Secure Encounter-based Social Networks: Requirements, Challenges, and Designs

Abedelaziz Mohaisen
University of Minnesota
Minneapolis, MN 55455, USA
mohaisen@cs.umn.edu

Eugene Y. Vasserman
Kansas State University
Manhattan, KS 66506, USA
eyv@ksu.edu

Max Schuchard
University of Minnesota
Minneapolis, MN 55455, USA
schuch@cs.umn.edu

Denis Foo Kune
University of Minnesota
Minneapolis, MN 55455, USA
foo@cs.umn.edu

Yongdae Kim
University of Minnesota
Minneapolis, MN 55455, USA
kyd@cs.umn.edu

ABSTRACT

In this paper we outline requirements, challenges, and designs for encounter-based mobile social networks, where relationships are based on a temporarily shared location. To illustrate the challenges we examine a recently proposed design, SMILE, against a set of functional and security requirements. We show that SMILE is vulnerable to several attacks such as impersonation, collusion, and privacy breaching, even though it was built with the explicit goal of resisting some of those attacks. With this in mind, we construct a flexible framework for secure mobile social networks, and describe how to use it in order to construct several networks which offer somewhat different security properties. Each of the designs is then examined against the ideal requirements where some are shown to outperform previous work.

Categories and Subject Descriptors

C.2.0 [**Computer Communication Networks**]: General – *Security and Protection*

General Terms

Security, Design

Keywords

Location-based services, Privacy, Trust, Social networking

1. INTRODUCTION

In the conventional model of social networks, users select their social contacts from a set of acquaintances. Despite their utility, these conventional networks support social networking in a very confined manner. Two users will only be able to establish a relationship in the social network if they know of, or are introduced to, each other. On the other hand, in an encounter-based social network the only requirement for establishing a connection is that you are in the same place at the same time as someone. Encounter-based social networks provide a good solution for systems that can't afford the high connection requirements of traditional social networking systems. For example, encounter-based systems can easily form the backbone of a powerful "missed connections" service.

CCS'10, October 4–8, 2010, Chicago, Illinois, USA.
ACM 978-1-4503-0244-9/10/10.

While encounter-based systems at first glance appear very similar to traditional systems, they present a dramatically different set of security concerns. Guarantees that are quite trivial in a traditional social network, such as ensuring you are communicating with the correct person, become problematic in an encounter-based social network. Additionally, requirements like anonymity, something not found at all in traditional social networks, need to be considered in an encounter-based network. In this paper, we describe the unique security challenges posed by encounter-based social networks. We examine SMILE [6], an encounter-based social network that utilizes mobile devices and a centralized meeting point. In SMILE, users exchange credentials via a mobile device to prove their presence at a location during a specific time. Users later interact with a central server which, with the aid of the supplied credentials, acts as a rendezvous point for the two users. SMILE presents itself as a secure system. In particular, SMILE claims to provide unlinkability to the two parties involved in an encounter from both the server and other users. This is, SMILE claims to guarantee that the server will not be able to link two users in the encounter settings by observing their encounter keys. However, SMILE fails to achieve many of its security goals. Additionally, SMILE also completely fails to consider how users authenticate the party they meet at the rendezvous point. After describing the limitations of SMILE, we construct a generic and flexible system which meets the security and functionality goals of an encounter-based social network.

2. REQUIREMENTS

In this section, we summarize some of the requirements of an ideal encounter-based social network. While these requirements are by no mean complete, they can be used as a guideline for evaluating potential designs.

2.1 Security requirements

- *Privacy:* the privacy of the two parties sharing an encounter must be protected, even from others in the vicinity who may also participate in the same encounter. In this case, privacy means that an external adversary, even one taking part in the encounter but is not one of the two users of interest, should not be able to conclusively determine that two users have made a connection. Note that we must take into account potential collusion between users in the encounter vicinity and the central "rendezvous" server.
- *Confidentiality:* information exchanged privately between the encounter parties should only be accessible to them.

- *Authenticity:* when two users of a network decide to make a connection, they should be assured that private messages indeed come from each other.

2.2 Functional requirements

- *Availability:* the infrastructure to exchange encounter information should be accessible by system users *most of the time*. Since the time at which encounter parties check for potential encounters associated with their activities could be any arbitrary time, the encounter-based social network is more sensitive to the availability than conventional social networks. Also, the availability in such system implies difficulty in disrupting the system by misbehaving users.

- *Scalability:* with typical social networks being large in size, any potential social network design, including those based on encounters, should scale to admit large number of simultaneous users. This requires a flexible design that minimize the dependence on a centralized server.

3. CHALLENGES

While it may appear that implementing the above requirements is straightforward, it presents a surprising challenge. Recently, Manweiler, Scudellari, and Cox introduced SMILE [6], an attempt to implement a subset of the above requirements. While they succeed in meeting the functional requirements, they fail to protect against a number of common security vulnerabilities, such as the "man-in-the-middle" attack, or MitM. In this section, we discuss some of these challenges and the approach taken in SMILE. Before getting into further details, we review the operation and claimed guarantees of SMILE.

3.1 Overview of SMILE

SMILE extends ideas from [5] and uses cryptographic construction to establish trust between individuals who shared an encounter. SMILE attempts to allow users equipped with mobile devices to build such trust relationships while preserving their privacy against potential attackers (e.g., the central storage server and other users). In SMILE, users who want to communicate with each other must prove that an encounter occurred. To do this, an interested person generates and passively broadcasts the "encounter key" to others within his communication range, and posts a hash of the encounter key, along with a message encrypted using the encounter key, to a centralized server. Other users in SMILE with the same encounter key may claim the encounter by simply looking up the hash of the key, which is used for indexing the encrypted message at the centralized server. Only the user with the correct key will be able to decrypt the message left by the first encounter party at the server.

In addition to the basic design, SMILE aims to provide two features: k–anonymity, and decentralized design. Anonymity is achieved by truncating hash values of keys so that a single user is obscured amongst k users with the same truncated value. Also, SMILE features a decentralized system that uses anonymizing networks of remailers for communication, claiming to provide k–anonymity by requiring each user to have at least k identifiers.

3.2 SMILE security revisited

We now examine exactly which of our previously-derived requirements SMILE meets. SMILE's availability and scalability are questionable, since the system depends on a centralized server that is easy to disrupt. (This problem is not unique to SMILE, but rather any design that uses a centralized online entity.) The security

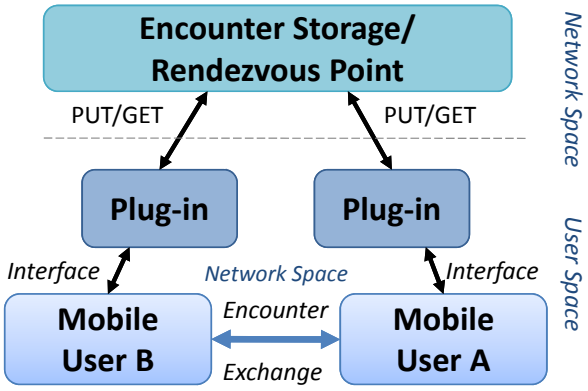

Figure 1: Architectural depiction of generic encounter-based social network.

guarantees of SMILE are also in question. While the confidentiality of encounter-related information is safeguarded by encryption, the privacy of users in SMILE may be breached in several different ways: first, SMILE is vulnerable to the simple, yet powerful, impersonation attack performed by an eavesdropper. Since no authentication is required or provided during key agreement, any user can eavesdrop on the encounter information at the encounter site and later claim to be the party of interest. This attack can further be extended to monitoring: if the adversary exchanges keys with the first user pretending to be the second, and repeats this with the other user, the adversary can monitor all messages passed between users. Second, SMILE is vulnerable to user collusion, an attack that is inherently valid in social interaction settings [4]. In particular, a few colluding users may possess enough information about the activities of other honest users (such as time-stamps, locations, and encounter keys) for the server to unmask users, determining the identities of communicating parties. Finally, the k–anonymity property in SMILE requires that each user know the number of other nearby SMILE users, which can be easily misrepresented by a simple Sybil attack [2, 8].

4. DESIGNS FOR SECURE ENCOUNTER-BASED SOCIAL NETWORKS

With the requirements outlined in section 2, we proceed to describe a generic design, that is in essence similar to the design of SMILE, but without the same security vulnerabilities. We divide the design into functional entities and describe potential attacks on various parts of the system. Finally, we show some instantiations of the generic design, each having different benefits and trade-offs.

4.1 Functional Components

The functional design of a typical encounter-based social network consists of three major components located at three different architectural layers: the user layer, the plug-in layer, and "the cloud," referring to the storage location of the encounters and private messages, and used by different encounter parties in the post-encounter phase. Storage components can be dynamically chosen using a plug-in architecture to support centralized services, distributed hash tables [7], or even Tor hidden services [3]. The three different layers are shown in Figure 1.

We have shown that simple unauthenticated key agreement during the encounter is vulnerable to a MitM attack. The only way to avoid this vulnerability is a visual authentication scheme, where users can recognize that they are communicating with the desired

party simply by looking at a picture. To provide user authentication, we assume each user to have a digital certificate signed by a trusted authority with sufficient information to identify users, including a photo of the user. It is not far-fetched to assume that future authentication tokens such as passports and driver licenses will be issued digitally, since cryptographic signatures make them more secure against malicious tampering than their physical counterparts. With that assumption, a user of an encounter-based system can broadcast a certificate with his or her picture and public key.

Here we face two potential design choices: do we require for immediate encounter key agreement between the two parties, or do we wait? Each approach has a benefit and drawback. Immediate generation of an encounter key requires manual selection of the target user. Delayed generation, on the other hand, requires no immediate action on the part of the user, but can potentially expose more user information during later communication. Both of these methods are discussed further below.

4.2 Immediate Key Generation

If a user is willing to manually select the picture of other users of interest while still at the encounter site, she can compose an encounter key, encrypt it to the selected user's public key using a non-malleable encryption scheme, and broadcast the resulting message. Each user in the vicinity will detect the transmission, and attempt to decrypt. However, only the target user will be able to decrypt the message correctly, learning the encounter key.This key will be used later to exchange private messages at the rendezvous point. This method prevents the rendezvous server and colluding adversaries at the encounter point from determining which two users are communicating. We can go a step further and use time-release encryption to hide the contents of the encounter key even from its intended recipient until the encounter is over, to ensure the users do not inadvertently give themselves away by using their devices at the same time.

4.3 Delayed Rendezvous

Devices will consistently broadcast their certificates, but will not require others users to immediately review their information. At a later time, the device user can look at the list of collected identities (and public keys), and select those with whom he or she wishes to communicate. As before, we will use non-malleable encryption to compose a message to the other user, but now the message must be stored at the rendezvous server in such a way that it is linkable to the public key of the user for whom it is intended. This may or may not be a problem, considering that only keys and faces are exposed, and not more personal components of users' identities.

4.4 Anonymity During Communication

In both schemes, users use the Tor network [1] to gain anonymity when posting encounter information or looking them up at the public server. Again, hashes of keys are used for indexing and unlike SMILE, our design is immune to impersonation since this is equivalent to forging a certificate, as any impersonator needs to own the corresponding private key of the public key of the user.

By limiting the amount of information exposed to the server, we limit chances of a malicious server to match users. The same functional guarantees are provided as in SMILE, which are limited by the centralized server.

5. CONCLUSION

In this work we show that existing designs for secure encounter-based social networks fail to meet their security requirements. We outline several requirements that an ideal encounter-based social networks need to satisfy in order to be secure. We further introduce generic framework, extended to several designs with the requirements in mind, and show that some of these designs outperform the other designs in literature.

Acknowledgement

This research was supported by the NSF under grant no. CNS-0917154 and a research grant from Korea Advanced Institute of Science and Technology (KAIST).

6. REFERENCES

[1] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the USENIX Security Symposium*, 2004.

[2] J. Douceur. The sybil attack. *Peer-to-Peer Systems*, pages 251–260, 2002.

[3] J. Lenhard, K. Loesing, and G. Wirtz. Performance measurements of tor hidden services in low-bandwidth access networks. In M. Abdalla, D. Pointcheval, P.-A. Fouque, and D. Vergnaud, editors, *ACNS*, volume 5536 of *Lecture Notes in Computer Science*, pages 324–341, 2009.

[4] M. Macy. Learning to cooperate: Stochastic and tacit collusion in social exchange. *The American Journal of Sociology*, 97(3):808–843, 1991.

[5] J. Manweiler, R. Scudellari, Z. Cancio, and L. P. Cox. We saw each other on the subway: secure, anonymous proximity-based missed connections. In *HotMobile '09: Proceedings of the 10th workshop on Mobile Computing Systems and Applications*, pages 1–6, New York, NY, USA, 2009. ACM.

[6] J. Manweiler, R. Scudellari, and L. P. Cox. SMILE: encounter-based trust for mobile social services. In E. Al-Shaer, S. Jha, and A. D. Keromytis, editors, *ACM Conference on Computer and Communications Security*, pages 246–255. ACM, 2009.

[7] P. Maymounkov and D. Mazières. A peer-to-peer information system based on the XOR metric. 2002.

[8] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao. Sybillimit: A near-optimal social network defense against sybil attacks. In *IEEE Symposium on Security and Privacy*, pages 3–17. IEEE Computer Society, 2008.