Equity-preserving Management of Privacy Conflicts in Social Network Systems

R. Marin[1], G. Piolle[1], C. Bidan[1]
SUPELEC / Inria
{regina.marin, guillaume.piolle, christophe.bidan}@supelec.fr

Abstract

Social Network Systems (SNSs) allow users to easily share data and link them to other users. For the sake of privacy, users may define policies over these data, which can trigger conflicts with other users. It happens that the resolution of these conflicts lead to unfair situations. We propose a characterization of this unfairness by introducing the concept of equity in SNSs. This allows us to define an equity-preserving conflict management algorithm. We evaluate this algorithm using the Gini coefficient, which is a standard metric for inequity in economics. This evaluation shows that our approach leads to better results than classical conflict resolution strategies.

1 Introduction

In 1997, the world watched the rise of Social Network Systems (SNSs) on the internet. They are usually described as "*a networked communication platform in which participants (1) have uniquely identifiable profile that consists of user-supplied content, content provided by other users, and system-level data; (2) can publicly articulate connections that can be viewed and traversed by others; and (3) can consume, produce, and interact with streams of user-generated content provided by their connections on the website*" [1]. SNSs, such as Facebook or Twitter, have attracted many users since their foundation, mainly because they allow and encourage publishing and sharing interests, news, hobbies, activities and documents (i.e., photos). Moreover, links can be made between documents and users by the means of *tagging*. The social relationships among users give semantics to these links, adding further value to this feature. In consequence, it happens that tagging is currently one of the most popular actions in SNSs [2]. However, it also worsens existing privacy breaches by exposing user identities and providing more means to reach documents.

Although many attempts have been made to define privacy [3], it is often based on a "right to be let alone" [4], associated with a "possibility to control the distribution and use of personal data" [5]. This is usually implemented through the specification of individual privacy policies. When several users are entitled to a form of control over the same data (for instance a picture of them), conflicts may arise as soon as two of them disagree about the permitted usages of the document. When resolving conflicts, it appears that many strategies lead to unfair situations [6], allowing a few users to gain advantage over others if their policies are more frequently enforced. In traditional human societies, the concept of equity is considered a basis for social justice and conflict resolution [7]. We believe that this notion could be of help in SNSs as well.

In this paper, we propose a characterization of unfairness in the enforcement of privacy policies by introducing the concept of equity in SNSs. We then present a novel conflict management mechanism based on an algorithm designed to maintain or restore this equity. The algorithm is then evaluated using the Gini coefficient [8] as a metric to measure the resulting degree of inequity in the SNS.

The paper is organized as follows. Section 2 presents the applicative context in which privacy conflicts arise. In Section 3, we introduce the notion of equity and present our equity-preserving algorithm. Section 4 describes our implementation and its evaluation. Section 5 compares our approach to some related works. We finally conclude and propose future research tracks in section 6.

2 Privacy Conflicts in SNSs

A Social Network System (SNS) is a social structure made up of a set of users organized in a social graph. A user usually has a profile as well as various resources. A user profile represents the identity of the user, often including her personal information: name, age, gender, birth day, and so on. Resources are the user's assets in the system, and can be pictures, videos, messages... The social graph represents social relationships among users in the SNS. We consider that the main objective of SNSs is to share data among users, and that uncontrolled access to these data may give rise to privacy issues. This is why we base our study on access control management.

2.1 Data Sharing in SNSs

SNS users share interests, news, hobbies, documents and activities. They can post notes and upload documents, like photos, on their web space, and tag friends in those documents. We understand tags as links between documents and users[1]. It appears that tagging is currently one of the

[1]In other contexts, people tagging may relate to non users, and tagging in general may relate to other concepts, like keywords describing a document.

most popular actions in SNSs [2], and derives from the idea of organizing and sharing resources efficiently on the Web.

In the context of data sharing and tagging, one may classify users as data holders, data subjects and data viewers, as shown in Figure 1. The *Data Holder* is the user whose webspace data is published on (i.e., the picture is on her "wall"); the *Data Subject* is a user to whom data is related (i.e., she is tagged in the picture); and the *Data Viewer* is the user requesting access to the shared document.

For instance, let us consider a scenario where five users (Alice, Bob, Charlie, Greg and Eve) interact within the same SNS. Alice uploads a document on her webspace. She becomes the *data holder* for this document, in which she tags both Bob and Charlie. Bob and Charlie therefore become *data subjects* of this document. Greg and Eve, for instance because of their position in the social graph, may request access to the document, hence becoming *data viewers*.

This scenario will be the base for the present study, in which we will focus on access control issues with respect to privacy.

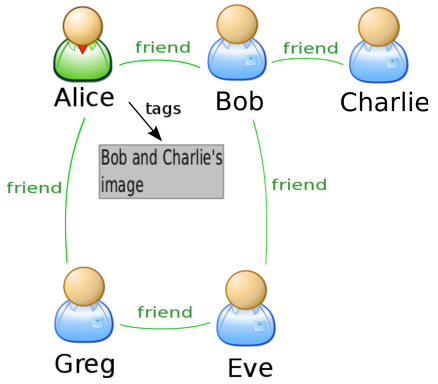

Figure 1: The SNS Scenario Specification

2.2 Expressing Privacy Policies in SNS

Traditionally, SNSs publish application-wide privacy policies which apply indistinctly to all users. In addition to those imposed policies, individual users are often given the opportunity to define their own policy, through a choice of settings and preferences (even though SNS providers may fail to properly enforce those user policies, or limit user empowerment through complex and abstruse privacy settings [9]).

The users' privacy policies usually rely on access control rules, allowing users to define their sensitive assets and limit access to them. In most SNSs, only the data holder has the right to define an access control policy over a published document. In our scenario for instance, only Alice would be in position to decide who can access the picture she has uploaded, and neither Bob nor Charlie have anything to say about who can see the picture. This may lead to privacy breaches for them. Moreover, even though the simple act of uploading a photo can represent a privacy violation, the fact that a person is publicly tagged in a picture can worsen the impact of the privacy breach, allowing other people to gain a searchable access to the user's identity. To take this situation into account, data protection regulation frameworks often provide data subjects with rights over the documents relating to them [10]. In the interest of privacy, we therefore choose to consider that any user tagged in a shared document may propose an access control policy for it.

In consequence, our version of the scenario presented in section 2.1 allows not only Alice (the data holder), but also Bob and Charlie (the data subjects) to specify privacy policies when the document is published on Alice's webspace. For instance:

- Alice may authorizes her "friends-of-friends" (including her friends) to be able to see her picture. All users in the scenario satisfy this condition;
- Bob has not defined a policy about pictures where he is tagged, therefore not imposing any restrictions on anyone;
- Charlie authorizes only his direct friends to see a photo where he is tagged. Bob does satisfy Charlie's policy, while Alice, Greg and Eve do not.

The variety in privacy policies reflects the fact that SNS usages are spread over a broad dynamic range of different user profiles (in terms of ages, nationalities, cultural backgrounds and so on). Indeed, studies have shown that users often have rich and complex photo-sharing preferences [11]. To take this diversity into account, several privacy policy languages have been proposed, with various levels of expressivity [12, 13]. In Enterprise Privacy Authorization Language (EPAL) for instance, a rule consists in a triple (user, action, data) to which the rule applies, a specific ruling (e.g., *deny*, *allow* or *do not care*), and a purpose constraining the actions of the data viewer (such as marketing or research). A rule may also contain activation conditions and associated obligations. For instance, a user allowing people to gather her profile information for research purposes may impose that she is informed of every usage of this information.

Since we have allowed every data holder and data subject to express a privacy policy on a shared document, we expose ourselves to the risk that those policies may not be consistent with each other.

2.3 Conflict Resolution Strategies

The multiplicity and heterogeneity of these privacy policies may lead to conflicting situations, which occurs when at least two users disagree about the outcome of an access request[2]. Two approaches can be taken to tackle this issue. The first one would be to merge individual policies in a single, centralized policy decision point in charge of resolving inconsistencies. Many formalisms have been proposed for this, some of them specific to privacy policy composition [14]. The second approach is to use a distributed policy decision point for each privacy policy, and to collect the

[2]It can be noted that a single policy may be self-conflicting when it includes inconsistent rules, but this issue lies outside the scope of this paper.

rulings in a decision aggregation point in charge of resolving the conflicts. In the second approach, the conflict is between rulings and not between policies. We have chosen this second approach because it allows users to keep their privacy policies secret, and because we expect the reasoning on rulings to be easier to design.

The example in the previous section implies a conflict if Greg, for instance, requests access to the document: Alice allows, Bob does not care, and Charlie denies. Of course, Bob is not involved in the conflict, which takes place between Alice's "allow" and Charlie's "deny". Once the conflict is identified, a resolution strategy must be chosen to determine which ruling to enforce [15].

A common way to eliminate a conflict in access control is to use what we call the *deny strategy*, in which a "deny" ruling will be given priority over "allow" rulings. In our example, Greg would be denied access because of Charlie's ruling. Conversely, the *allow strategy* gives priority to "allow" rulings. In our scenario, Greg would be granted access thanks to Alice's ruling. It can be noted that in most SNSs, allow strategy is the standard choice when resolving policy conflicts [16]. Another approach for deciding among multiple alternatives is to vote. What we call the *majority strategy* gives priority to the ruling supported by the majority of users. In our example, if Bob changes his mind and allows only his direct friends to access pictures where he is tagged, he will issue a "deny" ruling for Greg, thus triggering a majority in favour of this decision.

It is also possible to design a strategy based on social graph topology, for instance using the distance between the data subject issuing the ruling and the data holder. However, we did not follow this path because we did not want to introduce priorities between data holders and data subjects. Indeed, this kind of choice imposes to decide a trade-off between various rights, namely the data subjects' right to control their image, the putative photographer's copyright and the usage rights commonly attributed to the data holder in SNSs. We do not want to draw a line here.

All the strategies stated above can be used to resolve policy conflicts. Unfortunately, these strategies make decisions in a very static way, possibly allowing some users to take advantage over others. For instance, the classical deny strategy allows a user to prevent anyone from seeing any picture of her. Such a constant and systematic opposition can, in the long run, be considered highly unfair to other users because their privacy policy may never be enforced. We believe that a more dynamic resolution strategy could overcome this limitation.

3 Our Approach to Conflict Management

Prior to designing such a strategy, we need to characterize the issue at hand in a more clear and objective way. In this perspective, the notion of *equity*[3] seems to be a good candidate to qualify the intuitive fairness or unfairness of a situation.

3.1 The Notion of Equity in SNS

Equity has many meanings depending on the context: political, philosophical, societal or economical. The first necessary distinction is between equity and equality. The *Oxford English Dictionary*[4] defines the term "equity" as the quality of being fair and impartial, and "equality" as the state of being the same in quantity, size, degree or value. These definitions show that in human affairs, equity is not necessarily the same as equality, in the sense that giving the same amount to different people in different contexts can be deemed unfair, for instance. Our current work, stemming from the users' perception of the situation, focuses on equity rather than equality, even though equality may prove a useful tool for building equity.

In the social-economic context, due to the difficulty in achieving social justice and redistributing available resources in our society, Hayek believes that social justice is but a dream out of our reach [17]. However, Rawls defends equity as an ethical concept of justice or fairness, being the base of a society that cooperates and shares its benefits and burdens, thus creating a surplus which should be fairly distributed [7]. According to other opinions, for instance in the context of public health, the concept of equity is closely related to human rights principles [18]. The latter reference interprets equity as the absence of systematic disparities in health among groups showing different levels of underlying social advantages and disadvantages. These notions are related to our issue in the sense that they reason on situations where individual interests are conflicting and an acceptable solution must be reached.

It remains to characterize the notion of equity in SNSs, for instance in the special case of our scenario in section 2.1. Let us suppose, on the one hand, that Alice and Bob authorize only their friends to see the picture, and on the other hand, that Charlie has specified in his policy an interdiction, where nobody should be able to see the photos in which he appears. If the deny strategy is applied, we find ourselves in the unbalanced situation described at the end of the last section: Charlie always gets what he wants, while Alice and Bob never do. We consider that this situation is inequitable because Charlie sees its preferences prevailing over others', in the form of a more frequent enforcement of the policy he has chosen. In the light of this example, we choose to characterize equity in the context of SNSs as follows:

In the context of a Social Network System, a situation is said to be ***equitable*** *if and only if all the considered participants have seen their policies enforced or violated in the same proportion over past interactions.*

In the event of an inequitable situation, the SNS may try to restore equity by introducing some kind of compensation phase when resolving privacy conflicts, by influencing the enforcement of user policies.

[3]Since we wish to introduce a new notion in this field of study, we decide not to use the term "fairness", which already has a specific meaning in several subfields of computer science.

[4]http://www.oxforddictionaries.com/

3.2 An Equity-preserving Conflict Management Algorithm

We now propose our management mechanism and the development of an equity-preserving algorithm to tackle the issue of equitable enforcement of multiple user policies in social network applications. Figure 2 presents our algorithm. Our approach relies on three main steps: an equity evaluation (labelled A in the figure), a compensation phase (B) and an additional prioritization based on user preferences (C).

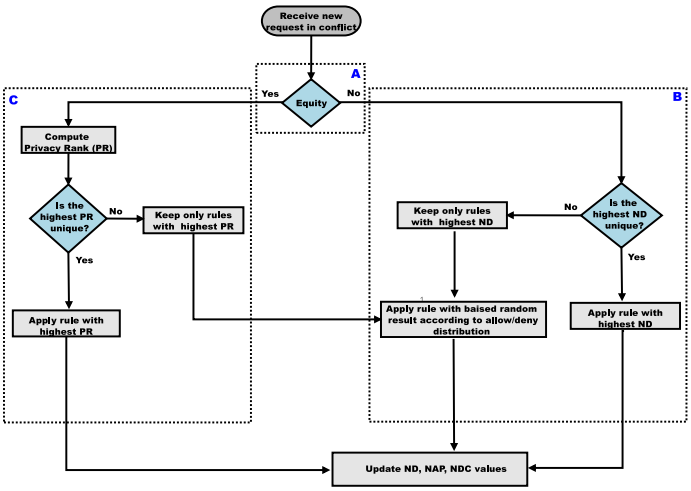

Figure 2: The Equity-preserving Conflict Management Algorithm

Users specify their privacy policies according to their privacy concerns and preferences regarding shared data. We assume that they can assign a ruling chosen among *allow*, *deny* and *do not care*. The ruling "do not care" is not an enforceable ruling. It allows more flexibility and expressivity, while being compatible with rich privacy policy languages such as EPAL [12].

The algorithm starts when a data viewer submits a request to access a document published by a data holder. Once all the necessary user rulings are collected, the algorithm identifies whether there is a conflict among them. If it is the case, then our conflict resolution algorithm is run.

Its first step is called *equity evaluation* (labelled A in the figure). It determines whether the current situation is equitable according to our definition. The algorithm needs to measure the number of times a given user has seen her policy rejected for enforcement (called *Number of Denials (ND)*) and the number of times this same user has been involved in an access request (called *Number of Interaction (NI)*). Note that each time a user is involved in an access request, we assume that it is a valid interaction, even in the absence of a conflict. The ratio between ND and NI (a "normalized" ND) is the proportion we base our equity decision on. It actually measures the proportion with which a user has seen her policy denied, and not enforced. It is equivalent, however, if we consider any "do not care" ruling to be always successfully enforced.

The equity evaluation we propose is simply an equality test over the normalized NDs of all users involved in the conflict[5]. If the equity evaluation result is "No", the situation is deemed inequitable according to our definition, because all users do not have had the same proportion of policy enforcement in the past. The SNS will then try to restore equity, and a *compensation phase* begins (section B of the figure). This part of the algorithm will try to favour a user who has the highest normalized ND, in order to lower this value for her and restore some equity.

At this point, the algorithm checks whether the highest ND is unique among the users in conflict (therefore excluding the ones having sent a "do not care" ruling). If there is a single user with the highest ND, then the algorithm will enforce her ruling. When this decision is taken, all users having sent the winning ruling will see their *Number of Successful Policy Enforcements (NAP)* incremented, all users having sent the losing ruling will see their ND incremented, and all users having sent a "do not care" ruling will see their *Number of Do Not Care (NDC)* incremented. The NI of all users is also incremented. If the highest ND is not unique, then the algorithm will keep the set of users with the highest ND and proceeds to the last step of the compensation phase, in which a ruling is chosen randomly, with a probability distribution identical to the proportions of "deny" and "allow" rulings in the set of highest-ND users. Once the ruling is chosen, all counters are updated as described sooner.

If at the first stage of the algorithm, the situation is deemed equitable, then a different branch is taken (section C of the figure). The decision must be made based on different criteria. The process we propose here intends to take into account the users' preferences and evaluations in terms of privacy, since this is the overall goal of the application. We have chosen to order rulings with respect to the privacy concerns expressed by the users, in the form of a *Privacy Rank (PR)*. From the technical point of view, a PR definition should satisfy the following properties:

- The higher the privacy risk associated to the data, the higher the PR. The variable embedding this notion in our application is the *Data Value (DV)*. DV defines the sensitivity associated with a resource. We consider that different users may have different privacy concerns regarding a same piece of information. We propose to use the following five levels of DV: 5 - extremely sensitive, 4 - very sensitive, 3 - sensitive, 2 - hardly sensitive, and 1 - not sensitive.

- The higher the risk associated to a particular access request, the higher the PR. This notion embeds the whole context of the access request, as well as the users and objects involved. The *Rule Value (RV)* variable embodies it. RV defines the importance of a specific rule within a user's policy. Rules may be marked as important when there is the necessity of a *strong allow* or a *strong deny*. Less important rules boil down to *weak denies* or *weak allows*[6]. Similarly to DV, RV has five

[5]We have tested broader implementations of equity, allowing for little differences between normalized NDs, but our tests shown that strict equality leads to far better results.

[6]While the need of a "strong deny" is pretty obvious, we feel that it is important to allow a user to express her need for the availability

levels: 5 - extremely important, 4 - very important, 3 - important, 2 - hardly important, and 1 - not important.

PR is an increasing function of both DV and RV. The implementation we have chosen is to define it as PR = DV × RV, although many other combinations could be used. In the context of our scenario, let us suppose that Alice becomes the new director in the company where Bob and Charlie are employed. Alice wants all employees to know who the new director is. Therefore, her ruling on this information is an "allow" and the associated rule value is set at level 4 or 5 (this is a "strong allow"), while the data value will be low, the information being somewhat public already. In another contrasting example, let us assume that Alice authorizes her direct friends to see her wedding picture. The data value of this document will be pretty high, since it may have a strong impact on the privacy of her family life. The rule value of the associated "allow" rule will be rather low if she thinks that it is not vital for her friends to see this picture, but the rule value of a "deny" rule regarding strangers would be quite high.

If the algorithm finds a unique highest PR, the corresponding ruling will be given priority and therefore enforced (and the metrics updated). If the highest PR is not unique, then the set of users with highest PR is kept and the algorithm switches to the last part of the compensation phase, with a random selection of the ruling biased by the distribution of "allows" and "denies" in the remaining set of users.

Of course, if no conflict is detected, then the consensual ruling is enforced and the metrics are updated as described sooner (in this case, no ND is incremented).

One may notice that the PR is only used when the initial situation is equitable. Thus, the resulting algorithm gives the priority to the equity between users on the privacy ranks defined by the users. We recognize that this may be questionable, but our main goal is to show that the notion of equity can help the conflict resolution. As part of future research work, we plan to modify our algorithm to take into account the PR during the equity evaluation and the compensation phase.

4 Implementation and Experimental Evaluation

We have developed a proof-of-concept prototype in Java as a demonstrator for our resolution algorithm. To evaluate the performance of our equity-preserving strategy, we compare its results with the ones obtained using the deny, allow, and majority strategies with respect to the Gini coefficient, which measures the fairness of a distribution.

4.1 Gini Coefficient

There is a real need in economics to measure various kinds of inequalities in the society, like the distribution of wealth or income over a population. Since 2008 in particular, the financial crisis has contributed to expose to the general public the income inequalities in our societies. A popular metric to characterize inequality in economics, as well as in biology, is the Gini coefficient [8].

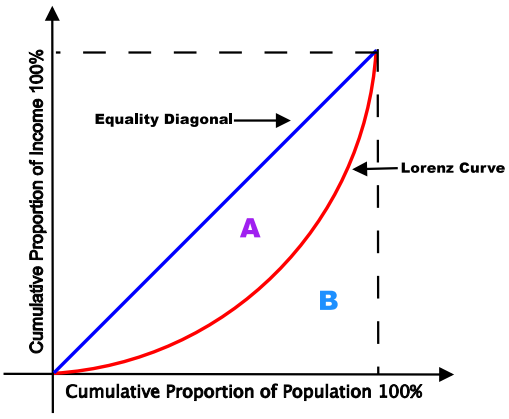

Figure 3: Graphical Representation of the Gini Coefficient

The Gini coefficient measures the degree of inequity of the distribution of a resource in a population. It ranges from 0 (perfect uniform equality) to 1 (perfect inequality), where smaller coefficients indicates a lower disparity in the distribution. The Gini coefficient is defined as a ratio between areas of two Lorenz curve diagrams (a Lorenz graph featuring the cumulative distribution function of wealth over a population). The first curve is the one of a perfectly equitable distribution (which boils down to the identity function), the second one is the actual distribution (a necessarily convex function on $[0, 1]$). If the area between the line of perfect equality and the actual Lorenz curve is A, and the area under the actual Lorenz curve is B, then the Gini coefficient is A/(A+B), as presented in Figure 3.

We have chosen the Gini coefficient as a metric to evaluate how our equity-preserving algorithm compares to other conflict resolution strategies. Since the Gini coefficient has more than fourteen alternative representations [20], we adopted in this work Brown's equation [21] for discrete distributions, as follows:

$$G = 1 - \sum_{i=0}^{k-1} (Y_{i+1} + Y_i)(X_{i+1} - X_i) \qquad (1)$$

where, Y_i is the cumulative proportion of resource, and X_i is the cumulative proportion of the population.

4.2 System Architectural View

In many social networking frameworks, a centralized approach is used for performing some or all of the security features, like access control. In this case, there is usually a single *Policy Decision Point (PDP)* to handle all data requests, and a single *Policy Enforcement Point (PEP)* to implement the decision (and these points are often the same entity) [22]. Nonetheless, it is believe that privacy policy management in SNSs requires a distributed architecture to cope with the individual policy specified by each user [23].

and publicity of a particular piece of information. Although this does not specifically increase any privacy level, it contributes to a higher control of the users over their information, a better user centring of the application and probably a more operational notion of free speech.

Therefore, if one aims at a distributed SNS, one needs to decentralize this process.

In the hypothesis of a distributed architecture, one may imagine that a data holder receives an access request and forwards it to the data subjects involved in the shared data. We consider that every user in our scenario has an individual PDP. Each PDP decides whether the requester has the right to access the data or not, according to the user's privacy policy, and sends the user's ruling ("allow", "deny" or "do not care") to the *Decision Aggregation Point (DAP)*.

The DAP is the architectural component that manages a set of decisions, resolves conflicts according to a selected strategy, and generates the final access control decision. The DAP aggregates all decision rulings provided by the relevant PDPs. In the current version of our model, the DAP is a centralized entity. We plan to distribute the process in the future, but this will need further work and an adaptation of our algorithm.

The decision of the DAP is sent to the PEP (possibly located where the data is stored, which may vary according to the SNS architecture), where the corresponding decision is enforced. If the answer is positive, the PEP grants access over data to the data viewer. Otherwise, access is denied. Figure 4 depicts an overview of this architecture.

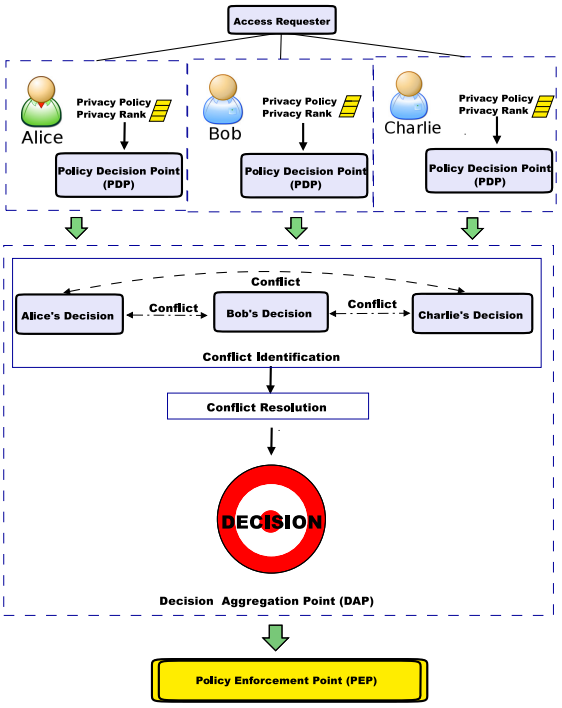

Figure 4: System Architectural View

4.3 The Experimental Setup

There were two experiments which are different between one another in terms of numbers of users, number of interactions and distributed resources.

4.3.1 Fictive

We have generated an artificial social network from a small graph of ten fictive users, simulating their behaviour regarding access control policies. For these first experiments, their PR was randomized. In each generated data access request, four users are involved. In order to validate that our algorithm improves equity in the system, we have artificially designed a drastically unbalanced population, where all interactions are in conflict, 10% of the users always answer "deny" and 90% of the users always answer "allow" to a user request. At each milestone, the Gini coefficient is calculated on an average of 100 runs, using NAP/NI as the distributed resource. We have used equation 1 to calculate the Gini coefficients for thirteen milestone iterations in our experiment (after 20, 50, 100, 250, 500, 700, 1000, 1500, 2000, 3000, 3500 and 4000 interactions).

This preliminary experiment outcome is presented in Figure 5. In the "deny-strategy", a single user always wins (and increases her NAP), while others always lose. The result implies a very high inequality, and thus a very high Gini coefficient. The bad result of the "deny-strategy" is not really surprising, since we designed the population in order to obtain this strong inequity where just a person in the society that always wins against everyone else. Indeed, deny-strategy reaches the upper bound of the Gini coefficient, which is given by $(n-1)/n$ [20].

In both "allow" and "majority" strategies, nine user always win (and increase their NAP) and one user always loses. These two strategies mechanically get significantly better results, and a pretty small (and constant) Gini coefficient.

Finally, the "equity strategy" from our proposal achieves the best results in this population, because it ensures that everyone reaches the same enforcement rate, leading to a quick convergence towards 0 for the Gini coefficient. In this basic example at least, no other resolution strategy achieves a better result.

One should note that even though the difference between the Gini coefficients of the equity strategy and of the other "efficient" strategies (all but deny) is only about 0.1, it is still a significant result when it comes to Gini coefficients. As an example, one can find a similar difference between the Gini coefficients measuring the 2005 income inequities in Sweden (Gini 0.250) and Egypt (Gini 0.344) [24].

This experiment simply confirms to us that the basic concepts of the algorithm are sound. However, more realistic experiments are now needed to evaluate the actual impact of our "equity strategy".

4.3.2 Facebook

To simulate a complex and realistic scenario, the first necessary step consists in obtain a large network structure. For that purpose we have used a Facebook social graph comprising 4,039 users, obtained from the Stanford University database[7]. For each user, we randomly generate one or two resources (e.g. photos), and according to research related to face recognition in online albums [29, 30], each resource is tagged and associated to two, three or four users, randomly chosen among friends and friends-of-friends. The Gini coefficient is calculated on an average of 100 runs, using the maximal inequity between normalized NAP and

[7]https://snap.stanford.edu/data/egonets-Facebook.html

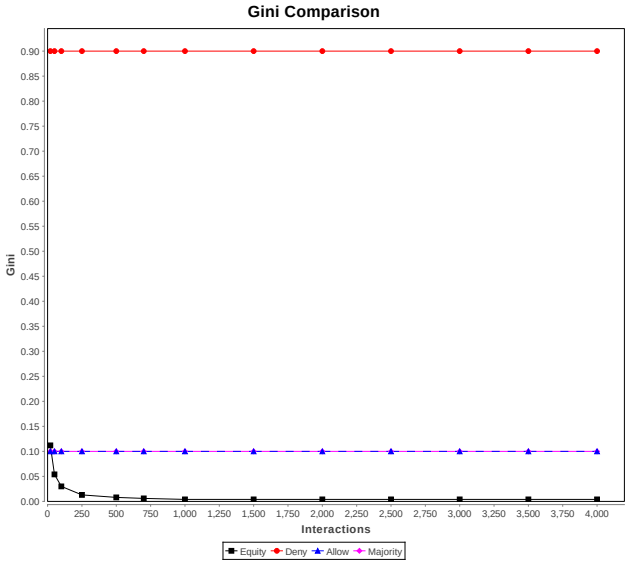

Figure 5: Strategy Results in a Very Unbalanced Population

normalized ND in order to analyze the worst case. We have used equation 1 to calculate the Gini coefficients for seventy eight milestone iterations in our experiment (between 20 and 40000 interactions). The population for our experiment has been designed considering that 20% of the users always answer "deny" and 80% of the users always answer "allow" to a user request, where all interactions are not necessarily in conflict.

In this experiment, the worst case was reached using the normalized ND, which represents the rejection of a user's policy. One of the reasons for high inequity at the beginning of the interaction for all strategies is due to the computation of the Gini coefficient run on few interactions (e.g. 20, 50, 100 ...), where only a very small fraction of the users have interacted and updated their metrics regarding to the population size. The experiment results are presented in Figure 6.

In the "deny" strategy, at each interaction the majority of users have their policy rejected (which increases their ND), while a few have their policy enforced. Then, the Gini coefficient decreases until 0.42 (a Gini coefficient between 0.4 and 0.5 indicates large inequity[8]).

In both "allow" and "majority" strategies, at each interaction a few users have their policy rejected (which increases their ND) while many users have their policy enforced. These two strategies lead to a very high inequality and thus a very high Gini coefficient of 0.79 for "allow" and 0.76 for "majority" (when the Gini coefficient reaches around 0.5, the inequity is considered extremely severe[9]).

Finally, our considered proposal called "equity strategy" achieves the best results achieving a Gini coefficient of 0.21, since it ensures a relatively uniform distribution without much concentration of policy rejection in the population. One should note that, even though the difference between the Gini coefficients of the equity strategy and deny strategy is about 0.2, a similar difference was measured in the 2005's income inequities in Sweden (Gini 0.250) and in Iran (Gini 0.430) [24]. We therefore consider it a significant improvement.

[8]http://www.marketwatch.com/story/china-refuses-to-release-gini-coefficient-2012-01-18

[9]http://www.marketwatch.com/story/china-refuses-to-release-gini-coefficient-2012-01-18

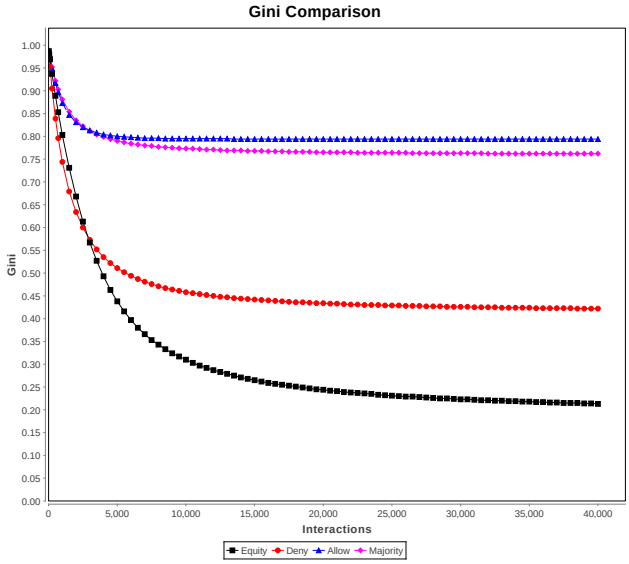

Figure 6: The Strategies' Outcomes Using the Facebook Population

5 Related Works

To the best of our knowledge, we are the first to provide and integrate the concept of equity in the field of policy conflict management, and in particular in the context of privacy policies. We could find no direct equivalent of our development of an equity-preserving algorithm to tackle the issue of equitable enforcement of multiple user policies in social network applications.

Other concepts present in our proposal, however, are far from new and can be found in several studies. It is the case, for instance, of the privacy risk evaluation embedded in the PR variable. In 2011, Hu et al. [26] have proposed a conflict resolution strategy based on the quantification of privacy risk and sharing loss from multiple users, in the context of a collaborative data sharing SNS. To address this issue, every *data controller* (that is every user issuing a policy over the shared data) defines a set of *trusted users* who can legitimately access the data. Any requester then needs to be a trusted user for all data controllers in order to access a content. This approach actually boils down to a strict implementation of the deny strategy. We find this approach too static, since friend behaviour change over time, and friendship links can be broken and re-created. We believe our approach to be more robust and flexible than this specific resolution strategy, since it is not built upon fixed trust relationships.

Squicciarini et al. focus on the collaborative management of privacy settings for shared content by using game theory [6]. Mainly, the Clarke-Tax mechanism promotes a collective policy that aggregates all individuals preferences into a single representative group. The expressiveness of this

model depends on the users' understanding of the Clarke-Tax mechanism, which significantly reduces its usability. We have chosen a very different approach, in which users are not encouraged to reveal their preferences in the form of privacy policies, allowing those policies to be fully dynamic, evolutive, context-dependent and even not fully consistent, depending on how users choose to manage them.

The Gini coefficient is a rather classic tool and has been used in several works on SNSs based on game theory [27, 28]. It is notably used to study how the structure of the network and its dynamics affect social welfare and inequality. Using it to measure inequities in the enforcement of privacy (or security) policies, however, seems to be an original line of work.

6 Conclusion and Future Works

In this paper, we examined the issue of the handling of conflicting privacy policies in the context of Social Network Systems (SNSs). We showed that traditional conflict resolution strategies could lead to inequities among users, some of them being able to see their policies enforced more often than other. We proposed a new, equity-preserving resolution strategy. We exposed the associated algorithm, its proof-of-concept implementation and showed through experimentation that this strategy did actually limit existing inequities in terms of policy enforcement.

Although the first validation experiments are satisfactory, this is only a starting research path. We have introduced an intuitive notion of equity in the field of policy conflict management, but the concepts and algorithms must be refined.

The first thing to do would be to elaborate more complex and realistic scenarios to put the current version of the algorithm to the test. We plan to generate access requests and ruling conflicts at a larger scale, in a multi-agent environment constructed from the actual topology of an existing social network (Facebook being an ideal candidate). User profiles, determining their responses to access requests, will need to be more varied, subtle and realistic, in order to get an idea of the true impact of our resolution strategy on the overall behaviour of a real social network.

The next priority will then be to amend the algorithm in order to make it more distributed, instead of relying on a single Decision Aggregation Point able to record the metrics (ND, NAP, NDC and NI) of all users. It should be made possible to perform the ruling aggregation of an access request at any node of the network, while measuring equity in a reliable way: users should not be able to fake their history metrics, corrupt the algorithm and manipulate the outcome rulings.

Another enrichment of the model would involve taking into account the notion of obligation (associated to a ruling), present in state-of-the-art policy languages and essential in data protection regulations.

If these various goals can be reached, we believe that the corresponding software component would be an interesting tool for applying the principles of *Privacy by Design* to the conception of new social networking software.

It should be noted that although our initial motivation was an improvement of the general level of privacy in such systems, the equity-preserving conflict resolution strategy that we have proposed has a broader application domain. Obviously, it can apply to other kinds of security policy conflicts (and surely enough, we only used very simple access control rules as examples). We also believe that this kind of algorithm can be of use in many multi-party decision taking scenarios in multi-agent systems, even outside the normative context. We suspect that it would be particularly useful in small agent societies, in which local inequalities cannot be compensated by the statistical expectations of a large number of interactions over a large number of entities.

7 Acknowledgement

This work has been partially funded by the grant ARED Presodis of the Bretagne region.

References

[1] N. B. Ellison and D. Boyd, Sociality through social network sites. In The Oxford Handbook of Internet Studies. Oxford University Press, 2013, pp. 151- 172.

[2] Y. Ding, E. K. Jacob, Z. Zhang, S. Foo, E. Yan, N. L. George, and L. Guo, Perspectives on social tagging.Journal of the American Society for Information Science and Technology, vol. 60, no. 12, pp. 2388 - 2401, 2009.

[3] A. F. Westin, Privacy and freedom. New York: Atheneum, 1967.

[4] S. D. Warren and L. D. Brandeis, The right to privacy. Harvard Law Review, vol. 4, pp. 193 - 195, 1890.

[5] G. Müller, Introduction of privacy and security in highly dynamic systems. Communications of the ACM, vol. 49, no. 9, p. 1013 - 1022, 2006.

[6] A. C. Squicciarini, M. Shehab, and F. Paci, Collective privacy management in social networks. in Proceedings of the 18th International Conference on World Wide Web (WWW09). New York, NY, USA: ACM, 2009, pp. 521 - 530.

[7] J. Rawls, A Theory of Justice. Harvard University Press: Atheneum,1971.

[8] C. Gini, Variabilit e mutabilit. Bologna, Italy: Cuppini, C., 1912

[9] P. Anthonysamy, A. Rashid, and P. Greenwood, Do the privacy policies reflect the privacy controls on social networks?. In SocialCom/PASSAT. IEEE, 2011, pp. 1155 -1158.

[10] The European Parliament and the Council, Directive 1995/46/EC of the european parliament and of the council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and

on the free movement of such data. in Official Journal of the European Communities, European Union, Ed., October 1995.

[11] P. Klemperer, Y. Liang, M. Mazurek, M. Sleeper, B. Ur, L. Bauer, L. F. Cranor, N. Gupta, and M. Reiter, Tag, you can see it!: Using tags for access control in photo sharing. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI12). New York, NY, USA: ACM, 2012, pp. 377 - 386.

[12] IBM Tivoli and World Wide Web Consortium, Enterprise privacy authorization language (EPAL 1.2), November 2003. [Online]. Available: http://www.w3.org/Submission/2003/SUBM-EPAL-20031110/

[13] Organization for the Advancement of Structured Information Standards. Extensible access control markup language (XACML 3.0), August 2010.[Online]. Available: https://www.oasis-open.org/committees/tc home.php?wg abbrev=xacml

[14] M. Backes, M. Duermuth, and R. Steinwandt, An algebra for composing enterprise privacy policies. In Proceedings of 9th European Symposium on Research in Computer Security (ESORICS). LNCS, vol. 3193. Springer, September 2004, pp. 33 - 52.

[15] S. Jajodia, P. Samarati, V. S. Subrahmanian, and E. Bertino, A unified framework for enforcing multiple access control policies. In SIGMOD Rec., vol. 26, no. 2, pp. 474 - 485, Jun. 1997.

[16] A. Yamada, T. H.-J. Kim, and A. Perrig, Exploiting privacy policy conflicts in online social networks. Carnegie Mellon University, Tech. Rep. CMU-CyLab-12-005, 2012.

[17] F. A. Hayek, Law, Legislation and Liberty: Rules and Order v. 3: A New Statement of the Liberal Principles of Justice and Political Economy. Routledge & Kegan Paul Books, 1973.

[18] P. Braveman and S. Gruskin, Defining equity in health. Journal of Epidemiology and Community Health, vol. 57, pp. 254 - 258, 2003.

[19] M. Backes, B. Pfitzmann, and M. Schunter, A toolkit for managing enterprise privacy policies. In Computer Security at ESORICS 2003. Lecture Notes in Computer Science, E. Snekkenes and D. Gollmann, Eds. Springer Berlin Heidelberg, 2003, vol. 2808, pp. 162 - 180. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-39650-5 10

[20] S. Yitzhaki and E. Schechtman, More than a dozen alternative ways of spelling gini. In The Gini Methodology. Springer Series in Statistics. Springer New York, 2013, vol. 272, pp. 11 - 31.

[21] M. C. Brown, Using gini-style indices to evaluate the spatial patterns of health practitioners: Theoretical considerations and an application based on alberta data. Social Science and Medicine, vol. 38, no. 9, pp. 1243 1256, 1994.

[22] A. Ahmad and B. Whitworth, Distributed access control for social networks. In 7th International Conference on Information Assurance and Security (IAS11), Dec 2011, pp. 68 - 73.

[23] R. Marin, G. Piolle, and C. Bidan, An analysis grid for privacy-related properties of social network systems, In 5th International Conference on Social Computing (SocialCom13), Sept 2013, pp. 520 - 525.

[24] R. Grabowski, S. Self, and M. P. Shields, Economic Development: A Regional, Institutional, and Historical Approach. M. E. Sharpe, 2007.

[25] A. Besmer and H. Richter Lipford, Moving beyond untagging: Photo privacy in a tagged world. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI10). New York, NY, USA: ACM, 2010, pp. 1563 - 1572.

[26] H. Hu, G.-J. Ahn, and J. Jorgensen, Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In Proceedings of the 27th Annual Computer Security Applications Conference (ACSAC11). New York, NY, USA: ACM, 2011, pp. 103 - 112.

[27] B. Ni, Y.-H. Chang, and R. Maheswaran, Social welfare and inequality in a networked resource game with human players. In IEEE International Conference on Social Computing / IEEE International Conference on Privacy, Security, Risk and Trust, pp. 967 - 970, 2013.

[28] Z. Li, Y.-H. Chang, and R. Maheswaran, Graph formation effects on social welfare and inequality in a networked resource game. In Social Computing, Behavioral-Cultural Modeling and Prediction, ser. LNCS, A. Greenberg, W. Kennedy, and N. D. Bos, Eds. Springer, 2013, vol.7812, pp. 221 - 230.

[29] M. Davis, M. Smith, J. Canny, N. Good, S. King, and R. Janakiraman. Towards context-aware face recognition. In Proceedings of the 13th annual ACM International conference on Multimedia, pages 483 - 486, New York, NY, USA, 2005. ACM.

[30] M. Naaman, R. B. Yeh, H. Garcia-Molina, and A. Paepcke. Leveraging context to resolve identity in photo albums. In Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries, pp. 178 - 187, New York, NY, USA, 2005. ACM Press.