Online Social Network Platforms: Toward a Model-Backed Security Evaluation

Erwan Le Malécot
lemalecot@nict.go.jp

Mio Suzuki
mio@nict.go.jp

Masashi Eto
eto@nict.go.jp

Daisuke Inoue
dai@nict.go.jp

Koji Nakao
ko-nakao@nict.go.jp

National Institute of Information and Communications Technology (NICT)
4-2-1 Nukui-Kitamachi, Koganei,
Tokyo, 184-8795 Japan

ABSTRACT

While presently enjoying a tremendous popularity among Internet users, Online Social Network (OSN) platforms have also recently increasingly come under fire for a number of security (and privacy) issues stemming from their usage. In an attempt to somehow formalize the study of such issues, we propose in this paper a conceptual model of the workings of a "typical" OSN platform as experienced by its users, putting a distinct emphasis on the resources published through such platforms for the pivotal role they have come to occupy, notably in regard to misuses. We then briefly discuss potential applications of that model, not only as a support to assess security properties inherent to the configuration of current OSN platforms, but also as a tool to further investigate practical attack scenarios against those (and their users).

Keywords

Social Network Services, Modeling, Information Security, Privacy, Graph Theory

1. INTRODUCTION

The Internet has since its early days been put to use to support social interaction, as exemplified by the quick rise to prominence and continuous success of the Email service. Likewise, the network of interlinked documents forming the World Wide Web can also be seen as an upshot of social activity, with hyperlinks manifesting existing relationships among websites [19]. Still, the advent of Web applications (cf. the so-called "Web 2.0" era) recently induced a profound shift on how people interact over the Internet, putting much more weight on social aspects than ever before. By empowering their users to effortlessly publish thoughts and resources online, such applications (e.g. blogging platforms, etc.) fueled the creation of large online communities. This in turn gave rise to constructs specifically designed to nurture such communities, thereby the birth and rise to prominence of Online Social Networks (OSN) platforms.

Under this moniker, we refer to those Internet-based platforms which build upon the digital expression of social relationships to put their users in relation with each other, and enable them to share resources accordingly. The last decade saw a tremendous growth in the usage of OSN platforms, with the most successful ones now catering to hundreds of million users (e.g. Facebook, Twitter). Alas, as with most success stories, this one is not exempt from its share of shortcomings. Indeed, OSN platforms have concurrently increasingly come under the spotlight for security and privacy implications associated with their usage [15], and for the sometimes insufficient take of their operators on such issues. That aspect combined with an ever growing interest from malicious individuals has led to a surge in attacks involving such platforms, putting at risks their users [16] but also third parties [2].

As a consequence, the operators of OSN platforms are now struggling to keep up with the mounting threat and, to deploy adequate defense mechanisms so as to thwart attackers attempting to exploit their platforms [20]. Given the variations among, but most of all, the core design similitudes shared by most present OSN platforms, we believe that a more abstract take on the issue at hands (i.e. the security of OSN platforms) could provide a way to better understand the inherent risks associated with such platforms, and therefore to help with the identification of likely attack patterns and the preemptive design of generic countermeasures. To this end, we develop a comprehensive conceptual model of the functioning of OSN platforms, tying the actions of end users and their effect on the internal data abstraction maintained by such platforms. One strength and contribution of our model is that we specifically integrate the notion of resources as "first-class" entities in it, making it fitting for security-oriented assessment. Indeed, when it comes to malicious activity involving OSN platforms, attackers are often either looking to get their hands on some target resources, or trying to exploit them for bad. To our knowledge, this is the first proposal of a model making such mechanisms directly explicit, and one of the first attempts to formalize the security review of OSN platforms.

The remainder of this paper is organized as follows. In Section 2, we give an overview of currently available OSN

PSOSM '12, April 17 2012, Lyon, France
Copyright 2012 ACM 978-1-4503-1236-3/12/04 ...$10.00.

platforms, detailing their modus operandi and deriving from it a list of characteristic features. In Section 3, we introduce our base model and describe how it can be extended to emulate a number of optional mechanisms proposed by some OSN platforms. In Section 4, we examine potential uses of our model with regard to the security of OSN platforms, investigating a few attack scenarios. In Section 5 we discuss related works, and finally in Section 6, we conclude on our work and consider future evolutions.

2. BACKGROUND

2.1 Current Landscape and Scope

Before delving into the construction of our model (and more security specific considerations), we begin with a brief review of existing OSN platforms. The first significant specimens of Internet-based platforms explicitly harnessing social mechanisms date back to the early 2000s, presumably spawned by the relative success encountered by the forerunner SixDegrees.com website launched in 1997. For a detailed synopsis of the history of OSN platforms from those early days to recent years, the reader can notably refer to the work of boyd et al. [5]. As often in the technological world, that history was shaped by a succession of rises and falls of a number of players in the field (e.g. Friendster, MySpace, etc.). That enduring quest for viability has incidentally led to a diversification of the services provided by OSN platforms, as well as of the audience they have come to target. As a result, there now exists a multitude of different thriving OSN platforms, ranging from very generic ones (e.g. Facebook, Twitter, Bebo, Orkut, Google+, etc.), to more targeted ones, for instance revolving around the establishment of business relationships (LinkedIn, XING, Ryze, etc.) or the collective enjoyment of specific hobbies (e.g. Flixter (movies), Flickr (photography), Last.fm (music), etc.).

A point of importance for our study is that, while certainly looking very diverse in appearance, almost all (if not all) of the above-mentioned OSN platforms that have come to dominate the market actually share the same core formula: centralized architecture plus Web-based interface for access (and tight integration with the rest of the World Wide Web). Given the privacy implications inherent to that design pattern (cf. the whole user data being left under the care of a single operator), efforts have recently emerged to try to break away from it by proposing more distributed solutions (e.g. the Diaspora project [14]), notably building up on the principles of peer-to-peer networks [7, 18]. That being said, in this paper we mainly restrict ourselves to the characterization of the predominant "Web-based centralized" family of OSN platforms. In the following section, we depict the modus operandi of a typical OSN platform, as derived from our experience and relevant literature on the subject.

2.2 Modus Operandi

Although a number of OSN platforms provide unrestricted access to large portions of the content they host (i.e. "public" content, in a likely bid to attract further users), one is usually required to go through registration to make full use of the services they advertise. And in practice, such registration more than often amounts to becoming a full-fledged "social entity" on the target OSN platform (i.e. only one type of account provided for all users and purposes). A typical registration goes as follows: the applying user is first requested to supply a token that attests to his identity (and individuality); it usually translates into providing a valid email address but can sometimes be more rigorous (e.g. request for a valid phone number, or a national identification number). He is then asked to pick up login credentials (username and password) for subsequent access to the platform, and to fill in a number of pieces of personal information to initialize his digital representation on the OSN platform (cf. his "profile"). Finally, he is usually encouraged to select a set of fellow registered users that he wishes to be put in relation with (cf. his initial "friends" or "contacts" on the platform).

From that point on, the user is eventually empowered to fully interact with the OSN platform, and to engage in the various social activities it supports. Given the large variety of OSN platforms that have come to be, it is rather difficult to provide a comprehensive compendium of those activities but for Web-based ones, it is fair to say that they mostly consist of the publishing and sharing of resources among registered users: diary entries, comments, photographs, etc. And of course, the user is expected to nurture and maintain his digital social circle, the list of his declared "friends". We then identify the following set of core functionalities with regard to the registered users of an OSN platform: they are enabled (1) to create a protected digital representation of themselves on the platform, (2) to define connections with other registered users based on social affinities, (3) to publish resources and to make them available to other users, and (4) to browse the resulting pool of content by connectivity. That enumeration can incidentally be regarded as our definition of what "OSN platforms" stand for in practice, definition from which we will eventually derive our model.

2.3 Abstraction and Motivation

Building on the above observations, we can then view OSN platforms as the mere combination of two basic parts: (1) the social graph formed by their registered users connected through the relationships that they specify among themselves, and (2), the pool of resources published by those users (which can also be seen as a graph in the Web context). While (1) has been the subject of rather extensive investigation (and modeling) since the advent of OSN platforms [22], (2) has yet to receive a comparable amount of formal scrutiny, as are the links that tie (1) and (2). Indeed, the "social aspect" associated with OSN platforms makes them a remarkable subject for the study of inter-relational behavior, published content then being often considered as no more than a by-product of such activity [21].

However, when looking at the recent state of OSN platforms, that content has progressively come to be given much emphasis in their design (e.g. content-centric user interfaces); and as such, has unfortunately come to become the focus of much of the malicious activity involving those platforms. This is notably evidenced by the increased presence of "spam content" on major OSN platforms [13], pushed in by attackers trying to mislead legitimate users into accessing "booby-trapped" resources (e.g. embedding malware or promoting phishing scams). Another substantial issue is manifestly the privacy implications brought by the aggregation of such an amount of personal information by OSN platforms, as that information may be susceptible to thief by willing attackers [4]. We will come back to those issues later on (cf. Section 4) but they bring forward the importance of understanding how resources published by users relate to the rest of the system (and to each other), which is one of the main motivations behind our model proposal.

3. MODEL

3.1 Definition

Our model can be decomposed into two complementary parts: first, an abstract depiction of the data structure exposed by OSN platforms to their users, and second, a rendition of those users and of the set of actions available to them in that framework. Regarding the structural part, we propose to represent OSN platforms as follows, aiming at the firm integration of both their "social graph" and "resource pool" components (cf. Section 2.2) into an unified abstraction: each OSN platform is portrayed as a vertex-labeled directed graph (thereafter referred to as its *operational graph*), the vertices of that graph standing either for identities registered by users on the platform (i.e. *identity nodes*, also thereafter referred to as *ID nodes*), or for published resources (i.e. *resource nodes*, also thereafter referred to as *RS nodes*). As for the arcs of the graph, they are associated with different semantic meanings depending on the label of their ending nodes (cf. Figure 1), namely:

- **(ID, ID) arc:** Such arc transcribes the existence of a *trust* kind of relationship between the identities mapped on its ending nodes, i.e. an arc pulled from ID node a to ID node b implies that *a is trusted by b* (or in more practical terms, that the user concealed behind identity a was accepted by the one concealed behind b as a "friend" on the platform).

- **(ID, RS) arc:** Such arc transcribes the existence of an *ownership* kind of relationship between the identity and the resource mapped on its ending nodes, i.e. an arc pulled from ID node a to RS node x implies that *a owns x* (or in more practical terms, that resource x was published through identity a on the platform).

- **(RS, RS) arc:** Such arc transcribes the existence of a *referral* kind of relationship between the resources mapped on its ending nodes, i.e. an arc pulled from RS node x to RS node y implies that *x refers to y* (or in more practical terms, that resource x embeds a mention to (e.g. an hyperlink), or part of resource y).

- **(RS, ID) arc:** As for (RS, RS) arcs, such arc transcribes the existence of a *referral* kind of relationship between the resource and the identity mapped on its ending nodes, i.e. an arc pulled from RS node x to ID node a implies that *x refers to a* (e.g. the tagging of identity a on resource/photograph x).

By construction, it ensues that the social graph of a given OSN platform is, in our model, materialized by the induced subgraph formed by the ID nodes of its operational graph. Then regarding the motivation behind our representation of resources as additional nodes grafted on to that social graph, it stems from the fact that in practice, identities and resources are indeed often presented as comparable objects to the end-users of the platform: a tissue of interlinked Web objects with identities assimilated to their corresponding "profile" Web pages, and resources to Web content linked from these and back (or if looking further "downward", a tissue of interlinked objects in a database).

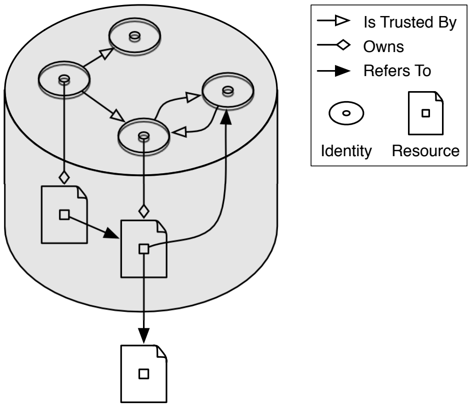

Figure 1: Representation of an OSN platform (grey container) in our model.

So, after having covered the structural facet of our model, what remains is for us to take care of the entities that interact with that structure. We will refer to those entities as the *consumers* of an OSN platform, and model them as active graph crawlers. By active we refer to their appointed ability to modify the constitution of the operational graphs they are interacting with. We then propose the following set of operations available to consumers, first relatively to "management":

- **`accessAs(Identity` i`):`** That operation specifies the *root* ID node i from which a consumer is to start crawling the operational graph, and by extension, the identity i under which subsequent operations are to be realized.

- **`register(Identity` i`):`** That operation enables a consumer to append the ID node i to the operational graph provided such node does not already exist; that ID node is initially disconnected (i.e. no edge to any other node).

- **`unregister(Identity` i`):`** That operation enables a consumer to remove the ID node i from the operational graph. It is to succeed if and only if that ID node is the root ID node for the identity currently assumed by the consumer (cf. `accessAs` through that identity).

- **`publish(Resource` r`):`** That operation enables a consumer to append the RS node r to the operational graph, r is to be connected to the consumer's current root ID node through an (ID, RS) arc (cf. ownership by the corresponding identity).

- **`unpublish(Resource` r`):`** That operation enables a consumer to remove the RS node r from the operational graph. It is to succeed if and only if there exists an (ID, RS) arc from the consumer's current root ID node to r (cf. `publish`, ownership by the corresponding identity).

- **`link(Node` n_1`, Node` n_2`):`** That operation enables a consumer to create an arc from node n_1 to node n_2, subject to the following conditions depending on their respective labels (cf. ID or RS):

  - **(ID, ID) arc:** If and only if n_2 is the consumer's current root ID node.
  - **(ID, RS) arc:** Unavailable (cf. operation conducted through `publish`).
  - **(RS, RS) arc:** If and only if n_1 is owned by the identity currently assumed by the consumer.
  - **(RS, ID) arc:** If and only if n_1 is owned by the identity currently assumed by the consumer.

  An additional condition is that the node linked from/to is to be *known* to the consumer (we will define what we mean by that terminology in the subsequent segment, cf. "crawling").

- `unlink(Node` n_1`, Node` n_2`):` That operation enables a consumer to delete an existing arc from node n_1 to node n_2, it is subject to the same conditions as the `link` operation as per the nature of the arc.

And second, relatively to "crawling":

- `followArc(Arc` a`):` That operation enables a consumer to crawl the arc a and learn about the node at its end, that node is then considered as *known* to the consumer.
- `getArcs(Node` n`):` That operation enables a consumer to retrieve the set of outgoing arcs from *known* node n.

Based on that description, the behavior of a consumer can be seen in our model as a succession of graph discovery phases through the use of the above-defined "crawling" operations (cf. starting from the initially only known root ID node determined through `accessAs`), and graph manipulation phases through the use of the remaining "management" operations – by design, those operations only being available on the discovered subgraph (i.e. the set of nodes known to the consumer). It then poses the problem of bootstrapping, that is the creation of the first outgoing arc from a freshly created ID node (cf. through `register`). On actual OSN platforms, this is solved by the recourse to an out-of-band channel that allows a consumer to make his root ID node known to another consumer, and request the creation of an arc from it to that other consumer's root ID node (cf. Section 2.2). We then assume the existence of such kind of out-of-band channel in our model (see also the notion of "public" resources/identities described in Section 3.2).

In a way, the operations that we defined for consumers embody the rules that dictate the construction of operational graphs. While a thorough analysis of the topological structure exhibited by such defined graphs (e.g. as derived from actual OSN platforms) is out of the scope of this paper, we expect them to share some characteristics with the social graphs they are "extending". For some insight into that matter, the reader can notably refer to the comprehensive analysis conducted by Mislove et al. of a set of large social graphs extracted from four major OSN platforms [17]. This would probably need to be collated with findings regarding the structure of large collections of hypertext documents [6] to account for the adjunct "resource graph" component introduced in our model.

3.2 Further Refinements

In this section, we put forward a few complementary pieces to our model, pieces that can optionally be brought in to render a set of extra mechanisms commonly found in "real-world" OSN platforms. One such mechanism is the possibility for consumers to freely access so-called "public" resources regardless of the connectivity of their assumed identities (cf. root ID nodes) in the operational graph. To provide for such behavior, we define an auxiliary identity – the "collective" identity – which is made to trust all the other registered identities (cf. arcs pulled from all the corresponding ID nodes to that "collective" ID node), and which in turn can be set as trusted by any identity/consumer wishing to be made "public" (i.e. made known to all, and by extension, make the resources he publishes accessible to all).

Another noteworthy artifact are the so-called "social applications" that build upon information extracted from OSN platforms to provide augmented services to their users. Usually, it involves those users conceding access rights to their accounts/identities on the relevant OSN platform. For that reason, we propose to portray such applications as "meta-ID nodes" fusing the root ID nodes corresponding to their users' registered identities. That concept of meta-node is used here to transcribe the fact that an application is enabled to interchangeably access (and operate on) the operational graph under the guise of any of its leased identities.

In a different area, OSN platforms also increasingly integrate the possibility for their users to leave "comments" on others' published content. What distinguishes those comments from "classic" resources is essentially the fact that a comment is automatically referred back to by the resource it refers to (i.e. the "commented" resource), allowing bidirectional navigation between the two. Our model can easily be extended to provide for such behavior, for instance by incorporating the following `tie` operation:

- `tie(Resource` r_1`, Resource` r_2`):` That operation enables a consumer to create an arc from *owned* RS node r_1 to *known* RS node r_2, a second arc from r_2 to r_1 being reciprocally created.

A user wishing to leave a comment on a resource would then just have to `publish` that comment and `tie` it to its target (provided that the conditions for `tie` to succeed are satisfied).

Last but not least, we warrant the creation of links to external resources (i.e. not belonging to the considered OSN platform) from internal ones. Indeed, very few OSN platforms can be considered as strict walled gardens and, as discussed earlier on, most of them rather integrate seamlessly with the World Wide Web, allowing cross-referencing through the use of hyperlinks.

3.3 Access Control?

As alluded to in the earlier definition of our base model (cf. `accessAs` and `register` operations), a major assumption made by parties interacting with an OSN platform is that only the entity (i.e. the consumer) that "registered" an identity is to be allowed to subsequently "assume" that identity on the platform. In other words, it implies the presence of mechanisms to authenticate the entities wishing to make use of the platform and enforce a set of access control rules accordingly. In Section 2.2 we saw that in practice such authentication was usually implemented by relying on textual credentials (cf. username and password) recorded upon the registration of an identity, and later requested upon access to the platform. Such type of admission access control mechanisms can be integrated in our model by analogously binding the use of `accessAs` to a prior use of `register` for the considered identity.

Another occurrence of access control mechanisms in OSN platforms is the often present option for a consumer to restrict access to the resources he publishes to a set of specific identities, the obvious motivation for doing so being to try to maintain a degree of privacy. One common policy is the possibility to restrict such access to his set of "friends", that is in our model the set of identities matching the ID nodes that are directly connected to that consumer's root ID node (i.e. "trusted by" him). More refined schemes can be devised as notably demonstrated by Carminati et al. [8], again exploiting the relative "location" of ID nodes in the social graph. Another work worth mentioning is that of Fong et al. who attempt to formalize, and generalize, such kind of access control mechanisms centered around a social graph [12].

Interestingly, those proposals also rely on models meant to transcribe the characteristic features of OSNs, features upon which the diverse access control schemes are then implemented. However, from our review, those models present the shortcoming of focusing exclusively on the notion of social graph (cf. resources being merely considered as static attributes attached to the nodes corresponding to their owners in the social graph, access being then exclusively mediated through those nodes) and therefore fail to address some of the challenges introduced by the intricate tissue of links woven between resources on actual platforms. For instance, considering the situation depicted on Figure 2 by means of our model, and supposing that a classic "friends only" policy is enforced on the traversal of the social graph: B can access A (cf. he is directly trusted by A) then X but C not being able to crawl the arc from B to A is prevented from accessing X by the policy which is probably the wish of A. However what if B publishes a resource Y and creates a referral link from it to X? That series of actions effectively makes X within the reach of C through a different path, inducing a discrepancy with the original policy. How to resolve that issue and to devise "consistent" policies in our model is left to future work.

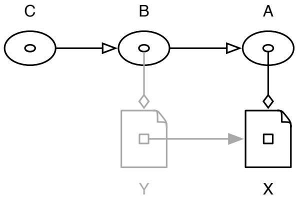

Figure 2: "Ambiguous" configuration in terms of access control, cf. case of access to semi-private RS node X via entry through root ID node C and following of the referral link from Y to X.

That discussion highlights a divergence in objectives: instead of a model crafted to provide an optimal solution to a specific issue, we rather intend our model to be as general, flexible and expressive as possible to encompass a large variety of cases and platforms (while providing for specialization when needed). In the following section, we further discuss the merits of our model by reviewing potential applications of it "in the field".

4. EXPLOITATION

4.1 Approach(es)

We envision two main ways of exploiting our model in relation with the security of OSN platforms: (1) a bottom-up approach focusing on the uncovering of potential attack patterns from the systematic review of the set of operations available to attackers, and (2) a top-down approach focusing on the investigation of observed concrete attacks through their decomposition into a sequence of elementary operations. In particular, (2) could be of use to gain some insight into how to devise mitigation mechanisms against such identified attacks (e.g. stricter access control rules). Moreover, the availability of a straightforward representation of consumers could also serve as a basis for the design of behavior-based evaluation mechanisms towards the automatic identification of the ones exhibiting malicious intents. In the followings, we provide some details as to how a number of security concepts can be expressed in our model.

4.2 Attack Scenarios

To begin with, we need to specify the framework in which our security review is to take place. We then advance the following attack model: we assume that malicious individuals can make use of any of the operations that we previously defined as available to consumers (cf. Section 3.1) to achieve their goals. It includes the option of accessing the operational graph of an OSN platform through identities registered by other consumers (i.e. identity theft, made possible by weaknesses in the access control mechanisms usually put in place), and of registering multiple identities for their own use (i.e. so-called sybil attack on the system [11]). As for the goals of attacks, their most obvious targets are assuredly the end-users of OSN platforms. In particular, we can identify the two following elemental threats vis-à-vis those users:

- **Privacy breach:** An attacker manages to get hold of a piece of information considered as confidential by operational standards; in our model jargon it can be expressed as an attacker managing to learn about, and follow an arc to a coveted resource.
- **Security breach:** An attacker manages to compromise the running environment of a targeted user; in our model jargon it can be expressed as an attacker managing to lure that consumer into following an arc to a previously arranged "booby-trapped" resource.

Regarding the notion of privacy, it is therefore intrinsically linked to the ability of an attacker to crawl the operational graph, and in particular, how large a portion of it. The strategy is then for an attacker to artificially extend his "reach". For instance, one option would be to randomly compromise a number of (well-connected) identities acquiring the ability to crawl from the associated root ID nodes. In a practical study on Facebook, Bonneau et al. showed that such a method could actually turn out to be very effective [4]. A more straightforward approach would be for an attacker to progressively create a "crawlable" path from a root ID node he added to the operational graph to the node he wishes to access (e.g. exploiting social engineering techniques to induce users to trust his assumed identity [3]).

Then, regarding the notion of security, the strategy would rather be for an attacker to increase the proportion of nodes of the operational graph being under his control (i.e. proportion of booby-trapped nodes), and, more to the point, to increase their connectivity so as to boost their chances from being reached by an unwary consumer.

4.3 Ramifications

From the preceding review, it ensues that the "dangerousness" of an attacker from the point of view of privacy can be correlated with the characteristics of the forest of (crawlable) directed trees springing from the root ID nodes under his control (i.e. arcs of the trees directed away from their roots). Similarly, his "dangerousness" in terms of security can be linked to the characteristics of the forest of trees rooted on the RS nodes under his control and this time directed toward those roots (i.e. aggregate of the paths leading to those nodes). Therefore, the study of those forests – as extracted from actual platforms, or artificially built through simulation – and of their characteristics should provide further insight on the latent capabilities of attackers given the "raw materials" at their disposal. For instance, one could attempt to measure how easy it is to expand the size of such forests depending on the operations made available on a target OSN platform (e.g. availability of the `tie` operation or not). On a different direction, once a RS node is identified as nefarious, the study of the portion of operational graph surrounding it could help uncover additional identities (i.e. in addition to the explicit owner) that might have been colluding to make that RS node more visible (e.g. creating links to it), and thus be equally nefarious.

4.4 Seek and Destroy

Finally, while the scenarios discussed so far mostly dealt with attackers either selectively expending the operational graph or leaving it untouched, in effect they also have the capability to break it down (cf. `unregister`, `unpublish`, `unlink`). Apparently, "real-world" OSN platforms have yet to fall victim to significant attacks exploiting such "destructive" operations but attackers might one day find reason enough to start doing so (e.g. to disrupt the availability of a target resource, or even, disrupt the overall usability of a target OSN platform by carefully segmenting its operational graph…). It is to be noted that OSN platforms based on peer-to-peer infrastructures tend, by their distributed nature, to be much more susceptible to such attacks (but advances are made toward their reinforcement [10]).

5. RELATED WORKS

As mentioned earlier on, we found only few attempts to formalize the security evaluation of OSN platforms through a model-based approach. In a position paper [1], Asher et al. discussed the value of such method and proposed a preliminary model based on the enumeration of the actions available to the users of an OSN platform, and therefore to its attackers. While reminiscent of our approach, they did not explicitly connect those actions to the exposed internal structure of OSN platforms, rather treating them as black-boxes. Moreover, they did not include actions geared towards the publication and management of resources into their model. Then, as part of their survey of security and privacy issues in OSN platforms [9], Cutillo et al. introduced a detailed layered model of the functioning of OSN platforms, going from their inner social graph abstraction to the interface exposed to users to interact with it, and down to the technical implementation of associated concepts. However, their model can be seen as rather descriptive in purpose, aiming at offering a comprehensive picture of the phenomenon but difficult to use for abstract reasoning.

6. CONCLUSION

In this paper, we proceeded with the formulation of an abstract model of the workings of an OSN platform, aiming at providing a ground for a more systematic evaluation of their security. Our model is composed of a stripped-down representation of the data abstraction maintained by OSN platforms (taking the form of an unified graph), combined with a transcription in that framework of the set of operations provided to their users for interaction. We then exposed how that model could be extended to make it more accurately represent some auxiliary mechanisms often offered by OSN platforms, and more to the point, discussed how it could be exploited for security assessment. In particular, we showed that, while rather concise, it indeed allowed us to express and characterize a number of practical attack scenarios. As for future work, we plan to make use of our model to conduct a more comprehensive generic survey of the security issues inherent to current OSN platforms, and then to evaluate the scenarios and mechanisms that we can derive from it through the use of concrete datasets extracted from a number of actual OSN platforms.

7. ACKNOWLEDGMENTS

We would like to thank all the anonymous referees who reviewed this manuscript (or an earlier version of it) for their very insightful feedback.

8. REFERENCES

[1] C. Asher, J.-P. Aumasson, and R. Phan. Security and Privacy Preservation in Human-Involved Networks. In *Proceedings of the 2009 Conference on Open Research Problems in Network Security (iNetSec'09)*, pages 139–148. Springer Boston, 2009.

[2] E. Athanasopoulos, A. Makridakis, S. Antonatos, D. Antoniades, S. Ioannidis, K. G. Anagnostakis, and E. P. Markatos. Antisocial Networks: Turning a Social Network into a Botnet. In *Proceedings of the 11th International Conference on Information Security (ISC'08)*, pages 146–160, Berlin, Heidelberg, 2008. Springer-Verlag.

[3] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All your Contacts are Belong to Us: Automated Identity Theft Attacks on Social Networks. In *Proceedings of the 18th International World Wide Web Conference (WWW'09)*, pages 551–560, New York, NY, USA, 2009. ACM.

[4] J. Bonneau, J. Anderson, and G. Danezis. Prying Data out of a Social Network. In *Proceedings of the 2009 International Conference on Advances in Social Network Analysis and Mining (ASONAM'09)*, pages 249–254, Washington, DC, USA, 2009. IEEE Computer Society.

[5] d. m. boyd and N. B. Ellison. Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication*, 13(1):210–230, 2007.

[6] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graph Structure in the Web. In *Proceedings of the 9th International World Wide Web Conference (WWW9)*, pages 309–320, Amsterdam, The Netherlands, The Netherlands, 2000. North-Holland Publishing Co.

[7] S. Buchegger, D. Schiöberg, L.-H. Vu, and A. Datta. PeerSoN: P2P Social Networking: Early Experiences and Insights. In *Proceedings of the 2nd ACM EuroSys Workshop on Social Network Systems (SNS'09)*, pages 46–52, New York, NY, USA, 2009. ACM.

[8] B. Carminati, E. Ferrari, and A. Perego. Enforcing Access Control in Web-based Social Networks. *ACM Trans. Inf. Syst. Secur.*, 13:6:1–6:38, November 2009.

[9] L. A. Cutillo, M. Manulis, and T. Strufe. Security and Privacy in Online Social Networks. In *Handbook of Social Network Technologies and Applications*, pages 497–522. Springer US, 2010.

[10] L. A. Cutillo, R. Molva, and T. Strufe. Safebook: A Privacy-Preserving Online Social Network Leveraging on Real-Life Trust. *Communications Magazine, IEEE*, 47(12):94–101, 2009.

[11] J. R. Douceur. The Sybil Attack. In *Revised Papers from the 1st International Workshop on Peer-to-Peer Systems (IPTPS'01)*, pages 251–260, London, UK, 2002. Springer-Verlag.

[12] P. W. L. Fong, M. Anwar, and Z. Zhao. A Privacy Preservation Model for Facebook-Style Social Network Systems. In *Proceedings of the 14th European Conference on Research in Computer Security (ESORICS'09)*, pages 303–320, Berlin, Heidelberg, 2009. Springer-Verlag.

[13] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: The Underground on 140 Characters or Less. In *Proceedings of the 17th ACM Conference on Computer and Communications Security (CSS'10)*, pages 27–37, New York, NY, USA, 2010. ACM.

[14] D. Grippi, M. Salzberg, R. Sofaer, and I. Zhitomirskiy. Diaspora Project. `https://joindiaspora.com/`, 2010.

[15] G. Hogben. Security Issues and Recommendations for Online Social Networks. *European Network and Information Security Agency Position Paper*, 80211(1):36, 2007.

[16] M. Huber, M. Mulazzani, and E. Weippl. Social Networking Sites Security: Quo Vadis. In *Proceedings of the 2010 IEEE 2nd International Conference on Social Computing (SOCIALCOM'10)*, pages 1117–1122, Washington, DC, USA, 2010. IEEE Computer Society.

[17] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and Analysis of Online Social Networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (ICM'07)*, pages 29–42, New York, NY, USA, 2007. ACM.

[18] F. Musiani. When Social Links are Network Links: The Dawn of Peer-to-Peer Social Networks and its Implications for Privacy. *Observatorio (OBS*)*, 4(3), 2010.

[19] H. W. Park. Hyperlink Network Analysis: A New Method for the Study of Social Structure on the Web. *Connections*, 25:49–61, 2003.

[20] T. Stein, E. Chen, and K. Mangla. Facebook Immune System. In *Proceedings of the 4th Workshop on Social Network Systems (SNS'11)*, pages 8:1–8:8, New York, NY, USA, 2011. ACM.

[21] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi. On the Evolution of User Interaction in Facebook. In *Proceedings of the 2nd ACM Workshop on Online Social Networks (WOSN'09)*, pages 37–42, New York, NY, USA, 2009. ACM.

[22] C. Zhang, J. Sun, X. Zhu, and Y. Fang. Privacy and Security for Online Social Networks: Challenges and Opportunities. *Network, IEEE*, 24(4):13–18, 2010.