Privacy and Security for Online Social Networks: Challenges and Opportunities

Chi Zhang and Jinyuan Sun, University of Florida
Xiaoyan Zhu, Xidian University
Yuguang Fang, University of Florida and Xidian University

Abstract

Online social networks such as Facebook, Myspace, and Twitter have experienced exponential growth in recent years. These OSNs offer attractive means of online social interactions and communications, but also raise privacy and security concerns. In this article we discuss the design issues for the security and privacy of OSNs. We find there are inherent design conflicts between these and the traditional design goals of OSNs such as usability and sociability. We present the unique security and privacy design challenges brought by the core functionalities of OSNs and highlight some opportunities of utilizing social network theory to mitigate these design conflicts.

Online social networks (OSNs) [1, 2] such as Facebook, MySpace, and Twitter enable people to stay in touch with their contacts, reconnect with old acquaintances, and establish new relationships with other people based on shared features such as communities, hobbies, interests, and overlaps in friendship circles. Recent years have seen unprecedented growth in the application of OSNs, with about 300 OSN systems collecting information on more than half a billion registered users [2]. As a result, OSNs store a huge amount of possibly sensitive and private information on users and their interactions. This information is usually private and intended for the eyes of a specific audience only. However, the popularity of OSNs attracts not only faithful users but parties with rather adverse interests as well [3, 4]. The diversification and sophistication of purposes and usage patterns of OSNs inevitably introduce privacy infringement risks to all OSN users as a result of information exchange and sharing on the Internet. It is therefore not surprising that stories about privacy breaches by Facebook and MySpace appear repeatedly in mainstream media [3–5].

Regardless of the purpose of an OSN, one of the main motivations for users to join an OSN, create a profile, and use different applications offered by the OSN is the possibility to easily share information with selected contacts or the public, and facilitate social interactions between the users of OSNs. Disclosing personal information in OSNs is a double-edged sword [6]. On one hand, information exposure is usually a plus, even a must, if people want to participate in social communities. Visibility of users' profiles and public display of connections (friend lists) are necessary for implementing core functionalities of OSNs such as social search and social traversal. On the other hand, leakage of personal information, especially one's identity, may invite malicious attacks from the real world and cyberspace, such as stalking, reputation slander, personalized spamming, and phishing [7]. Despite the risks, many of the privacy and access control mechanisms of today's OSNs are purposefully weak to make joining the OSN and sharing information easy. We believe that more effective and flexible security mechanisms are therefore required for the safety of OSN users as well as the continued thriving of OSNs.

In this article we present a general framework for assessing the security and privacy of current and next-generation OSNs. Whereas others have considered specific mechanisms for improving OSN users' security and privacy, such as finer-grained access control [8], we ask a more general question: What should be the security and privacy design goals for OSNs? When we evaluate these goals in the context of the usage of OSNs, we find inherent design conflicts between them and the traditional design goals of OSNs such as usability and sociability. To further complicate matters, the balance between security, privacy, usability, and sociability might differ depending on the particular purpose of the OSN in question. We also present a set of possible research directions for mitigating these design conflicts. Our framework and follow-on research will help provide a foundation for researchers and OSN providers to evaluate, understand, and address the unique security and privacy challenges created by OSNs.

Features of OSNs

Since the motivation of our research is to provide the advantages of OSNs without compromising privacy and security in OSNs, the desirable features of traditional OSNs and their possible implementations need to be discussed first.

The Main Functionalities of OSNs

Without doubt, there is vast diversity among the OSNs in all dimensions. However, it is widely accepted that all OSNs

This work was partially supported by the U.S. National Science Foundation under grants CNS-0916391, CNS-071645, and CNS-0721744, and the China 111 Project under Grant B08038. The work of Zhu was also partially supported by the China Fundamental Research Funds for the Central Universities under Grant No. JY10000901021.

 0890-8044/10/$25.00 © 2010 IEEE

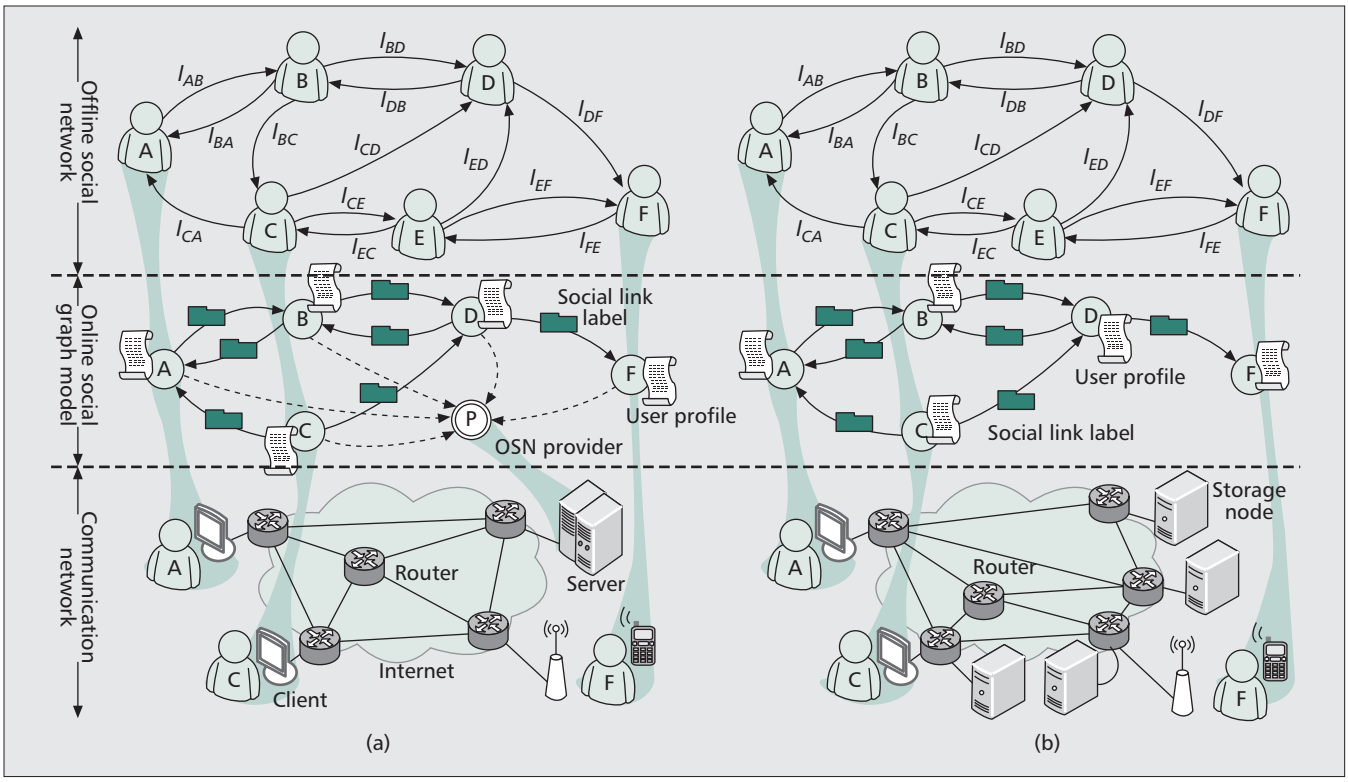

Figure 1. *A three-tiered description of online social networks: a) client-server architecture; b) peer-to-peer architecture.*

share some core features. According to Boyd and Ellison [1] and Wikipedia [2], an OSN is characterized as follows (our paraphrase).

An OSN is a digital representation of its users and (a subset of) their social connections/relationships in the physical or virtual world, plus networking services for messaging and socializing among its users. It provides a platform to:

- Allow users to construct digital representations of themselves (usually known as user profiles) and articulate their social connections with other users (i.e., lists of contacts)
- Support the maintenance and enhancement of preexisting social connections among users in the physical or virtual world
- Help forge new connections based on common interests, location, activities, and so on

Based on the above definition, an OSN should provide the following functionalities to facilitate users' self-representations and online social interactions.

Personal Space Management — An OSN should support a user to:

- Create/cancel an account (i.e., user registration or withdrawal)
- Create/edit user profile
- Upload/edit user generated contents (i.e., blog-like postings, status update)

Note that in most OSNs, a user's actions in her personal space will be reported to her social contacts automatically. As a result, uploading and editing user generated contents serve as an important communication primitive between a user and her contacts.

Social Connection Management — OSNs connect people by providing formal means for users to articulate their relationships with other users (e.g., friend lists). In addition to representing existing social ties, OSNs can be used to reestablish lost social connections if they are users of the same OSN. Social links can also be newly established, for instance, when users find they share common interests. Therefore, an OSN should support users to establish/maintain/revoke a social connection.

We use a simple labeled graph (i.e., *online social graph*) to model all data stored in an OSN, which is illustrated in Fig. 1. Here, vertices/nodes model individual users in an OSN, while edges/links model social connections between users. Each node maintains its corresponding user's profile, and each link is tagged with a social link *label*. Some OSNs allow their users not only to specify the types of social connections (e.g., friends or colleagues) but also to establish trust relationships, which express how much they trust the other members either with respect to a specific topic (topical trust) or in general (absolute trust). We can model all these descriptions and specifications on a social connection as the label of the corresponding social link in the online social graph. The collection of a particular user's profile and connected social links is called her *digital personal space* (personal space for short). The collection of personal spaces of all registered users of an OSN is called *digital social space*.

Means of Communication — Communicating with others is the central feature offered by OSN services. There are several possible channels of communication among users. First, one can leave public messages (e.g., blogging) in the form of texts, video, audio, photos, and so on using personal space. Moreover, one can send private messages, which are another form of asynchronous communication, as opposed to synchronous communication such as instant messaging readily available in almost all OSNs. An attractive feature of OSNs is that they are open to third-party applications. While many of these applications provide alternative means of asynchronous communications, there are also applications that allow users to engage in real-time activities such as online games.

Exploring Digital Social Space — OSNs are not restricted to interactions with existing contacts, and encourage users to establish new social connections (socialization). Consequently,

the support for searching or traversing digital social space is a crucial component of OSNs. Browsing a new user's personal space is authorized in two stages. There are two means by which the personal space of an unknown user may be reached [9].

Global Keyword Search (or Social Search) — The first means to find an unknown user is to conduct a global keyword search. A successful search would produce for the accessing user the search listing of a target user. A user may specify a *search policy* to allow only a subset of users to be able to reach her search listing through a global name search.

Social Graph Traversal (or Social Traversal) — A second means to reach a search listing is by traversing the online social graph. A user may traverse this graph by examining the friend lists of other users. More specifically, the friend list of a user is essentially the set of search listings of her friends. A user may restrict traversal by specifying a *traversal policy*, which specifies the set of users who are allowed to examine her friend list after her search listing is reached.

In order to achieve traditional goals of OSNs such as usability and sociability, we must support all functionalities mentioned above in designing new secure and privacy-preserving OSNs. Note that there also exist other goals such as efficiency and configurability. Our survey here only focuses on those that potentially conflict with OSN security and privacy.

System Architectures of OSN

There are two paradigms of implementing an OSN in the literature [10, 11], client-server architecture and peer-to-peer (P2P) architecture, which yield centralized and distributed systems, respectively.

Client-Server Architecture — Today's OSNs are centralized and web-server-based. All functionalities, like storage, maintenance, and access to OSN services are offered by the commercial OSN providers such as Facebook Inc., LinkedIn Corp., and XING AG. This traditional architecture has the advantage of being straightforward and easy to implement, while suffering from all the drawbacks of centralized systems. For example, any central entity can easily be a single point of failure,a single target for denial-of-service (DoS) attacks, and also a bottleneck for network performance.

P2P Architecture — There is a strong trend [10, 11] to design a P2P architecture for next-generation OSNs. It adopts a decentralized architecture relying on cooperation among a number of independent parties who are also users of the OSNs. Users' personal spaces are stored and maintained distributively. By supporting the direct exchange of information between devices, be it between users who have met before or between adjacent nodes of a city mesh network, a P2P architecture can take advantage of real social networks and geographic proximity to support local services when Internet access is unavailable. But for the P2P architecture, offering some functionalities (e.g., global search) of OSNs in a distributed manner is a challenging problem.

Figure 1 illustrates two kinds of architectures for OSNs, where the client-server architecture requires Internet connectivity for users to communicate via the remote central server, and the P2P counterpart supports communications via local connectivity since the role of the central server is distributed into each storage node. Note that for client-server architecture, in its corresponding online social graph model, a new entity called OSN provider is added, and each OSN user has a social (e.g., trust) relationship with her.

Privacy and Security of OSNs

To understand the unique challenges of balancing security and privacy with usability and sociability, we first review how the standard principles of network security, including confidentiality, integrity, and availability, extend to OSNs.

Privacy and Security Requirements for OSNs

Confidentiality or Privacy — Privacy is of paramount importance in OSNs, since the illegal disclosure and improper use of users' private information can cause undesirable or damaging consequences in people's lives. Privacy in the context of OSNs has several broad categories:

- User's identity anonymity — The protection of a user's identity changes across different types of OSNs. In college websites like Facebook, the use of real names to (re)present an account profile to the rest of the OSN is encouraged. There is no user identity anonymity because most applications in Facebook rely on connecting users' profiles to their public identities. In dating sites like Friendster, a weak pseudonymity is created by making only the first name of a participant visible to others, and not her last name. In pseudonymous-based dating websites like Match.com, the use of real names and personal contact information is discouraged. A random identifier is used to protect the public identity of a person.

- User's personal space privacy — The visibility of a user's profile also varies across different types of OSNs. By default, profiles on Friendster and Tribe.net are crawled by search engines, making them visible to anyone, regardless of whether or not the viewer has an account. OSNs like MySpace allow users to choose whether they want their profile to be public or friends only. Facebook takes a different approach: by default, users who are part of the same subnetwork can view each other's profiles, unless a profile owner has decided to deny permission to those in their subnetwork. For most OSNs, a friend list is visible to anyone who is permitted to view the profile, although there are exceptions. For instance, some MySpace users have hacked their profiles to hide their friend lists, and LinkedIn allows users to select hiding their friend lists.

- User's communication privacy — In addition to personal data disclosed in a user's digital space, an OSN user may also disclose personal information to the network operator or OSN provider using the network itself: data such as time and length of connections, location (IP address) of connection, other users' profiles visited, messages sent and received, and so forth. Therefore, additionally, communication privacy has to be met.

To sum up, a user's privacy requirement is twofold. First, unauthorized entities (i.e., who are not granted access to the private data) must not learn the content of the private data which reveals identifying information of the data owner (user). This aspect of data privacy implies data confidentiality and the owner's anonymity, and directly leads to the need for *access control*. Access to information on a user may only be granted by the user directly, and the access control has to be as fine-grained as each private information item has to be separately manageable. Second, unauthorized entities must not be able to link multiple private data files to profile the owner, indicating that the stored or transmitted private data should appear random and leak no useful information. This aspect is essentially the unlinkability requirement.

Authentication and Data Integrity — Although exceptions exist, previous research [1] suggests that most OSNs primarily support preexisting social relationships in real life. Therefore, in ideal situations an OSN is just a digital representation of an *offline social network* (also called *real-life social network* [RSN]). Since the data stored in an OSN can be modeled as an online social graph, there exists a one-to-one mapping from the RSN to the online social graph model (Fig. 1). Data integrity in the context of OSNs means that this consistency should be kept; that is, any attempt of deviating an online social graph model from its corresponding RSN is a kind of attack, and should be detected and corrected with appropriate mechanisms.

Obviously, there are two types of attacks on online social graphs: forging nodes/identities and forging social links/connections. Forging a node (e.g., identity theft) is a fundamental problem in OSNs and lies at the root of many other security problems. For example, an attacker can create fake profiles in the name of well-known personalities or brands in order to slander people or profit from their reputation. Also, an attacker can create multiple fake identities to disrupt digital reputation systems for OSNs [7]. Therefore, it is required that the communicating entities (e.g., OSN users) be assured of each other's legitimacy (i.e., the keys used for authentication are indeed assigned by the trusted authority) and authenticity (i.e., an entity is in possession of her claimed identity).

Others — Since some OSNs are used as professional tools to aid their members' businesses or careers, other security requirements, including *availability* (i.e., data published by users has to be continuously available) and *accountability* (i.e., misbehavior of users should be traceable), must also be satisfied.

Classes of Adversaries

No treatment of security is complete without a discussion of adversary model. For our purposes, the set of adversaries includes, but is not limited to:

- *Inside attackers*: Such adversaries primarily seem to be legitimate participants in the OSN but act in a malicious way in some cases; for example, a malicious OSN provider, a malicious third-party application provider, a malicious user of the OSN, or a malicious party that has access to the network infrastructure (e.g., an eavesdropper of a wireless link or a malicious ISP).
- *External attackers or intruders*: An intruder is not a legitimate participant in the OSN, but can perpetrate attacks on the OSN system or on the network infrastructures used by the OSN.

Design Conflicts

As mentioned earlier, inherent design conflicts exist between some security and privacy goals and traditional design goals such as usability and sociability of OSNs.

Social Space Exploring vs. Privacy

To support social search and traversal, disclosing some information about users' profiles and lists of contacts is inevitable. For social search, obviously, more personal data must be disclosed in order to support more efficient and accurate search in digital social space; thus, there is a trade-off between search capabilities and privacy. Specifically, higher search efficiency also means higher likelihood of potential privacy breaches.

For social traversal, public display of social connections also affects OSN users' privacy. First, social contacts themselves represent sensitive, private information that can be misused. For example, the fact that two professionals employed by rival companies are friends may trigger suspicion. Second, social contact information can be utilized to infer more private information. For example, suppose Bob's user profile is encrypted and only accessible to his friends while all users' friend lists are open in order to facilitate social traversal in OSNs. In this case Bob's private data can be indirectly inferred by adversaries. Intuitively, friends tend to share common traits. For example, elementary school classmates have similar ages and the same hometown, and members of a swimming club like swimming. Therefore, to infer someone's age or personal interests, we can check the values of these attributes of his classmates or club mates.

Note that on the other side disclosing social contact information also provides many new ways to protect OSNs. For example, the paths in an online social graph are useful to express trustworthy users: nearby users/nodes in an online social graph often deserve a higher level of trust. Also, paths in an online social graph can provide a basis for access control mechanisms suitable for OSNs, where users determine authorized users based on their distance to themselves in the social graph. Last, ensuring the receiver of an email that the sender is nearby in an online social graph can help avoid falsely flagging the email as spam [12].

Social Interaction vs. Privacy

To facilitate social interactions is the main functionality of all OSNs. However, it may lead to privacy leakage in a uncontrollable way for OSN users. For example, Bob is a high school teacher and a member of the gun club who wants to keep his personal and professional life separate. So he participates in two OSNs: Personal OSN (P-OSN) for his personal friends and Work OSN (W-OSN) for all his professional contacts. Alice and Carol are on his P-OSN and W-OSN friend lists, respectively. At the very beginning, Bob's P-OSN and W-OSN are completely disjoint. However, a few days later, Carol finds her lost-contact classmate Alice on the network, and Alice invites Carol to join the P-OSN. Carol wants to find more lost-contact classmates by examining Alice's friend list in the P-OSN, and finds that her colleague Bob is a gun collector. When Bob is aware of this change, lots of his personal information has already been made known to his contacts in the W-OSN.

In the previous example, Bob's identity information is disclosed in both his P-OSN and W-OSN. Even if Bob does not reveal any of his personal information, others may do so during their social interactions. We illustrate how this could happen with the following example. Suppose Bob uses the pseudonym 5473625 to protect his identity, and all his personal space information is encrypted and only accessible to his friends. Alice, a friend of Bob, may still reveal Bob's age, occupation, and even his real name during their interactions in the following ways:

- Alice may upload a photo taken with Bob to her online albums. She may also annotate the photo with "Bob and Alice at Kennedy Space Center" and link the photo to user 5473625's profile. This action would reveal user 5473625's full name, real identity (through face recognition), and the fact that Bob and Alice are friends.
- Suppose Alice also recommends Bob as "the best teacher in Orlando" on her own page, and thereby inadvertently reveals Bob's occupation and the city where he lives.

Thus, Alice may unintentionally reveal a great deal of information about Bob without his knowledge. In other words, Bob's efforts to protect his identity may easily be nullified by others' behavior during their social interactions. Moreover, it is difficult for Bob to detect occurrences of such privacy leakages due to their distributed nature.

Data Mining vs. Privacy

Data collected by OSN providers or data aggregators are an important source for social and marketing analysis, which may provide useful information on the evolution of a social community, collaborative problem solving, and so on. Additionally, they can also be used to optimize OSN services and customize them with respect to users' preferences and interests. However, potential conflicts between social data mining and OSN privacy requirements may arise. An adversary may intrude on the privacy of OSN users using the published OSN data and some background knowledge. Even after we hide users' identities by replacing the corresponding user names with meaningless random identifiers (i.e., node anonymization), it has been shown that [13], based on graph topology information, adversaries still can recover most users' identities. Therefore, in order to meet privacy requirements of OSNs, edge perturbation should be introduced in the graph anonymization procedure [14, 15]. Obviously, there is a trade-off between the quality of the result of data mining and the privacy requirements of OSN users. Although multiple graph anonymization techniques are proposed in the literature [13–15], our knowledge of this trade-off is still limited. It is desirable that these two research directions (i.e., data mining and OSN privacy preservation) find some common points in order to come up with a comprehensive design framework enabling the study of this trade-off in a quantitative way.

Client-Server vs. P2P Architectures

A client-server architecture exhibits several advantages over a P2P architecture in fulfilling traditional goals of OSNs. For example, OSN users are not restricted to interactions with friends they already have, but may re-establish lost social connections if the people are also on the OSN. When a user searches for old classmates in digital social space by global name search or social graph traversal, a centralized website can bring exposure more easily. Also, users want to find others with similar interests, location, or organization. By data mining, affinity groups can be found. For example, comparing users' book buying behaviors exposes indications for shared interests. Data mining is both more efficient and more effective in a central repository. However, with client-server architecture the complete data, directly or indirectly supplied by all users, is collected and stored permanently in the databases of the OSN provider, which potentially becomes a big brother capable of exploiting this data in many ways that can violate the privacy of individual users. This private data can also fall under the control of hackers. Given the reality of privacy breaches by centralized OSN providers, as exemplified by Facebook's beacon application [1, 4, 5], there is a motivation for giving control over data back to the users and not having one entity access all personal data of the users in the OSN. In a P2P architecture user privacy is strengthened by removing the big brother. Users' data is not stored in a central repository, but by OSN users themselves, who also carry out the tasks related to access control enforcement and privacy protection by encrypting their own data. It seems that in a centralized system, one could also store encrypted data centrally without loss of privacy. This is, however, not strictly true. First, OSN providers can forbid users to encrypt their data by refusing to provide any services. Second, OSN providers can cease to provide the service or change its terms at anytime. Third, OSN providers might still be able to infer who is related to whom based on accesses and correlations, such as to IP addresses.

A P2P architecture combined with an appropriate encryption scheme can provide better privacy protection in OSNs. However, for a P2P architecture, how to efficiently and effectively recreate all core functionalities of OSNs in a decentralized way is still an open problem. Taking data mining as an example, how to encourage cooperation among users to collaboratively collect online social graph information in a privacy preserving way needs to be further investigated.

Research Directions

Although completely eliminating design conflicts between the various goals for OSNs might be impossible, several directions deserve further research and exploration. We confine ourselves primarily to a high-level examination of these directions, some of which we plan to build on in future research.

Developing an Enriched Relationship Model for OSNs

In reality, social relationships are extremely diverse in terms of strengths and types of relationships. OSNs, on the other hand, often reduce these nuanced connections to simplistic binary relationships: *friend* or *not* [1]. This results in an inconsistency between the real-life social network and the online social graph model, and violates our security principle of keeping consistency between online and offline social networks. In order to capture multiple aspects of real-life social networks, we must extend the relationship model (or link model for an online social graph) to include, for example:

- *Type of relationships*, which can be roughly categorized into bidirectional relationships such as friend or colleague, and one-directional relationships such as fans or followers
- *Trust strength*, which expresses how much a user trusts other users either with respect to a specific topic (topical trust) or in general (absolute trust)
- *Interaction intensity*, which measures the quality and quantity of interactions between users

An enriched social relationship model can improve OSN privacy and security in multiple ways. First, we share different information with our friends and colleagues. A confusion of relationship types may cause embarrassment; therefore, all social relationships should be clearly articulated and treated accordingly in making privacy decisions. Second, trust relationships are the core information on which all security mechanisms are based. By their very nature, trust relationships among users are not equal. Traditional privacy polices based on binary trust relationships ignore the existing strength differences and treat them as equal. Therefore, they cannot provide fine-grained access control and may lead to privacy breaches. Third, interaction intensity can be used as a proxy for relationship quality for the purpose of making privacy decisions. In general, if a pair of users do not interact often, they only want to reveal a limited amount of information to each other [16]. The measurement of interaction intensity also introduces a way to characterize network dynamics.

However, the trade-off between accuracy and complexity in describing social relationships must be taken into account. Inaccurate and ambiguous descriptions will introduce security vulnerabilities, but evaluating and processing too complex descriptions may have computational costs that make the OSN infeasible in practice.

Protecting Online Social Graphs

A fundamental feature of OSNs is the online social graph that connects users. It collects the core information on which all the socialization services provided by OSNs are based; therefore, it should be primarily protected. In this subsection we show that trust relationships and connection patterns embedded in real-life social networks can be utilized to provide security mechanisms to protect the OSN and mitigate the attacks on online social graphs mentioned earlier.

Defense against Social Link Forging Attacks — A social link can be forged if a malicious user (Malory) successfully fools an honest user (Alice) into trusting him. However, the trust strength and interaction intensity introduced in the relationship model can effectively limit the impact of this kind of forged social links. Malory may also try to present as Bob, a good friend of Alice, in order to establish a strong social link with Alice. However, most people can tell quite easily if a friend's profile is faked, and a large proportion of OSN users also meet in person, allowing them to perform a *face-to-face* identity verification. The very nature of OSNs makes it difficult for attackers to forge social links.

Defense against Node Identity Forging Attacks — For many OSNs, it is easy for a malicious user to obtain multiple fake identities and pretend to be multiple distinct users in the OSN. If the OSN requires users to register with government-issued identity cards, the barrier against launching node forging attacks becomes much higher. However, users' privacy cannot be guaranteed with this centralized scheme. A more promising approach is to utilize the web of relationships embedded in a real-life social network to establish and verify user identity. The basic idea is that people are who they connect with, communicate with, or affine to. Your position in the social network defines your identity. Although it is not difficult to register under an alias, it is extraordinarily difficult to change one's friends and contacts. A node's position in the online social graph can be verified based on trust strengths of paths from the verifier to the claimant.

Although online social graphs can provide huge amounts of trust data to facilitate the design of defending mechanisms, personal relationships represent sensitive private information that can also be misused. The key point, therefore, is to find a privacy-preserving way to utilize the knowledge of online social graphs. One approach to mitigate the design conflicts is to try to find and utilize *qualitative properties* of real-life social networks. Unlike quantitative properties, qualitative properties are very malleable (i.e., applicable to any secure online social graphs) and reveal no information about individual users. For example, in [17], based on the fact that it is difficult to forge social links (i.e., attack edges) between honest nodes and Sybil nodes, Yu *et al*. propose a new defense scheme against Sybil attacks for OSNs. The basic insight is that if malicious users create too many Sybil nodes, the online social graph becomes strange in the sense that it has a small quotient cut, that is, a small set of social links (the attack edges) whose removal disconnects a large number of nodes (all the Sybil nodes) from the rest of the graph. On the other hand, real-life social networks do not tend to have such cuts (i.e., a qualitative property of OSNs). By utilizing a special kind of verifiable random walk in the online social graph and intersections between such walks, the small quotient cut can be identified, and the number of Sybil nodes can be accordingly bounded. We believe that more and more qualitative properties of social networks combined with privacy preserving cryptographic techniques [12] can be utilized to design novel security mechanisms for OSNs without compromising user privacy.

Conclusions

In this article we have discussed the security and privacy design issues on online social networks and pointed out a few research directions for mitigating the design conflicts between the various design design goals of OSNs. However, an ultimate solution will require experts from the social science and network security communities, industry, regulatory bodies, and all other relevant communities to collaboratively make decisions on both secure mechanisms and policies. This article is intended to provide a starting point for developing effective secure and privacy-preserving OSNs. We hope that this work will motivate OSN researchers and developers to move forward with more creative design of OSNs without compromising users' data security and privacy.

References

[1] D. M. Boyd and N. B. Ellison, "Social Network Sites: Definition, History, and Scholarship," *J. Comp.-Mediated Commun.*, vol. 13, no. 1, Oct. 2007, pp. 210–30.

[2] Wikipedia, "Social Network Service," 2010; http://en.wikipedia.org/wiki/Social_network_service

[3] S. B. Barnes, "A Privacy Paradox: Social Networking in the United States," *First Monday*, vol. 11, no. 9, Sept. 2006.

[4] R. Gross and A. Acquisti, "Information Revelation and Privacy in Online Social Networks," *Proc. WPES '05*, Alexandria, VA, Nov. 2005.

[5] B. Krishnamurthy and C. E. Wills, "Characterizing Privacy in Online Social Networks," *Proc. WOSN '08*, Seattle, WA, Aug. 2008.

[6] J. He and W. W. Chu, "Protecting Private Information in Online Social Networks," in *Intelligence and Security Informatics: Techniques and Applications*, H. Chen and C. C. Yang, Eds., Springer, 2008.

[7] ENISA, "Security Issues and Recommendations for Online Social Networks," Position Paper, Nov. 2007; http://www.enisa.europa.eu/doc/pdf/deliverables/enisa_pp_social_networks.pdf

[8] B. Carminati, E. Ferrari, and A. Perego, "Enforcing Access Control in Web-Based Social Networks," *ACM Trans. Info. Sys. Security*, vol. 13, no. 1, 2009, pp. 1–38.

[9] P. W. L. Fong, M. Anwar, and Z. Zhao, "A Privacy Preservation Model for Facebook-Style Social Network Systems," *Proc. ESORICS '09*, St. Malo, France, Sept. 2009.

[10] S. Buchegger *et al.*, "A Case for P2P Infrastructure for Social Networks — Opportunities and Challenges," *Proc. WONS '09*, Snowbird, UT, Feb. 2009.

[11] S. Buchegger and A. Datta, "PeerSoN: P2P Social Networking — Early Experiences and Insights," *Proc. SocialNets '09*, Nuernberg, Germany, Mar. 2009.

[12] G. Mezzour *et al.*, "Privacy-Preserving Relationship Path Discovery in Social Networks," *Proc. CANS '09*, Ishikawa, Japan, Dec. 2009.

[13] L. Backstrom, C. Dwork, and J. Kleinberg, "Wherefore art thou r3579x?: Anonymized Social Networks, Hidden Patterns, and Structural Steganography," *Proc. WWW '07*, Banff, Canada, May 2007.

[14] E. Zheleva and L. Getoor, "Preserving the Privacy of Sensitive Relationships in Graph Data," *Proc. PinKDD '07*, San Jose, CA, Aug. 2007.

[15] M. Hay *et al.*, "Anonymizing Social Networks," Tech. rep., Univ. MA Amherst, Mar. 2007; http://www.cs.umass.edu/~mhay/papers/hay-et-al-tr0719.pdf

[16] L. Banks and S. F. Wu, "All Friends are Not Created Equal: An Interaction Intensity Based Approach to Privacy in Online Social Networks," *Proc. WSPOSN '09*, Vancouver, Canada, Aug. 2009.

[17] H. Yu *et al.*, "Sybilguard: Defending Against Sybil Attacks via Social Networks," *IEEE/ACM Trans. Net.*, vol. 16, no. 3, June 2008, pp. 576–89.

Biographies

CHI ZHANG (zhangchi@ufl.edu) is a Ph.D. candidate in electrical and computer engineering from the University of Florida, Gainesville. His research interests include network and distributed system security, wireless networking, and mobile computing.

JINYUAN SUN (stellas@ufl.edu) is a Ph.D. candidate in electrical and computer engineering from the University of Florida, Gainesville. Her research interests include security and privacy in wired/wireless networks and application systems, mobile computing, and computer networks.

XIAOYAN ZHU (xyzhu@mail.xidian.edu.cn) received her Ph.D. from the School of Telecommunications Engineering at Xidian University, Xian, China, in 2009, where she is currently a lecturer. She was a research visiting scholar at the Wireless Networks Laboratory in the Department of Electrical and Computer Engineering at the University of Florida from 2008 to 2010. Her research interests include wireless networks, network security, and network coding.

YUGUANG FANG [F] (fang@ece.ufl.edu) received a Ph.D. in systems engineering from Case Western Reserve University in 1994 and a Ph.D. in electrical engineering from Boston University in 1997. He is a professor of electrical and computer engineering at the University of Florida and a Changjiang Scholar chair professor at Xidian University. He is currently serving as Editor-in-Chief for *IEEE Wireless Communications*. He is a member of the ACM.