Beyond-Privacy and Identity Spam:
What others say about us on the Federated Social Web

Position Paper

José M. del Álamo Yod-Samuel Martín Juan C. Yelmo

Universidad Politécnica de Madrid
{jmdela, samuelm, jcyelmo}@dit.upm.es

Abstract

Users' concerns about their digital identity have led to the spreading of privacy-management techniques, which put users back in control of what is done with their personal information. These techniques deal with information somehow originated in the users themselves, but what happens if someone else is talking about the user? Here we present two views of this issue: Beyond-Privacy (from the perspective of the affected user) and Identity Spam (from the perspective of the consumers of identity information), conjecturing their imminent relevance for the Federated Social Web.

Background and motivation: from privacy to beyond

Nowadays, users of the social web perform actions –either consciously or inadvertently– that reveal online different pieces of their **digital identity**. This information might consist of:

a) user-defined attributes such as profile information (e.g. name or address);
b) the social graph of the user (e.g. friends and relationships);
c) online activities the user carries out (e.g. attending an event or joining a group);
d) messages expressed by the user (e.g. a status or a micropost), etc.

Some **standardization efforts** have gained track to ease the sharing of this information for the benefit of the user and third parties using it as well. As a way of example, FOAF provides a scheme to define a user profile and OpenSocial allows sharing user attributes between two parties; XFN allows representing and qualifying user connections; ActivityStreams is aimed at providing a common data model and format to express user activities; Atom is commonly employed to format user messages, etc.

From the set of information released by the user, a third party may rebuild an ***Intended Profile*** i.e. *a partial but somehow accurate view of the user identity*. This intended profile can be used for different goals e.g. to provide users with contents and services tailored to their needs, preferences, interests and expertise; to improve the service delivery by adapting the contents to bandwidth or device capabilities; etc. It might not be a whole profile or even an exact one: Some information might be hidden to the third party due to privacy decisions, or the lack of some information might result into a misleading view.

Privacy can be kept and **managed** in this process. Users may give consent to reveal the information, know the context within which the information has been revealed and will be used, and have control over this information and what parties access it. If some of these principles (**consent, context or control**) are not kept, then privacy issues arise.

Tools and technologies exist for users to **assure** privacy is kept and properly managed. Web sites can declare their privacy policies using privacy policy expression languages such as P3P, and inform users about the dissemination and exploitation of their information. In turn, users can express their privacy preferences using suitable languages such as APPEL. User agent add-ons allow users to control the information they release e.g. by automatically enforcing their privacy preferences, or by authorizing the access to their personal information by third parties with initiatives such as OAuth or the Kantara Initiative User Managed Access – UMA. Auditing and trust systems provide users with hints on the Web site compliance with the stated policies; etc. They all try to place the **user** in sole control of the information release, thus deciding what information should be delivered, to which other parties, and under which terms.

However, what about if someone else (either acquaintance or stranger) publishes some information regarding the user? Traditional-privacy approaches might not apply since the information is not originated in the user. And even so, the consequences are somehow similar to those related to the release of traditional, privacy-manageable, user identity information. We have coined the term ***Beyond-Privacy*** to refer to this issue.

While that term refers to the issue from the perspective of the referred user, it might also affect other users consuming the information. When the information was part of the intended profile, consumers could at least rely on the information being originated in the user. However, when anyone can become a provider of other user's identity data, much noise is introduced, thus making the information quality decay. This problem is somehow similar to that of the spam, as it is difficult to know beforehand which messages we can trust and which other contain harmful contents, etc. Thus, we have coined the term ***Identity Spam*** to refer to it.

Beyond-Privacy: the perspective of the affected user

We have coined the term **Beyond-Privacy** to refer to issues faced when information referring a user is being released by other entities different from the referred user. We have analyzed this problem from different perspectives: the truthfulness of the information released and the motivations behind.

The **truthfulness** of the information released can vary from true to false. When true information is released, we talk about information ***leakage***. Leakage is usually linked to datacenter security, but here we refer to individuals leaking identity information they know about others. Major privacy concerns arise as the referred user might not be aware of this information being published.

On the other hand, an individual can ***fabricate*** new, false pieces of a user identity and publish them. This will result into **bogus profiles** being created, different from the intended one. Thus, when a third party tries to rebuild the user's identity from the information found, it might get a not only inaccurate but also (intentionally) misleading and false view of the user identity.

The referred user might be able to distinguish original from fake pieces of information. However, for the consuming users, the separation might be so blurred that most times it will be very difficult to discern between them.

We have categorized the **motivations** moving a third party to release pieces of other users' identity. So far we have identified five motivations described on the view of three different axes:

1. the benefit that the author of the information obtains;

2. the damage the release of the information inflicts on the affected user;
3. the closeness of the relationship between the author and the affected user.

The first type is what we named ***careless sharing***. It corresponds to two very close users (e.g. friends or relatives). There is no intentionality in the possible damage caused to the referred user. The information can be considered as truthful. We have not found that the author gets any specific benefit from leaking the piece of information; quite on the contrary, it might put the relationship in danger. There are loads of examples of this type ranging from publicly publishing compromising photographs to unfortunate comments revealing some sensitive details. A bad understanding of privacy policies by authors in current social networks usually leads to these situations.

Gossip is the second type: it is a primitive human behavior, given that reporting about other people's private life and behavior is somehow common. The closeness between the parties may vary, but the author has got some knowledge of the referred person. The motivation may vary as well, but usually there is no damage nor benefit motivators (just the pleasure in the activity). The information may vary from true information to pieces of fabricated information, being sometimes difficult to separate them.

The third category is ***libel***: the publication of a false statement that is damaging to a person's reputation. In this case, closeness and author benefit are not determining factors, but damaging the referred user is. By definition the information released is false and thus the referred person's privacy is not compromised but his or her reputation may be seriously damaged.

The enemy disguised as a friend, also known as ***frienemy*** or frenemy, is our fourth category. In this case there is a close relationship between the information author and the referred person. There is some intention in the release of the information (which is truthful) –either some damage to the referred person or some benefit for the author is sought, or even both.

We have labeled the last category as ***commensalism***. It involves the author getting some benefit by disclosing the information, even though no damage is sought for the referred person. We can differentiate depending on the closeness between the author and the referred person: when both are stranger, the author usually seeks easy, fast and cheap advertisement or publicity e.g. declaring that a celebrity is at my disco to increase my regulars. When they are closer the effects sought are quite the same but on a smaller scale e.g. appearing cooler to peers.

Many of these forms of *beyond-privacy* leverage on social networking, given that relating some information to a user's identity is as easy as linking it to his or her online identifier. For example, abouteveryone.com [1] is a website that allows people to anonymously comment on other users referring to them by their Facebook or Twitter identifiers. Another example was littlegossip.com [2], a similar website, now closed, which quickly became a place for slipping alleged sexual and drug-taking behaviors of students and classmates.

Identity Spam: the perspective of the consumer

The rational for **identity spam** being a problem is that if new pieces of identity information are fabricated and related to a user's identity, then it might become increasingly difficult to discern the real pieces from the fake ones. In the previous section, we have described how easy it is to create identity spam departing from users' declared identities at existing social networks. In this section we analyze some other drivers.

To understand our results we first have a look at email evolution. **Social networking and email** have evolved along similar paths. For some time, email providers used proprietary technologies; then they agreed on a set of protocols to interconnect. This benefitted all their users and increased the network value according to Metcalfe's law. The social web is following a similar evolution that may hopefully lead to realize the Federated Social Web. Unfortunately, interconnecting email providers also enabled one of the worse problems the Internet is facing: email **spam**, i.e the practice of sending unwanted e-mail messages in large quantities to an indiscriminate set of recipients. As a result, the amount of email spam is estimated to be more than 97% of all emails sent over the Internet. Several and complex reasons are behind email spam, among them, the social distribution of spam and the agnosticism of email providers.

We have identified that a ***social connection*** may become a driver for identity spam: the socially closer the sender of a message appears to be, the more credible and trustworthy the message seems for the recipient. This effect has been demonstrated by some of the most successful viruses that have flooded the Web. For example, “Iloveyou” email virus was such a success in the early 2000’s because the message was sent by a person the recipient knew. The same applies to some other malware such as ‘SPIM’ or the “Lechucd” virus in Instant Messaging tools. In social networking profiles, some applications have long been posting contents appearing as if coming from the profile owner, when actually being distributed without their knowledge[3].

The social web evolution into a **federated social web** might suffer from some of these issues. Although technologies such as OneSocialWeb allow authenticated connections between servers, it is still difficult to control what each user is saying about the others even in a single domain. As a matter of fact, some of these problems have already started with the upsurge of some social sites that offer information regarding a user but provided by someone else. We have already mentioned sites such as abouteveryone.com and littlegossip, where it is difficult to decide what is true and what is false about a person. In addition, new sites **aggregate** identity information from other sites without checking its truthfulness. For example, a site called 123people [4] describes itself as a people search tool. While some of the search results might indeed refer to identity attributes of the person being sought, many others are completely unrelated, thus providing inaccurate and misleading information, and contributing to identity spam.

Scenarios

As we have presented, real situations exist that demonstrate the topicality of these issues, and nothing suggests they are not going to reproduce in novel scenarios, such as those appearing in the Federated Social Web. Recapitulating, we encounter several **agents** involved in the distribution of a piece of identity-information, and each of them may rely on external services to perform specific tasks:

- the user authoring or disseminating the information, who may, in turn, depend on other entities or services to produce, store, publish or distribute information;
- the user referred to by the identity information;
- the user consuming that information, maybe relying on other entities or services to find, access, retrieve or consume it.

Depending on the **degree of distribution and control**, this may give rise to different scenarios:

a) In a **closed** scenario, all the agents use a shared set of services, and information is only accessible from a closed environment (e.g. a corporate Intranet).
b) In a **centralized** scenario, there is a central authority in charge of managing all the information, and the users account to it (e.g. a silo-alike, traditional social networking service).
c) In a **decentralized yet trustworthy** scenario, different services and entities may host information and user accounts, needing to interoperate in the process, notwithstanding they have trust relationships established (e.g. some proposals of distributed social networking services such as OneSocialWeb). Note that the information may also be distributed into different pieces stored by different entities.
d) In a **massively distributed** scenario, anyone may assume any of the roles, and trust relationships between agents are not warranted. This scenario corresponds to the present-day, provider-agnostic, Web and email, and we conjecture it will also become the most spread case of the Federated Social Web.

In between, different situations may partially combine characteristics of several scenarios, leading to a rich diversity that may potentially require different solutions.

UPM Position

Our position is that it should be possible and easy for individuals to manage any piece of online information concerning them, in a **dynamic yet trustworthy** manner. This way, the credibility of the social information found will increase, in addition to reputation problems being reduced, since non-asserted information might be filtered out. **Common or different solutions** may be necessary for different scenarios; nevertheless, they must provide individuals with **equivalent control** disregarding the **origin** of the information.

Conclusion

Beyond-Privacy and ***Identity Spam*** are two sides of the same coin appearing when others provide identity information about individuals. These problems are found in the current (social) web, but they are expected to rise along the Federated Social Web. First, because of its distributed and decentralized nature; and also, because 'identity' and 'social' are two factors increasing the damage power of spam.

Acknowledgements

This work has been developed in the context of the Verify project (social@dit.upm.es) at Universidad Politécnica de Madrid.

References

[1] http://www.abouteveryone.com/

[2] http://www.littlegossip.com/

[3] http://www.readwriteweb.com/archives/tagged_photos_on_facebook_new_source_of_marketing_spam.php

[3] http://www.123people.com/