Collaborative Privacy Management with Auctioning Mechanisms

Onuralp Ulusoy and Pınar Yolum

Abstract Online social networks enable users to share content with other users. Many times, a shared content, such as a group picture, may reveal private information about the uploader as well as others who are associated with the content. Ideally, protection of privacy in such cases would need to consider the privacy concerns of all relevant individuals. However, these concerns might conflict and satisfying one user's privacy needs could cause a privacy violation for others. This calls for computational mechanisms that can decide on the privacy policies of the content collaboratively. Accordingly, we propose an agent-based collaborative privacy management model for online social networks (OSNs). Agents represent OSN users and manage their privacy requirements on their behalf. We extend Clarke-Tax mechanism for auctioning to achieve fair handling of privacy settings and to tax the agents whose privacy settings are chosen. We evaluate our approach over multi-agent simulations and show that it produces privacy policies efficiently and more accurately than existing approaches.

1 Introduction

Privacy is one of the most important notions that should be taken into consideration while using online social networks (OSNs). Even a seemingly unimportant content sharing might have drastic effects on the life of the publisher, or in some cases, other people that are related to the content. With the tremendous worldwide usage of the OSNs, these violations might occur frequently. Hence, privacy resolution mechanisms are essential, especially for the published contents that concern multiple

O. Ulusoy (✉) · P. Yolum
Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands
e-mail: o.ulusoy@uu.nl

P. Yolum
e-mail: p.yolum@uu.nl

© Springer Nature Singapore Pte Ltd. 2021

T. Ito et al. (eds.), *Advances in Automated Negotiations*, Studies in Computational Intelligence 905, https://doi.org/10.1007/978-981-15-5869-6_4

people. Collaborative privacy management regulate such settings by enabling all the affected users to have their say in publishing or not publishing a content.

Reaching a consensus in collaborative privacy management is usually not an easy task, since people's privacy requirements can easily be in conflict. In real-life cases, the resolution requires much time and effort through mechanisms like negotiation, argumentation, and so on. However, current most widely used OSNs only enable their users to tweak with their own privacy settings while publishing contents. Since they can't provide collaborative solutions, people tend to resolve conflicts via different media [6] or in most cases, ignore others' privacy requirements and cause voluntary or involuntary privacy violations. In order to be able to handle them online, decision mechanisms should be in place. Recently, multi-agent agreement techniques have been used to address collaborative privacy management. Kekulluoglu et al. [4] and Such and Rovatsos [11] propose negotiation-based approaches that enable users to reach a consensus on how to share a content. Kokciyan et al. [5] use argumentation to enable one user to persuade the other into sharing with her own privacy constraints. These approaches have been successful but require heavy computations; that is, they can only be used when the entities can reason on its privacy policies and communicate with others intensively.

On a different line, Squicciarini et al. [9] propose a model where users enter auctions for deciding on a policy that requires collaborative management over a content. Each user creates bids based on how much she wants to see a content public or private. In that approach, users collect currencies by publishing content and tagging people that are related to content. These currencies are used in an auction, where users spend their currencies to convince other users to accept their policy, based on Clarke-Tax mechanism [1, 2]. However, with its current state, the system is open to abuse by the users, such that a single user's privacy can be ignored repeatedly when all others collaborate strongly. Further, since the bids are expected to be generated by users individually for each post, it is difficult to apply them in real-life online social networks.

Ensuring collaborative privacy management in a real-life scale OSN requires the system to scale to the tremendous amount of content being shared. To enable this, first, the operations expected from users should be handled automatically so that users do not need to think through the operations for each content. Second, the proposed mechanism should be easy to compute, because it will be repeated for each content separately. Third and most importantly, it should preserve privacy of the users fairly so that no user is left at a disadvantage.

Accordingly, this paper proposes PANO an agent-based collaborative privacy management system that uses ideas from the work proposed by Squicciarini et al. [9]. There are three main contributions of PANO: First, it employs agents for privacy management, where agents act on behalf of users to enforce their privacy constraints, so that heavy user involvement is reduced to minimum. The agents manage their users' privacy constraints and bid on behalf of them. Second, it contains a fair reward mechanism, which is protective against abuses, and at the same time encourages users to share content online. Third, it works with a group-wise currency system in auctions, where the agents cannot use the advantages they gain from the system against individuals. This disables an agent to abuse another agent's privacy. Our experimental evaluation shows that agents can indeed carry out this task successfully and help preserve their users' privacy with high accuracy.

The rest of this paper is organized as follows: Sect. 2.1 gives the necessary background for Clarke-Tax mechanism. Section 2.2 explains our agent-based approach PANO to handle privacy management. Section 3 presents the agent bidding mechanism and the proposed policy, bidding and robustness strategies to construct PANO. Section 4 explains the rules and process of the simulation implemented to evaluate our approach. Section 5 evaluates the success of PANO. Section 6 discusses our work in relation to existing methods in the literature and gives pointers for future work.

2 Auctioning Privacy

As a broad definition, privacy is the concept of individuals deciding on how much about themselves to be shared with the others. In OSNs, these decisions can be represented with privacy policies. Applying privacy policies when the information is solely related to an individual itself is an easy task, when the necessary tools are provided. However, a piece of information, e.g., a photograph content in an OSN, can be related to more than one individual. In such cases, the decisions of the individuals for the extend of how much to share may differ, resulting in conflicts. These conflicts require some resolution mechanism to define a generalized privacy policy with the goal to comply with every individual's privacy requirements. For this conflict resolution mechanism, we propose PANO, which employs Clarke-Tax mechanism with an agent-based approach.

2.1 Background: Clarke-Tax Mechanism

Clarke-Tax mechanism [1] provides an auction mechanism, similar to English auctions where participants bid for different, possible actions in the environment. The action that receives the highest total bids from the participants wins and is executed. Different from an English auction, participants who aid in the winning action to be chosen, i.e., that bid toward it, are taxed according to the value they put on it. This is achieved by subtracting the bid values of every single user from the overall values. If the subtraction of a single user's bid changes the overall decision, it shows that the user's bid on this action had a *decisive* value. Thus, the user is taxed with the difference of the actual action's score and the score of action to be taken if that user was not present in the auction.

In the context of collaborative privacy, Clarke-Tax mechanism is used to decide on how an image is going to be shared. Squicciarini et al. [9] consider three types of sharing actions: *no share*, *limited share*, and *public share*. We follow the same scheme here. When an image is about to be shared, all the relevant participants bid on

Table 1 Four user bids for sharing an image

Users	No share	Limited share	Public share
Alice	3	5	0
Bob	15	2	0
Carol	5	8	5
Dave	2	6	18

Table 2 Clarke-Tax mechanism example—decision and taxes

Values	No	Limited	Public	Taxes
Overall	**25**	21	23	–
Without Alice	22	16	**23**	1
Without Bob	10	19	**23**	13
Without Carol	**20**	13	18	0
Without Dave	**23**	15	5	0

these three possible actions. Table 1 shows an example of biddings of four users for deciding to share or not share a content. Users decide based on their own importance of the three actions. According to Table 1, it can be seen that Bob values the *No Share* action more than the others, while Dave values *Public Share* action the most.

According to the biddings of all users, Clarke-Tax auction mechanism decides on which action to take. Based on the bids from Table 1, no share action receives a total of 25, whereas limited share receives a total of 21, and public share a total of 23 points. Therefore, the no share action is chosen. Table 2 shows the resulting decision and applies the taxes according to the biddings in Table 1. According to the decided action, each user that bid for the decisive action is taxed.

Table 2 shows that Alice and Bob are taxed, since each of these users' absence in the auction causes the decisive action to be changed. When the scores of Alice is subtracted from the overall score, the decision of sharing the content receives the maximum score by 23, while not sharing gets a score of 22. This causes Alice to be taxed with a score of 1. As mentioned above, Bob bid a greater value for not sharing, and its absence from the auction also causes the action to be changed. Since the differences of the actions are much bigger in Bob's case (i.e., 13, obtained from the subtraction of *Public Share* and *No Share* scores), Bob is taxed with a much greater value. It is important to note that although the user is taxed, he gets the action to be decided what it values most, and prevents the content from being shared. The importance of good evaluation and truthfulness for bidding are crucial. For example, if Bob bid for not sharing with a rather big value, even though the decision was not that important to him, it would have paid the bid amount plus a great amount of tax. On the contrary, if he had bid for much less, then his privacy might have been violated. Hence, it is important to be able to create bids that reflect the true evaluations of the users.

2.2 Challenges

Auctioning with Clarke-Tax Mechanism is an efficient way of negotiation, since it has been shown that truthfulness is the best strategy for bidding [9]. The bidder who overvalues a decision to get its way can be taxed with a greater amount, because it changes the group decision by spending system currency way more than the other participants of the auction. This results in the participants bidding with truthful values, while trying to establish its own decision and not get taxed with a greater amount. Even with the notion of truthfulness, applying a pure Clarke-Tax approach still has some limitations that can result in abuse by the bidders or inflation in the currency used. Consider the following examples:

Example 1 Dave wants to share a photograph in which he is with Alice, Bob, and Carol. Alice and Carol do not want the picture to be shared. Dave has previously shared a lot of photographs that are unrelated to Alice, Bob, and Carol and gained substantially more currencies than the others. Since Dave can spend more than the others, he puts a high bid (e.g. >50 for the bids in Table 1), instead of the 18 that was bid for sharing in Table 1, to share the photograph, which bids of Alice, Bob and Carol cannot match. Therefore, the photograph is shared in the social network even though the majority of the people in the photograph did not want it to be shared.

Another drawback of the application of Clarke-Tax mechanism over OSNs is that it requires user involvement for every single auction. This could be necessary for some cases, but it could become a tedious task for greater number of contents. Also, unavailability of some users in auctions who are related to a content that has privacy conflicts could make the method get stuck, or decide on a semi-successful policy. To resolve these issues, an automated auction process where software agents represent the users and act on their behalf can be implemented.

Example 2 Alice, Bob, Carol, and Dave want to decide on sharing or not sharing a photograph they are in, in a social network. Carol shares a lot of photographs daily, has a lot of currency to spend, and does not want to join a Clarke-Tax auction, but still wants the photograph to be shared. Alice also doesn't mind if the photograph is shared or not, and doesn't enter the auction. Bob and Dave have a disagreement, and join into an auction, and Bob wins the auction for not sharing the photograph, by outbidding Dave. The final policy is decided for not sharing the photograph, even though majority of the people in the photograph wanted it to be shared.

3 Agent-Based Bidding

The pure Clark-Tax-based mechanism in Squicciarini et al. [9] requires user involvement for auctions. This could become a tedious work for the OSN users that shares a multitude of contents every day. Thus, we develop an agent-based approach, where each user is represented with an agent that maintains its user's privacy constraints, manages total currencies, and generates bids when necessary. In principle, understanding users' privacy constraints automatically is difficult. It would require the user behavior to be modeled and privacy constraints to be learned over time. There is a good body of literature on learning privacy constraints [8, 10]. Here, we assume that the user's agent is already aware of the constraints, either through learning or through elicitation.

3.1 Privacy Policy

Clarke-Tax auction mechanism depends on the decisions of the participants. These participants have a self-contained evaluation calculation to decide on the importance of different actions, and bids according to the result of the evaluation. If an agent can assess the importance of a content for its user correctly, and bid with neither excessive nor low amounts, it can help to get better decisive actions, while preserving the previously obtained scores of the user within the network.

PANO makes use of policies for the agents to compute the bidding evaluations. Agents have multiple policies that correspond to different actions, and in an auction, they correspond to these policies to place bids accordingly. In PANO, a policy P is a 5-tuple $P = \{a, n, p, q, i\}$, where a is the agent that the policy belongs to, n is the set of users in the network the policy is applied to, p is the conditions for the content types that the policy will be applied, q is the action of sharing or not sharing the contents when the policy is applied, and i is the importance of the policy, which is a value between 0 and 1. An example policy of Alice wanting to share photographs that contain scenery tag with friends, with 0.6 importance can be represented as P = *{Alice, friends, photograph[scenery], share, 0.6}*.

The success of the mechanism depends on how the final policy out of an auction satisfies the policies of the agents. The resulting policy of an auction should correctly assign the policy-applicable users of the network where there are no conflicts between the auction participants, and try to assign the rest as satisfactory as possible to protect the common good. Equation 1 measures how well the overall result found with PANO satisfies the n agents that enter the auction. Success is defined as the number of the users that the applied policy differing from the agents requirements (UPC: count of the users with unsatisfied policy for user u), divided by the entire set of users that were considered to share the content with (TNU: total count of users in auction participants' network).

$$\text{Success\%} = \left(1 - \frac{\sum_{u=1}^{n} UPC}{TNU}\right) * 100 \tag{1}$$

Example 3 Consider two agents in a network of 200 agents that enter an auction to decide how to share a picture. The first agent wants to share the picture with 140 people, and the second agents wants to share it with all. Assume that as a result, the picture is shared with 160 users that include the 140 users that the first agent preferred. The metric would result in $(1 - (20 + 40)/200) * 100 = 70\%$ success.

With the given policy notation, satisfaction of individual users can also be calculated. Equation 2 measures the user satisfaction after an auction, considering how well the outcome is aligned with the agent's policy and the importance of the policy. That is, while the satisfaction value for a single content can be computed with the Success metric, making use of importance values of policies can give us sensitivity levels (SL) of users for conditions of content types that are also represented in the policies. For example, if a user has importance level of 0.6 for a policy that is related to a condition of a content type, it can be assumed that the sensitivity level of the user for the same condition can also be given as 0.6. For multiple policies of the same user for the same condition, we take the average of importance levels of the related policies for SL value. Using the satisfaction metric for a single content as CS, and the sensitivity level of the content for the user as SL, we define the user satisfaction (US) metric for an agent with the formula below, where i is the content id from the previously policy applied contents.

$$US = \sum_{i=1}^{n}(SL_i * CS_i) / \sum_{i=1}^{n}(SL_i) \tag{2}$$

Example 4 Alice wants to share a photograph tagged with "bar," which also has other co-owners. She has a policy P for contents that are tagged with "bar," and policy P indicates Alice preferring contents tagged with "bar" to be shared with her "friends," with a sensitivity level of 0.7. According to the final policy decided with Clarke-Tax mechanism, the content is shared with 70 of her friends, while her policy was intended to share it with 100 people. The satisfaction of Alice for this content is 0.7. If Alice have shared a content before with a satisfaction of 0.6 where her policy had sensitivity level of 1.0, the combined user satisfaction is calculated with the user satisfaction metric as:

$$US_{Alice} = (0.6 * 1.0 + 0.7 * 0.7)/(1.0 + 0.7) = 0.64.$$

3.2 Preventing Abuses in the Auction Mechanism

The Clarke-Tax auctions are beneficial for decision making for multiple participants with different opinions, as they support truthfulness. However, the economic system and the currency used in the mechanism can allow abuses, as explained in Sect. 2.2. In order to prevent the system from facing malicious behavior by some users, some modifications are need to be made for earning the currency and spending it. The main modifications proposed in this paper to prevent abuses are the group-wise bid scoring, boundaries of the bids for the auctions and the balance between the income and expends.

Group-wise Spending: To prevent abuse of using currencies for irrelevant auctions with different users, earned currencies can only be used in new contents with same co-owners. Recall that there are three types of actions: not sharing with anyone, sharing with a limited number of people, and sharing with public. Limited number of people is decided from the conflicted share decisions of users. For example, when a co-owner of the content decides not to share it with a specific user and another co-owner wants to share the content with it, this user is added to the limited audience list. Since the conflicts, and the audience lists for actions are only related to the policies of the co-owners, spending pre-owned currency from previous auctions with different co-owners would give some participants an unfair advantage. This would result in some users cooperating to share trivial co-owned contents between themselves, and not spend any currencies for the auctions since the contents have no share value at all. This is prevented with group-wise spending, where the currency earned from auctions with some co-owners can only be spend in the future auctions with the same co-owners.

Boundaries: Clarke-Tax mechanism allows users to bid as much as the currency they hold. This free market approach economy adds a level of uncertainty to the auctions, since a participant cannot have a clear opinion about what others might bid. A person that has big earnings can bid with high numbers, and makes little of the taxes, since they can spend more than the others. Limitations to minimum and maximum bids allowed can be beneficial to prevent users that are richer in the currency from dominating the decisions. This also helps agents that participate in the auctions to have better evaluation functions, because they can have a better opinion about the other participants' bids, especially with prior knowledge about the others' characteristics and the context of the content.

With the notion of minimum-maximum boundaries, the balance between currency earnings and expenditures comes into consideration. Users earn currencies by being a co-owner of a content and spend them for decisions in auctions and taxes. Earning currencies equal to or more than the maximum expenditure for an auction would cause the economy to bloat and inflation to emerge. Therefore, earnings from being a co-owner should be less than one can spend for an auction of a content. We propose one to two balance, where the currency earned from a content should be half of the maximum boundary of an auction. Another benefit of this approach is to encourage agents to spend more wisely, since spending less for a relatively unimportant decision could help agents to spend more in the future decisions.

Consider Example 1 where Dave could spend 50 on a picture because he could afford it. With the proposed scheme, group-wise spending would only enable Dave to spend pre-owned currency earned from co-owned contents with the same group. Therefore, the notion of total owned currency becomes obsolete, replaced with group-wise currencies. In this situation, all the co-owners would earn the same amount from the co-owned contents, which would provide equality in auctions. Dave can only bid as much as the other co-owners, so he cannot take advantage by using earned currencies with contents co-owned with other people. The agent-based policy system can also solve the minority dominating the majority problem in Example 2, which is caused by the absence of co-owners in the auction. In this example, Alice and Carol, who do not join the auction which results in the only person, Dave, that doesn't want the photograph to be shared impose his policy on the others. With PANO, every user in the OSN is represented with an agent, which facilitates privacy policies to bid on behalf of their users. Therefore, even if the users are not available for an auction, their agents always are, bidding according to their owners' policies. With the help of automation, Alice and Carol will also be represented in the auction with agents. Since Clarke-Tax mechanism is based on truthfulness, with every co-owner present in an auction, the resulting action will be decided according to everyone's opinion, instead of a minority of the co-owners.

Example 5 Dave wants to share a photograph he is in with Alice, Bob, and Carol. Alice and Carol don't want the picture to be shared, and they want to do a Clarke-Tax auction to decide on the final policy. Dave shared a lot of photographs before, that are unrelated to Alice, Bob, and Carol, and gained bid scores from those contents. Since the mechanism only allows Dave to use currencies obtained from the contents that are related to all Alice, Bob, Carol, and Dave, the scores from unrelated content can't be used to gain advantage. The outcome of the auction is decided over who values their decision most, with the use of only the gained currencies from the previous shared contents of the same group, without any unfair advantage to any participant.

3.3 Bidding Mechanism

Agents bid on auctions based on their privacy policies. As explained in Sect. 3.1, policies have importance values. On top of this, agents also have privacy characteristics, which is the notion of how much privacy-aware an agent is, represented with a value between 0 and 1, named as characteristic coefficient. For a content in an auction, an agent checks its related policies and determines the set of social network users that it wants to share or not share the content with. The characteristic of the agents and the importance of related policies determine how much the agent wants to bid for an auction, according to the given actions of these policies. In Clarke-Tax mechanism, each type of action can receive a different bid from the participants, so even when an agent has conflicting policies (i.e., different policies of an agent preferring different actions for the same network users), it can place bid on conflicting actions, according to their evaluations. Agents should also consider how much currency they own, and place their bids accordingly (e.g., bidding small amounts when short on currency or bidding higher when have enough spendable currency).

In PANO, the bidding mechanism is a linear function that returns an integer value between the bidding boundaries for each action, namely sharing, not sharing and limited sharing, explained in Sect. 2.2. The function considers privacy policies, agent characteristics, the current score balance of the agent and the number of users that are in conflict against other agents in the auction to compute the outcome. For a content, an agent first considers its related policies to find *coefficient* of biddings. This coefficient (*PC*) is calculated as the mean average of the importance values of the policies. To spend the obtained currencies in a prudent manner, we consider both sharing and not sharing actions at the same time for the mean average, if the audience of both policies are the same. That is, we consider not sharing importance as negative, and sharing importance as positive values. For example, if an agent has two policies for a content; one sharing with importance of 0.8 and one not sharing with the importance of 0.2 for not sharing it with the same audience, the importance coefficient is calculated as the mean average of 0.8 and −0.2, thus 0.3 for sharing the content. This coefficient is multiplied by the privacy characteristic coefficient (CC) of the agent. The final parameter is calculated according to the bidding boundaries and the current score balance of the agent, namely spendable currency (SC). If the agent has two times of the maximum boundary (mb) as the balance, the bidding is made with the multiplication of the maximum boundary and the computed coefficient. If the currency is less, the coefficient is multiplied with the agent's owned currency (oc) divided into two, in order to save currency for the future auctions. Equation 3 gives the formula generating bids for sharing or not sharing where n is the number of related policies for an auction and $P_\alpha\{i\}$ is the importance value for policy α. Intuitively, an agent generates a high bid if the privacy importance of the image in question is high, the agent values privacy (CC) and it has the resources to spend on the bid (SC).

$$SC = \begin{cases} mb, & \text{if } oc/2 \geq mb \\ oc/2, & \text{otherwise} \end{cases}$$

$$\text{Bid Value} = \sum_{\alpha=1}^{n} (P_\alpha\{i\})/n \times CC \times SC \tag{3}$$

In addition to biddings for sharing or not sharing a content, an agent can also bid for the limited share action. Limited share action aims to share the content with the conflicting set of users, i.e. the users that at least one agent is willing to share the content and at least one agent does not want to share the same content with. For bidding to share with a limited audience, the agent checks how many of the users in the conflicting set of users it wants to share the content with, and simply calculates the ratio of the users fitting into its policies divided by the total number of users in the conflicting users set, namely as fitting ratio (FR). The agent bids half of the maximum boundary or half of its own score if it has less, multiplied with the computed ratio for the conflicting users and its characteristic. The formula for computing the bid values for limited sharing is given in Eq. 4, where n is the count of conflicting users considered for the policy, and β is the count of users in the conflicting list fitting into user's own share policies.

$$\text{Bid Value (Limited Share)} = \beta/n \times CC \times SC \tag{4}$$

4 Multi-agent Simulation

In order to evaluate the success of PANO, a multi-agent simulation environment is implemented, where agents represent users in a social network with privacy policies, and use Clarke-Tax auctions to decide on the share policies of co-owned contents. The policies are defined according to rules that regular social network users tend to rely on, and the simulation checks the success of final policy decisions with Clarke-Tax auctions over a chunk of co-owned contents.

4.1 Context-Based Privacy Constraints

Social networks users have privacy constraints for the contents they share in the network. These constraints can be represented by rules, which then can be used to model behavior of the users. Kekulluoglu et al. have done a user study to extract privacy rules of frequent users of online social networks [4]. In this work, there are seven decisive rules, which most of them had contextual content properties as a constraint. Two out of these seven rules were eliminated from the current work, since the simulation doesn't support location or event aspects of the social networks. The remaining rules, with relation types and contextual constraints, are presented below:

1. If the user: x is included in a photograph that depicts a depressed mood, don't share the photograph.
2. If the user: x is included in a photograph that is located in a bar then don't share the photograph with family members.
3. If the user: x is included in a photograph that is located on a beach don't share with work related people.
4. If the user: x is included in a photograph that is a mature content, don't share the photograph.
5. Do not share photographs with user: y.

In our policy notation, we define contextual constraints with conditions combined with content types. In the example policy given in Sect. 3, *photograph[scenery]* depicts that the content type is a photograph, and it's categorized as scenery. We represent the photograph categories presented in the rules above with such keywords, which can be obtained with automated photograph tags. In PANO, we assume that the photograph contents are already tagged, so we can assign them into policy conditions. The conditions can have multiple content types, and also multiple constraints for these content types, giving us flexibility in defining the policies. Since the rules in Kekulluoglu et al. [4] focus on not sharing contents in specific conditions, we would require some sharing rules that can create conflicts we aim to resolve in scope of this paper. Therefore, we generate some opposing rules to the ones above, where some agents would like to share the photographs in the same conditions. For example, when an agent has the third rule in their policy, we can generate a policy for another agent about sharing photographs tagged with beach with its friends. In this example, if both agents are co-owners for a photograph on a beach, the intersection of OSN users from work related people for *Agent #1* and friends for *Agent #2* will be in a conflicting situation, because the former do not want to share the photograph while the latter wants to share it.

4.2 Execution of the Simulation

PANO is implemented in Java. The simulation represents social network environments, with the aspects of users, contents, different relation types between users and agents. Agents have defined rules, based on the combinations of relations and contextual tags of photographs, as explained in Sect. 4.1.

As mentioned in Sect. 3, agents are only allowed to use the scores they obtained from earlier shared contents from the same co-owners. Another prevention for abuses is that each agent can only spend a total maximum of 20 scores for each content, so that none of the agents can overvalue it to force its own action. According to the characteristics of the agents, they can disperse their scores to different actions, and they can even decide not to spend any scores if they don't value the sharing or not sharing of the content. After all the bids are taken blindly from each user, the simulation calculates the overall decision and taxes of the agents (if any), and updates their currencies within the network, as explained in Sect. 2.1.

Algorithm 1: Main process of the simulation

```
load the social network;
set relationships;
add contents with contextual tags;
assign co-owners to the contents;
add user characteristics and rules;
foreach Content do
    foreach Co-Owner do
        determine the set of people to share or not share;
        if the decision for a person is differing then the other co-owners then
            remove person from other lists put the person on the conflicting list
        else
            if the person has no share decision & the person not exists in any lists then
                put the person on no share list
            if the person has share decision & the person not exists in any lists then
                put the person on share list
    foreach Co-Owner do
        place bids according to the final lists
    trigger Clarke-Tax mechanism;
    update currencies of the co-owners;
```

Algorithm 1 explains the main process of the simulation. The first five lines initialize the environment by loading the network, setting relationships between the users, adding contents with co-owner assignments and generating agent characteristics and rules. Then the simulation enters into a loop for each content that requires Clarke-Tax auction for deciding the final policy. This loop defines the share, limited share and no share lists of a content between lines 8 and 15. Using the formulas of bid calculations in Sect. 3, agents place bids on each action in lines 16 and 17, and the last two lines update scores according to the outcome of the auction, considering the placed bids and the taxes.

5 Evaluation

The main goal of a privacy management system is to maintain the balance between the system users' privacy requirements and fit the finalized privacy policies as much as it can, even with conflicting policies. We evaluate PANO using the overall system success and user satisfaction metrics presented in Sect. 2.2.

5.1 Correctness of Policies

The first evaluation of PANO is aimed to find out the success of the proposed modifications of the method to the Clarke-Tax mechanism, to prevent abuses and create more satisfactory privacy policies to the community. The success metric defined in Sect. 3 is used to measure the correctness of policies.

Hypothesis 1 Given a set of agents with various characteristics, PANO can assign collaborative policies more successfully than the native Clarke-Tax approach, regardless of the starting point and the level of abuse by the agents, while the higher abuse level decreases the performance of the former method.

Using the success metric, PANO is compared to the pure Clarke-Tax mechanism approach, where group-wise currencies or boundaries are not present. The evaluation is made with five randomly generated social networks of 1000 users, and 200 randomly generated photograph content with contextual tags related to the rules presented in Sect. 4.1. Each network was calculated with considering different starting points where some of the contents are considered shared before. These starting points are decided as 40, 80, 120 and 160 contents shared before, respectively. For representing the pure Clarke-Tax mechanism, two levels of abuse were used in comparisons to model the examples in Sect. 2.2. First level enabled some agents to bid more than the defined 0–20 range, but resulting in consuming all currency if used for every auction. Also, the agents that abuse to protect contents from sharing or try to disseminate to their entire social network are balanced, so the abuser agents can also

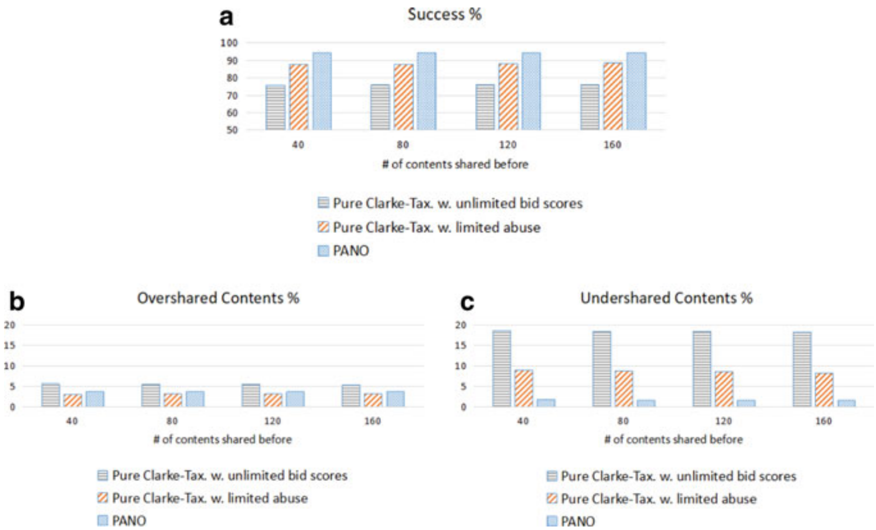

Fig. 1 **a** Percentages of the social network users that are correctly assigned within the final policies according to co-owners' characteristics. **b** Oversharing percentages. **c** Undersharing percentages

have conflicts. Second level of abuse were defined by giving unlimited currencies to the abuser agents, and only letting one-sided abuse for a single auction (i.e. having only agents that want to share the content or doesn't want the content to be shared). After running the simulations with these networks for three different models, the average of satisfactory policy metric for five networks, according to different starting positions is shown in Figure 1a. It shows that with PANO, correct assignments of users with created policies increase, with more than 90% (approximately 95%) success rate. Limited abuse level with pure Clarke-Tax had less than 90% success (approximately 88%), and unlimited currencies for abuser agents gave the lowest results with less than 80% (approximately 76%) success rate.

The same setup is also used for comparing oversharing and undersharing percentages of created policies between the two models. Oversharing covers the users a content is shared with, when some co-owners do not want to share the content with those users. In opposite, undersharing covers the set of users that some co-owners want to share the content with, but it is not shared with this set of users. The results of oversharing and undersharing evaluations are shown in Figure 1b and c, respectively. For oversharing metric, PANO and pure Clarke-Tax method with limited abuse almost performed the same (3.5%), with pure Clarke-Tax approach performing slightly better (less than a half percent). They both performed better than pure Clarke-Tax method with unlimited currency for abuser agents, which had about 5.6% oversharing percentage. The similar performances in oversharing is also investigated, and it has been seen that the abusing agents have more advantage with not sharing the content, since deny-overrides is the default strategy for equal situations. In some cases, limited abuse level had two conflicting abuser agents trying to outbid each other with the agent that prefers not to share always winning when the same bids are made, and oversharing percentages managed to be as low as PANO. For the undersharing metric, this causes both abuse levels with pure Clarke-Tax method to perform far worse than PANO, where PANO had about 1.5% undershare while two levels of abuse methods had about 8.8% and 18.5%, respectively.

In general, PANO performs better than applying pure Clarke-Tax mechanism with both abuse levels, considering all the different starting points. Since the evaluated contents were same for all starting points, the success rates were quite similar in four different setups. Even with the same contents, there is a slight increase in the performance when previously shared content number increases, indicating that with large number of pre-shared contents, Clarke-Tax mechanism will perform better.

5.2 User Satisfaction in Case of Conflicts

The second evaluation setup aims to calculate the user satisfaction metric described in Sect. 3 in a larger generated network. The main goal is to create conflicts between users for a large number of contents, and calculate ratios of intended/unintended audience for decided policies of these contents, using PANO and comparing it with the results of the application of a pure Clarke-Tax mechanism.

Hypothesis 2 Given that there is a conflict of policies between agents for every auction, PANO will ensure the user satisfaction of the majority by preventing some agents from abusing the mechanism continuously.

The user satisfaction metric in this evaluation has three measurement types. First is the satisfaction ratio, which shows the percentage of people that the policies of the co-owners matching the final policy, taking the sensitivity of the initial policies of the agents into account. The second and third metrics are the percentages of oversharing and undersharing, which are in the same concept as the first evaluation, but these also include the sensitivity levels of the agent policies.

The setup of the second evaluation is constructed with a 1000 agent network, where 100 agents owned contents and had approx. 250 friends each. 500 contents were generated with tags, and randomly assigned to 4 of the 100 agents within the network. Policies of the agents were also randomly generated with different sensitivity levels. For each content, at least one conflicting policy were assured with additional policy generation for some of the co-owners, to create 1 against 2, 1 against 3, or 2 against 2 conflicts. Clarke-Tax mechanism is run with PANO and the pure approach, where at least one agent was trying to abuse the mechanism to get decisions for sharing or not sharing the content according to its own policy. The user satisfaction, oversharing, and undersharing metrics were calculated at each 100 content mark according to equation (1), and these three comparisons are presented in Fig. 2.

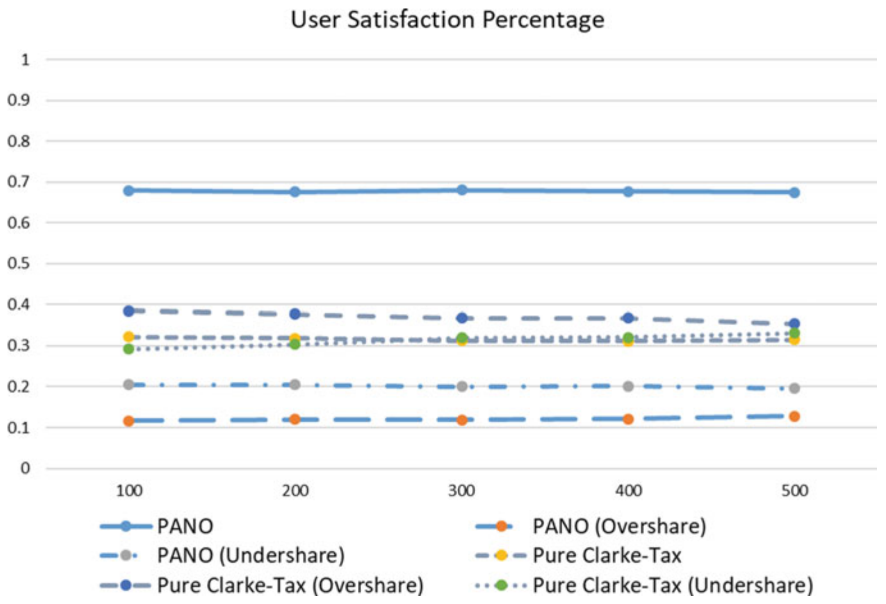

Fig. 2 Comparison of user satisfaction

6 Discussion

Applying multiple privacy preferences from several users for a single OSN content to define a single privacy policy is a major challenge, and it can even be more complicated when these preferences are conflicting. Several studies worked on collaborative privacy management, applying different methods with different approaches. Role-based access control, which was applied widely to other software or operating systems, is not fully suitable to OSNs, since it can't capture binary or context-dependent relationship features that are available in most OSNs. Fong [3] introduces Relationship Based Access Control (ReBAC) mechanism and provides a model to make it applicable to OSNs, where users can define their privacy constraints related to the relations that are available in OSNs, such as friends or colleagues. ReBAC is one of the earliest works which included user-related privacy constraints in access control decisions in OSNs, and provided a basis for many collaborative privacy management mechanisms.

Such and Rovatsos [11] propose a model, where predefined policies of the users are used to find out conflicts for contents and a middle ground is found with a negotiation mechanism. Such and Criado [10] extend that work by modeling user behavior and using a software mediator where some negotiation actions are made without direct user input. Kekulluoglu et al. [4] use multi-agent negotiation as well but apply a more comprehensive negotiation protocol. They further take into account incentives for agents. Kokciyan et al. [5] use argumentation to resolve privacy disputes. In that

Table 3 Comparison of desirable properties

Properties	Negotiation	Argumentation	Pure Clarke-Tax	PANO
Automation	✓	✓	✗	✓
Fairness	✗	✓	✗	✓
Concealment	✗	✗	✓	✓
Protection before exposure	✓	✓	✓	✓

work, OSN users were represented by software agents, where the agents have access to domain knowledge and infer semantic rules. With argumentation mechanism, agents can attack other agents' assumptions to convince the others to accept its users' privacy requirements. While this is successful, it requires agents to be able to reason semantically, which may not be possible in various environments.

Mester et al. [7] propose four desirable properties for a privacy management system, namely automation, fairness, concealment of privacy concerns, and protection before exposure. Automation property is the capability of the system to perform without human interference. Fairness property depicts the system's ability to provide equality to its members. Concealment of privacy concerns is about the coverage of the members' privacy policies; the more agents can hide their policies from the others, the more the property is satisfied. Protection before exposure is related to the privacy management method's application before the related entity is published in the system, causing possible violations. We compare our method with some other approaches for privacy management in terms of these properties in Table 3. The compared approaches are negotiation, argumentation, and pure Clarke-Tax approach with no modifications.

According to the comparison in Table 3, negotiation approaches for privacy management can be designed to be fully automated, even though some methods require human intervention. It can cover fairness to some extent, but it is still open to abuses by the negotiating agents. Since agents have to expose their privacy requirements for negotiation, they cannot offer the concealment of privacy concerns property. Negotiation protects the agents before exposure, but it might require multiple iterations over an algorithm to finalize a privacy policy. Argumentation is similar with negotiation for four of the properties, but it also provides fairness, since the agents are given the opportunity to attack others' arguments over multiple iterations to decide the final policy, thus defending themselves against unfair arguments. Pure Clarke-Tax mechanism provides a one-shot, simple algorithm, concealing privacy concerns of the agents, while being easy to compute. However, it is used by social network users, therefore with no automation, and is open to several types of abuses, as mentioned in Sect. 2.2. PANO is fully automated, where each user is represented with an agent, and agents act on behalf of their owners according to their privacy policies. Measures to prevent abuses to present a fair mechanism is explained in Sect. 3, which satisfies the desirable fairness property for such systems. prevent the system from having the first and the second properties.

Future work of the methods presented in this paper will focus on modeling user behavior and predicting the decisions of users with software agents. We plan to achieve this with a user study where participants can act according to their privacy concerns for co-owned contents and6 bidding with the same mechanism of PANO. With the results of the user study, we will cluster users according to their privacy concerns and behavior in OSNs. Using these clusters with a semantically related, hierarchical content categorization will be a guidance to implement efficient software agents, which can be used to represent OSN users more precisely with less amount of prior knowledge within our Clarke-Tax based Collaborative Privacy Management model.

References

1. E. Clarke, Multipart pricing of public goods. Publ. Choice **11**(1), 17–33 (1971)
2. E. Ephrati, J.S. Rosenschein, The Clarke tax as a consensus mechanism among automated agents, in *Proceedings of the Ninth National Conference on Artificial Intelligence, AAAI '91*, vol. 1 (AAAI Press, 1991), pp. 173–178
3. P.W.L. Fong, Relationship-based access control: Protection model and policy language, in *Proceedings of the First ACM Conference on Data and Application Security and Privacy, CODASPY '11* (ACM, New York, NY, USA, 2011), pp. 191–202
4. D. Kekulluoglu, N. Kokciyan, P. Yolum, Preserving privacy as social responsibility in online social networks. ACM Trans. Internet Technol. **18**(4), 42:1–42:22 (2018)
5. N. Kökciyan, N. Yaglikci, P. Yolum, An argumentation approach for resolving privacy disputes in online social networks. ACM Trans. Internet Technol. **17**(3), 27:1–27:22 (2017)
6. A. Lampinen, V. Lehtinen, A. Lehmuskallio, S. Tamminen, We're in it together: Interpersonal management of disclosure in social network services, in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11 (ACM, New York, NY, USA, 2011), pp. 3217–3226
7. Y. Mester, N. Kökciyan, P. Yolum, *Negotiating Privacy Constraints in Online Social Networks* (Springer International Publishing, Cham, 2015), pp. 112–129
8. A.C. Squicciarini, D. Lin, S. Sundareswaran, J. Wede, Privacy policy inference of user-uploaded images on content sharing sites. IEEE Trans. Knowl. Data Eng. **27**(1), 193–206 (2015). Jan
9. A.C. Squicciarini, M. Shehab, F. Paci, Collective privacy management in social networks, in *Proceedings of the 18th International Conference on World Wide Web, WWW '09* (ACM, New York, NY, USA, 2009), pp. 521–530
10. J.M. Such, Natalia Criado, Resolving multi-party privacy conflicts in social media. IEEE Trans. Knowl. Data Eng. **28**(7), 1851–1863 (2016). July
11. J.M. Such, M. Rovatsos, Privacy policy negotiation in social media. ACM Trans. Auton. Adapt. Syst. **11**(1), 4:1–4:29 (2016)