Reputation Management in Privacy-Enhanced E-learning

Mohd Anwar & Jim Greer
ARIES Lab
Department of Computer Science
University of Saskatchewan
Saskatoon, Canada
mohd.anwar@usask.ca

Abstract

Both privacy and trust are essential elements of an effective e-learning environment. Privacy provides a personal space to a member of an e-learning community, while trust is a crucial enabler for meaningful and mutually beneficial interactions that build and sustain collaboration (e.g. collaborative learning). Identity management (in the form of various degree of anonymity) is one technology-based approach to protect privacy. However, privacy-enhancing identity management (PIM) impedes trustworthiness assessment of an actor. We propose a guarantor-mediated reputation management together with a context-based identity management scheme (e.g. long-term or transaction term pseudonymity) to assess trustworthiness of actors without requiring specific knowledge of their identity. As part of such a reputation management system, a model for reputation transfer among the multiple pseudo-identities is presented where we deal with reputation in different contexts without compromising identity and thereby facilitating trust while preserving privacy.

1. Introduction

Today's e-learning systems facilitate a variety of tasks related to learning: supporting different learning scenarios (e.g. self-study or guided learning), authoring of learning objects, tutoring, communication, evaluation, annotation, administration, etc. These tasks necessitate interactions between various human actors as well as cooperation and collaboration among themselves. Through interactions, large amounts of personally identifiable information (PII) are transmitted, collected, and processed that could reveal personal details of an actor (e.g. learner, teacher, administrator, etc.). On the other hand, trust plays a major role in developing and sustaining cooperation and collaboration and thereby in building a community. Therefore, privacy is a natural concern at the same time that trust is an important factor in an e-learning environment.

Reputation is a factor in any online community where trust is important. In the real world, trust is developed through day-to-day activities where everyone gets to see and know each other on a regular basis. By contrast, in most e-learning environments, the (possibly pseudonymous) users are strangers whose interactions are limited to mostly written communications (synchronous or asynchronous). Any private information is susceptible to misuse if it is shared with a stranger. Following from the notion that it is important to establish the trustworthiness of entities and not so important to know their real world identity, reputation can be used as a measure of trustworthiness of an actor's future behavior.

Reputation is a longitudinal social evaluation on a person's actions. A good reputation is a return on a long term investment of good behaviour. In an e-learning community, actors sometimes assume numerous pseudo-identities (e.g. student, tutor, instructor) to allow them to explore different aspects of their persona, interests or hobbies. Transaction pseudonymity (i.e. a pseudonym used for a transaction) and anonymity cannot be effectively used because they do not allow linkability between transactions as required when building trust. In this paper, we propose a reputation management, particularly, a reputation transfer model that would allow reputation transfer among multiple pseudo-identities (e.g. pseudonyms) without letting anyone associate these pseudo-identities. As a result, this model facilitates both privacy and trust.

In the reputation transfer model, a trusted public actor, namely a guarantor, generates reputation based on its and other known and unknown witnesses' observations about an actor. In this model, actors receive a copy of their reputation and are given an opportunity to contest any portion of their reputation. Upon request, a guarantor vouches for actors to their partners by producing a certificate of recommendation based on their reputation. Most importantly, an actor may request the guarantor to attach their context-specific reputation to any of their pseudo-identities and enjoy an appropriate level of trustworthiness in that relevant context. For example, a student with a pseudonym Bob may have earned a reputation of good helper in a discussion group for programming courses. When Bob assumes a new pseudonym as Alice and wants to continue helping other students, this actor's reputation as a good helper should be

transferred from the Bob identity to the Alice identity. In this way, our model facilitates trust but preserves privacy.

In this paper, we presented an algorithm for reputation transfer among pseudo-identities without allowing linkability between a reputation transferring and a reputation receiving entity. A rudimentary implementation of this model is done in a socket based multiple Client/Server architecture using Java. The generic Public Key Infrastructure (PKI) support used in this implementation is provided by Bouncy Castle. Upon receiving simultaneous requests and after proving non-repudiation of the pseudonymous entities, a guarantor administers a secure reputation transfer between the two parties.

2. The Idea

Privacy can be seen as a basic need (according to Maslow's hierarchy of needs [2]) or a fundamental human right "to be left alone" [1]. Privacy is important to the members of online communities including e-learning as it gives them autonomy in their activities and provides them personal spaces in a public domain. Borcea et al. point out that privacy requirements are obviously important for e-learning, since they establish an unbiased environment [6]. On the other hand, trust is a word that people constantly use to mean different things in different circumstances, and in different scenarios (e.g. trust between parties, trust in the underlying infrastructure, etc.). According to Handy, trust is "a confidence in someone's competence and his or her commitment to a goal" [9].

Trust can be seen as a complex predictor of an entity's future behavior based on past evidence. As we always ponder if we could trust someone with our valuables, it is also crucial to calculate the trustworthiness of an actor to decide what piece of information would be safe with whom and in what context. Building up mutual trust is important for every communicative context. If trust is not present in a relationship a large amount of energy is wasted in checking up on the other's commitments and on the quality of their work.

Collaboration is an important part of learning, whether it is in classroom settings or in virtual settings. Mason and Lefrere state that, in e-learning, common goals and mutual benefits are discerned and pursued through collaboration [10]. It minimizes duplication of effort and stimulates innovation. Allan and Lawless point out that on-line collaboration can cause stress, and this stress is linked to the dependency of the collaborators on each other, and the level of their mutual trust [8]. An effective collaboration whether synchronous (e.g. chat, conferencing) or asynchronous (email, blogs, threaded discussions) depends upon trust.

Privacy and trust are circularly related. Nickel and Schaumburg suggest that the level of self-disclosure depends on the user's trust in a given website, which in turn is influenced by the level of perceived privacy offered by the website [3]. A person's privacy is their capacity to control the conditions under which their identity information will be shared, whereas trust is another party's ability to use knowledge about an individual to determine whether that individual should be allowed to perform some action. Since trust reduces the perceived risks involved in revealing private information, it is a precondition for self-disclosure [4]. Privacy and trust complement each other, and together they can make for a more stable e-learning community.

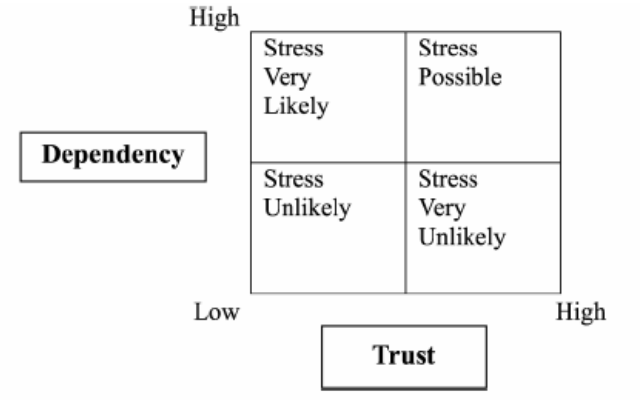

Figure 1. Stress due to online collaboration [8]

To provide an unambiguous description of how privacy is affected in an e-learning environment, we choose to define privacy in terms of identity. An identity is a dataset that holds information such as attributes (name, student number), traits (biometric information), and preferences (food choices, learning styles). An individual holds multiple partial identities in different contexts depending on the values of their context-relevant attributes (Figure 2). For example, graduate students hold multiple partial identities based on the roles they play: a student, a tutor, an instructor or a marker. In the context of teaching, one's student number is extraneous information whereas in the context of enrolling in a class, one's employee identification number is likely irrelevant. Therefore, we can say that a learner's or a teacher's privacy is their ability to limit the use of extraneous identifying information and to control the conditions under which their relevant identity information will be shared.

Since privacy is about protecting identity information, identity management appears to be a natural solution to privacy. Based on the amount of personal information disclosure, there are 3 approaches to identity management: anonymity (no identity information is disclosed),

pseudonymity (selective identity information is disclosed), or full identity (no restrictions to disclose identity information until it poses a threat). Anonymous actors in e-learning are strangers whose interactions are limited to certain selected written communications (synchronous or asynchronous). Although anonymity may ensure absolute privacy, it restricts trust because no longitudinal behaviour record can be associated with an anonymous actor. Since we don't live in an ideal world, full identity is a naïve scheme, and it totally disregards privacy.

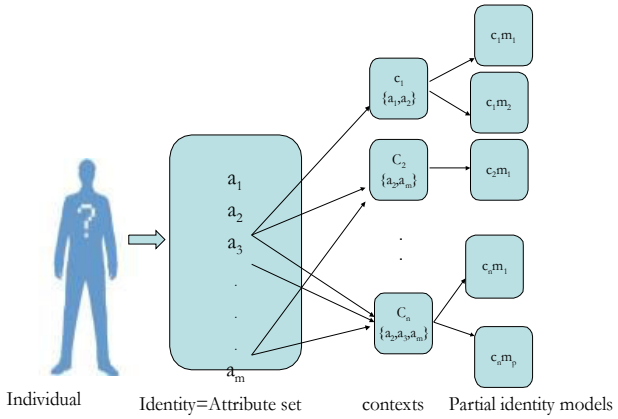

Figure 2. Context Sensitive Identity Model

In general, pseudonymity allows us to attach longitudinal behaviour records to an individual but it also throws a blanket of secrecy over the true identity and thereby helps protect privacy of the individual. A learner known as Student86 represents a partial identity of a particular individual. Pseudonymity supports reputation building as found in online auction services (e.g. eBay), discussion sites (e.g. Slashdot), and collaborative knowledge development sites (e.g. Wikipedia). Both good and bad reputation markers could be attached to the pseudonym based on observed actions. A pseudonymous user who has acquired a favorable reputation gains the trust of other reputable users. However, one of the problems with the pseudonym-based reputation is that it is a loss when a pseudonym must be thrown away (in case of identity-theft or slanderous attacks) and a new pseudonym has to be built from scratch [5]. Besides, when a pseudonymous actor joins a new community, they do not have any prior record to build up trust relationships with members of the new community. This problem can be addressed by allowing reputation transfer among pseudo-identities.

Although anonymity may ensure absolute privacy, it does not facilitate trust because an actor remains a stranger. There is no way to evaluate trustworthiness when an actor assumes anonymity. Johnson et al. note the following problems with anonymity [7]:

(a) it makes law enforcement difficult (tracking down and catching on-line law-breakers is difficult when their identity is unknown)

(b) it frees individuals to behave in socially undesirable and harmful ways (individuals seem to engage in behavior they wouldn't engage in if their identity were known);

(c) it diminishes the integrity of information since one can't be sure from whom information is coming, whether it has been altered on the way, etc.; and

(d) all three of the above contribute to an environment of diminished trust which is not conducive to certain uses of computer communication.

Yet if a favourable reputation provided by a trusted source could be associated with an anonymous user, many of the above problems are diminished. We claim that provided a pseudonymous chain of activity can be monitored, occasional uses of anonymity can be facilitated by having a trusted guarantor vouch for the context specific reputation of an actor using an anonymous identity and thereby effectively vouch for the actions of that anonymous actor.

Reputation management involves recording a person's actions and the opinions of others about those actions. In our model, a guarantor plays the role of a reputation manager by publishing and transferring reputation in order to allow others to make informed decisions about whether to trust about that person or not. An actor may choose to assume different pseudonyms for different purposes or at different times to: evade any disclosure of identity, regain privacy that was lost under some existing pseudonym, or disclose pseudo-identity appropriate in a context. For example, in a course discussion group, a shy student may want to be anonymous when asking some clarifications about an assignment whereas that student may want to be recognized as *BobTheWise* when helping other fellow students. However, by taking on a new pseudonym, the person loses their reputation which their past pseudonym has earned over a period of time. We propose a reputation-transfer mechanism to address this problem. In our model, we have the following four entities:

- Actor: An actor is a participant (e.g. student, tutor, instructor) in an e-learning environment, who takes part in various activities (e.g. chat, discussion) assuming anonymity or pseudonymity using their transaction-term or long-term pseudonyms.
- Reputation: In our model, reputation is the trustworthiness of an actor assessed over their past activities. For example, Alice may have worked in numerous collaborative course projects in the past. Based on her previous records, she could be trusted

as a hardworking participant. However her skills in programming assignments cannot be highly trusted.

- Guarantor: A guarantor is a public actor who is a trusted witness of the past activities of a pseudonymous actor. For example, since an instructor observes a student over a period of time, the instructor can serve as a guarantor of a student's reputation (e.g. good, bad, mediocre).
- KG: A trusted key generator that facilitates Public Key Infrastructure. This is a system component that will provide public/private key pair for the actors and the guarantor without knowing the purpose or usage of the key pairs.

In our reputation-transfer model, an actor picks a guarantor who would vouch for the actor and be credible in the community. The actor registers its pseudonym with the guarantor. The guarantor periodically evaluates the reputation of the actor based on their and other community members' observations. After each evaluation, a copy of the reputation is sent to the respective actor. The actor gets an opportunity to contest any misrepresentation of their reputation to the guarantor. The guarantor investigates the challenge and thereafter makes an appropriate adjustment to the reputation.

In the beginning, a key generator provides a set of keys to the guarantor. The guarantor publishes the public key so that any actor can send encrypted service request to the guarantor. An actor also acquires a key pair from the key generator and shares their public key with the guarantor so that a two-way encrypted communication may take place. Once public keys are exchanged, an actor makes a registration request to the guarantor by providing its identity profile (e.g. pseudonym) [step 2]. The guarantor generates a unique registration number for each of it clients and send it to them. This unique identifier is later used for authentication purpose.

The guarantor routinely generates reputations for its registered clients [step 3 & 4] and sends reputation to them. After a few iterations of step 5 & 6, a reputation certificate is finalized. Then the guarantor requests another key pair, namely the reputation key pair from the key generator. Using the reputation public key, RPub, the guarantor encrypts a reputation certificate and safeguards the public key so that no one else can tamper with the reputation certificate. The guarantor then sends the reputation private key, RPri and the encrypted reputation certificate to the actor [step 10]. The guarantor also generates a reputation digest, RD, which is an MD5 hash to uniquely identify each certificate and it can be used to check the integrity of the certificate.

An actor can use the private key to decrypt and view the reputation certificate. A reputation transfer is a two way process that has to be initiated by the transferring actor and followed by the receiving actor. The assumption is that both the transferring and receiving actors are just two pseudo-identities for one entity. The receiving actor also has to be registered with the guarantor. First, the transferring actor makes a transfer request to the guarantor providing the receiving actor's secret registration number and reputation digest that would authenticate the request of the transferring actor. Then the receiving actor makes a similar request by providing the transferring actor's registration number. Since both the actors know each others registration numbers that are provided to them through an encrypted communication, it is safe to assume that they are the same entity with multiple pseudo-identities.

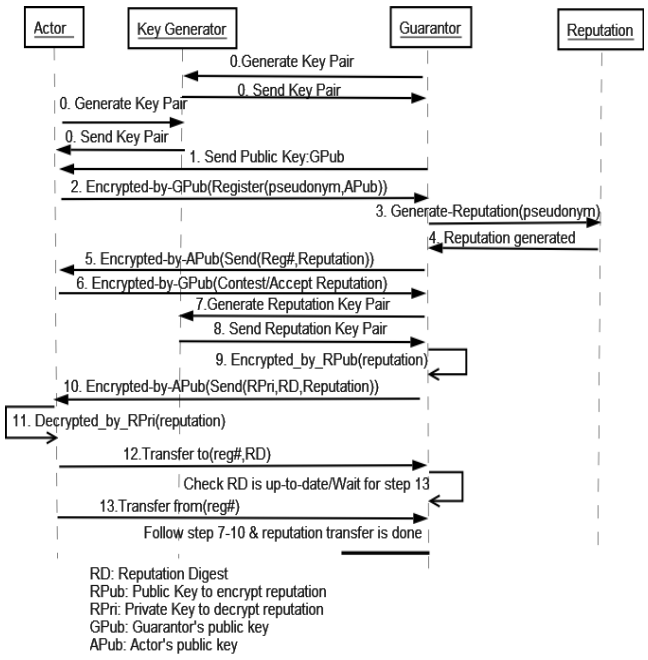

Figure 3. Reputation Transfer Model

The guarantor checks the reputation digest to make sure that a valid reputation transfer request is made. Upon verifying non-repudiation (a claimant is who they claim to be), the guarantor requests a new reputation key pair from the key generator. Then the guarantor encrypts transferring actor's reputation with the new public key and sends the corresponding private key to the receiving actor along with the encrypted reputation certificate [step 7-10]. As a result a reputation transfer takes place without letting other entities make a link between the transferring and receiving actors.

In summary, a pseudonymous actor can update the reputation of one pseudonym by transferring its reputation from another pseudonym. A guarantor vouches for an actor in two ways: i) responding to the queries about the actor, ii)

responding to the actor's reputation transfer request from one pseudonym to another. Since both the transferring and receiving actors are registered users of a guarantor, any bad acting can be traced and verified. All the communication between an actor and the guarantor takes place using each the other's public key. Moreover, the integrity of reputation can be checked using the reputation digest.

3. Implementation

Our implementation of this reputation transfer model has been done in the Java language using a Multiple Client/Server architecture. The Key Generator entity of the model is implemented using the RSA key pair generation algorithm provided by Bouncy Caste[1]. The implementation was first developed with JRE 1.4 and numerous methods of the java APIs failed to deliver their promised functionalities. Finally, the model was implemented using JRE 1.5 and java.security and javax.crypto APIs. The biggest challenge of this development was to build a string representation of an RSA key pair and converting the string representation back to the Java PublicKey and PrivateKey representation so that a key itself can be encrypted, serialized, decrypted, and de-serialized.

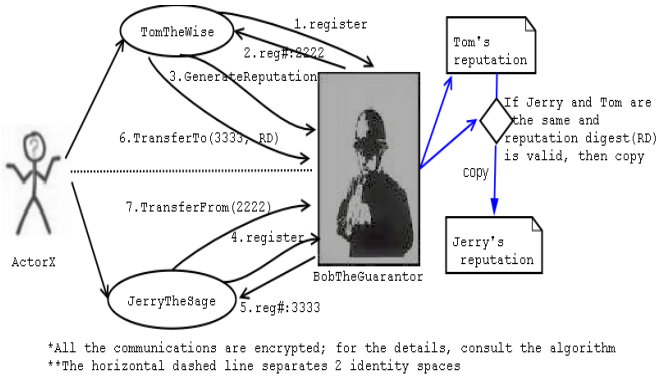

Figure 4. An implemented reputation transfer scenario

There are 4 different java classes in our implementation – Guarantor, Actor, Reputation, and Key Generator. The Guarantor is a server that can handle multi-threaded client requests. An actor is a client that takes part in encrypted communication, makes numerous service requests to the server (e.g. registration request, reputation transfer, etc.). Reputation represents reputation certificate that provides functionalities to generate and manipulate reputation. For this implementation, we represented reputation through a simple text file. The Key Generator provides the Public Key Infrastructure functionalities (e.g. generate public/private key, encrypt, decrypt, generate message digest, etc.).

The implementation is tested in the local host by creating a guarantor object and multiple actor objects. The guarantor and actors have separate repository to keep reputations. The reputation file is generated as a predefined text file. The secure transfer of reputation was the main focus of this implementation. The mechanism of dynamically generating reputation or contextualizing was not addressed. Figure 4 presents an implemented reputation transfer scenario where some actor (e.g. student, tutor, instructor), ActorX maintains two pseudonyms – *TomTheWise* and *JerryTheSage*. Both the pseudo-identities are registered with a guarantor, *BobTheGuarantor*.

These two pseudo-identities are not linkable by any third party, since the communications between the guarantor and each of the pseudonyms are encrypted using each other's unique public keys. Each of these pseudonyms receives a unique registration number (i.e. Tom: 2222 and Jerry: 3333). Reputation transfer request initiator, *TomTheWise* presents the registration number of the reputation receiver and the reputation digest (originally provided to it by the guarantor) to the guarantor. The guarantor awaits a similar request from the potential reputation receiver. The guarantor transfer *TomTheWise's* reputation to *JerryTheSage* only when they appear to be a registered user making simultaneous request and *TomTheWise* provides the most updated reputation digest. Empirical tests successfully show that the transferring aspect of the model works.

4. Conclusion

To realize the full potential of e-learning, legitimate privacy concerns must be identified and addressed. In this paper, an approach to address privacy protection and trust facilitation was explored. A mechanism to attach and remove reputation with a pseudonymous identity can help facilitate trust without the loss of privacy. For example, when Alice takes part in a discussion forum, her reputation as a friendly and knowledgeable user is all that matters to other participants. Reputation management can help attach a reputation marker to an anonymous or pseudonymous identity and thereby facilitate trust.

Since users assume many pseudonyms to represent many aspects of their identities, there is a need for reputation transfer among the pseudonyms without letting anyone link one pseudonym with the other. Privacy protection in reputation transfer requires that the transfer must occur without letting anyone observe such a transfer. We have developed and implemented a model by which this can be done with the aid of a trusted guarantor.

We are in the process of building a taxonomy for reputation and contextualizing reputation for the various

[1] www.bouncycastle.org

roles of the actors in an e-learning environment. We feel the need to build a reputation system to facilitate trust more effectively. This would involve the function of querying the reputation of a particular pseudonymous actor, generating reputation more effectively by fusing inputs from known and unknown witnesses along with the guarantor's first hand experience, and instructing an inquirer about what information pieces could be shared with the inquiring actor. For example if Bob maintains a reputation of a good mentor, it will not be an embarrassment to share with him one's learning weaknesses. When the system can help users to successfully identify potentially good helpers or collaborators, the system and users can work together to build an environment of trust.

5. References

1. I. Altman and M. Chemers, *Culture and Environment*. Stamford, CT.: Wadsworth Publishing Company, 1980.
2. A.H. Maslow, "Motivation and Personality," *Harper*, 1954.
3. J. Nickel and H. Schaumburg, "Electronic Privacy, Trust and Self-Disclosure in e-Recruitment," in *Extended Abstracts on Human Factors in Computing Systems*. CHI '04. Vienna, Austria: ACM Press, New York, NY, April 24 - 29, 2004, pp. 1231-1234.
4. J. L. Steel, "Interpersonal Correlates of Trust and Self-Disclosure," *Psychological Reports*, vol. 68, pp. 1319-1320, 1991.
5. D. Cvrcek and V. Matyas, "Pseudonymity in the Light of Evidence-Based Trust," in *Proc. of the 12th Workshop on Security Protocols*. Cambridge, UK: Springer-Verlag, April 2004.
6. K. Borcea, H. Donker, E. Franz, A. Pfitzmann, and H. Wahrig, "Towards Privacy-Aware Elearning," *Lecture Notes in Computer Science*, 2005.
7. D. G. Johnson and K. Miller, "Anonymity, Pseudonymity, or Inescapable Identity on the Net," in *Proceedings of the Ethics and Social Impact Component on Shaping Policy in the information Age*, T. Jewett and K. Miller, Eds. Washington, D.C., United States: ACM Press, New York, NY, May 10 - 12, 1998, pp. 37-38.
8. J. Allan and N. Lawless, "Stress Caused by Online Collaboration in e-Learning: A Developing Model," *Education & Training*, vol. 45, no. 8/9, pp. 564-72, 2003.
9. C. Handy, "Trust and the Virtual Organization," *Harvard Business Review*, vol. 73, no. 3, pp. 40-50, 1995.
10. J. Mason and P. Lefrere, "Trust, Collaboration, and Organisational Transformation," *International Journal of Training and Development*, vol. 7, no. 4, pp. 259-71, 2003.