

Establishing Global Policies over Decentralized Online Social Networks

Zhe Wang, Naftaly H. Minsky

Department of Computer Science

Rutgers University

Email: {zhewang, minsky}@cs.rutgers.edu

Abstract—Conventional online social networks (OSNs) are implemented in a centralized manner. Although centralization is a convenient way for implementing OSNs, it has several well known drawbacks. Chief among them are the risks they pose to the security and privacy of the information maintained by the OSN; and the loss of control over the information contributed by individual members.

These concerns prompted several attempts to create decentralized OSNs, or DOSNs. The basic idea underlying these attempts, is that each member of a social network keeps its data under its own control, instead of surrendering it to a central host; providing access to it to other members of the OSN according to its own access-control policy. Unfortunately all existing DOSN projects have a very serious limitation. Namely, they are unable to subject the membership of a DOSN, and the interaction between its members, to any global policy.

We adopt the decentralization idea underlying DOSNs, complementing it with a means for specifying and enforcing a wide range of policies over the membership of a social community, and over the interaction between its disparate distributed members. And we do so in a scalable fashion.

I. INTRODUCTION

An *online social network* (OSN) can be defined broadly as a *community* of people that interact with each other via some electronic media, which generally operates subject to a *policy* that may regulate the membership of the community, and the manner in which its members interact with each other. The policy of a purely social community is often informal, imprecise, implicit and only occasionally enforced. But such policy needs to be tightened for an OSN, because its membership can be larger than that of a purely social community, and its members tend to be less familiar with each other. Therefore, the policy of an OSN needs to be explicit and well defined, and it needs to be more strictly enforced, largely via computational means, so that it can establish desired regularities over the OSN.

Such policies are easily implementable via the conventional types of OSNs—such as the currently popular Facebook, Google+, and Twitter—because of their *centralized architecture*. That is each such OSN employs a virtually *central host*—which may be a centrally managed cluster of computers—that mediates all interactions between its members, subject to policies defined by the host. This central host also maintains the information supplied by the members of the community.

Unfortunately although centralization is a very convenient way for implementing OSNs, it has several well known drawbacks, which include: (a) lack of scalability; (b) the existence

of a single point of failure; (c) the risks to the security and privacy of information maintained by the central host; and (d) the loss of control over the information contributed by individual members. The first two of these drawbacks can be mitigated via very large, complex, and expensive infrastructures—like those used by Facebook and Twitter.

But the risk to security and privacy, and the loss of control over private information are harder to mitigate, because they are mostly the consequence of centralization itself. Indeed, maintaining the state of the membership of an OSN, and the history of interaction between members, under a single administrative domain makes it vulnerable to various malicious attacks. Such attacks can be mounted by insiders, say the programmer that maintains the software of the OSN; and by hackers from the outside, for whom the central repository of information is likely to be very lucrative.

Security seems not to be of much concern to the hundreds of millions of current users of Facebook, Twitter, and similar OSNs. But they are, or should be, of serious concern to other types of OSNs, whose members exchange more sensitive information—such as private medical and financial information; and information about the business of an enterprise, exchanged between its employees. We will consider examples of such OSNs in the following section.

Such concerns about centralized OSNs prompted several attempts to create decentralized OSNs, or DOSNs; such as LotusNet [1], Safebook [2], PeerSoN[3], and others. The basic idea underlying all these attempts to the decentralization, is that each member of the community in question should keep its data under its own control, instead of surrendering it to a central host, providing access to it to other members of the DOSN according to its own access-control policy.

Unfortunately all existing DOSN projects have a very serious limitation. Namely, they are unable to subject the membership of a DOSN, and the interaction between its members, to any global policy. This is a very serious limitation of the DOSN architecture, because, as pointed out above, an enforced global policy is generally essential for an OSNs, as it helps make it into a social community.

The Contribution of this Paper: We will adopt in this paper the decentralization idea underlying DOSNs, complementing it with a means for specifying and enforcing a wide range of policies over the membership of a social community, and over the interaction between its disparate distributed

members. And we shall do so in a scalable manner.

The rest of this paper is organized as follows. Section II introduces examples of OSNs for which security is critical, and would thus benefit from decentralization. Section III introduces examples of policies that are often essential for an OSN—particularly for the types of OSNs for which security tends to be critical—but which cannot be established under DOSN. Section IV provides a very brief outline of the LGI middleware, which serves as the basis for this work. Section V introduces our model of decentralized OSN—we call it OSC, for “online social community,” where the term “community” is meant to suggest two things: first, a decentralized nature, like that of purely social communities; and second, the existence of a shared policy, which characterize most social communities, and which under OSC is enforced. Section VI is an implemented case study that demonstrates how this abstract model can be used for a concrete application. The related works are discussed in Section VII. And we conclude in Section VIII.

II. EXAMPLES OF OSNS, FOR WHICH SECURITY IS IMPORTANT

We distinguish here between two types of OSNs: (1) *autonomous OSNs*, which are not bound by any outside authority; and (2) *bound OSNs*, which operate in the context of some organization, which has jurisdiction over it. We focus in this section on the security needs of these two types of OSNs, and on the risks to security that centralization poses to them. We will discuss both types of OSNs, but we will focus, here and in the rest of the paper, on the latter one.

A. Autonomous OSNs

Consider a set of physicians who form an OSN that enables them to consult with each other about various medical issues they confront. This OSN is to admit only qualified Medical Doctors as members, and may grow to be quite large if physicians all over the world join it. The information exchanged between members of such OSN is clearly very sensitive.

Having the consultation process mediated by a central host, and having the information exchanged between the physicians maintained centrally by this host can seriously compromise the security of both the doctors and their patients. The risk here is particularly serious because the host of such an OSN is likely to become a target for attackers, since the information maintained by it can be exploited for illicit financial gains.

There are many potential autonomous OSNs that exchange similarly sensitive information; such as an OSN formed by a certain type of workers for exchanging their views about their employers; an OSN used by students to exchange information about their teachers, and many others.

B. Bound OSNs

There is a growing realization[4] that OSNs that operate within an organization—such as manufacturing, commercial enterprises, medical centers, or even the military—can be beneficial for it. This seems to be particularly the case for OSNs that provide for micro-blogging, as is evident from the

recent purchase of the Yammer—a prominent micro-blogging OSN operating within organizations—by Microsoft, for \$1.2 Billion. We will have more to say about Yammer itself, but first we outline some of functional features one can expect from this kind of OSN.

Consider a large and geographically distributed enterprise E that provides a centralized micro-blogging OSN for its employees. Suppose that such an OSN, which we call B_E , distinguishes between groups of employees, enabling the members of each groups to communicate with each other. Such groups may be the following: (a) all the employees of E ; (b) the non-managerial staff of E ; (c) the managerial staff of E ; and (d) members of various task forces operating in E . Note that these groups may overlap partially, as a single employee may belong to several groups.

The Need for Security: We will distinguish here between two types of needed security. First, the information exchanged between the employees of enterprise E can carry sensitive information about the business of this enterprise. It is therefore important for this information not to be exposed to the outside, at least not on a large scale.

Second, one may need to prevent information exchanged between the members of a certain group of B_E from being accessible to anybody else, or to certain other groups. For example, suppose that enterprise E has several task forces that consult to other companies, some of which may compete with each other. And suppose that the members of each such task force form a group in B_E . It is obviously paramount for these subgroups not to have access to each other’s information.

The Risks to Security due to Centralization: There are two types of centralization to be considered, which we call strong and weak. Strong centralization is like the one practiced by Yammer, the Microsoft OSN that we mentioned above. Yammer provides services to a host of different enterprises—they currently claim to serve about 200,000 of them. Yammer supports policies that provide necessary separation between the various enterprises it serves. But the information belonging to all these enterprises is maintained centrally by the Yammer system. Such centralization of commercial and industrial information of many companies, is likely to attract attacks from the inside of Yammer, and from the outside.

A better approach would be to use an intramural Yammer-like OSN. This, weaker form of centralization, would be much safer than using Yammer. But if this system relies on a centralized database, it would still be vulnerable to breaches of security. Indeed, if all the information generated by the B_E is available to its software, then the rogue programmers of this OSN will have a fairly free access to all of it, disregarding the required boundaries between different groups.

III. A SAMPLE OF POLICIES THAT AN OSN MAY NEED TO ENFORCE

We illustrate here the type of *communal policies* that an OSN may need to establish. By “communal” we mean either global policy that is to govern all members of an OSN, or a policy that governs some subgroup of its members. All

the policies discussed here can be easily established by a centralized OSN, but none of them can be established under the DOSN architecture.

We distinguish here between three types of communal policies, and motivate some of them in the context of the B_E OSN, introduced in previous section. We show in Section VI how such policies can be realized in decentralized OSNs.

A. Membership Control

Control over membership is crucial to social communities, whether it is autonomous or bound. Such control may have several complementary aspects. We describe three of these.

First, one may require that to be a member of a given OSN one needs to authenticate itself via a specified kind of certificate. One may think that this policy can be established under the DOSN architecture by having every member of the DOSN in question require every interlocutor of his to authenticate itself in a specified manner. But DOSN has no way for ensuring that all its members behave in this way. (The inability of DOSN to enforce communal policies is even more obvious in the rest of our examples below.)

Second, one may require that to be a member of an OSN one needs to garner the approval of several current members.

Third, it is often important to establish some procedure for removing members from a given OSN. This can be done in many ways. For example, consider an OSN that has a member that plays the role of a *manager*. Now, let the manager be given the power to remove any current member x of the OSN, simply by sending it a message *remove*. Then x should lose its ability to interact with other members of the OSN.

B. Constraints on the Behavior of Members of an OSN

Sometimes one needs to impose constraints on what members can do. Such constraints may depend on the profile of individual members, and on the history of their interaction with others. We have just seen an example of such constraints: only a member that plays the role of manager can send the *remove* message to others. And any member that gets such a message must cease all communication with others.

More generally—but stated in the context of the B_E OSN—the type of messages that members are allowed to send, or the type of posts that they are allowed to issue, may depend on their roles in this OSN, which may be represented by their profiles. As another example, a member should be able to force an interlocutor to reduce the frequency of messages it is sending to it, or to cease sending messages to it altogether.

C. Global Access Control (AC) Policies

One of the intended consequences of decentralization under DOSN is that it enables each member to apply its own AC policy to its own data—e.g., to the set of posts it produced, which are maintained in its own database. The problem with this aspect of DOSN is that, unlike the case of Facebook or Twitter, a member of an OSN may not have the complete authority over the data it maintains. A case in point is a bound OSN, such as the B_E OSN introduced in Section II-B. The

posts being produced by the various members of this OSN really belong to the enterprise E , which thus has the ultimate authority about how they should be distributed. The enterprise may relegate to individual members the right to apply their own AC policies, *provided* that these policies conform to the global policy of an enterprise. For example, the global policy of the B_E may be that a group of members assigned to deal with the business of a given client-company can communicate only with each other, as long as they operate as members of that group—recall that under B_E , a single person may belong to several groups.

IV. THE LAW-GOVERNED INTERACTION (LGI) MIDDLEWARE—AN OVERVIEW

LGI is a middleware that can govern the interaction (via message exchange) between distributed *actors*, by enforcing an explicitly specified law about such interaction. We provide here a brief, and rather abstract, overview of LGI; focusing on what is the most relevant to this paper. A more detailed presentation of LGI, and a tutorial of it, can be found in its manual [5]—which describes the release of an experimental implementation of the main parts of LGI.

A. The Local Nature of Interaction Laws

Although the purpose of interaction laws is to govern the exchange of messages between different distributed actors, they do not do so directly under LGI. Rather, an LGI law \mathcal{L} governs the interaction of any actor operating under it, essentially by controlling its ability to send messages to others, and to receive messages from them. A law \mathcal{L} is local to each actor x operating under it, in that its rulings are based solely on the local state of x and on the event that occurs at it, and are completely independent of the coincidental state and events occurring anywhere else in the system. Such a law can be enforced locally, and thus very scalably, in a manner described in Section IV-B. Moreover, the locality of LGI laws has several other beneficial consequences, some of which will be pointed out in due course.

It should also be pointed out that although locality constitutes a strict constraint on the structure of laws, it does not reduce their expressive power. In particular, despite its *structural locality*, an LGI law can have *global effect* over what is called an \mathcal{L} -community, defined as the set of actors operating under a common law \mathcal{L} .

B. The Decentralized Law Enforcement Mechanism

Consider an actor x that chooses to operate under a law \mathcal{L} . It can do so by *adopting* a generic controller as its mediator, loading law \mathcal{L} into it. Once thus adopted, this controller is denoted by $T_x^{\mathcal{L}}$ —meaning that it operates under law \mathcal{L} , serving actor x —and the pair $\langle x, T_x^{\mathcal{L}} \rangle$, is called *agent* x and is referred to as an \mathcal{L} -agent—and sometimes simply an “agent”. This adoption, which signifies the *birth* of agent x , is one of the interactive events of LGI, so that the law in question has the possibility of refusing to be adopted by this actor, and can mandate some initialization for it, if it does not refuse.




Fig. 1. Interaction between a pair of agents, mediated by a pair of controllers under possibly different laws.

Note the fundamental difference between a bare actor and its agent: while the interactive behavior of an actor is unpredictable—unless its code is known—the interactive behavior of an \mathcal{L} -agent is known to conform to law \mathcal{L} .

Figure 1 depicts the manner in which a pair of agents, operating under possibly different laws, exchange a message. (An agent is depicted here by a dashed oval that includes an actor and its controller.) Note the *dual nature* of control exhibited here: the transfer of a message is first mediated by the sender’s controller, subject to the sender’s law, and then by the controller of the receiver, subject to its law. This dual control, which is a direct consequence of the local nature of LGI laws, has some important consequences. In particular, it facilitates flexible interoperation and it enables more sophisticated control than possible under many AC mechanisms that provide control only on the receiver side.

The overhead incurred by this kind of control turns out to be relatively small. In circa 2000 it was measured to be around 50 microseconds for fairly common laws, which is negligible for communication over WAN. This is one of the results of a comprehensive study of this overhead in [6].

Finally, we note that we replace the central host (of a social network) with a group of generic controllers, which are much simpler and designed not to keep any information they pass through. A generic controller needs to be trusted to enforce correctly any law it is adopted with. There are several ways for providing such trusted controllers as the TCB (Trusted Computing Base) of the system in question. In the case of a bound OSN, like our B_E example, we expect this to be done by the enterprise E , in the context of which B_E operates. This company could construct what is called a *controller service* (CoS) that maintains a set of well tested *controller pools*, each of which can host a number—it is usually in the hundreds—of individual controllers that can be used by arbitrary actors, upon request. For other types of OSNs one expects the CoS to be maintained by some commercial company that provides its services for a fee.

Note, therefore, that a controller $T_x^{\mathcal{L}}$ and the actor x that adopted it would run on different hosts. This would help prevent x corrupting its own controller.

V. A MODEL OF DECENTRALIZED OSN

We introduce here a model of decentralized OSNs that differs from the current approach to the decentralization employed under the current DOSN architecture, in that it enables the enforcement of communal policies over it. We call a specific OSN under this model an *online social community*

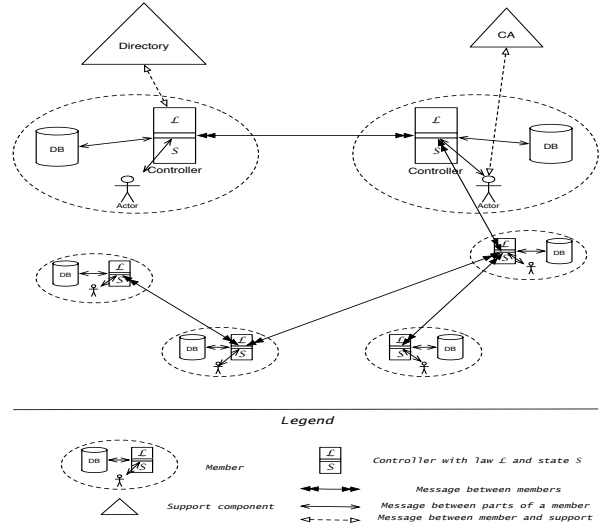


Fig. 2. The Anatomy of an OSC Community

or an OSC (or sometime simply a *community*), and we refer to this model itself as the OSC-model.

Now, a community C under the OSC model is broadly defined as a 4-tuple $\langle M, \mathcal{L}, T, S \rangle$, where M is the set of members of C ; \mathcal{L} is the policy that governs this community, which we call a *law* (an LGI-law, to be exact); T is a set of generic LGI controllers that serve as the middleware that enforces law \mathcal{L} ; and S is a set of components that support the operations of C , and is specific to it—this set is called the *support* of C , and it may be empty.

We now elaborate on this schematic definition of the OSC model by discussing the following aspects: (1) the anatomy of a community under OSC; (2) the launching of an OSC-community; and (3) the operations of a community. Note that an example of an OSC-community is described in Section VI.

A. The Anatomy of a Community Under OSC

We describe here the anatomy of a community C under this model by elaborating on its various components, and on the relations between them. This anatomy is depicted schematically in Figure 2.

The Set M of Members: An individual member m of a community C is a triple $\langle user, mediator, database \rangle$, where *user* is usually a human, operating via some kind of computational platform, like a smart phone; *mediator* is an LGI-controller that mediates all interactions between this member and the rest of the community—as well as between the other two components of the member in question—subject to law \mathcal{L}_C ; and *database*, which is an optional part of the member, is the private database of m that maintains information associated with this member, such as the set of Twitter-like micro-blog posted by m , or its Facebook-like page.

The Law \mathcal{L}_C of community C : The law endows an OSC-community with its overall structure and behavior. And the fact that the law can, in principle, be any well formed LGI law

(cf. Section IV) endows this model with great deal generality regarding the policy that can be enforced over a community.

The set T of LGI Controllers: Every user can create its own controller, using the software provided by the released LGI middleware. But if malicious corruption of controllers by their users is of concern, then it is better for the members of a community to adopt controllers created and maintained by a trusted *controller service* (CoS), so that they can authenticate each other as bona fide LGI controllers. For such a CoS to be trusted to provide genuine controllers, this service needs to be managed by a trusted organization. In particular, the CoS may be managed by the organization in the context of which the community is to operate—as in the case discussed in Section VI. For more about the security and trustworthiness of controllers see Section IV and [5].

The Support S : An OSC-community may require services of various components that are not themselves members of this community. Here are some examples of such components: (a) a certification authority (CA) used for the authentication the various members of the community; (b) a *naming service* that provides unique names of community members; (c) an index service for searching; and (d) a networking service for maintaining various networking structures of the community—more about which in Section V-D. It is worth pointing out that this set of support components may be empty for some communities.

B. The Launching of an OSC-Community

A specific OSC-community, C is launched by constructing its *foundation* and then having individual members join it. The construction of the foundation of a community C consists of the following steps: (a) defining law \mathcal{L}_C under which this community is to operate; (b) implementing the required support components; and (c) selecting, or constructing, a *controller-service* (CoS) for the use of this community.

Once the foundation of C is constructed, anybody can attempt to join it as a member, via the following three steps: (a) deploying its private database—if one is required by law \mathcal{L}_C —with an API required by this law; (b) adopting an LGI-controller, and loading law \mathcal{L}_C into it; and (c) providing this controller with a pointer to its database, if any. It should be pointed out that such an attempt to join a given community C may fail, if the conditions for joining imposed by law \mathcal{L}_C are not satisfied.

C. The Operation of a Community

Consider a member x of a community C sending a message m to another member y . The message first arrives at the controller of x , that operates under law \mathcal{L}_C . These controllers would then carry out the ruling of law \mathcal{L}_C , which can mandate the execution of any number of the following kind of actions: (a) change its own state in some way; (b) communicate with the database of x ; (c) send the message m , or some other message, to the controller of the original target of x ; and (d) send some other messages to the controllers of some other members, or to some of the support components of the

community. Among other things, this means that members of a community interact with each other via their controllers, and the controllers communicate with each other.

It is worth pointing out here that LGI provides an important *trust modality* which is critical to this model. This trust modality is called *law-based trust*, or simply *L-trust*, and can be introduced, broadly, as follows: any pair of interacting LGI-controllers can identify, cryptographically, each other as genuine controllers, and can identify the law, under which their interlocutors operates. One consequence of this is that the law \mathcal{L}_C of the given community C can be written so that members of C can interact only with other members of C . Now, *L-trust* can be defined as follows: *members of a community C can trust each other's interactive behavior to comply with their common law \mathcal{L}_C* . Another important observation about the behavior of a community under this model needs to be made: the ruling of a law for a given event that occurs at a controller depends on the state of this controller, which may be different for different members. This difference can come from some certificates submitted by the user to its controller, which may authenticate the role of the user in the organization in question. And the state may change dynamically in response to some interactive activity of the community. For example, the manager of the community under our B_E community, may be allowed by the law of B_E community to transfer its managerial baton to some other member, which would then be able to send *revoke* messages. In other words, *the members of a community C may not be equal under its law \mathcal{L}_C* .

D. Crowdsensing

We have already pointed out that some capabilities are easier to be provided via centralized OSN than via decentralized one. We have focused on the imposition of communal policies over an OSC in this paper. Another capability that is problematic under decentralized OSNs is the ability to analyze the networking relationship implicit in the community. Consider for example the *friend* relationship of Facebook, and let us examine its realization in an OSC.

It is easy to have each member of an OSC list his friends—we have done with a similar relation in Section VI—but it is very hard and expensive to analyze the entire friendship-graph, when this relation is recorded in such a distributed manner. Of course, such global analysis, which is central to Facebook, is not required for all kinds of OSN. But it is often required, and must be provided for.

A reasonable way for enabling global analysis of a network implicit in an OSC, is to maintain it explicitly in a central manner. That is, we maintain the friendship relation (or any other kind of relation between members) in a central place, as part of what we have called the *support* of an OSC, and then provide these components with various analysis tools. This is a reasonable solution under two conditions: (a) the relation in question is not, itself, highly sensitive from the privacy and security viewpoint; and (b) the central *network component* is not used too frequently, so it would not reduce substantially the scalability of the OSC in question.

More generally, an OSC may have several centralized support components, such as indices of various kinds. If these components are not used very frequently, they would not seriously undermine the scalability of it, due to the decentralization of its data and policy enforcement mechanism.

VI. A CASE STUDY

In this section, we describe the implementation of the B_E community, introduced in Section II. It has been implemented in the scale of more than two hundred users as a proof of concept. This community operates in the context of a large and geographically distributed enterprise E, providing a micro-blogging OSN for its employees, as a complement to its existing office systems. B_E enables the members of various groups of employees to communicate with each other. The groups could partially overlap, in the sense that a single employee may belong to several groups.

As described in Section II, there are two modes of communication in this community: publish/subscribe and direct message. And each post or message contains two parts: type and body. For both modes, there are certain global policies can be imposed to the community to control members' behaviors. We will discuss them in the following section.

Each member of the community holds a profile in its controller, as well as several internal states. A profile is a group of attributes of the user. There are mainly two types of attributes in profile. One type is the relatively stable attributes, like real name, position, etc. These items usually require users to provide certificates in order to get them in the profile by the rule of authentication. Since these attributes are stable, an index for searching is able to be built on them. Another type of attributes is the dynamic ones, such as interest, skill set, last ten posts, etc. Although these items don't require certificate, not all of them can be changed by the member arbitrarily. For example, subscriber list is handled by the subscription mechanism, reputation is maintained by controller according to the rates gotten from other members and an attendance attribute could be decided by the sign-in/sign-out time. We call these user-unchangeable dynamic attributes and the certified attributes together as controlled attributes, and the rest attributes as discretion attributes. The internal states are the states maintained by the controller for certain functions of the community. For instance, the frequency of publishing is used for preventing a member overwhelming the community by violently publishing posts.

The Law of the B_E Community:

The law \mathcal{B} of the B_E Community is used for regulating every aspects of the operations and behaviors of the community. We split it into several parts according to their functionalities. Due to lack of space, we only discuss the law of some functionalities of the communities. In Section VI-A, we discuss how a user becomes a member of the B_E community and its groups, how it configures its profile, and how a member is removed. Section VI-B shows the communication mechanism

and the imposition over it. We discuss other functionalities which are needed to be a complete OSN in Section VI-C.

A. Member Profile and Membership Control

To join the community, a member needs to adopt a controller under law \mathcal{B} . Rule $\mathcal{R}1$ allows a user to join the community by presenting a certificate from a CA run by the enterprise in question to prove that it is an employee. Once certificate is verified by the controller, the set of attributes in its profile will be inserted into the user's control state. An example of an attribute is $role(manager)$. There are two types of attributes in the profile: certified and uncertified. The certified attributes are the relatively stable ones, like real name, login ID, position, age, etc. These items can also be obtained by providing certificates after the adoption. Rule $\mathcal{R}2$ allows the user to join the group t_i by providing a group certificate. It will add an attribute t_i to its profile, as well as an access control filter which we will discuss in the next section.

```

 $\mathcal{R}1$ 
  UPON adopted(X, cert(issuer(ca), subj(X), attr(A))) :-
    do(+A).

 $\mathcal{R}2$ 
  UPON certified(X, cert(issuer(ca), subj(X),
    attr( $t_i$ ))) :-
    do(+ $t_i$ );
    do(+filter(group( $t_i$ ))).

 $\mathcal{R}3$ 
  UPON sent(X, addProfile(Attribute(Value)), X) :-
    if (  $\neg$  (Attribute in controlledAttributes) )
    then do(+Attribute(Value)).

 $\mathcal{R}4$ 
  UPON sent(X, updateProfile(Attribute(Value)), X) :-
    if (  $\neg$  (Attribute in controlledAttributes) )
    then do(-Attribute); do(+Attribute(Value)).

 $\mathcal{R}5$ 
  UPON sent(X, addFilter(Attribute(Value)), X) :-
    do(+filter(Attribute(value))).

 $\mathcal{R}6$ 
  UPON sent(X, #revoke#, Y) :-
    if(role(manager)@CS) then do(Forward);
    else do(Deliver(X, notAllow, X)).

 $\mathcal{R}7$ 
  UPON arrived(X, #revoke#, Y) :-
    update(certificateBlacklist);
    inform(certificateBlacklist);
    do(Quit).

```

Fig. 3. Law \mathcal{B} : Member's Profile and Membership Control

Another type of attributes is the dynamic ones. Although these items don't require certificate, not all of them can be changed by the member arbitrarily. For the discretion attributes, user can directly add some of them into its profile via Rule $\mathcal{R}3$ and update them via Rule $\mathcal{R}4$. Rule $\mathcal{R}5$ shows how a user sets up its subscription filter. Sometimes the user may not want to be subscribed by everyone. The existing way to do that in other social networks is to put the subscriber into blacklist, or we say to block specific user. This can only happen after somebody initiated or requested the subscription and needs to be done manually. However, our mechanism can prevent subscription by specifying a certain kind of attributes. User can use Rule $\mathcal{R}5$ to add the filter content into control state. Its controller will only allow the members who have the required attributes to subscribe to it. The following operations will be described in Section VI-B.

Finally, rules $\mathcal{R}6$ and $\mathcal{R}7$ regulate the removal of members from the community. Rule $\mathcal{R}6$ shows that only the manager role can remove a member from the community. Non-managers are not allowed to use the type *revoke* when sending messages. When the *revoke* message arrives at the member's controller, according to Rule $\mathcal{R}7$, the controller will directly terminate the connection to the actor and then put its certificate to the blacklist and broadcast to all controllers. Next time when another actor tries to use this certificate to adopt a controller, the controller will not verify it. The member has no way to avoid that. This rule guarantees that its participation in this community is seized immediately after the manager removed it and cannot get back again using the same certificate. This is just an example of how to handle the membership removal. Other methods, such as suspension, can also be supported.

B. Communication

For a member to subscribe to the posts from another member, the subscriber s sends a subscription request to the publisher p . When p receives the request, it will add s to its subscriber list unless such subscription is prohibited by the law, or if p itself blocks the subscription. When a post is published by a member, it will be automatically pushed to all its subscribers. Moreover, members can also send direct messages to each other, which can also be controlled.

We will show later in this section, how the communication is enabled and controlled. The control over communication has two complementary parts: global control and local control. The global control is imposed on every member of the community, but can be sensitive to the state of members, while the local control is discretionary to each member. We discuss both of controls below, and the according law later.

Global Control: The global control over P/S is imposed on both publishing and subscription. The control over publishing is on what types of posts members can publish. For example, only managerial staff can publish posts with type *management*. Upon publishing, a management post is allowed to be published only when the member has the attribute *role(managerial)* in profile.

The control on the subscription regulates who can subscribe to whom, and to which types of posts. Essentially, it is defined by a condition C on the profiles of the publisher and subscriber. An example of such global policies is that only the members from a same group can talk to each other. The problem is that there is no single place where these profiles can be evaluated because of the decentralization. To solve this problem, our law forces every subscription request to include the profile of the subscriber. And then it has the condition C to be evaluated and acted upon at the publisher side. This can be achieved by checking the profiles of the publisher and subscriber and rejecting the subscription request if the two members are from different groups.

The control over sending direct message is similar to the one over publishing. Certain types of direct messages are allowed to be sent only when the members have the required attributes

in their profiles. For instance, only the manager role can send the *revoke* message.

The control on receiving the direct messaging is different from the one on subscription. Whenever a member sends a direct message, its controller will append its profile to the message. Upon the arrival of the message at the receiver side, the controller of the receiver will not deliver the message if its profile does not satisfy the condition of receiving it.

Local Control: Sometimes, a member does not want to be subscribed by certain members, it can block the subscription requests from them. To achieve this, we introduce a profile attribute called *filter*. If a member adds a filter *filter(X)* in its profile, its cannot be subscribed by the member who has attribute X in profile. As described above, whenever a subscriber s sends the subscription request to a publisher p , it will be forced to attach its profile along with it. When the request arrives at the publisher's controller, s will not be added to p 's subscriber list if its profile has the banned attributes in filter. This rule of filter is just an example. More complex uses of the filter, such as OR or XOR, could also be achieved.

The Law: The rules of law \mathcal{B} that implements these provisions are defined in Figure 4 and described below.

```

 $\mathcal{R}8)$ 
UPON sent(X,publish(P),X) :- group( $t_i$ )@CS
if (typeof(P) == #management# and  $\neg$ 
role(manager)@CS)
then return;
updateProfile(lastTenPosts(P));
updateDB(P);
if(subList[group( $t_i$ )] = []) then return;
else forEach(subscriber in subList[group( $t_i$ )])
do(Forward(X,P,subscriber)).

 $\mathcal{R}9)$ 
UPON arrived(X,P,Y) :-
do(Deliver);
do(inform(X,P,Y)).

 $\mathcal{R}10)$ 
UPON sent(X,requestSubscribe(profile),Y) :-
do(Forward).

 $\mathcal{R}11)$ 
UPON arrived(X,requestSubscribe(profile),Y)
:- group( $t_i$ )@CS
if(filter(Attribute(Value))@CS and
Attribute(Value)@profile)
then do(Forward(Y,subscribeNotAllowed,X));
else do(updateSublist[group( $t_i$ )]);
do(Forward(Y,subscribeAllowed,X)).

 $\mathcal{R}12)$ 
UPON sent(X,M,Y) :-
do(Forward(X,M(profile),Y)).

 $\mathcal{R}13)$ 
UPON arrived(X,M(profile),Y) :-
if(group@CS == group@profile)
then do(Deliver(X,M,Y)).

```

Fig. 4. Law \mathcal{B} : Communication

In Rule $\mathcal{R}8$, when the user wants to publish a post to its subscribers, the controller will read local subscriber list and push the post to each of them. It will also update its attribute *lastTenPosts* of its profile in its control state. When the subscriber receives the post or message, according to the Rule $\mathcal{R}9$, controller will show the post to the user. In the meantime, to handle the situation that the user is not online, the controller can store the post or message into local file system, or use other ways to inform users, for example by

sending as email. Also, the controller can save the last several posts for the search functions. Note that if the post is of the type of management, only the managerial staff can publish it.

According to Rule $\mathcal{R}10$, any user can send a subscription request to any user. The controller attaches its profile to the request. In Rule $\mathcal{R}11$, when the request arrives at the user, the controller will check whether there is an access control filter in its control state. If there is not, it adds the request user to the subscriber list. If there are filters, it will examine whether this user satisfies by checking the required attributes of the profile. If the requester satisfies, the controller will add it to the subscriber list and send back the result to the request user.

If a member wants to send a direct message to a specific member, its controller will append its profile to the message, as shown in Rule $\mathcal{R}12$. When the message arrives at the controller of the destination member, it will check the group id in control state to see whether it matches the group id of the sender. If it does, the member is allowed to read the message. If the member belongs to a different group from the sender, the controller will discard the message, according to Rule $\mathcal{R}13$.

C. Other Implemented Functionalities

Due to lack of space, we do not cover all parts of the law for this community. However, following functionalities are very useful and important in forming a complete OSN. We discuss the general idea of them.

Naming and Addressing: When an agent joins a community, it must have a way of naming and locating other members of the community. After all, one joins a community only if one wishes to interact with some members. We employ a server, called secretary, that simply acts as a naming and locating service, negotiating with agents wishing to join the group in order for each agent to have a unique name within that group. More details about this mechanism are provided by [7].

Search: Search capability is also necessary for an OSN. It's relatively straightforward in the centralization, comparing to be achieved in a decentralized manner. The Distributed Hash Table, which is used by most of DOSN approaches, cannot do content search. The content search can be achieved via a gossip style search protocol—the search query initiator sends the query to its neighbors and then the neighbors forward the query to their neighbors. Both types of search, especially the DHT, are not secure and easy to be undermined[8], because they need the untrusted members substantially to carry out the search correctly. It's very vulnerable if there is no regulation imposed on each participant. We can make DHT secure by implementing it by law with similar technology discussed before, but this is beyond the scope of this paper.

VII. RELATED WORK

The concern about the security issues of centralized OSNs motivated several attempts to decentralize OSNs, creating several versions of what is called DOSNs. PeerSoN[3], Safebook[2], and LotusNet [1] are the main attempts among others. The basic idea underlying all these projects is that each member of the social networks keeps the data under its own

control, instead of surrendering it to a central host. This is a necessary measure of decentralization, but it is not sufficient. Because, as we explain in Section III, social network requires some global, or communal, policy to operate under. But none of the attempts known to us at the implementation of DOSNs provides any means for establishing such policies.

Moreover, all these attempts adopt the substrate of DHT to implement the p2p design. As we discussed in Section VI-C, DHT itself is not secure under the context of heterogeneous and distributed network and easy to be compromised. It's not able to defend some attacks if it cannot establish certain global policy to protect it. Furthermore, DHT is incapable of performing content search. Though some improvements or work-arounds are employed to provide limited content search, these are way off the basic requirement of an OSN.

VIII. CONCLUSION

This paper addresses the risks to privacy and security posed by conventional centralized OSNs. These risks, which are the consequence of centralization, seem not to be of concerns to most of the clients of OSNs such as Facebook and Twitter. But they are, or should be, of serious concerns to many other current and potential applications of OSNs.

Several recent attempts have been made to decentralize OSNs, by letting each member of such a network keep maintaining its own data. But this DOSN approach to decentralization is not able to establish any kind of regularity over the social network, which is necessary for both real life social community, as well as for OSNs.

We have introduced a decentralized architecture of OSNs, called OSC, for *online social community*, which is able to establish regularities concerning both the membership of OSC and the manner in which its members interact. The preliminary testing and experiments of our implementation show that our method is feasible and promising.

REFERENCES

- [1] L. M. Aiello and G. Ruffo, "LotusNet: Tunable privacy for distributed online social network services," *Computer Communications*, Dec. 2010.
- [2] L. A. Cuttillo, R. Molva, and T. Strufe, "Safebook: Feasibility of transitive cooperation for privacy on a decentralized social network." in *WOWMOM*. IEEE, 2009, pp. 1–6.
- [3] O. Bodriagov and S. Buchegger, "P2p social networks with broadcast encryption protected privacy," 2011, qC 20120126.
- [4] D. Zhao and M. B. Rosson, "How and why people twitter: the role that micro-blogging plays in informal communication at work," in *Proceedings of the ACM 2009 international conference on Supporting group work*, ser. GROUP '09. New York, NY, USA: ACM, 2009, pp. 243–252.
- [5] N. H. Minsky, *Law Governed Interaction (LGI): A Distributed Coordination and Control Mechanism (An Introduction, and a Reference Manual)*, February 2006, (available at <http://www.moses.rutgers.edu/>).
- [6] N. H. Minsky and V. Ungureanu, "Law-governed interaction: a coordination and control mechanism for heterogeneous distributed systems," *TOSEM, ACM Transactions on Software Engineering and Methodology*, vol. 9, no. 3, pp. 273–305, July 2000.
- [7] W. Zhang, C. Serban, and N. H. Minsky, "Establishing global properties of multi-agent systems via local laws," in *Environments for Multiagent Systems III, LNAI 4389*, D. Weyns, Ed. Springer-Verlag, 2007.
- [8] G. Urdaneta, G. Pierre, and M. V. Steen, "A survey of dht security techniques," *ACM Comput. Surv.*, vol. 43, no. 2, pp. 8:1–8:49, Feb. 2011.