

A Picture is Worth a Thousand Words: Privacy-aware and Intuitive Relationship Establishment in Online Social Networks

Delphine Christin, Tobias Freudenreich, Matthias Hollick
Secure Mobile Networking Lab
Technische Universität Darmstadt
64293 Darmstadt, Germany
{firstname.lastname}@seemoo.tu-darmstadt.de

ABSTRACT

In current online social networks (*OSNs*) such as Facebook, a new connection request usually only includes the name and photo of the requestor and possibly a list of mutual contacts. Given that the creation of a forged user profile is not too complicated, it is challenging to verify if the contact request is indeed genuine or from a forged account. Accepting a connection invitation from a forged profile might severely compromise the user's privacy, since the attacker gets access to a wealth of personal information and social relationships. In this paper, we present a novel and intuitive paradigm for secure establishment of friendship links in OSNs. We propose an interaction scheme that utilizes cues from snapshots captured using omnipresent smartphones to match and thus verify links in OSNs. We present a proof of concept implementation of our scheme using Android smartphones and embedding the same with Facebook. Finally, we show results from a user study with 25 participants, which demonstrates the intuitive and secure nature of our solution.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection; H.1.2 [Information Systems]: User/Machine Systems—*Human factors*

General Terms

Online Social Networks, Privacy, Context-based Interaction

1. INTRODUCTION

Social networking and the inclusion of user-generated content are among the latest success stories of the Internet. The social networking website *Facebook* nears 600 million users [13], who share diverse information including their attendance at upcoming events, up-to-date information about their status, pictures, and more. However, these contents may reveal sensitive information about the users, e.g., their favorite leisure activities, travel plans, or opinions. Their disclosure may lead to severe consequences ranging from dismissals [12] to safety threats [8]. On existing platforms, these data are protected against disclosure by access control mechanisms, generally realized in the form of contact lists. Other members can be added to these lists by sending them *contact requests*. The members can control what kind of content is disclosed to which of their contact lists using privacy preferences. Once the request is accepted, both members

can mutually access information posted on each other's profiles depending on their respective privacy preferences and a virtual relationship is established.

It has been demonstrated that members of online communities tend to accept contact requests obliviously [3]. This behavior may be driven by the simple curiosity of the members and their eagerness to augment their list of friends. Moreover, a user can base his decision to accept a contact request only on limited information, making the entire process susceptible to attacks. As depicted in Fig. 1, a Facebook friend request contains only the user name, the profile picture of the requester and, possibly, a list of mutual friends. Additional information such as date of birth, past and current affiliations or active 'networks' might also be found in the requester's profile. However, forgers may easily infer these data using multiple information sources (e.g., personal website or profiles in diverse online communities) to create fake accounts [3]. Based on this limited amount of information available in the friend request, the recipient may be unable to distinguish real accounts from fake ones. As a result, users may be tricked to authorize attackers to access their data, endangering their privacy and possibly incurring severe social and safety consequences.

We counter the risk of accepting friend requests from forged accounts by a mechanism that reliably verifies the identity of members who want to establish a virtual relationship. Instead of relying solely on virtual interaction, our approach creates unforgeable identity proofs based on intuitive physical interactions between the members. As the majority of new mobile phones come with embedded cameras, we base our interaction on this sensing modality. Two users located in physical proximity take a picture of a common subject. The subject could be a self-portrait of the two users as illustrated in Fig. 2 or any unique mutually agreeable object, which the users can recall at a later stage. The pictures complemented by spatiotemporal information are then uploaded by each user individually. Before establishing virtual links, these pictures are compared to verify the existence of a prior physical interaction between them and consequently, verify



Figure 1: Example of friend request in Facebook



Figure 2: Example of common subject

their respective identities. In the scope of this paper, we assume that the users establish friendship links with people whom they have met in person (e.g., at conferences, during holidays, or at social events), which is a fairly widespread scenario as shown in [6]. However, we will investigate expanding our scheme to include friend requests from virtual friends with whom no physical encounter has taken place in our future work.

Our contributions are as follows. In Section 2, we analyze the design space to identify interactions able to provide sufficient proofs of users’ identities and we discuss our design decisions. We describe our concept in Section 3 and analyze its resilience against identified attacks in Section 4. We detail the implementation of our proof-of-concept using Android phones and integrated in Facebook in Section 5. In Section 6, we evaluate the usability of our concept by performing a user study and we discuss related work in Section 7. We conclude this paper in Section 8.

2. SYSTEM MODEL AND DESIGN

We assume that two members of an online community, Bob and Carlos, who meet in person, want to establish a virtual relationship to share pictures. These pictures contain sensitive, personal information and shall remain private to both. Hence, to protect his privacy, Bob must ensure that the account presented by the system as Carlos’ account is not forged, before authorizing the access to his pictures. Reciprocally, Carlos must verify that Bob’s account is trustworthy. From this scenario, we can derive the need for a mechanism to verify the real identity of the users and to establish secure and trusted friendship links. To identify suitable interaction schemes supporting the verification of the user’s identity, we analyze the design space and discuss our design decisions.

2.1 Design Space Analysis

The design of the solution is guided by the following criteria: Reliability of the provided proofs of identity, privacy guarantee, and usability. Based on these criteria, we analyze the feasibility of the following selected classes of solutions: Cryptography-based solutions, text-based puzzles, visual contact as proof, and proofs based on physical proximity. The results of the analysis are summarized in Table 1.

Standard authentication protocols, e.g Kerberos [11], support authentication and identity verification, but they mainly provide guarantees at machine-level and not user-level and may appear obscure and complex for inexpert users. As the solution to design must be easy to comprehend, we discard interactions only based on cryptographic mechanisms.

Short messages containing puzzles (questions/answer) to verify the identity of the users could be exchanged in a text-based solution integrated in the existing *friend requests*, as

Table 1: Summary of the design space analysis

	Reliability	Privacy	Usability
Cryptography-based solutions	+	++	---
Text-based puzzles	-	++	-
Visual contacts	++	--	+
Physical proximity	++	++	++

the users can write personal messages. However, the puzzles must be carefully selected in order to prevent attackers from inferring the answers from publicly available sources of information. This scheme might be perceived as cumbersome by the users, as they need to adapt the questions to each person depending on the nature of their relationship.

Further, the establishment of each friendship link necessitates the resolution of two puzzles to ensure bidirectional identity verification. The security of this alternative is therefore determined by the selection of hard-to-solve puzzles, which negatively impacts on the usability and reliability if puzzles cannot be solved. As a result, users might recourse to trivial puzzles, rendering the mechanism inefficient.

Instead of using text-based proofs, visual contacts between the members via web cameras may be envisaged. Although this alternative promises to provide reliable evidence of the members’ identities, it requires both users to be online at the time of verification. Also, the immediate privacy may be endangered in case of a forger that may be able to temporarily establish a connection or to access the video streams. A preliminary mechanism may thus be introduced to provide sufficient identity proofs before establishing the connection without any privacy threat. However, this leads to the problem of the identity verification once again.

Alternatively to a purely remote, virtual exchange of proofs of identity, physical interaction of members may also be exploited. Such interaction requires physical proximity of the users, which, on a first glance, makes it less flexible compared with remote mechanisms. However, this alternative takes advantage of the fact that online social networks typically mirror real social links that are commonly established by direct interaction. Also, the scheme intrinsically guarantees the identity of both members, without endangering their privacy. Exploiting user proximity may hence support user-friendly interactions involving the users physically in the establishment of the relationships and increase their consciousness about disclosure issues, as detailed in [5].

2.2 Design Decisions

Based on the results of the above analysis, we have chosen to exploit physical proximity for securely establishing virtual relationships. In particular, it promises reliable proofs of identity without threatening the privacy of the users, while offering interesting interaction schemes that allow to improve user awareness and to increase usability. The physical proximity scheme also neutralizes remote attackers, since the physical presence is required to establish the relationships. Furthermore, despite the virtual nature of OSNs, most of the social links are established after offline meetings [6]. Our scheme consequently supports the establishment of most relationships. To not exclude people who never physically met,

our scheme may however be used in combination with existing mechanisms, such as friend requests.

As intermediary in the establishment process, we propose to use smartphones because of their ubiquity and their embedded sensors, which provide a large design space for physical interactions. Instead of using existing schemes, such as Bluetooth pairing or scanning the surrounding GSM base stations and Wi-Fi access points, we propose to utilize simple, yet intuitive interactions, where the users are physically involved to increase their awareness about the authentication process [4]. Preferably, the user should be familiar with the actions required, to increase the acceptance of the scheme. Among the possible sensor-based interactions, we selected camera-based interactions, as people are used to spontaneously take pictures of friends and relatives with their smartphones. Jointly taken pictures become the representation of the newly established relationships.

3. CONCEPT AND PROTOCOL

Considering the above discussions, we design an interaction mechanism and the corresponding protocol to establish secure and trusted relationships within online communities.

3.1 Overview

Assume that Bob and Carlos want to establish a friendship link. They initiate this relationship by taking a picture of a common subject with their respective devices, e.g., a snapshot of themselves or a common object. The pictures are stamped with time and location information. They are successively uploaded by Bob and Carlos to the server of the online community and serve as key for the pairing. The server searches its image database for matching images in given time intervals and in given locations. Once a match is detected, both Bob and Carlos are presented with the search results and the proposed pairing. They can then proceed to establish the relationship by confirming the pairing.

3.2 Protocol

Without loss of generality, we assume that Bob initiates the first step of the protocol illustrated in Fig. 3. Bob takes a picture P_B of Carlos and himself with his smartphone (Fig. 2). Then, Carlos repeats the same step with his own device and obtains the picture P_C including Bob and himself. Metadata including author's ID , timestamp T and location information L are generated and attached to the pictures P_B and P_C automatically. The location information does not require being the exact coordinates of the users, but must provide insights about their proximity by scanning, e.g., cell-IDs, and Wi-Fi access points. The smartphone initiates the communication with the server; we assume a secure channel is being established using the SSL/TLS protocol. The following communication between the server and the end devices are secured by the SSL/TLS protocol. Once the personal end devices are authenticated, the obtained pictures $P_B\{T_B, L_B\}$ and $P_C\{T_C, L_C\}$ are successively uploaded on the server. When Carlos' picture is uploaded, a search process is triggered to retrieve the picture formerly uploaded by Bob. To reduce the computation overhead and execution time, the search first exploits the metadata T and L stored with the pictures. Pictures taken within a predefined time interval and physical range (or sharing the same

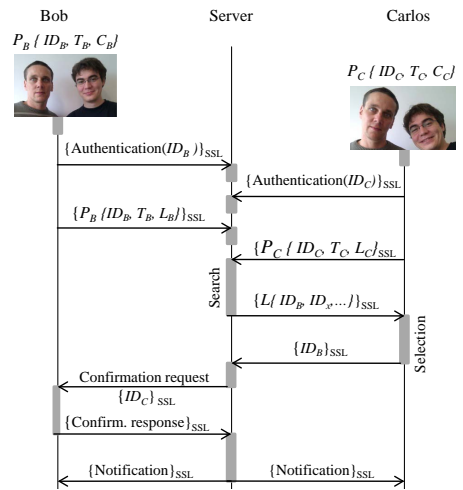


Figure 3: Protocol for establishing secure and trusted relationships

location-context) around the timestamp T_C and location L_C of Carlos' picture are selected. The time interval is selected according to an estimation of the time necessary for both users to perform the different tasks, while the physical range depends on the expected precision of the selected localization scheme. The pictures resulting from this first search iteration are then compared with Carlos' picture in order to recognize Bob's picture among them. Pictures showing the greatest amount of similarities with Carlos' picture are selected. The search process is achieved when the list of the resulting IDs $L\{ID_1, ID_2, \dots\}$ is securely transmitted and displayed on Carlos' personal end device. Scrolling the list, Carlos recognizes and selects Bob's identity ID_B among the proposals. The selection of Carlos is then transmitted securely and triggers a confirmation request sent to Bob. This confirmation request includes Carlos' identity ID_C . Bob verifies the data included in the confirmation request and his confirmation terminates the protocol by mapping both identities together and establishing a trusted relationship.

Without modifying the functionality of the above protocol, Bob and Carlos could also take arbitrary pictures involving the same elements/features of their surroundings. We selected self-portraits because they are a representation of the social context of the relationship, which is virtually established, and may increase the users' awareness.

4. SECURITY ANALYSIS

In our scenario, the main objective of the attacker is to gain access to data posted by the members on the sharing platforms. This objective may be motivated by multiple reasons including, e.g., simple curiosity, or willingness to harm the users. To reach this objective, our attacker may create forged accounts with Bob's and Carlos' identities and tempt to intervene in the process of establishing the relationship in one or both directions. To determine how secure our approach is, we analyze the possible consequences of the different attacks on the establishment of relationships.

The SSL protocol protects the data against manipulation during the transmission between the smartphones and the server. We further assume that state-of-the-art security




Figure 4: Examples of pictures (A, B, C, D) taken with our prototype

mechanisms protect the data stored on the server from access and manipulation by potential attackers. Attackers are therefore unable to modify pictures stored on the server in order to establish relationships with other users.

Attackers may conduct the equivalent of replay attacks by taking similar pictures as Bob and Carlos without being noticed. In this case, the account of the attacker appears in the list of IDs proposed to Carlos. If the attacker's account differs from Bob's account, Carlos notices the difference(s) and does not select the attacker's account. If the attacker created a forged account, two accounts with Bob's identity are contained in the list and Carlos can detect the multiple accounts. He can warn Bob of the existence of a second account and ask for his assistance to identify his actual account. If Bob is able to distinguish his account from the forged ones, the relationship between Bob and Carlos can be completed. Otherwise Carlos can abort the protocol to prevent the attacker from establishing of a relationship with him and therefore, protect his privacy.

Attackers may create a forged account of Carlos' account and take a similar picture as Bob and Carlos. When Bob uploads his picture, the attackers may jam the channel and then, upload their own picture from the forged account. Thus, the forged account appears in the list of IDs proposed to Carlos, who may not distinguish it from the actual one and select it. Yet, as Bob does not receive any confirmation request and notification, he notices that the procedure is erroneous and can inquire with Carlos. By introducing a delay between the establishment of the relationship and the activation of data access authorization, and an optional mechanism to revoke the established relationship, Carlos' privacy may not be endangered.

To summarize, our approach is resilient against the above attacks and enhances the privacy of members of OSNs, as only people met in person are authorized to access their information. Members can easily verify the identity of the contact requester and detect forged accounts, preventing them from accepting requests coming from forgers who aim to abuse their personal information and violate their privacy.

5. IMPLEMENTATION

We implemented our approach on Nexus One smartphones running the Android 2.1 operating system¹. Illustrated in Fig. 5, the user interface is designed to be simple, user-friendly and easy to comprehend by untrained users. The users log in with their Facebook account and a new entry (including the user's ID, last name, first name, and picture from his Facebook profile) is created at the first login in the MySQL database running on our server. The

¹A video demonstrating our prototype can be found at: <http://www.seemoo.tu-darmstadt.de/iwssi>

Table 2: Similarity results between each pair of pictures

Pictures	A	B	C	D
A	0	0,144	0,156	0,181
B	0,144	0	0,149	0,119
C	0,156	0,149	0	0,164
D	0,181	0,119	0,164	0

database also stores the uploaded pictures serving in establishing the friendship links. The search process is initiated by the server, which consults the database to retrieve the second element of the pair of pictures. The pictures are first selected according to their timestamps and their location information, which must be included in the predetermined time interval and location range. Only after this first processing, the content of the picture is analyzed. The image processing in our prototype is based on the *NaiveSimilarityFinder* application [14] comparing the pictures using spectral and spatial features. Higher precision algorithms using more complex operations or supporting face recognition may however be envisaged to scale to large sets of images, but their design remains out of scope of this paper. Table 2 presents the results of the comparison of each pair of pictures depicted in Fig. 4. To obtain these values, the pictures are first normalized and equally distributed squares are extracted from the pictures. For each square, the average of each RGB value is calculated. The similarity between the two pictures is then calculated by summing the Euclidean distance of each pair of squares. The pairs of pictures showing the lowest values present therefore the highest similarities, as e.g., pictures B and D. Despite the simplicity of the image processing algorithm, the pair matching was always successful during our experiments. Once the pictures are paired, the server a confirmation request including the personal information about the user.



Figure 5: Design of the user interfaces

6. EVALUATION OF USER STUDY

We performed a user study to compare our approach with the existing *friend requests* used in Facebook and to evaluate how potential users perceive the usability of our prototype. We submitted a questionnaire to 25 persons (7 female and 18

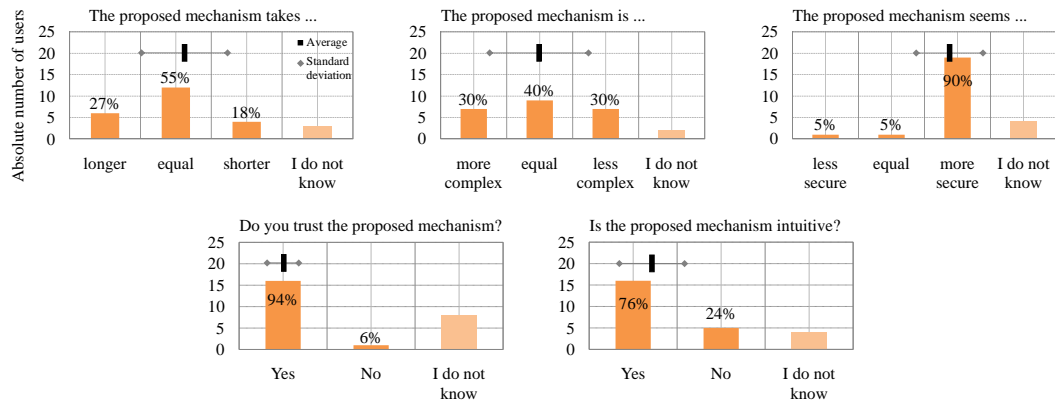


Figure 6: Evaluation of the user study showing the number of users per answer as well as the average answer and its standard deviation

male) at the TU Darmstadt and conducted a practical test. The user study was divided into three parts: (1) personal information and background on OSNs, (2) a practical test to add a new friend in Facebook, and (3) a practical test to establish a relationship using our prototype. The practical tests were conducted using a supervisor-based approach to guide the participants where necessary.

6.1 Personal Information

Our participants are primarily researchers (80%) working in 12 different research labs, mainly focusing on security. The remaining participants are undergraduate students (12%) and employees (8%). The users were between 20 and 37 years old, with an average age of 29. Due to their backgrounds, the participants are more familiar with OSNs than the average population. However, at the same time they are more suspicious about security and trust aspects. Moreover, 20% of our participants did not have any experience with OSNs. The repartition of memberships and the average corresponding amount of contacts for the remaining 80% is the following: 64% of users have an account at Facebook (with 74 contacts in average), 28% at Xing (78 contacts), 28% at StudiVZ (38 contacts), 16% at LinkedIn (21 contacts), and 12% at Wer-kennt-wen (65 contacts). StudiVZ and Wer-kennt-wen are German OSNs.

6.2 Adding a new Friend in Facebook

To investigate the ambiguity of the current Facebook member selection process, we created a forged Facebook account under the name of *Delphine Christin* and we completed it with information publicly available on her personal webpage including a profile picture and educational information. We also asked two coauthors of papers and her advisor to accept *friend requests* initiated by the fake account to create a realistic shortlist of friends. For the duration of the study, we disabled her primary Facebook account.

We asked the participants to search for *Delphine Christin* and add her as a friend in Facebook. Then, they indicated how many accounts with her name were displayed in the search results. 24% answered with 4 and 5 accounts respectively and 8% indicated 2 accounts. Only 44% of the participants gave the correct answer of 3 (precisely) matching accounts. The search results actually included 5 profiles with the user names: *Delphine Christin* (three times), *Christine Delphine*, *Christine Volpi*. We observe that the

majority of the participants did not notice the differences in the results, and may thus be easily influenced by attackers creating forged accounts using slightly modified user names.

We asked the participants to indicate how they identified the actual account of *Delphine Christin*. Multi-selection of answers allowed, among the proposed answers, 96% selected the profile picture, 28% the network and 28% the list of friends. We observe that the profile picture is the foremost selection criterion. However, finding a profile picture for a forged account is relatively easy and our group of users may therefore become victim of requests coming from forged accounts using a real profile picture of the users.

We concluded this task by inquiring if the participants were sure to have selected the correct account. 76% of the participants claimed to be sure, while 20% claimed the contrary. The remaining 4% were doubtful. We then explained them that the account they selected was a forged account created especially for the study purpose.

6.3 Adding a new friend using our Prototype

To evaluate the usability of our prototype, the participants next tested our proposed method to establish a virtual relationship with *Delphine Christin*, who acted as friend-to-be-established. In a supervised setting, each participant performed the respective steps required by the protocol with one of the smartphones, as described in details in Section 3.2.

The participants compared our approach with the default mechanism used in Facebook to search and add new friends, according to the following criteria: Duration of the process, complexity of the process, perceived security and trust. The repartition of the answers is detailed in Fig. 6. In summary, the participants found that our approach:

- Takes an equal amount of time
- Is perceived to be trustworthy and more secure
- Is perceived to be intuitive and not more complex

Asked, if they would adopt our approach to add new friends in online communities, 68% of the participants would prefer our approach, while 16% would prefer using the existing mechanism. The remaining of the participants (16%) were

indecisive. We finally offered the opportunity to the participants to add personal comments about our prototype. The majority of comments was positive including “I like the idea of interacting physically”, “I like it”, and “It’s cool”.

7. RELATED WORK

We propose a novel mechanism to establish trusted relationships between members of online communities. In the current state-of-the-art, these relationships are mainly initiated by friendship requests exchanged within the social networks, as e.g., the *friend requests* in Facebook. However, the latter mechanism does not provide any guarantee about the real identity of the account owner and the members may be easily tricked into befriending owners of forged accounts. In comparison, our approach eliminates this risk, as the utilization of physical proximity in combination with its connected context information provides a reliable and reciprocal verification of the member’s identities.

Our approach shares features with the hardware-based exchange of digital business cards, which often relies on physical proximity. Such exchange carries information that might be used to establish virtual relationships. Specialized devices, such as Poken [1], exchange personal information about their carriers when these enter in physical contact, while Bump [2] matches the accelerometer data of two smartphones bumped together to initiate the data exchange. In comparison, our approach does not require investment in special purpose hardware, or specific interactions, respectively. Moreover, these approaches only use physical interactions to initiate the information exchange protocol without providing any verification mechanism. To provide a similar level of security, the members would need to verify manually that the transmitted information was not forged, which may rapidly appear as burdensome to the users. To summarize, both mechanisms exploit physical proximity to exchange personal information, but the notion of secure and trusted establishment of relationship is not included.

Furthermore, our approach shares similarities with mechanisms to pair devices, such as presented in [7, 9, 10, 15]. The core idea to use sensor-based information is common, but we explore a novel dimension, as the nature of the paired subjects differs. We pair members of online communities via physical interactions, whereas device pairing focuses on the devices themselves. Based on simple and intuitive interactions already adopted by the users, our approach proposes a novel angle to combat forgers in online communities and establish secure relationships between their members.

8. CONCLUSIONS

Users of OSNs like to share personal details such as snapshots with their acquaintances and friends. Yet the basis for such exchange, the identification of friends, can be hit-and-miss in the presence of forged profiles. We have proposed an interaction scheme for reliably establishing friendship links in social networks. Departing from the common goal of utmost simplicity in user interface design, we designed the interaction to be natural yet explicit, thus increasing user awareness. In particular, we propose to leverage cues from snapshots taken in physical proximity, to match and verify links in OSNs. A secure protocol facilitates the interaction between users’ and the OSN. We have implemented a proto-

type system and have performed a user study. The obtained results are encouraging; users consider our approach to be more secure and trustworthy, while the perceived complexity stays the same with the existing schemes.

9. ACKNOWLEDGMENTS

Our thanks go to A. Reinhardt and S. Kanhere, as well as A. Oberle and the participants of the user study. This work was supported by CASED (www.cased.de).

10. REFERENCES

- [1] Poken. Online: <http://www.poken.com>.
- [2] The Bump Application. Online: <http://bu.mp>.
- [3] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks. In *Proc. of the 18th International World Wide Web Conference*, 2009.
- [4] M. K. Chong and H. Gellersen. Classification of Spontaneous Device Association from a Usability Perspective. In *Proc. of the 2nd International Workshop on Security and Privacy in Spontaneous Interaction and Mobile Device Use*, 2010.
- [5] D. Christin. Impenetrable Obscurity vs. Informed Decisions: Privacy Solutions for Participatory Sensing. In *Proc. of the 8th IEEE International Conference on Pervasive Computing and Communications*, 2010.
- [6] N. B. Ellison, C. Steinfield, and C. Lampe. The Benefits of Facebook “Friends”: Social Capital and College Students’ Use of Online Social Network Sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.
- [7] M. Goodrich, M. Sirivianos, J. Solis, G. Tsudik, and E. Uzun. Loud and Clear: Human-verifiable Authentication based on Audio. In *Proc. of the IEEE International Conference on Distributed Computing Systems*, 2006.
- [8] R. Gross and A. Acquisti. Information Revelation and Privacy in Online Social Networks. In *Proc. of the ACM Workshop on Privacy in the Electronic Society*, 2005.
- [9] R. Mayrhofer and H. Gellersen. Shake well before Use: Authentication based on Accelerometer Data. In *Proc. of the 5th International Conference on Pervasive Computing*, 2007.
- [10] J. M. McCune, A. Perrig, and M. K. Reiter. Seeing-Is-Believing: Using Camera Phones for Human-Verifiable Authentication. In *Proc. of the IEEE Symposium on Security and Privacy*, 2005.
- [11] B. Neuman and T. Ts’o. Kerberos: An Authentication Service for Computer Networks. *IEEE Communications Magazine*, 32(9):33–38, 2002.
- [12] B. Quinn. Virgin Sacks 13 over Facebook ‘Chav’ Remarks. Online: <http://www.guardian.co.uk>.
- [13] L. Rao. More Evidence That Facebook Is Nearing 600 Million Users. Online: <http://techcrunch.com>.
- [14] R. Santos. Java Image Processing Cookbook. Online: <http://www.lac.inpe.br/JIPCookbook/index.jsp>.
- [15] N. Saxena, J. Ekberg, K. Kostianen, and N. Asokan. Secure Device Pairing based on a Visual Channel. In *Proc. of the IEEE Symposium on Security and Privacy*, 2006.