

RelationGram: Tie-Strength Visualization for User-Controlled Online Identity Authentication

Tiffany Hyun-Jin Kim¹, Akira Yamada², Virgil Gligor¹, Jason Hong¹, and Adrian Perrig¹

¹ Carnegie Mellon University {hyunjin,gligor,jasonh,perrig}@cmu.edu
² KDDI R&D Laboratories Inc. yamada.akira@kddilabs.jp

Abstract. We consider the specific problem of how users can securely authenticate online identities (e.g., associate a Facebook ID with its owner). Based on prior social science research demonstrating that the social tie strength is a useful indicator of trust in many real-world relationships, we explore how tie strength can be visualized using well-defined and measurable parameters. We then apply the visualization in the context of online friend invitations and propose a protocol for secure online identity authentication. We also present an implementation on a popular online social network (i.e., Facebook). We find that tie strength visualization is a useful primitive for online identity authentication.

1 Introduction

Many real-world social interactions are based on various types of trust relations derived from strong social ties [4, 11–13]. As social interactions migrate from the physical to the online world, current systems do not provide many cues upon which users can base their identity authentication. For example, consider Facebook: how can a user be certain that a Facebook invitation is really from the claimed individual? As anyone can trivially set up a Facebook page with someone else’s photo, Facebook provides almost no help in ensuring correspondence between the online and physical identity [1, 2, 5], even fooling security-conscious individuals [14]. Furthermore, Irani et al. [7] recently propose reverse social engineering attacks in Online Social Networks (OSNs), where the attacker sets up fake accounts and lets the victim discover and contact the fake account.

Although at first glance Public Key Infrastructures (PKI) and Pretty Good Privacy (PGP) appear to enable users to link an online identity to an individual, these approaches have significant shortcomings. Despite the long existence of Certification Authorities (CAs), few users have personal certificates, which are cumbersome to obtain. Moreover, CAs have recently suffered from several attacks [3]. Unlike PKIs, PGP is a distributed approach based on the notion of “Web of Trust” enabling identity certification. However, PGP’s chains of trust are often unwieldy and offer limited security.

Personal recommendation systems may appear to address these issues, where a user would digitally sign a statement such as: “I trust that public key K_A really belongs to Alice, and I trust Alice to correctly validate other users.” In the context of PGP, users could specify how much they trust others to assist

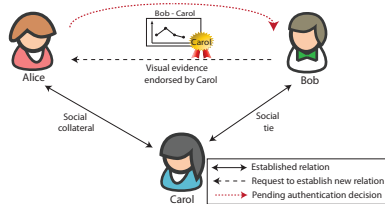


Fig. 1. Our approach for online identity authentication. Alice confirms Bob’s invitation based on a *RelationGram* – a visual evidence of Bob and Carol’s tie strength.

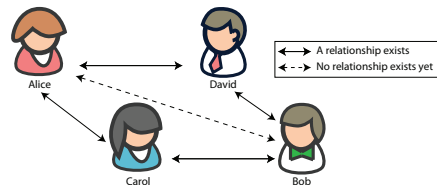


Fig. 2. An example of a trust graph. Bob wants to be Alice’s online friend, and Carol and David are their mutual friends.

validation of a chain of trust. Unfortunately, this approach suffers from the *distrust revelation problem*, defined by Kim et al. [8], where a polite or conflict-averse user does not want to publicly admit distrusting another individual, and thus specifies the untrusted user as trusted. Avoiding the distrust revelation problem is a core challenge we aim to address.

According to a social collateral model, users can accept online friend invitations from unknown senders based on explicitly endorsed recommendations made by social relations [9]. We extend the social collateral to provide different degrees of accountability for accepting friend invitations of unknown senders as follows: invitee Alice holds recommender Carol accountable for her actions such that Carol does not deceive Alice by creating or certifying bogus online identities. If Carol signs a bogus identity, the signature provides a non-repudiable statement, which may result in loss of the social collateral (e.g., loss of social relationship with Alice, loss of self-esteem, and feeling of guilt) for Carol. Preliminary evidence shows that social accountability can have a stronger deterrence effect than formal/legal punishment [6].

In this paper, we study how to enable users to authenticate OSN invitations to ensure that an invitation from an online individual is indeed tied to the correct individual. Our key idea is to derive tie strength between inviters and their mutual friends to represent real-world physical interactions, and provide it as evidence to empower users to authenticate online identities. More specifically, prior research indicates that in practice, tie strength can be represented using simple proxies such as frequency, reciprocity, and recency of communication, which we believe can be feasibly acquired by smartphones using call logs, emails, OSN comments, etc. Based on the simple proxies, we propose a *RelationGram* – a visualization of tie strength between an inviter and the invitee’s friend(s) by which the invitee can easily understand the degree to which her friends know the inviter before she makes her own context-dependent authentication decisions.

As shown in Fig. 1, Alice can authenticate Bob’s online identity using a *RelationGram* as follows. First, Alice personally knows Carol, and some level of social accountability exists between them. Second, Bob has sent a friend invitation to Alice and claims that Carol is a mutual friend. Before accepting the invitation, Alice wants to validate that Carol has a strong tie with Bob. Thanks to the *RelationGram*, visualizing tie strength between Bob and Carol, along with

Carol’s digital signature of the visualization and Bob’s public key, Alice gains evidence and endorsement implying the strong tie. Hence, the combination of Carol’s social accountability to Alice and the strong tie between Bob and Carol results in Alice authenticating Bob’s online identity.

2 Problem Definition, Adversary Model, Assumptions

In this paper, we explore how to provide the evidence of social distance between users through a simple visualization that is endorsed in the form of a digital signature. Our goal is to help users correctly authenticate online identities using endorsed visualizations of social tie strength.

A challenge then is to accurately capture aspects of tie strength among OSN users and visually represent it to convey social proximity to other OSN users.

Relevance with respect to social parameters. Every individual is unique and has different criteria in judging social distance. Hence, it is important to carefully select *relevant* parameters which accurately convey tie strength.

Robustness. Tie strength represented using social parameters must be robust against active attackers who attempt to claim close social proximity to others. Also, tie strength must be difficult to inflate due to social pressure (i.e., *distrust revelation problem* [8]), because users do not want to publicly admit that they do not trust another user.

Usability. It is crucial that OSN users can correctly interpret the visualization of relevant social parameters and easily understand social tie strength.

We consider an adversary whose goal is to manipulate social parameters for measuring tie strength such that he can claim to have a strong tie to a victim’s friend. When the adversary deceives the victim who accepts the friend invitation, he can successfully gather sensitive personal information of the victim and possibly her friends.

We assume that trusted friends of a user do not misbehave due to their social accountability. Furthermore, we consider an attacker who compromises a user’s account to be orthogonal to the issues we address in this paper.

3 Interpersonal Tie Strength Visualization

Bob wants to be Alice’s online friend. Bob already knows Alice’s friends who have social collateral with her. Rather than verifying any evidence provided by Bob (since Alice has not met Bob in person), we want to help Alice make a decision based on the evidence provided by her mutual friends who are socially accountable to her. In Fig. 2, these mutual friends would be Carol and David.

Using this scenario, Fig. 3 depicts a RelationGram, a visualization from which Alice can deduce appropriate social relationships between Bob and her friends, Carol and David. Below are the 7 parameters that the RelationGram visualizes:

Interaction Frequency. This parameter is displayed on the y-axis without a detailed scale, and we assume multiple types of communication (e.g., phone calls, emails, OSN comments). Note that we also display a normalized average line of the communication frequency (i.e., the normalized average frequency of

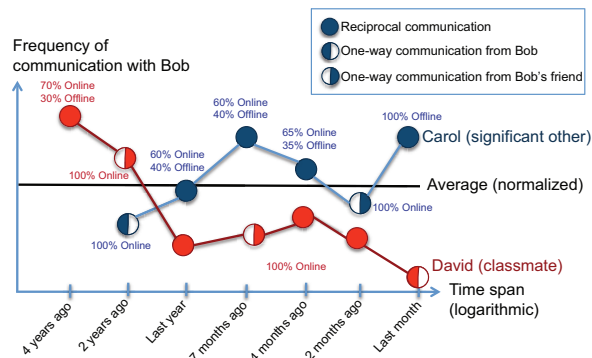


Fig. 3. RelationGram: a tie strength visualization from a trust graph.

communication between Carol and all her other friends). The same line also represents the average frequency of communication between David and all his other friends. Based on this average frequency, Alice can distinguish how frequently her friends interact with Bob as compared to their other friends, such that Alice can evaluate their tie strengths in a fair and unbiased manner.

Communication Reciprocity. This parameter is shown using the amount of shade on each plotted dot. For example, a fully-colored dot implies that the interaction is reciprocal, and a half-colored dot implies that the interaction is one-way where the originator is based on the side of the color as shown in Fig. 3.

Recency of communication. The x-axis to the right represents a more recent time span than the left side. Fig. 3 confirms that Bob has communicated with Carol more recently (last month) than with David (2 months ago).

The existence of more than one mutual friend. This parameter is depicted by the number of graphs in the same plot. In this case, Alice can infer that two of her friends are also friends with Bob.

Length of relationship. The x-axis represents a logarithmic timeline which captures the length of Bob’s social relationships with Carol and David.

Relationship status. Fig. 3 display the relationship status (e.g., classmates, family, etc.) between the inviter and the mutual friends, and this relationship label is assigned by the mutual friend. As a result, the inviter has no control over the relationship label.

Communication type. The RelationGram labels the composition of *online* and *offline* communication frequencies, where online communication may include emails, OSN conversations, etc., and offline communication may include phone conversations, physical interactions, etc. The “100% online” labels will help indicate individuals who have only established a relationship over purely online means. As a result, Alice may be able to infer, with higher confidence, that Carol’s graph indicates a strong tie with Bob with both online and offline interactions in Fig. 3.

We entrust full disclosure control to users such that the users themselves can decide to either reveal or protect their own graphs depicting their interactions with a particular set of friends. For extensive explanation of visualized parameters, including how they are computed, as well as security and privacy analysis of the RelationGram, please refer to our technical report for details [10].

4 Authenticating Online Friend Inviters

We introduce Indirect Friend Authentication (IFA) that can be used to authenticate an online friend inviter who is a friend of the invitee.

For this application context, we assume that people use smartphones for communication, as a greater number of smartphones are being sold.³ We further assume that every user can use a smartphone to generate a public-private key pair, measure the parameters to represent tie strength, and automatically communicate with cloud application providers. Cloud application providers may be similar to Google which provides a backup service for contact information on phones, and we trust them for the availability of user information.

4.1 Indirect Friend Authentication

Alice receives an online invitation from someone named Bob, and this invitation indicates that they have two mutual friends: Carol and David.

Our Indirect Friend Authentication (IFA) protocol helps Alice authenticate Bob by leveraging two mutual friends. In a nutshell, the IFA protocol presents evidence that reflects the interpersonal tie strength between Alice’s friend(s) and Bob in a RelationGram as explained in Section 3. Based on the visual evidence and the strength of social ties with her friends, Alice can exercise sound judgment when accepting Bob’s invitation.

Evidence Generation. Bob and Carol mutually agree to disclose the information that reflects their social tie, and so do Bob and David. There are different ways of gathering information to represent these parameters. For example, Bob’s and David’s phones can automatically detect and record the duration of a meeting, the call history between them, exchanged SMS text messages, Facebook posts, etc. Furthermore, OSNs can analyze information about their online message exchanging behavior, photos in which both are tagged together, etc. Note that these are the optional features that users opt-in for usage and users with privacy concerns may decline to use our protocols. When all the information representing tie strength is properly gathered on David’s phone, it would sign the visual graph of their tie strength from David’s perspective, sign Bob’s public key, and hand it over to Bob. (With David’s permission, this process is transparent to David.) Carol does the same for Bob. Thanks to Carol and David’s release of the visual graphs, Bob has the evidence implying his social relations with Alice’s friends and he inserts the graphs into the invitation.

Evidence Verification. When Alice receives the invitation from Bob, she has an option of seeing the RelationGram to determine the tie strength between Bob and her friends. Alice first verifies Carol’s signed graph and Bob’s public key using Carol’s public key that she can retrieve from her own phone or from the cloud application provider. She also verifies David’s graph in the same manner. When Alice successfully verifies that the graphs are generated by her real, trusted, and accountable friends (e.g., by verifying their digital signatures), she

³ As of Q3 of 2012, 56% of all mobile consumers in the U.S. own smartphones (http://blog.nielsen.com/nielsenwire/online_mobile/nielsen-tops-of-2012-digital/).

may decide to accept Bob’s invitation based on his strength of social ties with her friends. However, it is possible that the authentication fails or the graphs do not convey strong ties, possibly due to some abnormal interaction conditions (e.g., David could have recently relocated, reducing his interaction frequency with Bob and limiting the communication medium to Facebook only). We emphasize that the visualization is one type of available evidence for users to make better tie strength evaluation, and the IFA protocol recommends that Alice gathers other evidence before accepting Bob’s invitation.

For the security analysis of IFA, please refer to our technical report [10].

5 Implementation

We have implemented the IFA protocol in the context of Facebook friend invitations. Fig. 4 shows the architecture and the flow of our protocol to validate Facebook friend invitations.

“Do I Really Know You?” is an integrated Facebook web application such that (1) users can access their friends’ invitations and present visualizations in a seamless manner and (2) the visualizations can be displayed on any smartphone with a web browser.

We have implemented our application using three types of APIs that Facebook provides: GraphAPI, OldREST API, and Facebook Query Language (FQL). This application seeks permission from users to retrieve posts and comments from their walls.

When a user invokes “Do I Really Know You?”, it gets a token which enables this application to access the Facebook database on behalf of the user. The application then queries the database according to the user’s policy. First, the application retrieves a list of pending friend invitations (via Facebook’s *notifications.get* API). With at least one invitation, the application queries information

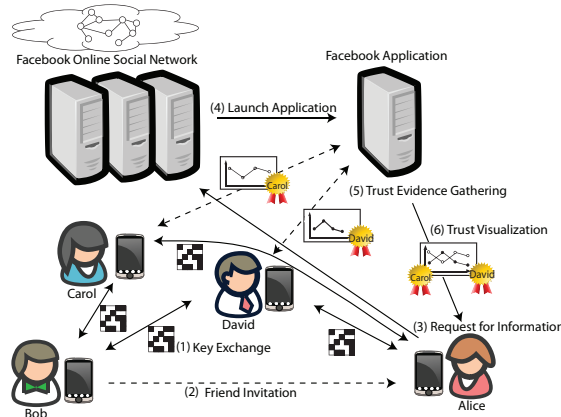


Fig. 4. Architecture of IFA protocol on Facebook. This diagram illustrates the process of IFA on Facebook as follows: (1) Alice exchanges keys with her friends Carol and David. (2) Bob wants to be Alice’s friend. (3) Alice requests a RelationGram. (4) Facebook launches our application. (5) Our application obtains RelationGrams from Alice’s friends after endorsement. (6) The application returns the endorsed RelationGram to Alice. (7) Alice verifies the tie strength between Bob and her friends (Carol and David) using the endorsed RelationGram.

about the inviter and the mutual friends (via *friends.getMutualFriends*). Then, the application retrieves a stream of wall information (via *stream.get* query with *limit=0* as a parameter).

When there are more than three mutual friends, this application prompts the user to select the “best” friends with whom he wants to infer the inviter’s tie strengths. Based on the comments from the selected mutual friend’s Facebook wall, this application computes the number of comments between each mutual friend and the inviter, and plots the interaction frequency for a RelationGram on the web browser.

We conducted an online user study with 93 participants to verify how much OSN users understand tie strengths between their own friends and the invitation senders, and whether our visual approach provided more convincing evidence to accept invitations as compared to the current OSN approaches. Participants found RelationGrams to be relevant with social parameters, robust against attackers, and easy to use. They also appreciated the visualizations to make informed authentication decisions. Please refer to the technical report for the evaluation results [10].

6 Discussion and Future Work

A first question is how usable such a system would really be, or whether it would be such a burden on the user that its utility would be negated. Although further research is needed, several points indicate that the burden would be minimal. Existing systems could automatically collect interaction information without burdening users by aggregating email, SMS, Google+, etc. exchanges. Smartphones could also collect information about people users physically meet, through the use of voice recognition or by detecting the proximity of the other party’s smartphone. Generation of evidence, endorsement (i.e., digital signature), and distribution to friends could also be automated. A minor burden would be configuration, where a user can decide which tie strength visualizations to share with others. This could occur through an opt-in process, where a user could add friends whose tie strength information could be shared.

Another important question is on incentives: would users really have incentives to share their tie strength visualizations, and how can privacy concerns be dealt with? In our user studies, it was clear that users seemed eager to obtain such information to validate online invitations with confidence. Although further studies are needed, we believe that people’s inherent altruism that explains Internet phenomena such as Wikipedia would also encourage users to share their tie strength visualizations, because little burden is required on their part, and they can help their friends to befriend each other with more safety.

7 Conclusion

Online user behavior is faced with an uncomfortable trade-off: should we really accept unauthenticated friends’ invitations that might represent impersonation attempts to deceive; or should we deny them at the cost of losing potentially valuable relationships and become socially isolated? Currently, there is no secure and usable mechanism that would enable us to resolve this dilemma.

Our online identity authentication system implements a simple identity authentication logic in a visually compelling manner that is consistent with mental models derived from real-life experience. That is, it enables a casual user to authenticate online identities in a safe and easy-to-use manner.

References

1. Sophos Facebook ID Probe. <http://www.sophos.com/pressoffice/news/articles/2007/08/facebook.html>.
2. L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks. In *Proceedings of WWW*, 2009.
3. Economist. Duly notarised. <http://www.economist.com/blogs/babbage/2011/09/internet-security>, Sept. 2011.
4. M. S. Granovetter. The Strength of Weak Ties. *The American Journal of Sociology*, 78(6):1360–1380, 1973.
5. N. Hamiel and S. Moyer. Satan Is On My Friends List: Attacking Social Networks. In *Black Hat Conference*, 2008.
6. Q. Hu, Z. Xu, T. Dinev, and H. Ling. Does Deterrence Work in Reducing Information Security Policy Abuse by Employees? *Communications of The ACM*, 84(6):54–60, 2011.
7. D. Irani, M. Balduzzi, D. Balzarotti, E. Kirda, and C. Pu. Reverse social engineering attacks in online social networks. In *Detection of Intrusions and Malware and Vulnerability Assessment (DIMVA)*, July 2011.
8. T. H.-J. Kim, L. Bauer, J. Newsome, A. Perrig, and J. Walker. Challenges in Access Right Assignment for Secure Home Networks. In *Proceedings of USENIX Workshop on Hot Topics in Security (HotSec)*, Aug. 2010.
9. T. H.-J. Kim, V. Gligor, and A. Perrig. Street-Level Trust Semantics for Attribute Authentication. In *Proceedings of the 20th International Workshop on Security Protocols*, April 2012.
10. T. H.-J. Kim, A. Yamada, V. Gligor, J. I. Hong, and A. Perrig. RelationGrams: Tie-Strength Visualization for User-Controlled Online Identity Authentication. Technical Report CMU-CyLab-11-014, Carnegie Mellon University, 2011.
11. D. Krackhardt. The Strength of Strong Ties: The Importance of *Philos* in Organizations. *N. Nohria and R. Eccles (eds.), Networks and Organizations: Structure, Form, and Action*, pages 216–239, 1992.
12. D. Z. Levin and R. Cross. The Strength of Weak Ties You Can Trust: The Mediating Role of Trust in Effective Knowledge Transfer. *Management Science*, 50(11):1477–1490, 2004.
13. R. Reagans and B. McEvily. Network Structure and Knowledge Transfer: The Effects of Cohesion and Range. *Administrative Science Quarterly*, 48(2):240–267, 2003.
14. T. Ryan. Getting in Bed with Robin Sage. In *Black Hat Conference*, 2010.