Trustworthy Information: Concepts and Mechanisms

Shouhuai Xu[1], Haifeng Qian[2], Fengying Wang[3], Zhenxin Zhan[1], Elisa Bertino[4], and Ravi Sandhu[1]

[1] University of Texas at San Antonio (UTSA)
[2] East China Normal University (work done while visiting UTSA)
[3] Shandong University of Technology (work done while visiting UTSA)
[4] Purdue University

Abstract. We used to treating information received (from recognized sources) as trustworthy, which is unfortunately not true because of attacks. The situation can get worse with the emerging shift of information sharing paradigm from "need to know" to "need to share." In order to help information consumers make the "best" decision possible, it is imperative to formulate concepts, models, frameworks, architectures, and mechanisms to facilitate information trustworthiness management in distributed and decentralized environment. In this paper we initiate a study in this direction by proposing an abstraction called *information networks* as well as two supporting mechanisms called *provenance digital signatures* and *optimal security hardening of information network*.

1 Introduction

Suppose Alice receives a piece of information. To what extent should she trust it? Like most users, she will likely treat it as trustworthy, especially when the information is digitally signed by a peer whose public key is known to her. However, the presence of attacks such as malware and malicious insiders makes such trusting behavior questionable. This example scenario suggests the need for "trustworthy information" or "information trustworthiness management", which has been a missing piece of traditional approaches to data and information sharing. This paper is a significant first step towards addressing the problem.

Our Contributions. First, we propose the concept of "information trustworthiness management" in the context of information networks. Information networks are an abstraction we use to capture the "flow" or "movement" of information that is often in the format of messages with respect to some appropriate network protocols. Second, we formulate the abstraction of "trustworthiness graph" with respect to a piece of information. In a trustworthiness graph, a node represents a principal (e.g., a user or organization), and an arc represents the transmission of messages (containing information) between the nodes. Many mechanisms are needed for managing trustworthiness graphs. We investigate two useful mechanisms: (1) We identify a new kind of cryptographic primitive we call "provenance digital signatures" which goes beyond the standard concept of digital signatures by preserving the history of a message (i.e., a subgraph or subnetwork associated with the message). (2) We identify the need of optimal security hardening. We show that the algorithmic problem in question is NP-hard, but has a good approximation algorithm.

L. Chen et al. (Eds.): WAIM 2010, LNCS 6184, pp. 398–404, 2010.

Discussion. From a conceptual perspective, the concept of information networks as well as the resulting notion of trustworthiness graphs are reminiscent of the well-studied notion of "information flow" in computer security [7,19], which aims to ensure information secrecy and integrity. However, there are important differences. In particular, we focus on the notion of information trustworthiness, which is a broader requirement than secrecy and integrity and is, as mentioned above, about embedding security into information management at the very beginning of information lifecycle (rather than treating any given information as trustworthy). The differences are essential because, for example, an authorized user can insert a bad data item into a database, which is not trustworthy but disseminated to many other participants without violating the information-flow policy in question. While the present paper is in parallel to the recent active research that aims to inject provenance at the Operating System level and at the DataBase level (see, for example, [8,20,16,18,1] and the references therein), it is orthogonal to these studies because we take the perspective of the network that poses a new challenge that nodes in the network (e.g., OS/DB) can be compromised.

From a mechanism perspective, our notion of provenance digital signatures moves a step beyond the recent proposals for securing provenance information. Specifically, [9] investigated provenance integrity in the setting of file systems with an emphasis on the total-order operation-chain provenance of atomic objects (e.g., files). This study was extended to deal with the partial-order operation provenance of compound objects such as records and tables in database systems [22]. Compared with the integrity protection mechanisms presented in [9,22], provenance signatures offer stronger security definition and better efficiency. To see this, we note that the scheme presented in [22], when used in our setting (i.e., network), has the drawbacks of linear increasing of signature size and vulnerable to the "peeling off" attack. Intuitively, our scheme achieves strictly stronger security because signatures are always aggregated together (i.e., attacker does not "see" the individual signatures). Moreover, the "peeling off" attack has another security consequence on those schemes, which our scheme does not suffer. Specifically, a malicious user *by itself* can arbitrarily duplicate any prefix of a received signature so as to fake multiple incoming paths, which could manipulate (e.g., amplify or deflate) the trustworthiness of message it sends to downstream nodes.

From a cryptographic perspective, provenance signatures bear some similarity to (sequential) aggregate signatures [14,3]. However, we consider partial-order aggregation (i.e., with respect to information networks) whereas the later considers total-order aggregation. As such, our security definition is strictly stronger.

Other Related Work. This paper is inspired in part by a recently proposed framework for "trustworthiness-centric information sharing" [21]. The notion of trustworthiness graph was somewhat inspired by the concept of social networks [10,11].

2 Information Network and Trustworthiness Graph

We start with a simple example (Figure 1), which captures what is needed for evaluating information trustworthiness: from which principals an information item is originated; which principals have manipulated and forwarded an information item; which algorithms have been used to process an information item. Specifically, $P_1, \ldots, P_6$ in the

graph represent principals, and the arcs indicate how information has moved in the information network. Specifically, suppose P_1 enters message M_1 into the system at time T_1 and P_2 enters message M_2 into the system at time T_2. At time T_3, P_3 receives M_1 from P_1 and processes M_1 to produce M_3. At time T_4, P_4 receives M_1 and M_2 from P_1 and P_2, respectively, and processes them to produce a new message M_4. At time T_5, P_5 receives M_3 and M_4 from P_3 and P_4, respectively, and processes them to produce M_5. Finally, P_6 receives M_5 at time T_6.

The goal is to help the nodes evaluate the trustworthiness of the messages. For this purpose, we first define the notion of information networks in the context of information sharing and their associated trustworthiness graphs, and then show how the goal can be achieved while requiring what support. In what follows, we assume a simple notion of time domain, that is, we assume that the domain is represented as the pair $(\mathbb{Z}; \leq)$, where each element of $\mathbb{Z}$ is referred to as a time instant and $\leq$ is a total order on $\mathbb{Z}$. Also given $T, T' \in \mathbb{Z}$, $[T, T']$ denotes a time interval starting at time instant T and ending at time instant T'. Such model is very simple but it is adequate for the purpose of the present paper. More advanced notions of time in distributed systems have been investigated [12,15] which we will consider as part of our future work.

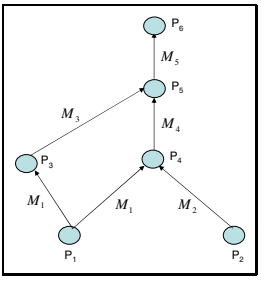

Fig. 1. Motivating example of provenance digital signatures

Definition 1. (information network) *Let* $[T_1, T_2]$ *be a time interval and* $\mathsf{V}([T_1, T_2])$ *be a set of principals (users, organizations) which exchanged information during* $[T_1, T_2]$. *An information network over* $[T_1, T_2]$ *and* $\mathsf{V}([T_1, T_2])$, *denoted as* $\mathsf{G}([T_1, T_2])$, *is a pair* $(\mathsf{V}([T_1, T_2]), \mathsf{E}([T_1, T_2]))$, *where* $\mathsf{E}([T_1, T_2])$ *is the set of edges. An edge* $(u, v) \in \mathsf{E}([T_1, T_2])$ *if* $u \in \mathsf{V}([T_1, T_2])$ *has sent a message to* $v \in \mathsf{V}([T_1, T_2])$ *during* $[T_1, T_2]$.

Definition 2. (trustworthiness graph of an information network) *Let* $[T_1, T_2]$ *be time interval and* T *be a time instant, where* $T_1 \leq T \leq T_2$. *A trustworthiness graph* $G(T) = (V(T), E(T))$ *at time* T *is defined as* $\mathsf{G}(\mathsf{V}[T_1, T]) = (\mathsf{V}([T_1, T]), \mathsf{E}([T_1, T]))$ *with the following annotations. If* $(u, v) \in E(T)$ *we say that* u *is an "upstream" node of* v *and* v *is a "downstream" node of* u. *Moreover, each* $(u, v) \in E(T)$ *is annotated with a pair* $(w_T(u, v), \theta_T(u, v))$, *where* $w_T(u, v) \in [0, 1]$ *is* v*'s trustworthiness evaluation of* u *at time* T *(e.g., based on the trustworthiness of information it has so-far received from* u*), and* $\theta_T(u, v) \in [0, 1]$ *is a threshold specified by* v.

Definition 3. (most/least trustworthy path) *Given a trustworthiness graph* $G(T) = (V(T), E(T))$ *with annotations and a path* $p = (v_1, \ldots, v_\ell)$, *we can define the trustworthiness of path* p *as (for example)* $W_T(p) = \prod_{i=1}^{\ell-1} w_T(v_i, v_{i+1})$, *which is a real number in the interval* $[0, 1]$. *For a given pair of nodes* $(u, v) \in V(T) \times V(T)$, *let* $P_T = \{(u, \ldots, v)\}$ *denote the set of paths from* u *and* v. *We say that path* $\bar{p} \in P_T$ *is (one of) the most trustworthy if* $W_T(\bar{p}) = \max\{W_T(p) : p \in P_T\}$ *and path* $\underline{p} \in P$ *is (one of) the least trustworthy if* $W_T(\underline{p}) = \min\{W_T(p) : p \in P_T\}$.

3 Provenance Signatures

Let us recall the motivating example illustrated in Figure 1. At first glance, one may think that this can be achieved by letting each P_i sign the respective subgraph using a standard digital signature mechanism (as in a standard PKI) such that a new signature is attached to last signature. However, this approach has two major drawbacks. (1) The size of the resulting signatures linearly increases with the number of signers. (2) A dishonest principal could manipulate the provenance by, for example, dropping some signers who signed in the past. For example, assume that P_1, P_2, P_3 and P_4 are honest and that M_1 and M_2 are sent through private channels such that no other principals have access to them. A dishonest P_5 could convince P_6 that M_5 is derived from M_3 received P_3, M_1 received from P_1, and M_2 received from P_2 (e.g., by simply setting $M_5 = M_3||M_4$). This "peeling off" attack is undesirable at least for the sake of crediting, which may serve as incentive for information sharing. The problem cannot be solved by requiring that each signer signs the identity of next principal in the graph because, for example, P_1 may broadcast M_1 to a large population of principals, and may not know in advance that P_4 will receive or accept it.

As mentioned above, the scheme presented in [22] suffers from these problems. As we will see, our provenance signatures overcome the above two drawbacks simultaneously. The basic idea is to utilize some cryptographic aggregation technique to ensure that (1) the size of the resulting provenance signatures (not including the messages as they are needed anyway) is constant and independent of the number of signers, and (2) attacks like the above "peeling off" one are prevented. While our use of cryptographic aggregation technique can be viewed as a generalization of the recently investigated notions of multi-signature, (sequential) aggregate signatures, ordered multi-signatures [5,13,2,3], the generalization is non-trivial. This is because we have to ensure (extra) unforgeability with respect to the graph structure, which makes the security proof complicated. In contrast, existing schemes cannot assure unforgeability with respect to graph structure and are actually subject to simple attacks (e.g., two corrupt signers can arbitrarily manipulate their "positions"). Moreover, we must allow multiple aggregations of a single signature on the same message (e.g., P_1 sends the same M_1 to both P_3 and P_4). Now we describe our provenance signature scheme.

$\mathsf{Setup}(1^k)$: Generate a bilinear group $\mathbb{G}$ with order $2^{k+1} \geq p \geq 2^k$ and an associated bilinear pair $e(\cdot,\cdot) : \mathbb{G} \times \mathbb{G} :\longrightarrow \mathbb{G}_\mathsf{T}$. Return $pp = (e, \mathbb{G}, \mathbb{G}_\mathsf{T}, H)$, where $H : \{0,1\}^* \rightarrow \mathbb{G}$ a random oracle.

$\mathsf{Keygen}(pp)$: Randomly choose $x \stackrel{R}{\leftarrow} \mathbb{Z}_p$ and output a pair of private and public keys $(sk = x, pk = X = g^x)$.

$\mathsf{GraphCom}(pp, \mathsf{loc}, \{G_\lambda\}_{\lambda \in \mathcal{R}})$: On input of public parameters pp, local information string $\mathsf{loc} = (\mathsf{id}_i, m_i, t_i)$ for a local message m_i of trustworthiness t_i, and incoming $|\mathcal{R}|$ provenance subgraphs $\{G_\lambda\}_{\lambda \in \mathcal{R}}$ where $\mathcal{R} \subset \mathcal{S}$, P_i generates a new message $\overline{m}_i = \mathsf{alg}_i(m_i, \{G_\lambda\}_{\lambda \in \mathcal{R}})$ of trustworthiness $\overline{t}_i = \mathsf{tru}_i(t_i, \{G_\lambda\}_{\lambda \in \mathcal{R}})$, where the specification of algorithms alg_i and tru_i is application-dependent and beyond the scope of the paper. Finally, P_i outputs a provenance subgraph $G_i = ((\{G_\lambda\}_{\lambda \in \mathcal{R}}), a_i)$ for its newly produced message $\bar{m}_i$, where $a_i = (\mathsf{id}_i, \mathsf{alg}_i, \overline{m}_i, m_i, \mathsf{tru}_i, \overline{t}_i, t_i)$ is the "end" node in G_i. Note that if $\{G_\lambda\} = \emptyset$, then $((G_\lambda), a_i) = (a_i)$.

$\mathsf{PSign}(pp, sk_i, \mathsf{loc}, \{\Sigma_\lambda\}_{\lambda\in\mathcal{R}})$: on input of public parameters pp, local information $\mathsf{loc} = (\mathsf{id}_i, m_i, t_i)$ of message m_i of trustworthiness t_i, a private key sk_i of P_i, and provenance signatures $\{\Sigma_\lambda\}_{\lambda\in\mathcal{R}}$ on respective provenance subgraphs $\{G_\lambda\}_{\lambda\in\mathcal{R}}$ received from P_i's upstream nodes belonging to $\mathcal{R} \subset \mathcal{S}$, P_i executes $\mathsf{PVrf}(pp, \Sigma_\lambda)$ to verify the individual provenance signatures Σ_λ. If any verification fails, abort; otherwise, set $G_i \leftarrow \mathsf{GraphCom}(pp, \mathsf{loc}, \{G_\lambda\}_{\lambda\in\mathcal{R}})$, set $\omega \leftarrow H(G_i)^{x_i}$, and output $\Sigma_i = (G_i, \sigma_i)$ where $\sigma_i = \omega \prod_{P_\lambda\in\mathcal{R}} \sigma_\lambda$.

$\mathsf{PVrf}(pp, \Sigma)$: given parameters pp, provenance signature $\Sigma = (G, \sigma)$, the algorithm parses G to obtain $\{G_i | i = 1, \cdots, \ell\}$ and the signers' identities $\{\mathsf{id}_i | i = 1, \cdots, \ell\}$, and returns 1 if the following equation holds and 0 otherwise: $e(g, \sigma) \stackrel{?}{=} \prod_{i=1}^{\ell} e\left(X_i, H(G_i)\right)$.

Theorem 1. *If the $\mathcal{BLS}$ signature is $(\mathrm{T}', q_s, \varepsilon')$-secure under a chosen-message attack, our provenance signature scheme is $(\mathrm{T}, q_p, \varepsilon)$-secure where*

$$\varepsilon \geq \varepsilon', \quad q_p = q_s \quad \textit{and} \quad \mathrm{T} \leq \mathrm{T}' - (q_s + 1)N \cdot \mathrm{T}_e, \tag{1}$$

where q_s, q_p are the numbers of the queries to the $\mathcal{BLS}$ signing oracle and the **PSigning** *oracle, respectively, and T_e is the time cost of exponentiation computation.*

4 Optimal Security Hardening

Given a trustworthiness graph, we want to identify the most "influential" K nodes so as to harden their security. This optimization problem turns out to be related to the *independent cascade model* in the context of social networks [10]. However, there is an important difference. In the model by Kempe at al. [10], the influence passes each *arc* with a probability that is independent of any other arc. In our model, a piece of information passes each *node* with a probability that is independent of other nodes. We now show that the optimal security hardening problem is NP-hard. The reduction is based on the NP-complete Set Cover problem.

Theorem 2. *The optimal hardening problem for trustworthiness graphs is NP-hard.*

We now show that the optimal security hardening problem also has a certain submodular structure. A function $f(\cdot)$ mapping sets to $\mathbb{R}^+$ is said to be submodular if it has the so-called *diminishing returns* property: for all $v \in V$ and all $A \subseteq B$ it holds that

$$f(A \cup \{v\}) - f(A) \geq f(B \cup \{v\}) - f(B).$$

This can be exploited to prove that the greedy algorithm produces a solution within an approximation factor of $(1 - 1/e)$, where e is the base of the natural logarithm [17].

Theorem 3. *The function $\sigma(\cdot)$ incurred by optimal hardening is submodular.*

5 Conclusion

We presented the concept of "information trustworthiness management" in the context of information networks and the abstraction of "trustworthiness graph". We investigated a new cryptographic primitive we call "provenance digital signatures" and investigate the issue of optimal security hardening in information networks.

Acknowledgement. This work was partially supported by AFOSR MURI and State of Texas Emerging Technology Fund.

References

1. Archer, D., Delcambre, L., Maier, D.: A Framework for Fine-grained Data Integration and Curation, with Provenance, in a Dataspace. In: TaPP 2009 (2009)
2. Bellare, M., Namprempre, C., Neven, G.: Unrestricted aggregate signatures. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 411–422. Springer, Heidelberg (2007)
3. Boldyreva, A., Gentry, C., O'Neill, A., Yum, D.: Ordered multisignatures and identity-based sequential aggregate signatures, with applications to secure routing. In: ACM CCS 2007 (2007)
4. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, p. 514. Springer, Heidelberg (2001)
5. Boneh, D., Gentry, C., Lynn, B., Shacham, H.: Aggregate and verifiably encrypted signatures from bilinear maps. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 416–432. Springer, Heidelberg (2003)
6. Dai, C., Lin, D., Bertino, E., Kantarcioglu, M.: Approach to Evaluate Data Trustworthiness Based on Data Provenance. In: Jonker, W., Petković, M. (eds.) SDM 2008. LNCS, vol. 5159, pp. 82–98. Springer, Heidelberg (2008)
7. Denning, D.: A Lattice Model of Secure Information Flow. CACM 19(5), 237–243 (1976)
8. Halevy, A., Franklin, M., Maier, D.: Principles of dataspace systems. In: PODS 2006 (2006)
9. Hasan, R., Sion, R., Winslett, M.: The case of the fake picasso: Preventing history forgery with secure provenance. In: FAST 2009 (2009)
10. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the Spread of Influence through a Social Network. In: ACM KDD 2003 (2003)
11. Kossinets, G., Kleinberg, J., Watts, D.: The Structure of Information Pathways in a Social Communication Network. In: ACM KDD 2008 (2008)
12. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. CACM 21(7), 558–565 (1978)
13. Lu, S., Ostrovsky, R., Sahai, A., Shacham, H., Waters, B.: Sequential Aggregate Signatures and Multisignatures Without Random Oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 465–485. Springer, Heidelberg (2006)
14. Lysyanskaya, A., Micali, S., Reyzin, L., Shacham, H.: Sequential Aggregate Signatures from Trapdoor Permutations. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 74–90. Springer, Heidelberg (2004)
15. Mattern, F.: Virtual time and global states of distributed systems. In: Workshop on Parallel and Distributed Algs. (1989)
16. Muniswamy-Reddy, K., Macko, P., Seltzer, M.: Making a Cloud Provenance-Aware. In: TaPP 2009 (2009)

17. Nemhauser, G., Wolsey, L., Fisher, M.: An analysis of the approximations for maximizing submodular set functions. Mathematical Programming 14, 265–294 (1978)
18. Reilly, C., Naughton, J.: Transparently Gathering Provenance with Provenance Aware Condor. In: TaPP 2009 (2009)
19. Sabefield, A., Myers, A.C.: Language-Based Information-Flow Security. In: IEEE JSAC (2003)
20. Spillane, R., Sears, R., Yalamanchili, C., Gaikwad, S., Chinni, M., Zadok, E.: Story Book: An Efficient Extensible Provenance Framework. In: TaPP 2009 (2009)
21. Xu, S., Sandhu, R., Bertino, E.: TIUPAM: A Framework for Trustworthiness-Centric Information Sharing. In: IFIPTM 2009 (2009)
22. Zhang, J., Chapman, A., Lefevre, K.: Do You Know Where Your Data's Been? — Tamper-Evident Database Provenance. In: SDM 2009 (2009)