Safebook: A Privacy-Preserving Online Social Network Leveraging on Real-Life Trust

Leucio Antonio Cutillo and Refik Molva, Institut Eurécom
Thorsten Strufe, TU Darmstadt

ABSTRACT

Online social network applications severely suffer from various security and privacy exposures. This article suggests a new approach to tackle these security and privacy problems with a special emphasis on the privacy of users with respect to the application provider in addition to defense against intruders or malicious users. In order to ensure users' privacy in the face of potential privacy violations by the provider, the suggested approach adopts a decentralized architecture relying on cooperation among a number of independent parties that are also the users of the online social network application. The second strong point of the suggested approach is to capitalize on the trust relationships that are part of social networks in real life in order to cope with the problem of building trusted and privacy-preserving mechanisms as part of the online application. The combination of these design principles is Safebook, a decentralized and privacy-preserving online social network application. Based on the two design principles, decentralization and exploiting real-life trust, various mechanisms for privacy and security are integrated into Safebook in order to provide data storage and data management functions that preserve users' privacy, data integrity, and availability. Preliminary evaluations of Safebook show that a realistic compromise between privacy and performance is feasible.

INTRODUCTION

Social networking services (SNS), like *facebook*, *LinkedIn*, and *orkut*, are a predominant service on the web today. Catering for a broad range of users of all ages, and vast differences in social, educational, and national backgrounds, they allow even users with limited technical skills to publish personal information and communicate with ease. In general, the online social networks (OSNs) that are stored for this purpose are digital representations of a subset of the relations that their participants, the registered persons or institutions, experience in the physical world. Spanning all participating parties through their relationships, they model the social network as a graph. However, the popularity and broad acceptance of social networking services as platforms for messaging and socializing attract not only faithful users, who are trying to add value to the community, but parties with rather adverse interests, be they commercial or plain malicious, as well.

The main motivation for members to join an OSN, create a profile, and use the different applications offered by the service is the possibility to easily share information with selected contacts or the public, for either *professional* or *personal* purposes. In the first case, the OSN is used as a facility geared toward career management or business goals; hence, SNS with a more serious image, like XING and LinkedIn, are chosen. As members in this case are aware of the professional impact of the OSN, they usually pay attention to the content of the data they publish about themselves and others. In the case of more private use, they share more personal information like contact data, personal pictures, or videos. Other members in the shared pictures can be marked (*tagged*), and links to their respective profiles are created automatically.

The core application used by the members of SNS is the creation and maintenance of their contact lists, which describe the members' milieu and maps them into the digital OSN graph.

Through informing members automatically on profile changes of their contacts, SNS thus help users to stay up to date with news of their contacts and very often the popularity of users is measured in the number of contacts to which their profile links.

These properties of the services have led to the definition of boyd and Ellison [1], according to which *social network sites* or *online social network services* are: "*web-based services that allow individuals to (1) construct a public or semi-public profile within a bounded system, (2) articulate a list of other users with whom they share a connection, and (3) view and traverse their list of connections and those made by others within the system*."

This definition, however, leaves aside an additional service that becomes apparent when observing the use of SNS: the communication of

This work in parts has been supported by the SOCIALNETS project, grant no 217141, funded by the EC FP7-ICT-2007-8.2 for Pervasive Adaptation.

0163-6804/09/$25.00 © 2009 IEEE

members through direct, sometimes instant, message exchange or annotation of profiles (by comments or recommendations). Additionally, SNS typically enable a wealth of third-party applications featuring advanced interactions between members ranging from simple *poking* of another member or support for *interest groups* for a common topic to *likeness* testing with other members and the exchange of virtual *gifts*.

Storage, maintenance, and access to the OSN and their services are offered by commercial providers, like Facebook Inc.,[1] LinkedIn Corp.,[2] Google Inc.,[3] XING AG,[4] and the like.

Analyzing the OSNs with respect to their security properties and the privacy of their users, some obvious threats become apparent. Generally, a wealth of personal data on the participants is stored at the providers, especially in the case of OSNs targeting non-professional purposes.

This data is either visible to the public, or, if the user is aware of privacy issues and able to use the settings of the respective SNS, to a somewhat selected group of other members. As profiles are attributed to presumably known persons from the real world, they are implicitly valued with the same trust as the assumed owner of the profile. Furthermore, any actions and interactions coupled to a profile are again attributed to the assumed owner of this profile as well. Different studies have shown that participants clearly represent the weak link for security in OSNs and that they are vulnerable to several types of social engineering attacks [2–5]. This is partially caused by a lack of awareness regarding the consequences of simple and presumably private actions, like accepting contact requests, tagging pictures, or acts of communication like commenting on profiles or leaving wall posts. However, the usability of privacy controls offered by the SNS, and, finally and most important, inherent assumptions about other participants and trust in other profiles, which are actually a desired characteristic, certainly add to the problem.

However, analyzing the privacy problems in current OSNs, it becomes apparent that even if all participants were aware and competent in the use of SNS, and even if a concise set of privacy measures were deployed, the OSN would still be exposed to potential privacy violations by the omniscient service provider: the complete data, directly or indirectly supplied by all participants, is collected and stored permanently in the databases of the providing company, which potentially becomes a "Big Brother" capable of exploiting this data in many ways that can violate the privacy of individual users or user groups. The importance of this privacy exposure is underlined by the market capitalization of these providers, which ranges from US$580 million (acquisition of Myspace through the news corp. in 2005) to US$15 billion (Facebook Inc., according to the investment of Microsoft in 2007) [6].

In consequence, we consider the protection of private data in OSNs a pressing topic, which current providers are not likely to address. In this article we suggest an SNS called Safebook[5] that is specifically designed to prevent privacy violations by intruders, malicious users, and OSN providers alike. Safebook is mainly characterized by a decentralized architecture relying on cooperation among peers in order to prevent potential privacy violations due to centralized control. In addition to the description of Safebook, this article presents:

- A multilayered model of social networking services
- A security analysis of threats and attacks in online social networking

The next section states the security objectives for OSNs. We then analyze the security requirements of current SNS and present Safebook, our new approach to a privacy-preserving SNS. We conclude with a summary and an outlook in the final section.

Security Objectives in OSNs

In the context of OSNs, we generally identify three main security objectives, *privacy*, *integrity*, and *availability*, which come in slightly different flavors than in traditional systems.

Privacy

In accordance to previous studies [7, 8], we assume the protection of the user's privacy to be the main objective for SNS. Privacy not only encompasses the protection of personal information, which users publish on their profiles, presumably accessible by their contacts only. Additionally, communication privacy has to be met. Hence, none but directly addressed or explicitly trusted parties may have the possibility to trace which parties are communicating. Furthermore, details of messages have to be hidden, so only the requesting and responding parties should know one another's identity and the content of the request. Finally, disclosure of information about a third party to some member that is not explicitly trusted by the third party, without the consent of the latter, has to be prevented. In summary, privacy calls for the possibility to hide any information about any user, even to the extent of hiding their participation in the OSN in the first place. Moreover privacy has to be met by default; that is, all information on all users and their actions has to be hidden from any other party internal or external to the system, unless explicitly disclosed by the users themselves.

Requiring explicit disclosure directly leads to the need for *access control*. Access to information on a user may only be granted by the user directly; the access control has to be as fine-grained as the profile, and each attribute has to be separately manageable.

Integrity

As part of integrity, the user's identity and data must be protected against unauthorized modification and tampering. In addition to conventional modification detection and message authentication, integrity in the context of OSNs has to be extended: parties in an OSN are not arbitrary devices, but real, unambiguously identifiable persons. The creation of personae — bogus accounts, cloned accounts, or other types of impersonation — in traditional SNS is easy to

Privacy does not only encompass the protection of personal information, which users publish at their profiles, presumably accessible by their contacts only. But additionally, communication privacy has to be met.

[1] www.facebook.com

[2] www.linkedin.com

[3] www.orkut.com

[4] www.xing.com

[5] www.safebook.us

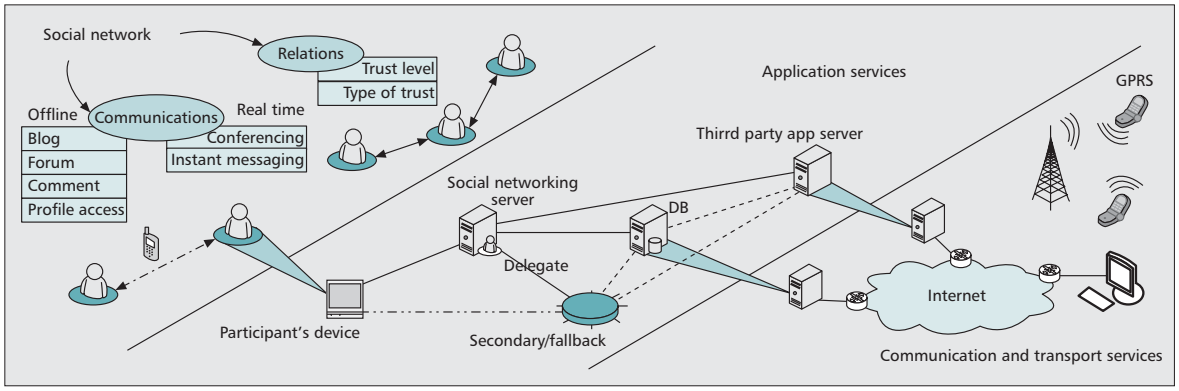

Figure 1. *OSN levels: three architectural layers of social networking services.*

achieve. However, users have a strong inherent trust in an OSN, and it has been shown that this combination may lead to new kinds of vulnerabilities [2, 3, 5]. In consequence, the authentication has to ensure the existence of real persons behind registered OSN members. Identity checks do not necessarily have to be performed by a centralized service; however, all identification services have to be trusted by all participants.

Availability

Since some SNS are used as professional tools to aid their members' business or careers, data published by users has to be continuously available. Availability of user profiles is consequently required as a basic feature, even though considering recreational use, the availability of some content may not seem a stringent requirement. In OSNs, this availability specifically has to include robustness against censorship, and the seizure or hijacking of names and other key words. Apart from availability of data access, availability has to be ensured along with message exchange among members.

Security Analysis of OSNs

First of all, we shall sketch a model for SNS, to get an overview on the aim and possible implementation schemes of SNS.

SNS can be divided into three different levels (Fig. 1):

- A **social network** (SN) level: The digital representation of members and their relationships
- An **application services (AS) level**: The application infrastructure, managed by the SNS provider
- A **communication and transport** (CT) level: Communication and transport services as provided by the network

The SN level provides each member with a set of functions corresponding to social interactions in the real life, like finding friends, accessing profiles, commenting, and the like.

To implement these functions, the SN level relies on the AS level. This second level includes the infrastructure managed by the SNS provider, together with basic services to create the SN service, such as web access, storage, and communication. Common strategies to enhance availability for these are redundancy and delegation: both for organizational reasons, or if a server faces failures or other inabilities to provide a service, it may delegate requests to secondary servers. Data storage and retrieval, indexing of content, management of access permissions to data, and node join or leave are implemented in a centralized or decentralized, distributed fashion on the AS level.

The AS level, on the other hand, relies on the transport and internetworking protocols and infrastructures, implemented by the CT level.

Based on this architecture of OSN, we define an attacker as one of the following:

- A malicious member on the SN level
- A malicious service provider on the AS level
- A party that has and misuses access to the infrastructure at the CT level (an eavesdropper with a local, or a malicious ISP with possibly even a global, view)

Other than these inside attackers that primarily seem to be legitimate participants in the system but act in a malicious way in some cases, there may be external attackers, or *intruders*. An intruder can perpetrate attacks at one or more of the SNS levels.

After defining the different levels of SNS, we shall characterize major attacks on SNS.

Privacy: The protection of a member's identity is one of the key aspects that still need to be addressed in current OSNs. In *identity theft*, for example, a malicious member or service provider acquires the credentials of authorized users and acts on their behalf with full access to the profile, relations, and communication traces. Due to the inherent trust in other profiles, plain *impersonation* by creation of a *clone* of the targeted profile[6] may suffice to be able to establish trust relationships with parties on a victim's contact list by simply sending new friendship requests. *Profile porting* attacks, in which the attacker creates a profile under the victim's identity in an OSN where the victim is not present, are more difficult to detect. However, with most existing accounts being unprotected, profile porting poses a valid threat. The collection of existing data is the basis of *profiling* attacks, data aggregation that gives an attacker the possibility to

[6] http://www.nature.com/news/2009/090423/full/news.2009.398.html

guess the value of a potentially huge set of usually disclosed properties, such as the victim's social security number, income bracket, potential interest in some product, and so on. They additionally supply potential attackers with the knowledge needed for *secondary data collection*, as from the data published on an OSN they may easily be able to guess the social security number (or, e.g., the *Foedselsnummer* in Norway), which often acts as a key to accessing personal information from a wide range of different sources. Matching the profiles of a person in both a professional and a more informal OSN for analysis and comparison of the content published in both is another obvious and frequent type of secondary data collection.

Moving from data to communication privacy, a series of other threats arises. A malicious SNS provider or, to some extent, a malicious member with the appropriate set of privileges can be able to perform *communication tracking* and reveal who is talking to whom. The problem becomes relevant, and much more difficult to solve, at the CT level with an omniscient ISP.

Another series of attacks on privacy is *profile harvesting*, in which an attacker, a malicious participant, or an SNS provider gathers data on the participants on a large scale for purposes that the victims have not considered, intended, or foreseen. More sophisticated harvesting comes in the form of *image retrieval*, possibly even in association with automated *face recognition* algorithms for further profiling.

Integrity: The above mentioned impersonation threats are due to a basic shortcoming: none of the current major OSNs is able (or cares, in many cases) to ensure that a profile is associated with a single real person. *Faked profiles* are a common phenomenon resulting from this shortcoming, as well as clones or ported profiles. Such impersonation paves the way for *Sybil attacks*, which aim at creating fake identities, as well as *defamation* and *ballot stuffing* attacks that aim at forging the reputation for a person using the system or disrupting digital reputation systems.

Availability: Several types of *denial of service* (e.g., to cover a victims profile or selected data, or to disrupt the possibility to communicate with a victim) are possible in SNS. A centralized OSN obviously is vulnerable to *censorship* through the SNS provider. However, distributed SNS, which are implemented as decentralized, possibly peer-to-peer (P2P), systems or follow other types of service delegation, may be vulnerable to a series of attacks from these domains. *Black holes*, *selective forwarding*, and *misrouting* are serious threats in this case.

Table 1 gives an overview of the relationship between the stated attacks and the involved security objectives. Some attacks breach several objectives, but still primarily focus on or mainly exploit a vulnerability to one of these objectives, in which case they are attributed to only this objective. Attacks that are not mainly related to a single security objective, such as collusions, have to be countered by a number of measures regarding different objectives, and they are hence attributed to more than one objective.

In conclusion, it becomes apparent that current SNS are still vulnerable to different attacks on all three levels by either insiders (legitimate parties) or outsiders (intruders). In the following section we describe Safebook, a new approach to decentralize SNS, to convey this approach as an alternative solution to open vulnerabilities that are unlikely to be fixed by current SNS providers.

A Decentralized OSN

Some of the security and privacy exposures analyzed in the previous section could be addressed through the enhancement of existing OSN applications, by integrating various security and privacy mechanisms. However, the privacy of users' data is at risk due to the central storage and management and hence threatened by potentially malicious service providers or unintended access following short-sighted publication,[7] security breaches, or plain misconfiguration of the OSN. It inherently cannot be ensured with centralized server-based architectures on which all existing OSNs rely. Peer-to-peer architectures seem to offer a suitable alternative to the centralized approach as the basis for a decentralized OSN, avoiding the all-knowing service provider. As a major drawback, P2P systems suffer from a lack of a priori trust, thus creating the need for cooperation incentives. We thus suggest a decentralized OSN based on a P2P architecture whereby basic security and privacy problems as well as the lack of a priori trust and incentives are addressed by leveraging on real-life trust between users, such that services like data storage or profile data routing are performed by peers who trust one another in the social network.

We suggest a decentralized OSN based on a P2P architecture whereby basic security and privacy problems as well as the lack of a priori trust and incentives are addressed by leveraging on real life trust between the users.

Safebook: Security Based on Real-Life Trust

Safebook consists of a three-tier architecture with a direct mapping of layers to the OSN levels depicted in Fig. 2 as follows:

- The user-centered social network layer implementing the SN level of the OSN
- The P2P substrate implementing the AS services
- The Internet, representing the CT level

Each party in Safebook is thus represented by a node that is viewed as a host node in the Internet, a peer node in the P2P overlay, and a member in the SN layer.

The nodes in Safebook form two types of overlays:

- A set of ***matryoshkas***, concentric structures in the SN layer providing data storage and communication privacy created around each node
- A P2P substrate providing lookup services

In addition to these nodes, Safebook also features a *trusted identification service* (TIS), providing each node unambiguous identifiers: the *node identifier* for the SN level and a *pseudonym*.

Each Safebook component plays an essential role since it implements a particular set of countermeasures against the threats presented earlier.

Matryoshka: Matryoshkas are concentric rings of nodes built around each member's node in

[7] http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1450006

Security objectives	Privacy	Integrity	Availability
Attacks			
ID theft	x	x	x
Profile cloning	x	x	
Profile porting	x	x	
Secondary data collection	x		
Profiling	x		
Communication tracking	x		
Face recognition	x		
Image retrieval	x		
Harvesting	x		
Fake profiles		x	
Sybil	x	x	x
Ballot stuffing		x	
Defamation		x	
Censorship			x
Collusion	x	x	x

Table 1. *Attacks vs. security objectives in online social networks; primarily affected objectives are highlighted.*

order to provide trusted data storage, profile data retrieval, and communication obfuscation through indirection. Each matryoshka thus protects the node in its center, the *core*, which on the SN layer is addressed by its node identifier. The nodes in the matryoshka are connected through radial paths on which messages can be relayed recursively from the outermost shell to the core and vice versa. All paths are based on trust relationships akin to the social network; thus, each hop connects a pair of nodes belonging to users linked by a trust relationship in real life. The innermost and outermost shells of a matryoshka have a specific role: the innermost shell is composed of direct contacts of the core, and each of them stores the core's data in an encrypted form. Hence, they are called the *mirrors*. Every node in the outermost shell acts as a gateway for all data requests addressed to the core, and is thus called an *entrypoint* (Fig. 3). All requests to a core are addressed using its node identifier. Real-time communication is responded to by the core itself; any kind of offline communication can be served by one of its mirrors as well. While the number of mirrors and entrypoints in each path is fixed, the number of nodes between them is variable, thus leading to paths with variable length on the same matryoshka.

[8] http://xlattice.sourceforge.net/components/protocol/kademlia/specs.html

P2P system: In order to provide a location service to find entrypoints for a user's matryoshka, the nodes create a P2P substrate. Currently, this substrate resembles a KAD,[8] and the pseudonyms are used as identifiers for the DHT. The searchable and registered keys are the hashed properties of the participating members and their node identifiers. Unlike the path across a matryoshka, the communication through the P2P layer does not rely on trusted links. However, using pseudonyms still protects members from privacy violations based on node identification and tracing through the untrusted P2P links.

TIS: The TIS ensures that each Safebook user gets at most one unique identifier in each category of identifiers. Based on an out-of-band identification procedure, the TIS grants each user a unique pair of a node identifier and a pseudonym, computed as the result of a keyed hash function on the set of properties that uniquely identify a party in real life, such as full name, birth date, birth place, and so on. Even if at first glance a centralized trusted third party service such as the TIS seems to contrast the purpose of decentralization as pursued by Safebook, the TIS, even though a centralized service provider, does not pose a privacy threat as it cannot trace users or their messages; nor can it peek into their private data. Moreover, the TIS can be implemented in a distributed and offline fashion.

Operations

Safebook implements different OSN operations:
- Account creation
- Data publication
- Data retrieval
- Contact request and acceptance
- Message management

Account Creation — In order to join Safebook, a new member $\mathcal{V}$ has to be invited by one of its real-life friends $\mathcal{A}$ that must already be a registered user. $\mathcal{V}$'s account is then created in the two steps of identity creation and matryoshka creation.

Identity creation: After $\mathcal{A}$'s invitation, $\mathcal{V}$ provides the TIS with its identity property set $name_v$, together with a proof of owning it. This credential request also contains the public keys $\mathcal{P}_v^+$ and $\mathcal{I}_v^+$ belonging to two keypairs $\mathcal{P}$ and $\mathcal{I}$, which are generated by $\mathcal{V}$ itself. The TIS then computes the node identifier of $\mathcal{V}$ and its pseudonym by applying two different keyed hash functions with two different unknown master keys to $name_v$. At this point, the TIS sends $\mathcal{V}$ back its pseudonym P_v and node id I_v together with the certificates $Cert(P_v, \mathcal{P}_v^+)$ and $Cert(I_v, \mathcal{I}_v^+)$ associating the peer and member identifiers of $\mathcal{V}$ to its public keys $\mathcal{P}_v^+$ and $\mathcal{I}_v^+$, respectively. The pseudonym keypair $\mathcal{P}$ is used to guarantee integrity and confidentiality of all messages exchanged in Safebook, as in each hop every message is signed using the sender's pseudonym private key and encrypted using the receiver's pseudonym public key, while the node id keypair $\mathcal{I}$ is used to guarantee the same properties for end-to-end communication between members.

It becomes evident that even if a valid member $\mathcal{V}$ repeats the account creation operation

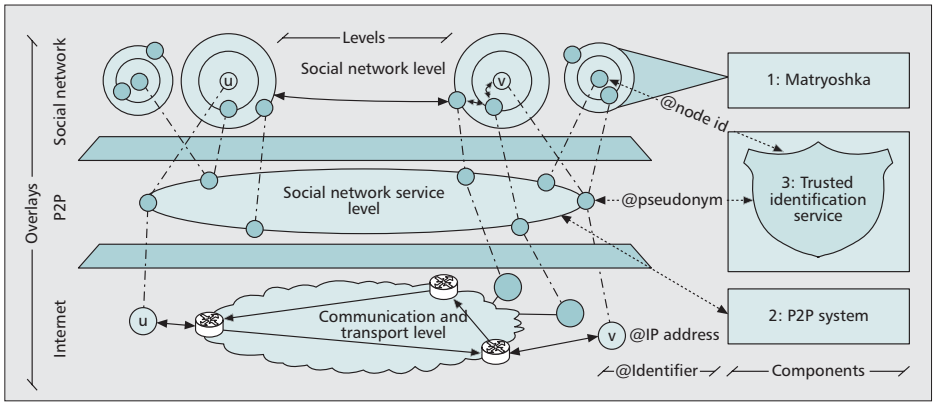

Figure 2. *Safebook overlays (left) and main components (right).*

multiple times, it will always receive the same pseudonym and node identifier, since they are a function of $\mathcal{V}$'s identity itself. Moreover, $\mathcal{V}$ cannot claim ownership of an identity that is not its own, since it would not be able to prove this fact. Identity proof is an out-of-band process that relies on real-life mechanisms to ascertain the identity of a potential member, such as a face-to-face meeting between a user and the representation of the TIS, or relying on existing tamper-proof schemes such as a passport or ID card. According to this fact, Sybil and impersonation attacks are not possible, as $\mathcal{V}$ cannot manipulate its node id or its pseudonym.

Once $\mathcal{V}$ gets its identifiers, it can join the P2P system by using $\mathcal{A}$ as a bootstrapping node and start the matryoshka creation process.

Matryoshka creation: $\mathcal{V}$ has only $\mathcal{A}$ as a contact to start with, so it sends $\mathcal{A}$ a request for path creation containing the distributed hash table (DHT) lookup keys it wants to register, a time to live (ttl), and the number of members to whom $\mathcal{A}$ should forward the request, hereafter called the *span* factor. $\mathcal{A}$ then selects between its friends a number *span* of next hops and forwards them this registration message. This process is recursively done until the ttl expires: the receiving node $\mathcal{D}$ registers the lookup key in the P2P system together with its reference $@d$ and starts acting as an entrypoint for $\mathcal{V}$.

Matryoshkas provide for privacy based on hop-by-hop trust, as all nodes in each matryoshka are only aware of their direct neighbors. As soon as $\mathcal{V}$ has created its matryoshka, it can publish its profile (Fig. 4).

Data Publication — The data managed in SNS can be generalized to:

- Profile information
- Contact relations
- Messages

The profile information is the part of the data each user intends to publish. To guarantee fine-grained access control, it is organized in atomic attributes for which particular access policies can be set. Contact relations represent a member's real-life relations and can be seen as the friend list of the user. As the strength of a relation is not the same for all links [9], in Safebook each user associates a particular trust level to each of its contacts. This level is used to select closely related contacts that primarily will store the published data. Finally, personal messages or comments on profiles can be exchanged between members. In case of comments, the receiver has the right to publish or discard them.

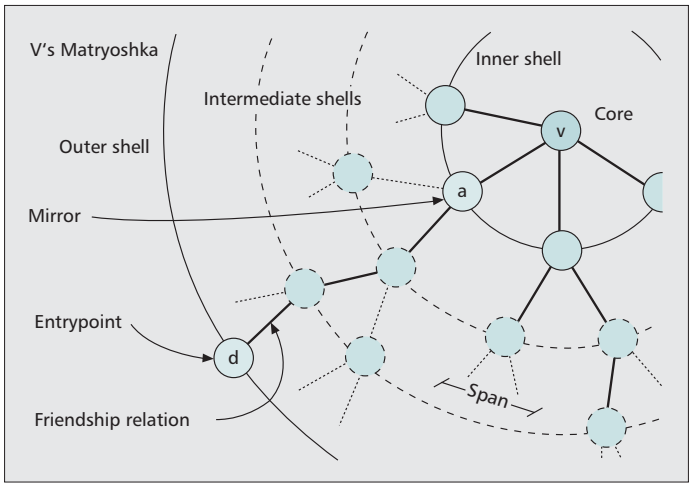

Figure 3. *Matryoshka structure.*

To guarantee privacy, data in Safebook can be private, protected, or public: in the first case the data is not published, in the second case it is published and encrypted, and in the third case it is published without encryption. All the published data of a member $\mathcal{V}$ is replicated to its mirrors, the nodes in the innermost shell of $\mathcal{V}$'s matryoshka.

Data Retrieval — The lookup of $\mathcal{V}$'s data through member $\mathcal{U}$ starts with a recursive query in the P2P system: according to the DHT structure, the node responsible for the lookup key responds with the entrypoint list building $\mathcal{V}$'s outer shell. Consequently, $\mathcal{U}$ can request that one of $\mathcal{V}$'s entrypoints forward the request through $\mathcal{V}$'s matryoshka until a mirror is reached. $\mathcal{V}$'s encrypted data then reaches $\mathcal{U}$ through the inverse path (Fig. 5).

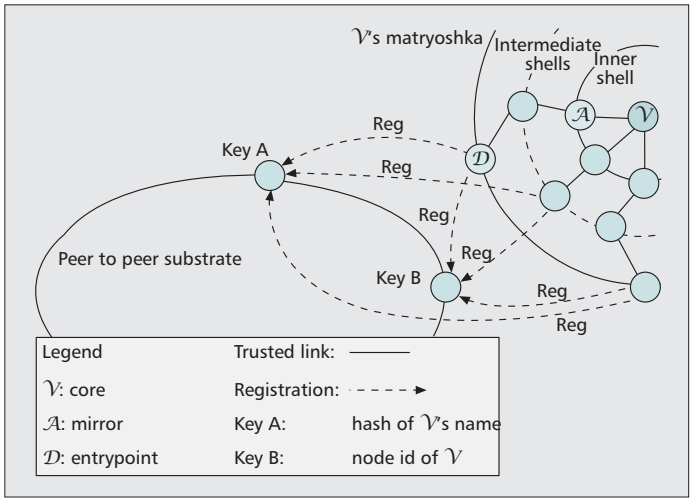

Figure 4. *Entrypoint registration in the P2P substrate.*

The protocol of Safebook uses recursion to hide the source of requests. Additionally, the addressing and routing, for both P2P lookup and data retrieval using the matryoshkas, are based on the pseudonyms of nodes. Attackers consequently have no means to identify a source of a request for some content, as there is no way to distinguish between generated and forwarded requests. Since the mapping between the pseudonym of a node and its identifier is only known to the TIS and direct connections (*friends*) in the matryoshka, which are trusted by the node, no private information can be derived from it either. Finally, communication tracking is not possible, as a malicious node would always have to be the first hop for all requests to the matryoshka of a certain node in order to be able to link the pseudonym of the sender to its real identity.

A preliminary feasibility study conducted with a previous and less well performing approach [10] showed that data retrieval performs well, even though the messages are forwarded along multiple hops in the overlay.

Contact Request and Acceptance — A member $\mathcal{U}$ that wants to add another member $\mathcal{V}$ to its contact list sends a contact request message following the same steps as in the data request case. Assuming $\mathcal{V}$ accepts $\mathcal{U}$ as a new contact, $\mathcal{V}$ associates with $\mathcal{U}$ a certain trust level (known by $\mathcal{V}$ and nobody else) and sends it back an opportune key that will enable $\mathcal{U}$ to decrypt the selected parts of $\mathcal{V}$'s published encrypted data.

Message Management — Offline messaging, such as wall posts, recommendations, and other annotations to a profile, is implemented using the steps of retrieving some member's data, decrypting the shared parts, annotating some content, and sending this data back, signed with the key bound to the annotator's node identifier and encrypted with the public key bound to the receiver's node identifier. On reception of this updated message the receiving mirror advertises it to the other mirrors and to the addressed node that finally can choose to sign and republish or discard it.

Real-time messages, like chats, are forwarded to and handled by the core solely and responded to with an error message if the core is offline.

Hence, in Safebook, only members with appropriate privileges can access and update the profiles of other members. Entity and data authentication are provided through common signatures and encryption schemes.

Related Work

While a series of studies [3–5] has investigated privacy and security exposures of current OSNs, several other articles suggest solutions to these exposures in various directions combining cryptography with advanced distributed computing techniques.

The approach in NOYB [11] mitigates the existing problems by cryptographic means thanks to the application of substitutions according to secret dictionaries. Public profiles, which still may be stored in a centralized OSN, are thus made useless to anybody lacking access to these dictionaries. Whereas some of the contents of the profiles are protected, this is not the case for relations between users as expressed by contact lists or message exchange.

Yeung *et al*. [12] propose using the existing World Wide Web Friend-Of-A-Friend representation of people and their relations as an OSN. Conventional content and friendship relations are stored in the user's personal space hosted by a server, the choice of which is left at the discretion of the users. While access control for user data can be efficiently ensured based on articulated policies, the system does not protect the identity of users.

Persona [13] offers flexible and fine-grained access control for user data by combining attribute-based encryption with traditional public key cryptography. Users are identified by public keys they exchange out of band while creating OSN links, while data confidentiality and privacy are ensured through encryption. Users have to trust a Firefox extension to interact with Persona and can also create multiple identities.

The related work closest to Safebook is probably PeerSon [14]. PeerSon achieves decentralization thanks to an external P2P system, OpenDHT, and ensures access control through encryption. Whereas it represents a fully distributed OSN, PeerSon leverages on an untrusted P2P system and thus offers weaker privacy protection than Safebook.

Although not designed originally for the purpose of social networking, darknets and related P2P systems [15–17] aim at anonymous communication through hop-by-hop encryption among trusted users, as in Safebook. Unfortunately, such systems suffer from delays that could be prohibitive for an OSN.

Conclusion and Future Work

This article outlines a new approach to the design of online social networks that addresses privacy problems known in existing social network applications. Potential access to the private

data of users, such as profiles and contact lists, and possible misuse of such information by the providers of social networking services is viewed as the highest privacy exposure. In order to ensure users' privacy in the face of such potential exposure, the suggested approach adopts a decentralized architecture relying on cooperation among a number of independent parties that are also the users of the online social network. The second strong point of the suggested approach is capitalizing on the trust relationships that are part of social networks in real life in order to cope with the problem of building trusted and privacy-preserving mechanisms as part of the SNS. The result of these design principles is Safebook, a decentralized and privacy-preserving SNS. Various mechanisms for privacy and security are integrated into Safebook in order to provide data storage and data management functions that preserve privacy, data integrity, and availability. The current design and prototyping of Safebook raise an interesting trade-off between privacy and performance. While increasing the number of hops through trusted links increases privacy, it severely affects lookup and communication delays. A preliminary evaluation of Safebook shows that a realistic compromise between privacy and performance is feasible. Fine tuning of the performance models and simulation results also helps determine critical design parameters such as obfuscation layers and data replication factors. Furthermore, the underpinnings of Safebook can serve as a model to tackle various problems that were left unsolved in the area of secure communications. Thus, a decentralized approach relying on social links can shed new light on hard problems of the past such as anonymous communications, secure routing, or cooperation enforcement in self-organizing systems.

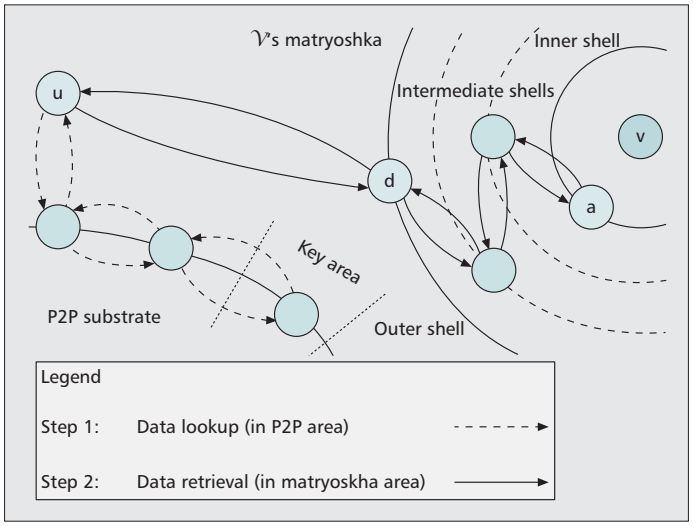

Figure 5. *Data lookup and retrieval.*

References

[1] d. m. boyd and N. B. Ellison, "Social Network Sites: Definition, History, and Scholarship." *J. Comp.-Mediated Commun.*, vol. 13, no. 1, 2008.
[2] Sophos, Aug. 14, 2007; http://www.sophos.com/pressoffice/news/articles/2007/08/facebook.html
[3] Black Hat; http://www.blackhat.com/html/bh-usa-08/bh-usa-08-archive.html
[4] T. N. Jagatic *et al.*, "Social Phishing," *Commun. ACM*, vol. 50, no. 10, 2007, pp. 94–100.
[5] L. Bilge *et al.*, "All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks," *18th Int'l. W3C*, 2009.
[6] M. Arrington, "Modeling The Real Market Value Of Social Networks," 2008; http://www.techcrunch.com/2008/06/23/modeling-the-real-market-value-of-social-networks/
[7] R. Gross and A. Acquisti, "Information Revelation and Privacy in Online Social Networks," *ACM Wksp. Privacy Elect. Soc.*, 2005, pp. 71–80.
[8] d. m. boyd , "Facebook's Privacy Trainwreck," *Convergence: Int'l J. Research into New Media Tech.*, vol. 14, no. 1, 2008, pp. 13–20.
[9] W. X. Zhou *et al.*, "Discrete Hierarchical Organization of Social Group Sizes," *Proc. Royal Soc. B: Bio. Sci.*, vol. 272, no. 1561, 2005, pp. 439–44.
[10] L.-A. Cutillo, R. Molva, and T. Strufe, "Safebook: Feasibility of Transitive Cooperation for Privacy on a Decentralized Social Network," *World Wireless, Mobile, Multimedia Net.*, 2009.
[11] S. Guha, K. Tang, and P. Francis, "NOYB: Privacy in Online Social Networks," *Online Social Net.*, 2008, pp. 49–54.
[12] C. M. A. Yeung *et al.*, "Decentralization: The Future of Online Social Networking," *Future Social Net.*, 2009.
[13] R. Baden *et al.*, "Persona: An Online Social Network with User-Defined Privacy," *ACM SIGCOMM*, Barcelona, Spain, Aug. 2009.
[14] S. Buchegger *et al.*, "PeerSoN: P2P Social Networking," *Social Net. Sys.*, 2009.
[15] M. Rogers and S. Bhatti, "How to Disappear Completely: A Survey of Private Peer-to-Peer Networks," 2007.
[16] I. Clarke *et al.*, "Freenet: A Distributed Anonymous Information Storage and Retrieval System," *Design Issues Anonymity Unobservability*, 2000, pp. 46–66.
[17] K. Bennett and C. Grotho, "GAP — Practical Anonymous Networking," *Privacy Enhancing Tech.*, 2003, pp. 141–60.

Biographies

Leucio Antonio Cutillo (cutillo@eurecom.fr) is a Ph.D. student at EURECOM, Sophia Antipolis, France. He received his M.S. in computer engineering from the Polytechnic of Turin in 2008, his Diplôme d'Ingénieur in communication systems from EURECOM, and his Master of Research in image and geometry for multimedia and life modelization from TELECOM ParisTech in 2007. He actually works in the Networking and Security Department dealing with security and privacy concerns in distributed systems under the supervision of Prof. Refik Molva.

Refik Molva (molva@eurecom.fr) is a professor at EURECOM. His research interests are the design and evaluation of protocols for security and privacy in self-organizing systems. He has been Program Chair or General Chair for security conferences such as ESORICS, RAID, SecureComm, and IEEE ICC, as well as various security workshops. He is an area editor for *Computer Networks Journal*, *Computer Communications Magazine*, and *Pervasive and Mobile Computing Journal*. He worked in the Zurich Research Laboratory of IBM as one of the key designers of the KryptoKnight security system.

Thorsten Strufe (strufe@cs.tu-darmstadt.de) is professor for peer-to-peer networks at Technische Universität Darmstadt, Germany. His research interests lie in the areas of decentralized distributed systems and security, with an emphasis on network analysis and the construction of resilient systems. Recently, he has focused on studying privacy and security in online social networks and possibilities to provide social networking services through P2P technologies. Previously, he took a post as senior researcher at EURECOM and at TU Ilmenau, working on resilient networking technologies.