PANO: Privacy Auctioning for Online Social Networks

(Extended Abstract)

Onuralp Ulusoy
Department of Computer Engineering
Bogazici University
Istanbul, Turkey
onuralp.ulusoy@boun.edu.tr

Pınar Yolum
Department of Information and Computing Sciences
Utrecht University
Utrecht, Netherlands
p.yolum@uu.nl

ABSTRACT

Online Social Networks (OSNs) enable their users to share content with their connections. Shared contents over OSNs raise privacy concerns, since they tend to contain personal information of users. More importantly, a single content, e.g, a photo of a group of people, can potentially contain private information of multiple users, which become available without their consent. Ensuring that all relevant users' privacy requirements are met is important but difficult since the requirements can easily be conflicting. Hence, mechanisms to resolve privacy disputes are needed. Accordingly, this paper proposes an agent-based collaborative privacy management model, where agents represent users and manage their privacy requirements. When an image is about to be shared, the relevant agents enter an auction and bid on behalf of their users about how private the considered image is. The bids are processed with a modified version of Clarke-Tax mechanism that achieves fair handling of privacy settings and taxes the agents whose privacy settings are chosen.

KEYWORDS

Privacy; Online Social Networks; Auctioning

ACM Reference Format:
Onuralp Ulusoy and Pınar Yolum. 2018. PANO: Privacy Auctioning for Online Social Networks. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Online social networks are widely used systems that enable their users to share content online. As of 2017, there are almost two billion users over the world that use at least one of the social networking applications [1]. With that many users, the amount of content that is being shared online is tremendous. The shared content tends to reveal personal information about the user that owns the content as well as others that are affiliated with the content. For example, a picture of a group of friends do not only reveal information about the uploader but also about all others that can be identified in the picture. Hence, when preserving the privacy, it is not enough to consider just the uploader, all others that might be affected should be taken into account. This leads to the concept of *collaborative privacy management*, where all affected users should get a chance to influence how a content is going to be shared [4, 5, 7].

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

Accordingly, this paper proposes PANO, an agent-based collaborative privacy management system that builds on the work proposed by Squicciarini *et al.* [7]. 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.

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. 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 photo 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.

Clarke-Tax mechanism [2] provides an auction mechanism 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. Participants who aid in the winning action to be chosen are taxed according to the value they put on it. If the exclusion 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.

Auctioning with Clarke-Tax Mechanism is an efficient way of negotiation, since it has been shown that truthfulness is the best strategy for bidding [3, 7]. Even with the notion of truthfulness, the approach used in [7] still has some limitations that can result in abuse by the bidders or inflation in the currency used.

3 AGENT-BASED BIDDING

Current Clark-Tax based mechanism in Squicciarini *et al.* [7] 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 [6, 8]. Here, we assume that the user's agent is already aware of the constraints, either through learning or through elicitation.

3.1 Privacy Policy

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 photos that contain scenery tag with friends, with 0.6 importance can be represented as $P = \{Alice,friends,photo[scenery],share,0.6\}$.

The success of the mechanism depends on how the final policy out of an auction satisfies the policies of the agents. 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).

$$Success\% = 1 - \frac{\sum_{u=1}^{n} UPC}{TNU} \tag{1}$$

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. 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.

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

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 Section 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 and boundaries of the bids.

Group-wise Spending: 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 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. Limitations to minimum and maximum bids allowed can be beneficial to prevent users that are richer in the currency from dominating the decisions. With the notion of minimum-maximum boundaries, the balance between currency earnings and expenditures come into consideration. We propose one to two balance, where the currency earned from a content should be half of the maximum boundary of an auction.

3.3 Bidding Mechanism

Agents bid on auctions based on their privacy policies. As explained in Section 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. 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).

4 FUTURE WORK

Our current evaluation of the system is based on multiagent simulations. First direction we want to pursue is to model real user behavior and predict the decisions of users with software agents. This is planned to be achieved by gathering data with a user study that applies a game model, where participants can act according to their privacy concerns for co-owned contents, bidding with the same mechanism of PANO and clustering users according to data that is already present in the social networks. Using these clusters, with a semantically related, hierarchical content categorization will be a guidance to implement efficient software agents for our model. Second direction is to derive various bidding strategies for agents such that they can learn from the outcomes of their actions and make trade-offs between privacy and actual bids.

ACKNOWLEDGMENTS

This work is supported by the Turkish Ministry of Development under the TAM Project number DPT2007K120610 .

REFERENCES

[1] Dave Chaffey. 2017. Global Social Media Research. (2017). http://www.smartinsights.com/social-media-marketing/social-media-strategy/new- global-social-media-research.

[2] Edward Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33.

[3] Eithan Ephrati and Jeffrey S. Rosenschein. 1991. The Clarke Tax As a Consensus Mechanism Among Automated Agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence - Volume 1 (AAAI'91)*. AAAI Press, 173–178.

[4] Dilara Kekulluoglu, Nadin Kökciyan, and Pınar Yolum. 2018. An Argumentation Approach for Resolving Privacy Disputes in Online Social Networks. *ACM Transactions Internet Technology* (2018). To appear.

[5] Nadin Kökciyan, Nefise Yaglikci, and Pınar Yolum. 2017. An Argumentation Approach for Resolving Privacy Disputes in Online Social Networks. *ACM Trans. Internet Technol.* 17, 3, Article 27 (June 2017), 22 pages.

[6] Anna C. Squicciarini, Dan Lin, Smitha Sundareswaran, and Joshua Wede. 2015. Privacy Policy Inference of User-Uploaded Images on Content Sharing Sites. *IEEE Transactions on Knowledge and Data Engineering* 27, 1 (Jan 2015), 193–206.

[7] Anna C. Squicciarini, Mohamed Shehab, and Federica Paci. 2009. Collective Privacy Management in Social Networks. In *Proceedings of the 18th International Conference on World Wide Web (WWW '09)*. ACM, New York, NY, USA, 521–530.

[8] Jose M. Such and Natalia Criado. 2016. Resolving Multi-Party Privacy Conflicts in Social Media. *IEEE Transactions on Knowledge and Data Engineering* 28, 7 (July 2016), 1851–1863.