Security and Trust through Electronic Social Network-based Interactions

Patrik Bichsel*, Samuel Müller‡, Franz-Stefan Preiss*, Dieter Sommer* and Mario Verdicchio†
*IBM Research GmbH, Rüschlikon, Switzerland, *Email:* {*pbi, frp, dso*}*@zurich.ibm.com*
†Università degli Studi di Bergamo, Bergamo, Italy, *Email: mario.verdicchio@unibg.it*
‡ETH Zürich, Information Security Group, Zürich, Switzerland, *Email: smueller@inf.ethz.ch*

***Abstract*—The success of a Public Key Infrastructure such as the Web of Trust (WoT) heavily depends on its ability to ensure that public keys are used by their legitimate owners, thereby avoiding malicious impersonations. To guarantee this property, the WoT requires users to physically gather, check each other's credentials (e.g., ID cards), to sign the trusted keys, and to subsequently monitor their validity over time. This trust establishment and management procedure is rather cumbersome and, as we believe, the main reason for the limited adoption of the WoT. To overcome this problem, we propose a solution that leverages the intrinsic properties of Electronic Social Networks (ESN) to establish and manage trust in the WoT. In particular, we exploit dynamically changing profile and contact information, as well as interactions among users of ESNs to gain and maintain trust in the legitimacy of key ownerships without the disadvantages of the traditional WoT approach. We see our proposal as an effective way to make security and trust solutions available to a broad audience of non-technical users.**

I. Introduction

The Public Key Infrastructure (PKI) paradigm aims at assigning asymmetric cryptographic keys to entities, such as people or organizations, to allow for several security features, including secret communication, rights delegation, and access control. A key distribution mechanism comes with a *trust infrastructure* that enables the establishment of the authenticity of the binding between a public key and a person. To achieve this result, mechanisms are prescribed to associate metadata with the public key, which allow users to identify the key's owner. One example of such a trust infrastructure is the Web of Trust (WoT) [1]. In a WoT, trust between users is ensured through mutual key signing. In particular, a user verifies the metadata by checking a physical witness, like an ID card or a driver's licence, at gathering events called *key signing parties*. The metadata attached to a user's key thus not only consists of personally identifying information (PII) but also of signatures from people that the user has met at an event and that were willing to personally guarantee that she is the legitimate owner of that key. These procedures are unfortunately far from simple, and thus do not provide an appealing and straightforward way to increase communication security for a wide audience.

The aim of this work is to exploit the widespread and successful Electronic Social Network (ESN) mechanisms for personal information exchange to support the trust establishment and management processes prescribed by the WoT. In this work, we rely on the WoT definition of trust: the confidence in the fact that the PII attached to a public key corresponds with the real identity of the possessor of the key. Our aim is to simplify the processes that establish and manage trust by leveraging the information provided by ESN users in the form of personal data attached to a user's profile and in the form of interaction happening through the ESN. Such information can be exploited by other users to identify the physical person controlling that particular profile in the social network, and thus, to trust that she is the legitimate owner of her public key. Indeed, to achieve this purpose on the sole basis of the WoT mechanisms, users must go through the afore-mentioned cumbersome procedures. The main idea of this work is that by linking an ESN profile to a WoT certificate, one is enabled to exploit the ESN-based information exchange mechanisms to establish trust in people.

The paper is organized as follows: Section II motivates our work and gives an overview of the underlying concepts; our solution is detailed in Sections III and IV, which show how ESNs can support trust establishment and management, respectively; implementation guidelines are provided in Section V; the possibility to exploit multiple ESNs is described in Section VI; Section VII discusses related work; finally, we conclude and give hints on future research in Section VIII.

II. Motivation

The WoT has not seen a widespread adoption, and we believe that the complexity of its trust establishment process is the main reason. Moreover, an established WoT is static and its maintenance requires significant effort. For example, if a user loses her private key, she needs to re-establish all her trust relations, which means that she has to re-prove her identity to all the contacts that had trusted her before the key loss. The second establishment entails the same investments (i.e. face-to-face meetings) as the creation of the original trust relations.

Our proposal for tackling these shortcomings is based on the idea that the interactions over an ESN can work as a valid substitute for the WoT-prescribed key signing parties. The underlying assumption is that mimicking the behavior of a user over a long period of time in an ESN is as hard as forging an identification document. An important characteristic

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 216483 for the project PrimeLife.

of ESN-based interactions is that they do not consist of a single interaction, but take place over long temporal intervals. This allows for a fine-grained assessment of the ESN members a user interacts with. For example, it is easy to search for a detailed set of personal data which might allow for an episodic impersonation of a user, while it is much more difficult to impersonate her for a longer time span in an ESN as this possibly entails live chatting sessions, message exchanges or uploading of pictures.

In our view, as the information exploited to build a trust relation comes from an ESN, the traditional WoT methods should be integrated into the ESN management system to simplify the necessary key signing processes. After certain conditions have been met (e.g., a number of interactions over a given timespan), the ESN could query a user whether she currently believes that her communication partner is actually the person indicated by the relevant personal data. If the user confirms that, the key signing process could be handled transparently. Our approach prescribes the exploitation of the information exchange over an ESN to establish a trust infrastructure.

The ESN-based approach is beneficial also to the management of the trust infrastructure. Let us again consider the case of a user losing her private key. Assuming that the user can still authenticate herself towards the ESN, it is sufficient to convince the other users that she has lost her key and that they should sign the newly produced key. The users prompted to sign the new key can re-authenticate the person through the ESN in a simple and fast way. This method clearly mitigates the investments of a user after a key loss while maintaining an acceptable level of security.

These considerations shed light on the main advantage of the proposed approach: ESNs provide a simple way for a wide audience, which already exploits its interaction mechanisms, to achieve security and trust properties that traditionally rely on far more complex mechanisms. The following sections show how this result can be achieved.

III. Trust Establishment

In our model, the establishment of trust consists of two steps: (1) *trust assessment*, by which a user u evaluates the confidence she has in the identity of another user v, and, in case this confidence is enough for u to believe that the identifying metadata belongs to v, (2) *trust declaration*, whereby u makes her trust explicit. Therefore, when we say that u trusts v, u has assessed and declared her trust in v.

In the following we define our concept of trust, we describe trust assessment and declaration as performed intuitively by a user and we elaborate on trust assessment through ESN interaction. Finally, we depict how trust assessment can be simulated on the basis of ESN data to allow for meaningful suggestions to the users.

A. Definition of Trust

We consider a trust infrastructure E, a set of ESNs $W = \{w_1, w_2, \ldots\}$, a set of users $U = \{u_1, u_2, \ldots\}$ as well as a set of attribute names $A = \{a_1, a_2, \ldots\}$. The trust infrastructure comprises a set of users $U_E \subseteq U$, a set of attributes $A_E \subseteq A$ as well as a (possibly partial) function $a_{E,u}$, mapping the attributes in A_E onto concrete values for $u \in U_E$. Similarly, an ESN w comprises a set of users $U_w \subseteq U$ and a set of attributes $A_w \subseteq A$. Let $F_w \subseteq U_w \times U_w$ be a set of symmetric friendship relations between users in w: if a user $u \in U_w$ is in a friendship relation with a user $v \in U_w$, we have $(u, v) \in F_w$, as well as $(v, u) \in F_w$ because of the symmetry of F_w. Every user $u \in U_w$ has a profile $P_{w,u}$ that contains (1) a (possibly partial) function $a_{w,u}$, mapping the attributes in A_w onto concrete values for u, as well as (2) the set of u's friends in w, defined as $F_{w,u} = \{v \mid (u, v) \in F_w\}$.

We now formalize our concept of trust. A user u trusts another user v when she has *enough confidence* about the fact that v is indeed the person she claims to be according to her trust infrastructure attributes, i.e., that the values $a_{E,v}$ indeed correspond with the values of v's PII. More formally, let $T \subseteq U_E \times U_E$ be a trust relation, where $(u, v) \in T$ means that u trusts v and is denoted as $u \twoheadrightarrow v$. Moreover, let $T_u = \{v \mid (u, v) \in T\}$ be the set of users that u trusts. Trust relations are, in contrast to friendship relations, not symmetric, i.e., $u \twoheadrightarrow v$ does not entail $v \twoheadrightarrow u$. Trust relations are not transitive and we assume trust propagation to be handled by the trust infrastructure (e.g., the WoT). For now, we consider trust to be binary. We will introduce different trust levels in Section IV-B. Where clear from the context, we will omit the indexes that identify the ESN. Note that T is not ESN-specific since we consider the trust relations to be publicly available through the trust infrastructure (like the WoT) the ESNs rely on.

B. Trust Assessment

The confidence of u that a given set of attributes belongs to v is strongly evidence-based, i.e., the more evidence u gathers, the higher her confidence about v's identity becomes. The evidence indicating that a user v is indeed who she claims to be takes various forms, such as information, credentials, or characteristics v proves to have as well as actions v performs. Information that v has could be the content of a previous conversation with user u. A valid passport is an example of a credential v might possess, and a recognized voice or style of expression are characteristics that v might prove to have in a phone or chat conversation. An action that v performs might be responding to an e-mail sent to her mail account.

Different pieces of evidence contribute differently to the confidence in a user's identity. For example, the presentation of a valid passport is, due to the high trust in the issuer, considered a much stronger identity evidence compared to a membership card from the local gym. Therefore, the presentation of the passport leads to a higher increase in confidence than showing the membership card. Such increase is very subjective as different users may ascribe different values to the same piece of evidence. In the above mentioned example, people who know the very strict identification procedures at a particular gym will value the relevant membership card more than people who are not familiar with these procedures. In addition, the level of confidence necessary for trusting another user is also a very subjective factor.

In fact, there may also be counter evidence for a user v's identity, i.e., evidence that indicates that a user v is *not* the one she claims to be. However, the only factor we consider for decreasing a user's level of confidence is time, i.e., the level of confidence decreases gradually with time in case u and v not having any interaction.

Our formalism includes $c_u(v) \in \mathbb{R}$, denoting u's level of confidence that the values $a_{E,v}$ correspond to the values of v's PII, and the threshold $t_u \in \mathbb{R}$, indicating how much confidence u needs to consider another user as trusted. Note that because of the before-mentioned considerations on the subjectiveness of these concepts, it is not possible (not even for u herself) to frame either $c_u(v)$ or t_u into an absolute quantitative scale.

C. Trust Declaration

Let $w \in W$ be an ESN and $u, v \in U_w$ be users. As soon as $c_u(v) \geq t_u$ holds, u may make this explicit by adding v to her set of trusted users T_u. In order to do so, however, we require $u \in U_E$ and $v \in U_E$. To ensure that u can identify v in both the ESN w and the trust infrastructure E, we require a dedicated attribute which is mapped to the same value (e.g., the hash value of v's public key) in $a_{w,v}$ and $a_{E,v}$.

D. Trust Assessment through ESN Interaction

The growing list of data managed by common ESNs comprises, in addition to the profile attributes, friends lists, blog entries, messages, comments, pictures and relevant tags, videos, status messages, etc. These data serve as the evidence that is necessary to perform trust assessment. We also regard the mere interaction between u and v in the ESN, such as conversations, tagging of pictures involving the other user, or commenting on the other user's content, as evidence for their identities. For example, consider u assessing trust in a user v who introduces herself as a former work colleague. The profile, including photos showing common friends, seems legitimate to u. Not yet fully convinced, u engages in a chat conversation with v talking about a past joint event. This information, together with v's writing style increases u's confidence level significantly such that $c_u(v) \geq t_u$ holds and u declares her trust, i.e., $u \twoheadrightarrow v$.

To keep trust relationships up to date, the ESN might notify a user u about the possibility of user v having reached u's trust threshold and propose to establish a trust relationship. In addition, the ESN should prevent u from taking unwise decisions like declaring trust in a user she did not assess the attributes closely. To do so, the ESN must be enhanced with trust simulation mechanisms.

E. Simulating Trust Assessment

The model illustrated in Section III-B describes the trust assessment as it is *intuitively* performed by a user. Thus, a way for the ESN to assist the user in her trust decisions is to simulate her assessment process. The challenge for the ESN is to provide an appropriate calculation model for estimating the confidence levels as well as the trust threshold. To make this explicit, in the following we only consider the simulated level of confidence and trust threshold, denoted as $\hat{c}_u$ and $\hat{t}_u$, respectively. As the confidence level is driven by the available evidence, the different types of evidence accessible to the ESN need to have assigned appropriate values that gradually increase the level of confidence.

We prescribe the trust threshold of a user to be initially determined by and automatically adjusted according to her trust declarations towards other users. The interaction that leads user u to declare trust towards user v will be used by the ESN as an estimate of the confidence level needed by u to trust other users she interacts with at a later stage. We provide details of these mechanisms in Section IV-B.

The advantage of this approach is avoiding the burden of elaborating a computable definition of trust, or, more precisely, enumerating and dealing with all conditions that cause persons to trust each other. An exhaustive list of such factors is very hard to compile, as many of them are very subjective.

Many factors influencing a person's trust lie outside an ESN, e.g., phone calls, work meetings, dinner parties, and thus cannot add up to whatever metrics the social network relies on to register the electronic interactions among its users. These considerations show why the user's autonomous decisions must play a fundamental role in any ESN-based solution to support a trust infrastructure.

IV. Trust Management

The trust establishment process enables a user to build trust relations. As important as the establishment is the management of such relations, for two reasons. Firstly, social relations are very dynamic and so is the evidence that an ESN uses when suggesting to build a trust relation. This triggers the need to describe how changes in the communication frequency affect existing trust relations. Secondly, trust comes with a propagation effect: the confidence one has in a very trusted friend of a very trusted friend is clearly higher than the confidence in a complete stranger with no attachments. Both aspects will be discussed in this section.

A. Dynamics of Relations

Relationships in real life change over time for various reasons such as people getting new jobs or hobbies. The change in a relationship can also be observed in an ESN. In particular, the evidence used for the calculation of $\hat{c}_u(v)$ can show the modifications in how close u and v are. Let us discuss how the trust between u and v evolves, while $\hat{c}_u(v)$ changes its value. As developed in Section III-B, the level of confidence can increase as well as decrease. In the case of an increase, the user goes through a trust establishment process. Thus, whenever the condition $\hat{c}_u(v) \geq \hat{t}_u$ holds, the ESN proposes to u to enter in a trust relation with v.

We propose that a decrease of $\hat{c}_u$ below $\hat{t}_u$ should not trigger the ESN to suggest a change to an already established trust relation. This is because our trust definition is a statement that

u once had the possibility to assess the correspondence of the values of $a_{E,u}$ with v's PII. Thus, we propose that only a change in metadata attached to a key might trigger the ESN to suggest a new trust assessment among users having an established trust relation but not enough interaction to fulfill $\hat{c}_u(v) \geq \hat{t}_u$.

B. Trust Levels

Some trust infrastructures rely on a more detailed model in which not only the fact that u trusts v ($u \twoheadrightarrow v$) is formalized, but a degree is also assigned to such relation. For example, the WoT includes two *trust levels*, namely, *marginal* and *full*, where the latter indicates stronger trust. Those trust levels are exploited in the trust propagation process. Let us show how to integrate such trust levels in our formalism.

We assume that the trust infrastructure provides a totally ordered list of $n + 1$ increasing trust levels $R = \{r_0, \ldots, r_n\}$, corresponding to trust relationships that increase in strength. In the case of the WoT, we have $R_{WoT} = \{marginal, full\}$. We denote a trust relation of level r_j between u and v as $u \overset{r_j}{\twoheadrightarrow} v$. Let $T_{i,u} = \{v \mid u \overset{r_i}{\twoheadrightarrow} v\}$ be the set of users that u trusts with level r_i.

Let us remind that the ESN's estimate of u's level of confidence $\hat{c}_u(v)$ increases with the amount of interaction u has with v, and that when the condition $\hat{c}_u(v) \geq \hat{t}_u$ is reached, this may cause u to make her trust explicit to $u \twoheadrightarrow v$. Similarly, we define *trust level thresholds* that, once reached by the estimated confidence, can cause u to increase the level of a trust relation accordingly. In particular, these trust level thresholds are used by an ESN to make proposals for advancing trust relations to a higher level. For every $r_i \in R$ we define a threshold $\hat{t}_{i,u}$ and $\hat{t}_{0,u} = \hat{t}_u$, i.e., the threshold of the lowest trust level is equal to the user's general trust threshold. As soon as condition $\hat{c}_u(v) \geq \hat{t}_{i,u}$ is reached, the ESN proposes to assign the trust level r_i to $u \twoheadrightarrow v$.

In accordance with the computational model for a user's general trust threshold, we also prescribe the values of the trust level thresholds to be defined by the user's behavior.

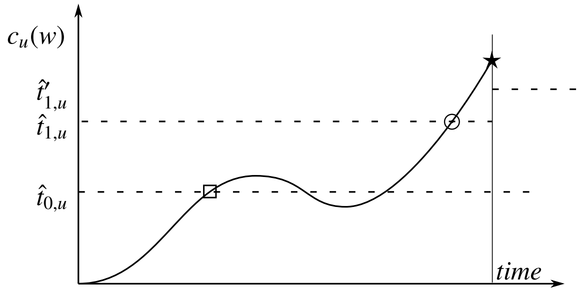

Fig. 1. Level of confidence $c_u(w)$ between users u and w. We assume that u accepts the proposal of the ESN to add v to $T_{0,u}$ at $\square$. The proposal to add w to $T_{1,u}$ is rejected (at $\odot$) but u manually assigns w the next trust level (at $\star$). Thus, the ESN adapts $\hat{t}_{1,u}$ to $\hat{t}'_{1,u}$.

Let us illustrate how trust level thresholds are set in more detail. The basic idea is that the ESN establishes these thresholds on the basis of the user's past behavior, in terms of the interaction she previously needed before assigning another user a certain trust level. At the beginning, a user u has no trust relations nor any trust level set. In particular, we assume that all trust relations of u start at a very high level in order not to propose trust relations too early. Due to her interaction with v through the ESN, the level of confidence increases and at a certain instant the user declares that trust has been established: $u \overset{r_0}{\twoheadrightarrow} v$. The ESN records u's current $\hat{c}_u(v)$ and uses the value to make a first estimate of u's lowest trust level $\hat{t}_{0,u}$. Further interaction with v increases $\hat{c}_u(v)$ and at a certain instant u assigns the next trust level to v, resulting in $u \overset{r_1}{\twoheadrightarrow} v$. Again, the value of $\hat{c}_u(v)$ is used by the ESN as an estimate for $\hat{t}_{1,u}$.

The ESN uses these estimated trust level thresholds to assist u in finding an appropriate amount of interaction before entering trust relations. So, as soon as interaction with w makes u's level of confidence reach the previously established $\hat{t}_{0,u}$, the ESN will propose u to grant w trust level r_0. Should u accept this proposal, $\hat{t}_{0,u}$ is confirmed. In another case u might refuse the proposal of granting w trust level r_1, waiting for more trust-building interaction to take a decision. Thus, $\hat{t}_{1,u}$ needs some adjustment. If $\hat{c}_u(w)^*$ is the level of confidence at which u finally grants w trust level r_1, the new trust level threshold $\hat{t}'_{1,u}$ can be set as follows: $\hat{t}'_{0,u} = \alpha \cdot \hat{t}_{0,u} + \beta \cdot \hat{c}_u(w)^*$, where α and β are parameters that weigh the contribution of the gap to the new threshold estimation (e.g., $\alpha = \beta = 0.5$, see Fig. 1).

A situation like the following, $\hat{t}'_{i,u} \geq \hat{t}_{j,u}$ with $r_i < r_j$, in which the gap is so wide that the new threshold for trust level r_i is above the threshold of level r_j, although $r_i < r_j$ needs to be handled with special care. All thresholds that are affected, that is, are overtaken by the newly estimated threshold $\hat{t}'_{i,u}$, must be adjusted. Let r_k be the first trust level whose threshold is above $\hat{t}'_{i,u}$. The estimation of trust level r_k should not be changed, as there has not been any significant clue on its inadequacy. All thresholds of levels between r_i and r_k must be then repositioned in the interval $\hat{t}_{k,u} - \hat{t}'_{i,u}$. The repositioning can be performed with a uniform distribution of the new thresholds in the interval, or by keeping the proportions between the relative positions that the old thresholds occupied with respect to $\hat{t}_{i,u}$ and $\hat{t}_{k,u}$. Should $\hat{t}'_{i,u}$ be above also the maximum trust level $\hat{t}_{n,u}$, then all thresholds between level r_i and r_n can be shifted accordingly. Analogous considerations hold for changes that lower the trust level thresholds.

One might object against the use of the trust levels established with a user v to perform estimations involving another user w: any motivation leading u's decisions with respect to v regards v, and v only. However, this position entails that user u should decide individually for the trust levels of all the users she connects to through the ESN, which is clearly an undesirable burden for most users, while the use of estimated trust level thresholds allows for the decision process to be supported by the ESN.

C. Trust Propagation

Trust propagation is important in scenarios where a user u wants to use another user x's public key which she has not signed. In fact, the trust propagation enables users to benefit not only from their direct trust relations but also from a larger network of people they might gain confidence in. In addition to the standard trust propagation mechanisms as the one used in the WoT or relevant improvements as described in [2], we propose to use the additional ESN information to create more confidence. For example, adding the information about the friendship and communication frequency between two people in the chain of the trust propagation may improve to a user's confidence.

V. Architecture and Implementation Guidelines

Let us now focus on the technical aspects of our approach. The description of the architecture relies on concepts related to standard WoT principles. Our aim is to shed light on the added value provided by the ESN.

A. Key Generation

The key generation process consists of (1) the generation of a private/public key pair, (2) the binding of PII metadata to the public key resulting in a certificate, and (3) the publication of the certificate to a publicly accessible repository.

Security considerations allow for the first step only to take place on a device trusted by u. Still, the ESN can enable u to initiate the key generation process on a user-trusted device via the ESN itself. The collection of the metadata as well as the publication of the certificate can be entirely performed by the ESN. For the generation of the certificate, however, a device holding the user's private key is needed. The ESN can facilitate this process with a specific request to the device. Finally, the ESN must add a reference to the certificate of u's profile.

B. Key Signing

When building trust relations as described in Section III-D, users finally need to declare their trust. Given u willing to express $u \twoheadrightarrow v$, from a technical perspective, this entails that u signs v's public key together with the relevant attributes and uploads the resulting certificate to a publicly accessible repository. Thus, u confirms the binding of v's attributes as stated in the certificate and allows other people, who are not necessarily part of an ESN, to access this information.

The ESN can facilitate the key signing process by automatically comparing v's ESN attributes with the ones given in her self-signed certificate. In case of a mismatch the trust conditions are not met and u is warned. Signing v's certificate must rely on a device trusted by u as in the case of key generation, and the ESN can provide an analogous support.

C. Key Management

Key management turns out to be a complex task due to the following issues: (1) the availability of *all keys in* T_u to u on all devices that u uses even if only temporarily (e.g., a computer at an Internet cafe), (2) the availability of the *user's key pair*, especially her private key, from all her devices (devices not owned by u are excluded here), and (3) the *correct usage* of all keys, i.e., renewal of the own key, timely revocation, and refraining from the use of expired keys.

The ESN-based approach allows for optimization compared to the WoT approach in all those aspects. Firstly, u's *key ring* (the set of the public keys of the user u in T_u) can be downloaded transparently by the ESN whenever u connects with a new device. Thus, this problem boils down to a connectivity problem. Secondly, the portability of u's private key and thus of the possibility of executing transactions such as decryption or key signing is more critical. One solution is to let the user have the key on a portable device (e.g., a smart card). Another possibility is to use a group signature scheme where the user might register several devices which can all execute the transactions traditionally requiring the private key. Note, that the current WoT implementation does not allow for group signature keys to be used. Thirdly, correct usage again boils down to a connectivity problem as whenever the user u is online, the ESN can update revocation lists or the expiration of u's key. Thus, we propose that the ESN provides a mechanism to allow for a timely replacement of a key coming close to its expiration. The renewal itself could consist of a proof of possession of the old private key and the generation of a new key pair.

The revocation of public keys in a PKI is a known issue. However, our proposal allows for improvement in that area by using short key life-cycles with automatic ESN-based renewal that involves the user's host.

VI. Integration of multiple ESNs

So far, we focused on the benefits of a single ESN to the WoT. Let us now further potential of our approach by considering scenarios in which the WoT trust infrastructure is connected to multiple ESNs.

An issue arises in the task of establishing and managing trust when having several ESNs as opposed to one. Given that users u and v have profiles in a number of ESNs w_k where $k \in K = \{1, \ldots, \ell\}$, all ESNs w_k estimate $c_u(v)$ individually as $\hat{c}_{w_k,u}(v)$, whereas u's perception would be better modeled by $\sum_{k \in K} \hat{c}_{w_k,u}(v)$. A solution to this problem would be the communication of the respective confidence levels between the involved ESNs, either directly or via u's host. However, this seems unrealistic as ESNs currently do not allow for automatic information flow outside their own network.

Alternatively, ESNs could indirectly infer information by observing changes in the WoT. Let us assume that $\hat{c}_{w_m,u}(v) \geq \hat{t}_{w_m,u}$ holds for $m \in K$. ESN w_m then asks u whether v should be added to T_u. Should u accept the ESN's proposal, u would issue a certificate on v's public key. This change in the WoT can be noticed by all the other ESNs, which can infer that (1) u uses at least another mechanism (e.g., another ESN) to assess the trust in v, and (2) such mechanism has been used more frequently with respect to the interactions with user v. The second statement relies on the assumption that interactions affect the establishment of trust in all ESNs in a similar way. When a change in the WoT shows that user v has gained trust in some other network, an ESN can adjust its current estimated level of confidence by adding u's threshold, i.e., $\hat{c}_{w_k,u}(v) = \hat{c}_{w_k,u}(v) + \hat{t}_{w_k,u}$, for all $k \in K \backslash \{m\}$, to factor in interaction between u and v that is sufficient to insert v in T_u that has taken place in some other ESN. Expanding this example to a trust model with n levels is straightforward.

VII. Related Work

We agree with Hogben [3], where he hints at the possibility to establish trust by means of the information provided by an ESN. Our work goes further though and illustrates in more detail how such information can be exploited to actually build a trust relationship.

In [4], an alternative methodology to the WoT is discussed. A PKI with limited dimensions is required to establish trust in URIs. Trust is modeled after a transitive relation, and such transitivity is considered to be sufficient to ensure that the proposed mechanism is effective. Any aspects involving people, the essential component of a key signing party, is left out.

Bootstrapping an open ESN is discussed in [5], where trust aspects are not taken into account, but the focus is on the concept of security, interpreted here as the possibility to have protected personal information, as opposed to public and available to any member of the ESN.

Several works aim at exploiting the users' behavior in ESNs to establish trust relations. In [6], a policy-based approach is proposed, where access to private ESN data is granted on the basis of the ESNs by which the requester is linked to the data owner, and the interaction frequency on those channels. Although we share the authors' approach in considering dynamic aspects of users' behavior over time, we part from their effort in the following respect: such aspects are considered only in the process of writing access control policies to private data (e.g.: "only people who commented on my blog at least 10 times in the last 2 weeks can see the pictures"), while in our view, the behavior of users is continuously assessed to update the trust relationship with them.

In [7], a trust metric is proposed that not only takes third-party opinions into account, but also considers 'aging' as a factor that weakens a previously established link, unless it is refreshed with new interactions or mediated opinions. From the authors' perspective, the trust built by means of iterated interaction is to be interpreted as a measure of how reliable is the information provided by the ESN users. Our work, instead, is closer to the basic concepts of the WoT, as we are more focused on exploiting such exchanges to assess the link between an ESN entry and the person that created it.

Zhang et al. in [8] provide a more complex model of trust, where 'trust rating' to choose interaction partners is distinguished from a 'reliable factor' that assesses the believability of their acquaintances' assertions. Nevertheless, the evaluation of the link between the ESN entry and the real person is once again not considered. This work is strongly related to Goldbeck and Hendler's [9], where the authors aim at providing a trust rating system that allows for the creation and evaluation of links between people who are not directly connected in an ESN.

VIII. Conclusions and Future Work

We presented an approach that integrates widely-popular electronic social networking technology with the hitherto unsuccessful Web of Trust paradigm to overcome the main obstacle to WoT adoption: cumbersome establishment and management of trust in the legitimacy of key ownership. Our proposal leverages available ESN profile and contact information, as well as interactions between users for establishing and managing a sufficient degree of trust for use in the WoT. Thereby, trust is established to a large extent by assessing ESN user behavior, requiring little to no extra effort on the side of the participating users. Moreover, our approach addresses key revocation and key renewal, two common key management problems of the traditional WoT. Overall, combining ESNs with the traditional WoT paradigm has the potential to provide security and trust solutions to non-technical users in a largely transparent manner.

To assess the practical feasibility of our ideas, it would be instructive to implement our model on top of existing ESN platforms. An implementation would also provide the basis for collecting empirical data to fine-tune the parameters of our model. A further area of future work concerns the study of privacy issues associated with the progressive integration of communication devices and disparate data sources as implied by our ideas.

References

[1] P. R. Zimmermann, *The official PGP user's guide*. Cambridge, MA, USA: MIT Press, 1995.

[2] R. Haenni and J. Jonczy, "A new approach to PGP's web of trust," in *EEMA'07, European e-Identity Conference*, Paris, France, 2007.

[3] G. Hogben, "Security issues in the future of social networking," in *W3C Workshop on the Future of Social Networking*, 2009.

[4] H. Story, "FOAF+SSL: Creating a web of trust without key signing parties," [online; 12 April 2009], Sun Microsystems, Tech. Rep., 2009, http://blogs.sun.com/bblfish/entry/more_on_authorization_in_foaf.

[5] ——, "Building secure, open and distributed social network applications," [online; 12 April 2009], Sun Microsystems, Tech. Rep., 2008, http://blogs.sun.com/bblfish/entry/building_secure_and_distributed_social.

[6] A. Passant, P. Karger, M. Hausenblas, D. Olmedilla, A. Polleres, and S. Decker, "Enabling trust and privacy on the social web," in *W3C Workshop on the Future of Social Networking*, 2009.

[7] V. Carchiolo, A. Longheu, M. Malgeri, G. Mangioni, and V. Nicosia, "An approach to trust based on social networks," in *Proceedings of the 8th International Conference on Web Information Systems Engineering (WISE 2007)*, ser. LNCS, vol. 4831. Springer, 2007, pp. 50–61.

[8] Y. Zhang, H. Chen, and Z. Wu, "A social network-based trust model for the semantic web," in *Proceedings of the 3rd International Conference on Autonomic and Trusted Computing (ATC 2006)*, ser. LNCS, vol. 4158. Springer, 2006, pp. 183–192.

[9] J. Goldbeck and J. Hendler, "Inferring binary trust relationships in web-based social networks," *ACM Transactions on Internet Technology (TOIT)*, vol. 6, no. 4, pp. 497–529, 2006.