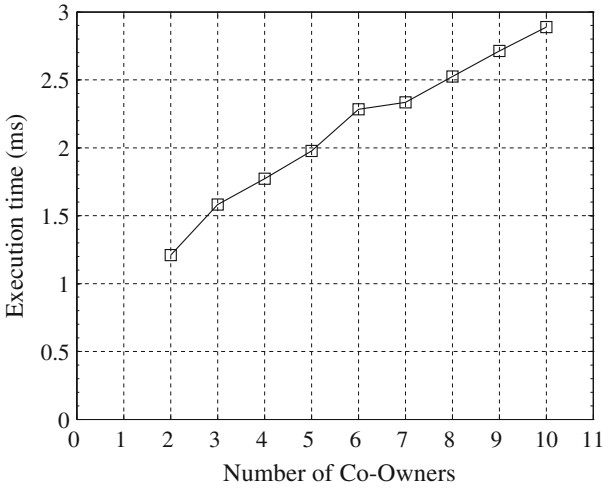

Fig. 4 Auction execution times

collective enforcement of privacy policies is transparent to the user.

8.2 Cross-site enforcement of co-ownership

Co-ownership could involve users who are not social network members or who belong to different sites. In order to guarantee a fair approach, also this class of potential owners should be involved in the auction process, so as to have a role in the decision process. To support this, level of interoperability several design and technical challenges arise. First, a major issue is how to control and communicate the non-social network members about the potential ownership. Clearly, contact information of some sort should be provided. A second important concern deals with the numeraire. Clearly, non-member users do not have any images posted (or very few, if they are already being tagged in other images with the focus social network); therefore, the system may not be fair to them. Additionally, some architectural challenges arise, such as where to store the data for non-member users and how to retrieve it as needed. To address the above concerns, we focus on the possibility of supporting cross-site interactions, that is, to allow users to take part of an auction where they appear as stake holders, even if the other participants belong to different social network domains. Our solution relies on a cross-site manager that allows information mapping among distinct sites.

Our solution relies on a cross-site manager that allows information mapping among distinct sites. The cross-site manager can be implemented as using the Google Open Social API initiative[26], which denes a common API for social applications across multiple websites by generating a mapping between the APIs and a general common API. We chose to use the OpenSocial API to utilize the API mappings to enable the interaction between social sites. The main functions of the Open Social API are to validate a user ID and the image and parse these values to the original web- site where the auction should take place. The validation involves identifying the prole owner and sending a version of the user ID. We use the make Request function in the gadgets API to manage the authorization and key mappings of the different user IDs. The Open Social APIs functions are used whenever an image shared policy is to be created, and the co-owner is possibly be part of a different site. A query is sent to verify whether the co-owner belongs to the linked site. Again, this is carried out by the make Request function. If the co-owner is found, then the user ID is sent out to the database storing the mappings of the different IDs of a single user on the different SNs. This additional database is used for matching users proles, that is, to map User U1 in the SNs A as User U2 in the SNs site B. The ID mapping function constitutes a mapping engine that sends out the corresponding user ID to the original website hosting the image. This engine, along with the database, constitutes our trusted reference monitor. The friends (or other such equivalent relationships) of the user for each and every website are imported every time a policy is being enforced using the dataRequest.PeopleRequest functions.

9 PrivateBox user study

Participants Participants were recruited from a large U.S. university community (staff, students, faculty, and the community at large), including users and non-users of social network sites. One hundred and twenty-two participants (out of 440) agreed to participate in our study (27.7% response rate). Three participants were excluded as they completed less than 25% of the questionnaire items; all other participants completed at least 90% of the items. The average age of the respondents was 23.79 (SD = 5.69) years old (Range: 18-39). Participants were asked to indicate any social networks they were a part of 99.2% indicated Facebook, 18.6% indicated MySpace, 13.6% indicated Linkedin, 3.6% indicated Orkut, and less than 1% indicated Friendster and LiveJournal. Less than 1% reported that they were not part of any social networks. In terms of network usage frequency, 93.1% of the respondents accessed social network sites at least once a week, with 75.9% of reporting that they were daily users. Participants were also asked to indicate how many minutes they were willing to spend in configuring privacy settings. The responses ranged from 0 to 60 minutes, with an average of 13.23 min (SD = 12.34).

Procedure Participants were shown an animated sequence describing the concept of shared ownership using the PrivateBox application for Facebook, the auction approach, and

three auction scenarios illustrating various features of the auction. Participants did not interact with the application directly, as the goal of the study was not to examine the interface or assess the intention of use. In the study, we first presented a usage scenario to illustrate the various features of the PrivateBox application and the auction. The following concepts were elaborated on in the scenario:

- *PrivateBox*: The participants were shown snapshots of the PrivateBox application, explaining how photos can be added to the application and how users can request co-ownership of pictures that they appear in.
- *Auction*: The credit system and the auction process were described to the participants, indicating that the winning bids would determine who was able to view the co-owned pictures. Three simulated auction users (Alex, Tom, and Bob) stated their privacy preferences for individual pictures. The participants watched the users submit bids for privacy settings under 3 different scenarios. The three scenarios included a range of behaviors from the users to show a variety of auction outcomes.

After learning about the various features of PrivateBox and the auction approach, participants were asked to complete a web-based questionnaire assessing their social network usage and attitudes related to privacy. The animated sequence took approximately 6 min to view after which they were directed to complete the survey. Completion of the survey took an additional 8–10 min.

Measures The following measures were considered as part of the study.

Concern with privacy was measured using 3 items ($\alpha = 0.72$) rated on a Likert scale (5-point rating scale, where 1 = strongly disagree and 5 = strongly agree). An example item is 'I have had concerns about the privacy of my data on Social Networks.' *Frequency of social network use* was measured on a frequency rating scale (1 = never; 2 = once or a few times a day; 3 = once or a few times a week; 4 = once or a few times a month) with the item 'How often do you access Social Network Sites?'

Minutes willing to spend configuring privacy was measured with a unrestricted free-report response based on the question 'how many minutes are you willing to spend in configuring privacy settings?'

Comprehensive profile was measured on a Likert Scale (5-point rating scale, where 1 = strongly disagree and 5 = strongly agree) with the item 'I have a comprehensive Social Network profile.'

Understanding of co-ownership was measured using 3 items ($\alpha = 0.69$) rated on a Likert scale. An example item is 'I have trouble understanding co-ownership of privacy settings.'

Table 1 Descriptive statistics of study variables

Variable	Mean	St. Dev.
Subscales		
Concern with privacy	3.91	0.89
Understanding of co-ownership	3.86	0.83
PrivateBox as a privacy-enhancing technology	3.98	0.68
Usefulness of the auction	3.94	0.70
Understanding of the auction	3.71	0.70
Fairness of the auction	3.76	0.76
Single items		
Frequency of application use (4pt. scale)	3.77	(0.54)
I have a comprehensive social network profile	3.83	(0.97)

PrivateBox understood as a privacy-enhancing technology was measured using 3 items ($\alpha = 0.76$) rated on a Likert scale. An example item is 'Using the PrivateBox application would enhance my control over my photos.'

Usefulness of the auction was measured using 3 items ($\alpha = 0.82$) rated on a Likert scale. An example item is 'The auction is a useful way to express my privacy settings.'

Understanding of the auction was measured using 3 items ($\alpha = 0.81$) rated on a Likert scale. An example item is 'The auction is easy to understand.'

Fairness of the auction was measured using 3 items ($\alpha = 0.78$) rated on a Likert scale. An example item is 'The auction is a fair approach to decide the privacy settings of shared pictures'.

Results Table 1 presents the descriptive statistics of the study variables. The main purpose of the study was to assess participants understanding of shared privacy settings and their understanding of the auction process. The subscales contained in the questionnaire show that participants understood the idea of co-ownership (M = 3.85, SD = 0.83) and additionally thought that shared privacy settings would enhance users control over data (M = 3.98, SD = 0.68). Furthermore, participants indicated that the auction approach was both useful (M = 3.93, SD = 0.70) and fair (M = 3.74, SD = 0.76).

In addition to assessing participants' understanding of the auction approach and whether the participants viewed the approach as fair, we were interested in understanding the factors affecting users' attitudes toward the idea of the auction. Specifically, we examined whether understanding of the auction approach is predicted by participants concern with privacy, understanding of shared privacy, usefulness of the auction approach, frequency of social network use, age and minutes willing to spend to configure privacy settings. We conducted an exploratory least squares multiple regression analysis, regressing understanding of our approach

simultaneously to all the possible predictors. Not surprisingly, the most significant predictor of the understanding of the auction approach was the usefulness of the auction ($\beta = 0.382$, $p < 0.01$). Participants who found the auction approach useful indicated a better understanding of the auction approach. Understanding of co-ownership also predicted understanding of the auction approach ($\beta = 0.245$, $p < 0.05$). Participants who reported greater understanding of shared privacy also reported greater understanding of the auction approach. Finally, participants' age also predicted the understanding of the approach ($\beta = -0.206$, $p < 0.05$). Specifically, the younger participants were, the more they understood the auction process. Frequency of application use, concern with privacy, and minutes willing to spend to configure privacy settings were not significant predictors of auction understanding.

Furthermore, we were interested in understanding why users thought the auction approach was fair. We conducted a least squares multiple regression analysis, regressing fairness of the auction approach on age, concern for privacy, understanding of co-ownership, and usefulness of the auction approach. Participants' understanding of co-ownership was the only significant predictor of fairness of the auction approach ($\beta = 0.367$, $p < 0.01$). If participants understood co-ownership, they viewed the auction approach more fairly. Age, concern for privacy, and usefulness of the auction approach were not significant predictors of fairness of the auction approach. Users who did not understand the co-ownership had a hard time understanding the need of such an approach. These results indicate that one of the key components to an auction-based approach is whether or not users understand the concept of co-ownership which is essentially the reason behind the need of an auction. Understanding of co-ownership significantly predicted both understanding and fairness of the auction approach. Given the positive feedback we obtained when assessing the users' understanding of co-ownership, we believe the PrivateBox represents a good implementation of the auction.[4]

10 Manipulating the auction-based privacy system

Generally speaking, the problem of voting becomes complicated when some of the voters try to manipulate the outcome by expressing false preferences. Voters might cast bogus votes to help a preferred alternative win or to avoid an undesired alternative being chosen. The Gibbard–Satterthwaite theorem[24,60] proves that any choice rule that satisfies no dictatorship and universal criterion will be susceptible to manipulation when there are three or more alternatives. The Clarke-Tax approach, which is at the core of our solution, is shown to alleviate the problem of manipulation, since it reduces the incentive for a voter to manipulate the voting outcome. There are, however, some aspects of concern, which we now briefly discuss.

As any voting schema, our system requires that the privacy of the bidders is guaranteed, and that the voting actions are verifiable. While such requirements are inherently met by the design of our solution (the design of the PrivateBox is such that users do not see the others' bids and their actions are logged for verification), it is still possible for malicious users to bypass the auctions' rules. For example, malicious users may try to manipulate the system to prevail in the collective decisions in at least two possible ways: (1) first, they can refuse to add certain stakeholders who most likely will not share their same privacy opinions and (2) second, they can try to gain large amount of illegitimate numeraire to make their bids pivotal, preventing honest users' privacy preferences from being respected fairly.

Concerning the first issue, we begin with highlighting that even without considering coercive measures, we notice that our design is such that (according to equation (1)) every user has strong incentives in allowing as many potential co-owners as possible. Therefore, limiting the number of co-owners is actually a discouraged behavior that users will unlikely take. Second, in order to actually prevent these cases from occurring, it is possible to extend the proposed approach as follows. Id-tags can be validated using facial recognition techniques. Although this may add some overhead to the system, automated detection of potential stakeholders does not allow other co-owners and/or stake holders to deliberately leave out other entitled users from the stakeholders' list, making it unlikely that they will be able to participate in the auction. Automated tagging is provided by several lightweight tools[40], and their integration in social network platforms has been already successfully tested [62]. Forcing users to actually identify the potential players in the auction game addresses the issue of stakeholders' being denied to join the privacy setting procedure.

Issue (2) leads to a general observation related to the fairness of the auction. Ensuring that each and every auction is fair is a controversial issue, since it is not possible to ensure fairness while maintaining an incentive-based system. There are, however, some extensions that can be investigated to mitigate this issue.

Due to the possibly sensitive nature of certain shared content, users should be given a chance to always express privacy preferences on content that is related to them. However, in the current implementation, it is safe to assume that if users choose not to collect numeraire, it can be assumed that they do not have a strong interest in bidding in an auction. On

[4] We believe better results could be obtained by improving the prototype design, and having users interact with the application, however, user studies related to the prototype design are beyond the scope of this work.

the other hand, it is plausible that some users may prefer not to upload images or other documents to their SN profile for privacy reasons. Hence, one possible approach to avoid situations where users feel forced to expose data in order to gain numeraire is to increase the ways such numeraire can be gained, allowing users to earn more credits based on other SN interactions (such as posting comments, tagging content etc). Further, adjustments to the approach used to calculate the maximum bid within an auction can be taken. One adjustment that could be made is limit bids up to a certain numeraire, for example, by limiting bids to 150% of the lowest numeraire available among all the co-owners. While users with a lower numeraire can still be outbid, they are now given some bidding power (especially in light of the fact that multiple users may bid on the same option, leading to a high collective numeraire), leaving the incentive system still valid. Further, by extending the ways users can gain numeraire, upon the request of participating in an auction, they can gain the credits required to compete.

11 Conclusion

In this paper, we discussed a novel model for privacy management within social networks, where data may belong to many users. We presented a theoretical representation of the collective privacy management problem and proposed a solution that builds upon well-known game theoretical results. We implemented a tool prototype hosted in Facebook and carried out a performance analysis and a user study evaluation. The issue tackled in this work is quite challenging, and our solution, although promising, has room for improvement and can be extended in several ways.

On the policy language front, we will explore more sophisticated and flexible policies, such as network-based policies, to include predicates related to the users' geographic locations. Further, we plan on taking into account time dependency of privacy requirements. As users' privacy preferences may change over time, our approach should be able to reflect this if necessary. Currently, users can choose to undo the auction and update the privacy preferences. However, more sophisticated and effective approaches are required to incorporate temporal constraints in the policies themselves. To this extent, we will investigate further the implications of our approach in case of revocation or leave of some co-owners. On the architectural side, we would like to improve the design of Private Box to make it more user-friendly and evaluate its predicting capabilities using the inference techniques discussed in the paper. Finally, we plan to extend our analysis concerning the systems manipulation by elaborating on colluding users.

Acknowledgments 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.

References

1. Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. IEEE Secur. Priv. Mag. **3**(1), 26–33 (2005)
2. Acquisti, A., Gross, R.: Imagined communities: Awareness, information sharing, and privacy on the facebook. In: Proceeding of Privacy Enhancing Technologies, pp 36–58. Springer (2006)
3. Bartal, Y., Gonen, R., Nisan, N.: Incentive compatible multi unit combinatorial auctions. In: Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge, ACM. pp. 72–87 (2003)
4. Beaver, D.: 10 billion photos. http://www.facebook.com/note.php?note_id=30695603919, October (2008)
5. Bonneau J., Preibusch, S.: The privacy jungle: On the market for data protection in social networks. In the eighth workshop on the economics of information security (WEIS 2009) (2009)
6. Borgatti, S.P., Everett, M.G.: A graph-theoretic perspective on centrality. Soc. Networks **28**(4), 466–484 (2006)
7. Carminati, B., Ferrari, E.: Privacy-aware collective access control in web-based social networks. In DBSec, pp. 81–96 (2008)
8. Carminati, B., Ferrari, E., and Perego, A.: Rule-based access control for social networks. In OTM Workshops (2), pp. 1734–1744 (2006)
9. Chen, L., Den, X., Fang, Q., Tian, F.: Condorcet winners for public goods. Ann. Oper. Res. **137**, 229–242 (2005)
10. Clarke, E.H.: Multipart pricing of public goods. Public Choice **11**, 17–33 (1971)
11. Clarke, E.H.: Multipart Pricing of Public Goods: an example. In public price for public products, urban inst (1972)
12. Cormen, T.H., Leiserson, C.E., Rivest, R.L.: Clifford Stein Introduction to Algorithms (3rd ed). MIT Press, Cambridge (2009)
13. Cox, I.J., Kilian, J., Leighton, T., Shamoon, T.: Secure spread spectrum watermarking for images, audio and video. In: Proceedings of International Conference on Image Processing, IEEE. pp. 243–246 (1996)
14. Davis, M., Smith, M., Canny, J., Good, N., King, S., Janakiraman, R.: Towards context-aware face recognition. In: Proceedings of the 13th Annual ACM International Conference on Multimedia, ACM, pp. 483–486 (2005)
15. Enterprise, C. F., Josang, A., Pope, S.: Auscert conference 2005. In in Asia Pacific information technology security conference, AusCERT2005, Austrailia, pp. 77–89 (2005)
16. Ephrati, E., Rosenschein, J.-S.: The Clarke-tax as a consensus mechanism among automated agents. In national conference on artificial intelligence, pp. 173–178 (1991)
17. Ephrati, E., Rosenschein, J.-S.: Voting and multi-agent consensus (1991)
18. Ephrati, E., Rosenschein, J.S.: Deriving consensus in multi-agent systems. J. Artif. Intell. **87**, 21–74 (1996)
19. Facebook. Facebook web site. http://www.facebook.com/
20. Felt, A.: Defacing Facebook: A security case study. Technical report. University of Virginia, Charlottesville (2007)
21. Felt, A., Evans, D.: 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)
22. Gates, C.: Access control requirements for Web 2.0 Security and Privacy. In IEEE Web 2.0 privacy and security workshop (2007)

23. Geambasu, R., Balazinska, M., Gribble, S.-D., Levy, H.-M.: Homeviews: peer-to-peer middleware for personal data sharing applications. In SIGMOD conference, pp. 235–246 (2007)
24. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica **41**(4), 587–601 (1973)
25. Gollu, K. K., Saroiu, S., Wolman, A.: 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)
26. Google open social api available at: http://code.google.com/apis/opensocial/
27. Greenberg, J., Mackay, R., Tideamn, N.: Some limitations of the Groves-Ledyard Optimal mechanism. Public Choice **29**(2), 129–137 Springer (2005)
28. Gross, R., and Acquisti, A.: Information revelation and privacy in online social networks. In workshop on privacy in the electronic society (2005)
29. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? a game-theoretic analysis of information security games. In World Wide Web Conference pages 209–218 (2008)
30. Herlocker, J.: Evaluating collaborative filtering recommender systems. ACM Tran. Inf. Syst. **22**(1), 5–53 (2004)
31. Groves, T.: Incentives in teams. Econometrica **41**, 617–631 (1973)
32. Hart, M., Johnson, R., and Stent, A.: More content - less control: access control in the Web 2.0. In IEEE Web 2.0 privacy and security workshop (2007)
33. Hobgen, G.: Security issues and recommendations for online social networks. ENISA. Pos. Paper N. 1 (2007)
34. Jameson, A.: More than the sum of its members: challenges for group recommender systems. In Working Conference on Advanced Visual interfaces, ACM, (2004)
35. Jiang, J., Conrath, D.: Semantic similarity based on corpus statistics and lexical taxonomy. In: Proceedings of ROCLING X Sep (1997)
36. Josang, A., Zomai, M. A., Suriadi, S.: Usability and privacy in identity management architectures. In: ACSW '07: Proceedings of the Fifth Australasian Symposium on ACSW Frontiers, pp. 143–152, Darlinghurst, Australia, Australia, (2007). Australian Computer Society, Inc
37. Krishna, V.: Auction Theory. 1st edn. Academic Press, Elsevier (2002)
38. Linden, G., Smith, B., York, J.: Amazon.com recommendations: item-to-item collaborative filtering. IEEE internet computing, pp. 76–80, January/February (2003)
39. Lenhart, A., Madden, M.: Teens, privacy & online social networks. Pew internet & American life project, 18 April (2007)
40. Lowensohon, J.: Facebook's auto-tagging features could be tip of tagging icerberg. CNET News. http://news.cnet.com/8301-17939_109-10004835-2.html. August (2008)
41. Maliki, T. E., Seigneur, J.-M.: A survey of user-centric identity management technologies. In SECUREWARE '07: Proceedings of The International Conference on Emerging Security Information, Systems, and Technologies, IEEE Computer Society, Washington, DC, USA, pp. 12–17 (2007)
42. Mannan, M., van Oorschot, P.-C.: Privacy-enhanced sharing of personal content on the Web. In WWW, ACM, pp. 487–496 (2008)
43. Mas-Colell, A., Whinston, M.D.: Micro-Economic Theory Chapter 23. Oxford University Press, Oxford (1998)
44. Mathes, A.: Folksonomies: cooperative classification and communication through shared metadata. http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html (2004)
45. Miller, G.A.: Wordnet: a lexical database for english. Commun. ACM **38**(11), 39–41 (1995)
46. Naaman, M., Yeh, R.B., Garcia-Molina, H., Paepcke, A.: Leveraging context to resolve identity in photo albums. In: Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital libraries, pp. 178–187, ACM Press (2005)
47. Ellison, C.L.N.B., Steinfield, C.: Benefits of Facebook "Friends": social capital and college students' use of online social network. J Comput Mediat Commun-Electron (2007)
48. McCarthy, J., Anagnost, T.: MusicFX: An arbiter of group preferences for computer supported collective workouts. In: Proceedings of the 1998 Conference on Computer-Supported Cooperative Work, pp. 363–372, (1998)
49. McCarthy, K., Salam, M., Coyle, L., McGinty, L., Smyth, B., Nixon, P.: group recommender systems: a critiquing-based approach. IUI 2006: international conference on intelligent user interfaces, pp. 267–269. ACM Press (2006)
50. McCarthy, K., Salam, M., McGinty, L., Smyth, B.: CATS: A synchronous approach to collective group recommendation. In: Proceedings of the Nineteenth International Florida Artificial Intelligence Research Society Conference, Melbourne Beach, FL (2006)
51. Minr, S., Magnusson, B.: A model for semi-(a)synchronous collaborative editing. In: Proceedings of the Third Conference on European Conference on Computer-Supported Cooperative Work, pp. 13–17 (1993)
52. Norberg, P.-A., Horne, D.-R., Horne, D.-A.: The privacy paradox: personal information disclosure intentions versus behaviors. J. Cons. Aff (2007)
53. O'Connor, M., Cosley, D., Konstan, J., Riedl, J.: PolyLens: A recommender system for groups of users. In: Proceedings of the Seventh European Conference on Computer-Supported Cooperative Work, Kluwer, Dordrecht (2001)
54. Newman, M.-E.-J.: Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. Physical Review E **64**(1), 016132+ (2001)
55. Pirro', G., Seco, N.: 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)
56. Ray, P.: Independence of irrelevant alternatives. Econometrica **41**, 987–991 (1973)
57. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., and Riedl, GroupLens, J.: an open architecture for collaborative filtering of netnews. In ACM conference on computer supported cooperative work. ACM, Chapel Hill, NC, pp. 175–186 (1998)
58. Manuel Romero Salcedo: Dominique Decouchant, structured cooperative authoring for the World Wide Web, computer supported cooperative Work 6(2–3):157–174 (1997)
59. Rosenblum, D.: What anyone can know: the privacy risks of social networking sites. IEEE Secur. Pri. **5**(3), 40–49 (2007)
60. Satterthwaite, M.A.: Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J. Econ. Theory **10**, 187–217 (1975)
61. Spiekermann, S., Grossklags, J., Berendt, B.: E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. In EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce. ACM. pp. 38–47 (2001)
62. Stone, Z., Zickler, T., Darrell, T.: Autotagging facebook: social network context improves photo annotation, computer vision and pattern recognition workshops, pp. 1–8 (2008)
63. Sun, C., Ellis, C.: Operational transformation in real-time group editors: Issues, algorithms, and achievements. In conference on CSCW, ACM, pp. 59–68, Seattle (1998)
64. Varian, H.R. (2002) System Reliability and Free Riding. In Economics of Information Security. Kluwer Academic Publishers, pages 1–15
65. Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. J. Financ., p. 8–37 (1961)
66. Vidot, N., Cart, N.M., Ferrić4, J., Suleiman, M.: Copies convergence in a distributed real-time collective environment.

In: Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work, ACM, pp. 171–180 (2000)
67. Wang, C., fung Leung, H.: A secure and private Clarke-tax voting protocol without trusted authorities. In: Proceedings of 6th International conference on Electronic Commerce, ACM, pp. 556–565, New York, NY, USA (2004)
68. Watson, J. (2008) Strategy, an introduction to game theory. Second Edition, Norton Publisher
69. Wu, X., Zhang, L., Yu, Y.: Exploring social annotations for the semantic Web. In World Wide Web conference, ACM, pp. 417–426 (2006)
70. Yao, M.Z., Rice, R., Wallis, E.K.: Predicting user concerns about online privacy. Am. Soc. Inf. Sci. Technol. **58**(5), 710–722 (2007)