Facilitating Trust in Privacy-Preserving E-Learning Environments

Mohd Anwar and Jim Greer

Abstract—This research explores a new model for facilitating trust in online e-learning activities. We begin by protecting the privacy of learners through identity management (IM), where personal information can be protected through some degree of participant anonymity or pseudonymity. In order to expect learners to trust other pseudonymous participants, we realize that a reliable mechanism is needed for managing participants' reputations and assuring that such reputations are legitimately obtained. Further, because participants can hold multiple identities or can adopt new pseudonymous personas, a reliable and trustworthy mechanism for reputation transfer (RT) from one persona to another is required. Such a reputation transfer model must preserve privacy and at the same time prevent linkability of learners' identities and personas. In this paper, we present a privacy-preserving reputation management (RM) system which allows secure transfer of reputation. A prototypical implementation of our reputation transfer protocol and the successful experimental deployment of our reputation management solution in an e-learning discussion forum serve as a proof of concept.

Index Terms—E-learning environments, trust, reputation, reputation management, identity management, privacy.

1 Introduction

Trust relationships among colearners are important for collaboration activities in e-learning environments. A trust relationship may need to be developed between two unknown learners who find themselves working together. The meaning of trust differs from one context to another. For example, when Bob seeks help from his colearner Alice on a math assignment, Bob may trust Alice's competence as well as her willingness to help. On the other hand, when Alice shares her frustration about the math course with her colearner Jill, Alice may trust that Jill will not disclose these feelings to the course instructor. Now if Jill wants to maintain a trust relationship with Alice, she will act according to Alice's expectation and not publicize Alice's feelings about that course to others. In aforementioned trust examples, one thing is common: reliance on the counterpart is central to trust. The paper deals with this aspect of trust. Therefore, to engage in and maintain a trust relationship, users need to do two things: 1) assess the trustworthiness of the counterpart, and 2) act according to the degree of trustworthiness expected of each other.

An expectation of trust has impact on and is influenced by the expectation of privacy. In a trust relationship, an individual's (e.g., Alice's) requirement for privacy may be diminished by expectations of trust (e.g., Alice's expectation of trust from Jill); or an individual may forfeit privacy to gain trust. Privacy risk is minimized when a trust-based disclosure decision is made. However, misplaced trust poses severe threats to privacy. Privacy and trust are equally desirable in a learning environment. Privacy promotes safe learning, while trust promotes collaboration and healthy competition, and thereby, knowledge dissemination.

Reputation appears to be one effective source for measuring trust. Reputation is a contextual and longitudinal social evaluation on a person's actions. In traditional face-to-face academic settings, trust is developed through day-to-day activities where everyone gets to see each other on a regular basis and thus grows to know one another. By contrast, an e-learning environment may bring the possibly pseudonymous users together through chat, message board, threaded discussion, online conferencing, email, blogs, etc. Research has shown that it is both unnecessary and privacy threatening to divulge a user's real identity in most online-learning related activities [1], [2]. Therefore, the trustworthiness of a pseudonymous entity needs to be estimated without the full knowledge of a real-world identity. We investigate how reputation can effectively be used as a predictor of a pseudonymous user's future behavior, which is actually a prediction of trustworthiness.

Identity management (IM) has been shown to offer an effective solution to privacy [3], particularly in the learning domains [1], [2]. In such a privacy-enhancing identity management scheme, each user participates in a context by assuming a context-specific partial identity and potentially many different identifiers or pseudonyms. Besides for privacy reason, learners may use multiple identities in open learning environments (e.g., OpenLearn) for different learning purposes. The trustworthiness of a pseudonymous user can be computed by measuring reputation on various aspects of trust pertinent to the underlying context. However, a proper reputation assessment is disrupted when an individual acts under multiple partial identities. Since the partial identities and pseudonyms offered by the

- *M. Anwar is with the School of Information Sciences, University of Pittsburgh, 135 N. Bellefield Ave., Pittsburgh, PA 15260. E-mail: manwar@pitt.edu.*
- *J. Greer is with the Department of Computer Science, University of Saskatchewan, 176 Thorvaldson, 110 Science Place, Saskatoon, SK S7N 5C9, Canada. E-mail: jim.greer@usask.ca.*

Manuscript received 16 Mar. 2010; revised 20 Dec. 2010; accepted 18 Apr. 2011; published online 5 May. 2011.
For information on obtaining reprints of this article, please send e-mail to: lt@computer.org, and reference IEEECS Log Number TLT-2010-03-0025.
Digital Object Identifier no. 10.1109/TLT.2011.23.

privacy-enhancing identity management solutions are not linkable, the complete assessment of reputation can easily be disrupted by switching and shedding of pseudonyms: reputation earned over a pseudonym is unusable with the shedding of that pseudonym or switching to another pseudonym.

This paper is about building a privacy-preserving reputation management (RM) system that performs two major reputation assessment tasks: 1) contextual (i.e., partial identity-based) reputation assessment and 2) reputation transfer (RT) across and merger among partial identities so as to support comprehensive assessment of reputation. The crux of privacy preservation lies in ensuring that task 2 maintains nonlinkability of partial identities. In other words, reputation transfer or merger process should not allow an observer to link partial identities involved in the process. As a result, the presented system measures trust while supporting an identity-management-based solution to privacy. Our contributions are as follows:

- *Relationship between identity management and reputation management.* We define reputation as a component of an identity, and consequently, we establish the relationship between identity management and reputation management.
- *Reputation assessment in learning environments.* We propose a contextual reputation assessment technique within a learning environment.
- *Supporting trust while preserving privacy.* We face the challenge of supporting trust while preserving privacy, and devise a privacy-preserving reputation management solution to address this challenge.
- *Implementation.* As a proof of concept, we implement and evaluate our solution in an online learning environment.

This paper is organized as follows. Section 2 describes trust and privacy issues apparent in learning environments. Section 3 discusses the relationships between identity and reputation management. In Section 4, we discuss supporting trust in learning environment through reputation assessment. Section 5 presents the challenges and techniques of supporting trust while preserving privacy. Section 6 presents our reputation management system. Section 7 describes related work and Section 8 concludes and describes future work.

2 Trust and Privacy Issues in E-Learning

Many assumptions about privacy in a traditional classroom do not apply to online learning—whether it is an online offering of a course or an online community of practice. A traditional classroom represents a close group where learners get to know each other. Yet some information is private including precise grades or confidential conversations. In contrast, e-learners become acquainted with one another by means of looking into each others' profiles. A profile is a self-constructed identity model presented under some label, popularly known as pseudonym. An e-learner may construct many such profiles depending how they want to present themselves in many different contexts. For example, an e-learner may want to position herself differently to her colearner peers than to her instructors, or might want to share more personal information with her project team than with the members of other project teams. Since each of the profiles consists of a different subset of personal information, they represent partial identities. To e-learners, privacy is about the autonomy of presenting themselves differently in different contexts.

In a traditional classroom, learners do not enjoy the same freedom of presenting themselves so differently in different contexts as do e-learners. In a traditional classroom, an observer can easily construct an identity model of another learner. As a result, unlike e-learning, a self-constructed identity model of a learner may not be well accepted by another learner in a traditional classroom setting. However, the lack of privacy is compensated by greater degree of trust in a traditional classroom. E-learners are often strangers whose interactions are limited to certain selected written communications (synchronous or asynchronous). Any private information is prone to misuse when shared with a stranger. It is also hard to engage in a trust relationship with a stranger. With a certain degree of familiarity, one can form an opinion about another person's trustworthiness. While in a traditional classroom, physical presence works as the guarantor of authenticity, in e-learning a learner needs to worry about the authenticity of their peers or instructors.

We observe the need for privacy and trust in the following popular learning activities:

- **Peer tutoring.** Peer tutoring is a widely practiced learning method. The main idea behind forming an online community of practice is peer tutoring. A learner needs to trust the competence and benevolence of their peer tutors. In a tutoring activity, a tutee shares her weakness with an expectation that her privacy will be preserved. A privacy breach may put the tutee in a disadvantageous or embarrassing situation. Privacy and trust concerns can easily demotivate learners from participating in peer-tutoring activities.
- **Peer reviewing.** Online portfolios are becoming increasingly common to engage learners in peer reviewing and assessment. These portfolios contain various sensitive information such as tests and test scores, projects, and self-reflections. Accessibility to an e-portfolio has privacy implications. Learners need to decide who they should trust with their e-portfolio items.
- **Learning object selection.** The selection of a suitable learning object requires making a trust decision of a sort. This trust may involve trusting (the reliability of) a learning object, trusting (the competence of) the author of the learning object, or trusting (the competence or authoritativeness of) the recommender of the learning object.
- **Collaboration.** Trust is essential to successful collaboration among learners [4], [5]. Online collaboration can cause stress depending on the level of the collaborators' mutual trust [6]. 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 works. In a learning environment, various key relationships of recommender-recommendation seeker, peer-peer, helper-helpee, and

mentor-mentee are formed based on mutual trust. Privacy concerns are inherent in a collaborative environment. The privacy concerns in collaborative systems originate from individuals' desire to control how one is perceived by another [7].

- **Group learning.** Group learning in the form of discussion forum, or reading group, offers valuable learning experience to learners. A group functions well when each member trusts each other and respects each other's privacy. An online learning system should facilitate a trust- and privacy-preserving learning environment.
- **Evaluation.** Confidentiality is very important in the learner assessment and evaluation process. Sometimes, learners experience various biases such as gender, ethnic, or connectedness (more connected to the evaluator). Biases in learner evaluation can be prevented through privacy-preserving techniques [8]. In a trust relationship, learners' confidence can grow regarding the fairness of evaluation.
- **Role playing.** Role playing is an effective technique for exploring complex social issues in certain courses (such as Sociology). Safety is an essential condition for authentic role playing. When a learner plays a controversial role, the learner may run the risk of being stigmatized or feel embarrassed. For example, when talking in favor of same-sex marriage, a learner may fear to be ridiculed. Learners' safety can be assured through trusting and privacy preserving learning environments.
- **Personalization.** Personalization of learning objects increases the motivation and interest of learners [9]. As a result, in recent time, we have witnessed an increasing volume of research and development efforts to offer personalized e-learning. Trust has been identified as a prerequisite [10] and a consequence of good personalization practice [11]. Anwar et al. define key characters of an e-learning environment that offers personalization together with trust and privacy [1].

3 Relationships between Identity Management and Reputation Management

An identity is a representation of an individual through a data set that holds information such as attributes (e.g., name, student number), traits (e.g., biometric information), and preferences (e.g., food choices, learning styles) [1]. A partial identity is a context-dependent identity model which is often published through user profiles. Each partial identity can be presented with many different identifiers or pseudonyms. An individual's behavior is manifested by a set of actions (or interactions) that the individual performs.

When an observer monitors someone's behavior with full knowledge of their identity, the person being monitored does not enjoy any privacy. On the other hand, when behavior is observed while the identity of the person being observed is not known (e.g., in the case of anonymous behavior), the person being observed enjoys privacy. In the former case, the observer can easily attribute some characteristics to the person being observed. In the latter case, the observer can still monitor a stranger. However

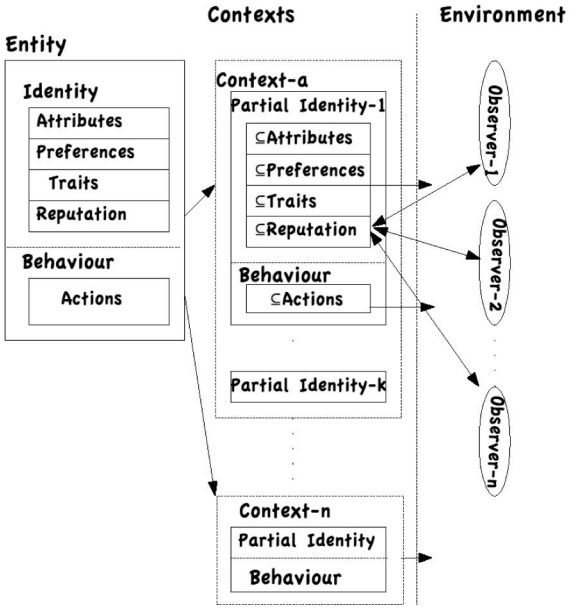

Fig. 1. A contextual notion of identity and behavior.

since the observer cannot identify the person being observed, the stranger enjoys a degree of privacy. Even though rigorous analysis of behavior may reveal the real-world identity of a person, without identity information one cannot make high probability association between identity and behavior of the person. With similar motivation, privacy models of k-anonymity [12] and l-diversity [13] make identification harder in released person-specific records. Therefore, we separate the data set representing a person into two proper subsets: identity and behavior. For example, when seeking help, Bob may only know Alice's identity. Or, Bob may have watched Alice's behavior without knowing her identity.

Even though identity and behavior are separable, a person's identity attributes (or partial identity attributes) may include information about reputation earned over their behavior (cf. Fig. 1). An advantage of carrying reputation with identity is that it allows an individual to establish a trust relationship fairly easily. Separation of identity from behavior allows us to observe someone's behavior without compromising their privacy. Since reputation is an evaluation on one's behavior, we argue that, the longitudinal study of just the behavior part of a person could sufficiently assess reputation of the person in a given context. Essentially, such a longitudinal study would require classifying behaviors by contexts, and for each context, accumulating observers' ratings of the suitability of such behaviors.

Since action should not be judged out of context, reputation is contextual. For example, a graduate student in a researcher role may not carry as prominent a reputation as he might in a tutor role. Since partial identities represent various aspects of one's projected self, each partial identity can draw a contextual boundary of an individual's actions, and therefore, each partial identity can serve as a context for reputation as well. Therefore, users may need to manage reputation that stems from actions taken under their respective partial identities.

The primary goal of identity management is to achieve information parsimony (and thereby privacy) by partitioning a user's identity into multiple partial identities according to their participations in various communicative contexts (e.g., my peer helpers need not know my class standing). We take the view that one of the challenges that identity management seeks to address is impression management (i.e., desire to be perceived by others in different ways in different contexts) [14], one of the important purposes of privacy preservation [15], [16]. In different contexts, users need to convey different impressions in accordance with their needs. In our running example, Alice may want to convey a different impression to Bob (from whom she is seeking help) than what she might convey to her confidant Jill. Conveying a certain impression may also require conveying a certain reputation. For example, Bob has to maintain and convey a reputation of high competence to convey the impression of a capable potential helper. Therefore, proper impression management can be supported through incorporating reputation management within identity management.

4 Supporting Trust in Learning Environments

Trust is contextual, and trustworthiness (measured by reputation) is assessed against an identity. For example, "Bob identity" may be trusted for his math competence, however, may not be trusted for his benevolence toward his peers as a math helper. We propose that **user-to-user trust**, during collaborative learning activities, be realized in two forms: **trust about a purpose** and **trust in a partner** (partner's identity) for which the partner's trustworthiness needs to be assessed.

- **Trust about purpose.** In e-learning, each context explicitly or implicitly manifests some purpose for its participants. For example, a math discussion forum context may have a purpose of offering peer tutoring in math. Within the math forum context, there could be more granular contexts like an Algebra thread or Calculus thread for the purpose of peer tutoring the respective topics. This form of trust is based on the expectation from the purpose of a context. For example, Alice may highly trust the Math Forum to find an effective helper in Calculus.
- **Trust in partner.** This form of trust considers the trustworthiness of a partner in a given context. For example, in a Calculus course, Alice may be considered as a trusted peer helper. Trust in partners may need further consideration of the roles of, and relationships with, the transacting partners. Some roles convey more trust than others. For example, an instructor role may convey a higher degree of trust. However, not all instructors are equally trusted by learners. A learner may trust one instructor over another based on their perceived relationship or reputation.

To facilitate assessment of accurate reputation, a system is needed that would: be able to prove itself unbiased and trustworthy, be able to judge individuals' behavior in light of context, recency, completeness, etc., allow individuals to contest or update their reputation, and help individuals manage their reputation across their partial identities. To this end, this paper presents a guarantor-mediated reputation management system, where the guarantor plays the role of a judge (who possesses above-mentioned qualities) with automated tool support for reputation management.

One important challenge for establishing reputation for a pseudonymous learner is foreseeable: it is a loss when a partial identity needs to be forsaken (e.g., in case of identity theft or slanderous attacks) and a new partial identity has to be built from scratch [4], or when a learner wants to have multiple pseudonyms for the same role (e.g., role-relationship pseudonym [17]). Besides, when a pseudonymous learner joins a new community of learners, they do not have any prior record from which they can build up trust relationships with members of the new community. **This problem can be addressed by allowing reputation transfer across partial identities.**

Though anonymity does not support building of reputation, sometimes a pseudonymous actor needs to act anonymously (e.g., doing peer evaluation, reviewing paper of a colearner). Yet if a favorable reputation provided by a trusted source could be associated with an anonymous user, the user could enjoy appropriate credibility. For example, because of negative bias, a specific editor may never pick Bob as a reviewer of a journal. With anonymity, a high competence score associated with Bob's anonymous reviewing may attract the same journal editor to want to work with him. If 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.

E-learning systems are different from many other online communities in that learners typically have more trust in the system and have long working relationships with one system [1]. As a result, the system can play the role of an acceptable reputation guarantor. With the aid of automated privacy-enhanced reputation management tools (e.g., reputation evaluation, reputation transfer/merger), instructors in a traditional learning setting or an elected senior member of a community of practice, can also play the roles of guarantors and adjudicators of learners' reputation. Since an instructor in a class or a senior member in a community of practice is accountable for the well being of their respective communities, their guarantor roles, along with automated reputation management tool support, will empower them to carry out their responsibilities.

In a high-risk or low-trust environment, we may need to require multiple guarantors to work together to address any bad acting. We realize that users may be able to defeat our reputation management system by colluding with the guarantor(s). However, this is an inherent problem in or a limitation of any reputation system, in general any system that uses any type of third party information. One way to address the collusion problem is to ensure the credibility of any trusted third party involved [18]. Since our guarantor-mediated reputation system is situated in a learning environment, we assume that none could be more credible to learners than an instructor. Therefore, an instructor playing a guarantor role is unlikely to collude with some learners to game the reputation system. In a similar way,

senior members of a community of practice or moderators of a discussion forum are expected to play the role of guarantor. In this work, we do not consider any threat model where these guarantors are involved since we perceive them as facilitators of these communities.

5 Challenges and Techniques to Support Trust

Trust can be seen as a complex predictor of an entity's future behavior based on past behavior. In our daily life, we always deliberate whether we could trust someone with something. Likewise, it is also crucial to calculate the trustworthiness of a user to decide what piece of information would be safe with whom and in what context. People are not likely to reveal confidential information about themselves to an untrustworthy party.

Trust plays a major role in reducing privacy concerns. If the evidence is provided to the users that the data they disclose will be treated as defined, then this can potentially enhance trust of users in a data processing environment of the service providers (SPs). For example, the learner needs assurance that the service provider will only use his/her private information, such as name, address, credit card details, preferences, and learning behaviors in a manner expressed in the policy provided for the e-learning system users.

In order to maintain privacy, a user faces the biggest challenge of making a trust-based decision at the time of sharing personal information. In a well-understood context, a user can relatively easily understand privacy implications of trusting another user (e.g., disclosing their identity to another user). For example, a learner can have different privacy expectation from a peer-tutoring context than from an evaluation context. Therefore, contexts draw boundaries of trust and privacy. A pseudonymous user, who has acquired a favorable reputation, gains the trust of other users.

The solution to privacy through maintaining partial identities in different contexts can be less appealing due to the fact that reputation earned over a partial identity is unusable across other partial identities. Since the partial identities and pseudonyms offered by the identity management solutions are not linkable, the complete assessment of reputation can easily be disrupted by switching and shedding of pseudonyms: reputation earned over a pseudonym is unusable with the shedding or switching of that pseudonym. Although a mechanism for reputation transfer across partial identities of an entity may address this problem, it may pose the threat of linkability to privacy: by observing a reputation transfer, an observer may be able to link the transferor identity with the transferee identity. Therefore, reputation aggregations/transfers across multiple partial identities have to happen unobservably and securely. Such a transfer has to restrict any undue advantage for bad acting (e.g., recurring merger of a bad reputation with a good reputation).

To facilitate reputation-based trust (i.e., trust is associated with the reputation of an actor) in the online domain, we need to support complete assessment of reputation across partial identities. As a result, a secure and privacy-preserving reputation transfer model is developed to transfer/merge reputation across contextual partial identities.

Assessment of reputation across partial identities in a privacy-preserving manner involves 1) assessing reputation from behavior analysis of a user under each of their partial identities, which we term *partial reputation* and 2) transferring/merging reputation of a user across their partial identities in similar contexts, while preserving nonlinkability of these partial identities.

In the RT model, a pseudonymous user can update the reputation of one partial identity by transferring its reputation from another partial identity, effectively merging reputation across partial identities. Though anonymity does not support building of reputation, sometimes a pseudonymous user needs to act anonymously. For example, in a course discussion group, a shy student, Bob may want to be anonymous when conversing with peers about some research ideas, whereas that same Bob may want to be recognized as *BobTheHelper* when helping peers. Yet if a favorable reputation provided by a trusted source could be associated with an anonymous user, the user could enjoy appropriate credibility. For example, despite anonymity, a high competence score associated with Bob's anonymous identity may attract other students to converse with him.

In the RT model, a guarantor (an appropriate public trustee) vouches for a pseudonymous user in two ways: 1) responding to the queries about the user's reputation, and 2) responding to the user's reputation transfer request from one partial identity to another. The reputation is generated as a reputation point average (RPA) on a 0 to 5 scale, 0 representing unknown rating or lack of input and 5 representing the best rating. Depending on subjective judgement, a user may consider any lower arbitrary value in the 0-5 scale as bad rating. The guarantor assesses reputation for its registrants (i.e., pseudonymous users) by aggregating ratings submitted by their transacting partners.

To provide a solid foundation for the empirical study of trust, Schoorman et al. [19] observe three characteristics of a trustee appearing often in the literature: ability, benevolence, and integrity. For learners, reputation is a mechanism for ascertaining the trustworthiness of colearners, analogous to those in eBay (e.g., integrity of the seller) and to those in Wikipedia (e.g., authority/competence of the contributor). Therefore, using trust as a scale to find a suitable recommender, peer, helper, and mentor, a learner should be able to find out the status of each participant in an e-learning environment: Is someone really the expert or well-intentioned peer that they claim to be? One can also decide whether trust can replace the need for privacy: Can one confide in their peers? Most importantly, in assessing reputation of a learner, their behavior has to be evaluated (when the knowledge of their identity is inconsequential) by their transacting partners.

We consider reputation evaluation as a process of aggregating observers' opinions on the performance of individuals against the expectations of their roles in similar contexts.

6 Reputation Management

Due to the observed relationship of identity and reputation management (see Section 3), we offer a standard mechanism for reputation assessment across partial identities. As a result, reputation management involves reputation assessment and reputation transfer or merger. We have deployed

our reputation management system in the $iHelp$[1] (see [20] for iHelp architecture) Discussion Forum, which acts as an online forum for students at the University of Saskatchewan to converse asynchronously with one another, with subject matter experts, and with their instructors. Based on the requirement of the course, a discussant can have as many as three types of partial identities: user-level, context (or category)-level, and role-level. Context and roles are defined by the course designers. Both context- and role-level partial identity types can further be categorized into group-scope and individual-scope. Based on their group memberships, group-scope identities are created for discussants by the system. The system provides user interface facilities for creating individual-scope partial identities. The system also provides a user-level identity based on the true identity of a discussant. The system allows discussants to create as many additional user-level partial identities as they like. The remainder of this section explains how the reputation assessment and reputation transfer (across partial identities) components have been implemented and evaluated.

6.1 Reputation Assessment

We implemented a mechanism for reputation assessment for an actor along the dimensions of competence, benevolence, and integrity. What a particular dimension represents in a given context is specified through a list of features. A list of dimension-relevant features is presented to a rater to capture the rater's opinion along the respective trust dimension. Each feature carries certain weight (strength), according to which it contributes to the relevant dimension. In the $iHelp$ implementation, anyone who is authorized to read a posting (excluding the poster) is eligible to rate a posting. Each rating contributes to the overall reputation of the poster. Finally, the weighted sum of all the relevant ratings is averaged to calculate reputation along a respective dimension. The three dimensions of reputation are calculated on the following features: insightful, timely, informative, well-written, constructive, and relevant. These features are qualities of learners desirable in learning activities. Our contention is that it will help participants to articulate on the postings (i.e., poster's behavior), not on the posters (i.e., poster's identity).

This feature-based assessment of reputation can be employed for personalized reputation assessment. A user may define a dimension of trust on their own by choosing a list of features and/or their respective weights for measuring a specific dimension of reputation. Given that $Features_d$ is a set of features chosen for a dimension d, the system can compute the dimension d of trust using the formula

$$R_{d \in \{Competence,\ Benevolence,\ Integrity\}} = \frac{\sum_{f \in Features_d} \sum Rating_f \times Weight_f}{number\text{-}of\text{-}observations}.$$

We have classified these features based on their expected impacts (i.e., real weights in the range [0, 1]) on determining the level of competence, benevolence, and integrity of a poster in an e-learning discussions context. In our implemented system, weights on features have been empirically assigned. For example, in determining competence of a poster, an insightful or an informative posting has been assigned twice as much impact as a well-written posting. Reputation of an identity for a specific dimension (e.g., competence) is estimated by averaging the weighted sum of relevant features. For example, in calculating competence, the following formula is used:

$$\begin{aligned} R_{competence} &= \Big(\sum Rating_{insightful} \times Weight_{insightful} \\ &\quad + \sum Rating_{informative} \times Weight_{informative} \\ &\quad + \sum Rating_{well\text{-}written} \times Weight_{well\text{-}written} \Big) \\ &\quad /number\text{-}of\text{-}observations. \end{aligned}$$

In $iHelp$ Discussion Forum, a poster's reputation is contextualized by their group identities or individual partial identities.

6.2 Reputation Transfer across Pseudonyms

With the persistent use of a pseudonym (a partial identity), the attribution of reputation markers to the pseudonym takes place. A pseudonymous user cannot, on their own, transfer or merge reputation across their multiple pseudonyms, yet such ability is highly desirable. Let us consider scenarios from an e-learning discussion forum where users can participate using individual identity or group identity. With a group identity, all the members of the group are represented. For example, all the students in peer-helper role can be grouped into one identity with a pseudonym "peer-helper." Ratings on a posting made by a user using a group identity should contribute to the reputation of that group identity as well as to the reputation of the group member's (poster's) individual identities. This is a trivial example of a need for reputation transfer from a group identity to an individual identity.

Let us consider another scenario from the e-learning context, where an identity expires and reputation from the expired identity needs to be transferred to an existing identity. Anwar and Greer observed that contexts in the e-learning domain are hierarchical and proposed the notion of contextual identity [21], [17]. As a context expires, the reputation of an identity under that context may need to be propagated back to its parent context resulting in a backward propagation of reputation (reputation transfer) from the innermost context to the outermost context. For example, in the outermost context, a person becomes a student for the purpose of attaining a degree. In the innermost context the student is evaluated in an assignment of a course, the student's mark in that assignment is propagated to its parent context of the course and the course grade is eventually propagated backwards to the outermost context contributing to achieving their degree.

There is another variation of reputation transfer, which we call reputation merger. It is a process where reputation of two partial identities (involved in reputation merger) are updated by each other or aggregated to reputation of a new partial identity. A reputation merger can be viewed as a two-way reputation transfer between two identities or two

1. http://ihelp.usask.ca.

one-way transfer between each of the identities and a new third identity, which is the case when two partial identities are merged into a new partial identity. We anticipate two scenarios of transfer or merger: 1) a user requests transfer or merger and the system obliges with the mediation of a guarantor, 2) the system automatically performs transfer or merger based on the decision of the guarantor. In our system, reputation earned on any partial identity is merged with reputation of all other partial identities of a user within the same context.

Unfortunately, a privacy concern is inherent in reputation transfer. Observing a transfer of reputation from one identity to another, an observer can easily link two identities involved in the reputation transfer, failing an identity-management-based solution [22] to privacy. Therefore, a pseudonymous actor needs a privacy-preserving mechanism for the transfer or merger of their reputation across their multiple pseudonyms. Such a mechanism has three objectives: 1) provide cryptographically secure reputation transfer protocol, 2) restrict Bad Acting, and 3) restrict link-ability of partial identities.

6.2.1 Secure Reputation Transfer Protocol

In the secure reputation-transfer protocol, a user registers its pseudonym with a guarantor who would vouch for the user and be credible in the community. The guarantor periodically evaluates the reputation of the user based on their and other community members' observations. After each evaluation, a copy of the reputation is sent to the respective user. The user 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 RT model, there are the following four entities:

- Actor. An actor is a user (e.g., student, tutor, instructor in an e-learning environment), who takes part in various activities (e.g., chat, discussion) assuming their various contextual partial identities.
- Reputation. Reputation measures trustworthiness of a user 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 may not be highly trusted.
- Guarantor. A guarantor is a "public" user who is a trusted witness of the past activities of a pseudonymous user. For example, since an instructor observes a student over a period of time, the instructor can serve as a guarantor of a students reputation in a traditional e-learning context.
- Key Generator (KG). A trusted key generator that facilitates Public Key Infrastructure. This is a system component that will provide public/private key pairs to the users and the guarantor without knowing the purpose or usage of the key pairs. The steps of reputation transfer model are detailed in the table found in Appendix, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TLT.2011.23.

In summary, in the RT model (see the figure found in Appendix available in the online supplementary material), a pseudonymous user can update the reputation of one pseudonym by transferring its reputation from another pseudonym. A guarantor vouches for a user in two ways: 1) responding to the queries about the user, and 2) responding to the user's reputation transfer request from one pseudonym to another.

6.3 A Proof-of-Concept Implementation

The prototypical system incorporating the RT model has been implemented through a client (for users) and a multithreaded server (for guarantor) suite written in Java language. The Key Generator entity of the secure reputation transfer protocol is implemented using the RSA key pair generation algorithm provided by Bouncy Castle. The model was implemented using JRE 1.5 and *java.security* and *javax.crypto* APIs. The system manages reputation for three different generic roles that are present in an e-learning community: helper, peer, and lurker. The system allows a user to perform any of the following four tasks: **register** (i.e., register a pseudonym with a guarantor), **evaluate** (i.e., rate a user), **transfer** (e.g., transfer/merge reputation across pseudonyms), and **query** (e.g., query reputation of a pseudonymous user).

- Register. A user registers with a guarantor entity of the system. The communication between a user and a guarantor is cryptographically secure. At the time of registration, a user provides their pseudonym (partial identity) and context (reputation context for which the user wants to be evaluated for reputation). Upon registration, the user receives two pieces of information to be kept secret: 128-bit unique registration number and a digest (MD5 hash) for reputation. For any change in reputation, the system generates a new digest.
- Evaluate. Any user can evaluate others (i.e., pseudonyms) against the features specific to the role of the user being evaluated on a scale of 0 to 5. Additionally, an evaluator may write comments in support of their evaluation.
- Transfer. Reputation transfer is a two way process that has to be carried out by both the pseudonyms—transferor and transferee. First, the transferor and then the transferee authenticate themselves by providing their respective contexts, registration numbers, and reputation digests. Reputation from one pseudonym can be transferred to a new pseudonym, or reputation of one pseudonym can be merged with the reputation of the other pseudonym. Reputation merge takes place incrementally by combining each rating transaction of a pseudonym one-by-one to the aggregate rating of the other pseudonym and vice versa. Though the end result of the merge is two pseudonyms with the same reputation, their reputations are different on each time step of the merge. There is a little time delay induced in between each step to give the impression that there could have been another transaction (evaluation) taking place.

Fig. 2. A screen shot of reputation Window in iHelp Discussion Forum.

- Query. A user may query reputation about another user (corresponding pseudonym). A reputation summary, which is an aggregation of collected ratings against context-relevant features, is displayed in the following format: *Feature* | *Score* | #*Trans* (i.e., number-of-ratings).”

6.4 Evaluation

This section reports on two studies: 1) value of reputation management system in e-learning and 2) validating the implementation of the RT model. The study 2 [23] was designed to see whether the system facilitates secure reputation transfer/merge across multiple pseudonyms.

6.4.1 Value of Reputation Management System

Methodology. The system was used in an experiment to support online course discussions of 35 students (19 female and 16 male) in an intensive six-week undergraduate course on Introduction to Sociology. The study was done in two phases: 1) in the first phase, the class made 173 postings using the original version of iHelp Discussions (without reputation management system), and 2) in the next phase, they made 302 postings using a version of iHelp Discussions with reputation management system features.

The system allowed the participants to create multiple role- and relationship-level identities, provided awareness support of contexts and identities, and enabled them to rate others and query others’ as well as their own identity-specific reputation (a screen shot of reputation Window in iHelp Discussion Forum is shown in Fig. 2). In each phase, the participants (students and the instructor) discussed topics under 11 contexts (chosen by the instructor of the course as per the course objectives), each addressing 11 different social and behavioral questions. The goal of the discussion is to collaboratively find answers to different social phenomena (e.g., Dating Older man, Spitting on the Ground, Eye-contact on elevator, etc.). Prior to each phase of the study, users were trained to use the system. At the end of the second phase, 25 participants (of the 35 who used the system) took a post-use online survey to share their use experience and their attitudes toward reputation-based trust.

Results. The usage data reveal that every participant received reputation ratings on their posts and that 43 percent of the participants checked their own or others’ reputation. On an average, each participant received 12.5 ratings. 31 percent of the participants consulted self-reputation. We realize that the need for reputation or trust in the study is not as critical as it is in an online setting where there is no bodily presence to act as a trust guarantor. Since the participants of this study were classmates, they were already involved in trust relationships. However, it was observed that those who cared about trust measures (based on the survey) used the trust and reputation features of the system more extensively. The postuse survey reveals that 28 percent of learners used the system to identify trustworthy peers. 36 percent of learners valued postings based on posters’ reputation while 40 percent found that reputation management system helped them identify trustworthy postings (see Table 1 for details).

Discussions. In this study the guarantor role is automated by the system. The system transferred a participant’s reputation earned using a group identity (i.e., while a group identity is used to make a posting) to all of her individual

TABLE 1
User Survey Response

item	%-of users
System Helped Identifying Trustworthy Peers	28%
Valued Postings Based on Posters’ Reputation	36%
System Helped Me Identify Trustworthy Posting	40%
System Facilitates Trust	60%
Replied More Often to Posters with Good Reputation	28%
Paid More Attention to Posters with Good Reputation	36%
Rated Postings with a Purpose to Reward/Discipline	28%
More Open when Replying to Posters with Good Reputation	28%

partial identities within the same context. 22 percent of postings (66 of 302 postings) are made using group identities. Also, reputation is transferred among partial identities within the same context. Even though 43 percent of users (lower than our expectation) were interested in seeking out reputation information, every user was interested in managing their identities—switching identities in different contexts. They engaged in this identity switching activity because they felt that identity linkability was not going to be a problem—that is, they implicitly trusted the security of the reputation management system. Perhaps those who were seeking out more reputation information were indeed checking up on how well the reputation mechanism preserved their privacy. We plan to conduct future study in an environment, where the need for reputation or trust is naturally higher so that we can fully understand the impact of our system.

6.4.2 Validating RT Model

Methodology. For validating the RT model, the system was initialized to generate multiple instances of four types of events (reputation evaluation request, reputation transfer request, reputation merge request, and null requests) in some random order for n pseudonyms representing m actors. At multiple time steps during the simulation, the system (the component representing the guarantor) was queried for the latest reputation of each of the $n \times m$ registered pseudonyms and the query results are logged. A version of this simulation was run for $n = 4, m = 2,$ and reputation update actions were logged accordingly. These logs were then provided to a security attack-defense expert to attempt to deduce types of events might have occurred based on an analysis of the reputation score patterns over various time steps. The expert was also asked to see whether he could distinguish among or determine instances of reputation transfer, reputation merge, and normal updates of reputation ratings.

Results. The simulation performed three transfers and seven merges of reputations across four pseudonyms of two actors. Although the data set was relatively small, the expert could not make any definitive conclusions that would identify which pseudonyms corresponded to the same actor. Our expert suspected that four mergers or transfers of reputation occurred. The one merger hypothesis in which the expert was most confident was totally incorrect. Two of our expert's suspected mergers or transfers actually did correspond to real mergers or transfers, but the expert entirely missed eight of the merger/transfer events. Our expert correctly had a suspicion that one transfer and one merger (of the 10) had occurred, but he could not be sure. Out of these two correct hypotheses, the expert could not confirm conclusively about any of the mergers or transfers.

We could say that these correct guesses are no more than random luck. With an increase in the number of actors or pseudonyms, it becomes even harder to guess about any reputation transfer or merge. Therefore, we could say that our system supports reputation transfer with privacy preservation.

6.5 Restricting Bad Acting in Reputation Transfer

The RT model provides mechanisms for restricting bad action in reputation transfer.

- The integrity of reputation can be checked using the reputation digest, a 128-bit "fingerprint" of reputation information generated through the calculation of MD5 hash.
- Since both the transferring and receiving pseudonyms are registered to the guarantor, any bad acting can be traced and verified by the guarantor.
- To restrict the taking of undue advantage from recurring merger of a bad reputation with a good reputation, a history of already merged ratings is kept and compared before entertaining a new merge request.
- The model also supports rollback of reputation to recover from bad acting.

6.6 Restricting Link-Ability of Partial Identities

Since linking of partial identities results in unintended disclosure defeating the purpose of partial identities, the transfer of reputation among the pseudonyms or update of reputation because of new ratings has to happen without letting anyone link one pseudonym with the other. Privacy protection in reputation transfer further requires that the transfer must occur without letting anyone recognize such a transfer. In the RT model, non-observable and nonlinkable reputation transfer is done by means of following techniques:

- Use of public key infrastructure ensures secure reputation transfer channel so that an observer cannot snoop a reputation transfer or identify two pseudonyms involved in the process of a reputation transfer.
- One pseudonym's reputation (i.e., aggregated ratings) is incremented one-by-one by each rating transaction of the other pseudonym and vice versa allowing longitudinal increase or decrease in reputation to make transfer indistinguishable from reputation update by a new rating.
- A random time delay is induced between each of the increments to make reputation transfer indistinguishable from reputation update by a new rating, which may not happen in a continuous succession of a short burst.
- A time delay proportional to the amount of activities takes place in the system is induced between updates of reputation so that multiple partial identities of an individual are not linkable because of one reputation update triggering changes of reputation of multiple pseudonyms.

While our approach offers mechanisms for restricting linkability of partial identities, the limitation of our approach is that if an attacker continuously changes the ratings she assigns to various identities and observes the results for a long time, then the attacker might be able to link identities. However, unlike a financial institution, stakes of doing so is low in a learning environment. Furthermore, we believe that the guarantor can address these attacks through routine auditing and proper mediation.

7 Related Work

Trust issues on the web have been around since the inception of the web. Trust is a word that people constantly use to mean different things in different circumstances. For example, "confidence in someone's competence and his or her commitment to a goal" [24] or "the choice to expose oneself to a risk toward one's counterpart, in the expectation that the counterpart will not disappoint such expectation" [25]. Our work is motivated by Handy [24]. In the literature, trust is identified in different forms relating to: whether access is being provided to the trustor's resources, the trustee is providing a service, trust concerns authentication, or trust is being delegated [26]. Even though all the stated forms of trust may take place in e-learning, our work mainly targets on user-to-user trust that relates to the trustee providing services. For example, in peer-help scenario, a learner is providing help to another learner. Learner-to-learner trust relationships can be used to address many different issues in learning environments. For example, Carchiolo et al. exploited trust relationships among peers to select suitable learning resources [27].

Policies and Reputation are two common ways of determining trust [28]. Policy-based trust approaches are widely used in security and access control. Our work integrates reputation (reputation is calculated on three dimensions) with policies (guarantor vouches for credentials based on reputation) in determining trust.

7.1 Trust and Privacy

Trust and privacy are interrelated constructs—disclosure of personal information depends on trust [11]. Since trust reduces the perceived risks involved in revealing private information, it is a precondition for self-disclosure [29]. On the other hand, trust invokes the threat of privacy violation, identity theft, and threat to personal reputation [30]. In policy-based trust, privacy loss from credential disclosure is addressed through trust negotiation [31], [32]. This paper supports privacy while facilitating reputation-based trust.

Privacy awareness becomes very important in a collaborative environment. The primary desire for privacy control in collaborative work settings comes from the desire of "impression management" [7]. Furthermore, since high reputation creates positive impression about a user, we take the view that reputation management also contributes to "impression management." Individuals with good reputation are usually trusted and valued in a relationship. Privacy in the form of anonymity could diminish trust. All the points below may contribute to an environment of diminished trust, which is not conducive to certain uses of computer communication [33]: 1) anonymity makes law enforcement difficult; 2) it frees individuals to behave in socially undesirable and harmful ways; 3) it diminishes the integrity of information since one cannot be sure who information is coming from, whether it has been altered on the way, etc.

7.2 Trust Models

Marsh addresses the issue of formalizing trust as a computational concept in his PhD dissertation [34]. In his model, trust is treated as a subjective and mathematical entity, and it is computed using a subjective real number arbitrarily ranging from -1 to $+1$. In the work of Golbeck and Hendler, trust is treated as a measure of uncertainty in a person or a resource [35]. Specifically, they suggested an algorithm for inferring trust by polling ratings from one's trusted neighbors in a social network. In both of the models [34], [35], reputation is synonymous with the measure of trust. We use reputation to measure trust for e-learning because of the following reasons: reputation is more of a social notion of trust [35], and reputation-based trust works well because of small world web [36] effect.

The use of more formal methods for reputation assessment of a site or of a user are also common on the web. The eBay rating system tries to use customers' positive and negative feedback ratings as a measure of a seller's reputation.[2] Epinions, a consumer review website, also allows customers to rate the transactions with sellers, and maintains a more explicit trust rating system.[3] The PageRank algorithm [37] used by the Google search engine, is also a trust metric of a sort. It uses the number of links coming into a particular page as votes for that site.

The three most common types of trust solutions found in the literature are as follows: 1) based on digital certificates and signatures (e.g., X.509, PGP), 2) based on one's own past experience, and 3) based on the recommendations from third parties. In the first case, trust measure is binary—one party is authenticated to be trustworthy or not. On the other hand, trust built by experience or recommendation is referred as reputation-based trust and it is of "nondiscrete" nature, for example, the interuser trust we seek to capture in this paper could be defined as a value between 0 and 1. Certificate-based trust vouches for the certificate holder's identity, whereas we are interested in modeling trust that would vouch for behavior.

One interesting approach of assessing reputation is the federated reputation model of Agudo et al. [38]. In this work, the authors propose that an Identity provider (IDP) not only authenticates users to different service providers, but also collect information from the SPs about the reputation of given users and a reputation manager inside IDP maintains reputation of users. Pingel and Steinbrecher propose an interoperable reputation system to serve multiple online communities with the assumption that intercommunity and within community agreement on appropriate contexts for exchanging reputation [39]. Our work treats trust, reputation, and identity to be contextual and allows transfer and merge of reputation among partial identities within the same context in a unlinkable and secure way.

8 Conclusion and Future Work

The expectations of trust and privacy among the users of e-learning systems affect learning activities and learning outcomes. A naively constructed privacy-enhanced learning environment offers isolated personal learning spaces, which allow learners to be sometimes frustrated, overwhelmed, or dissatisfied with learning objects or instructors. In this

2. http://www.ebay.com.
3. http://www.epinions.com.

paper, an approach to address privacy protection and trust facilitation is explored. Reputation is an effective means to measure trust in e-learning environments. A mechanism to evaluate and attach reputation to a pseudonymous identity can help measure trust without the loss of privacy. For example, when Alice takes part in a discussion forum, her reputation as a friendly and knowledgeable user may be 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 need to assume multiple nonlinkable partial identities to protect their privacy, there is a need for reputation transfer among the partial identities. Privacy protection in reputation transfer requires that the transfer must occur without letting anyone easily observe such a transfer or be able to link two partial identities querying reputation. Besides, reputation is contextual and needs to be assessed within a context for accuracy. A solution has been developed and implemented by which privacy-preserving and contextual reputation assessment can be done with the aid of a trusted guarantor. The system can help learners to successfully identify potentially good helpers or collaborators.

8.1 Future Work

Even though our work is geared toward e-learning, the problem of nonlinkability disrupting reputation assessment and vice versa is not peculiar to e-learning. This is a limitation of identity management-based solution to privacy. Therefore, our solution has broader applications, and we expect to apply our solution in other domains like e-business, where both privacy and trust are important. Since our work shares similar over-arching goal of privacy-enhanced trust management with other research efforts like PICOS project,[4] our work can be expanded to facilitate reputation-based trust while supporting privacy-preserving identity management in online communities.

In order to better analyze the impact of our system on the users' experience, we plan to conduct a large-scale study in an online environment where there is no existing trust relationship among users. Furthermore, we plan to look more deeply into privacy trust tradeoff issues. A user may choose to trade their privacy for a corresponding gain in their partner's trust. In an asymmetric trust relationship, the weaker party must trade this privacy loss for a trust gain, which is required to start interaction with the stronger party [40]. For a privacy trust tradeoff, we would like to build a heuristic tool that would help users with answers to various privacy and trust related questions, such as

- How much privacy is lost by a user when disclosing the given data?
- How much does a user benefit from a particular trust gain?
- How much privacy should a user be willing to sacrifice for a certain amount of trust gain?

4. http://www.picos-project.eu.

References

[1] M. Anwar, J. Greer, and C. Brooks, "Privacy Enhanced Personalization in E-Learning," *Proc. Int'l Conf. Privacy, Security, and Trust,* 2006.

[2] K. Borcea, H. Donker, E. Franz, A. Pfitzmann, and H. Wahrig, "Towards Privacy-Aware eLearning," *Proc. Privacy Enhancing Technologies,* pp. 167-178, 2005.

[3] R.E. Leenes, "User-Centric Identity Management as an Indispensable Tool for Privacy Protection," *Int'l J. Intellectual Property Management,* vol. 2, no. 4, pp. 345-371, 2008.

[4] J. Mason and P. Lefrere, "Trust, Collaboration, and Organisational Transformation," *Int'l J. Training and Development,* vol. 7, no. 4, pp. 259-271, 2003.

[5] C. Haythornthwaite, "Facilitating Collaboration in Online Learning," *J. Asynchronous Learning Networks,* vol. 10, no. 1, pp. 7-23, 2006.

[6] J. Allan and N. Lawless, "Stress Caused by Online Collaboration in E-learning: A Developing Model," *Education and Training,* vol. 45, nos. 8/9, pp. 564-572, 2003.

[7] S. Patil and A. Kobsa, "Privacy in Collaboration: Managing Impression," *Proc. First Int'l Conf. Online Communities and Social Computing,* 2005.

[8] E. Aimeur, H. Hage, and F.S.M. Onana, "A Framework for Privacy-Preserving E-learning," *Proc. Joint iTrust and PST Conf. Privacy, Trust Management and Security,* vol. 238, pp. 223-238, 2007.

[9] E.T. Bates and L.R. Wiest, "Impact of Personalization of Mathematical Word Problems on Student Performance," *The Math. Educator,* vol. 14, no. 2, pp. 17-26, 2004.

[10] A. Kobsa and J. Schreck, "Privacy through Pseudonymity in User-Adaptive Systems," *ACM Trans. Internet Technology,* vol. 3, no. 2, pp. 149-183, 2003.

[11] P. Briggs, B. Simpson, and A.D. Angeli, "Personalisation and Trust: A Reciprocal Relationship?" *Designing Personalized User Experiences in eCommerce,* pp. 39-55, Kluwer Academic, 2004.

[12] L. Sweeney, "k-Anonymity: A Model for Protecting Privacy," *Int'l J. Uncertainty Fuzziness and Knowledge-Based Systems,* vol. 10, pp. 557-570, Oct. 2002.

[13] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam, "L-Diversity: Privacy Beyond K-anonymity," *ACM Trans. Knowledge Discovery from Data,* vol. 1, p. 3, Mar. 2007.

[14] E. Goffman, *The Presentation of Self in Everyday Life.* Anchor-Doubleday, 1961.

[15] S. Patil and A. Kobsa, "Privacy as Impression Management," Technical Report UCI-ISR-03-13, Inst. for Software Research, Univ. of California, Dec. 2003.

[16] M. Raento and A. Oulasvirta, "Designing for Privacy and Self-Presentation in Social Awareness," *J. Personal and Ubiquitous Computing,* vol. 12, no. 7, pp. 527-542, 2008.

[17] M. Anwar and J. Greer, "Implementing Role- and Relationship-Based Identity Management in E-Learning Environments," *Proc. 14th Int'l Conf. Artificial Intelligence in Education,* 2009.

[18] T.D. Huynh, N.R. Jennings, and N.R. Shadbolt, "Certified Reputation: How an Agent Can Trust a Stranger," *Proc. Fifth Int'l Joint Conf. Autonomous Agents and Multiagent Systems (AAMAS '06),* pp. 1217-1224, 2006.

[19] D.F. Schoorman, R.C. Mayer, and J.H. Davis, "An Integrative Model of Organizational Trust: Past, Present, and Future," *Academy of Management Rev.,* vol. 32, no. 2, pp. 344-354, 2007.

[20] J. Greer, G. McCalla, J. Vassileva, R. Deters, S. Bull, and L. Kettel, "Lessons Learned in Deploying a Multi-Agent Learning Support System: The I-Help Experience," *Proc. Int'l AI and Education Conf. (AIED '01),* pp. 410-421, 2001.

[21] M. Anwar and J. Greer, "Role- and Relationship-Based Identity Management for Private yet Accountable E-Learning," *Proc. Joint iTrust and PST Conf. Privacy (IFIPTM '08),* 2008.

[22] M. Hansen, A. Schwartz, and A. Cooper, "Privacy and Identity Management," *IEEE Security and Privacy,* vol. 6, no. 2, pp. 38-45, Mar./Apr. 2008.

[23] M. Anwar and J. Greer, "Enabling Reputation-Based Trust in Privacy-Enhanced Learning Systems," *Proc. Ninth Int'l Conf. Intelligent Tutoring Systems (ITS '08),* June 2008.

[24] C. Handy, "Trust and the Virtual Organization," *Creating Value in the Network Economy,* pp. 107-120, Harvard Business School, 1999.

[25] N. Luhmann, "Familiarity, Confidence, Trust: Problems and Alternatives," *Trust: Making and Breaking Cooperative Relations,* D. Gambetta, ed., pp. 94-107, Oxford Univ., 2000.

[26] T. Grandison and M. Sloman, "A Survey of Trust in Internet Applications," *IEEE Comm. Surveys and Tutorials*, vol. 3, no. 4, pp. 2-16, Oct.-Dec. 2000.
[27] V. Carchiolo, D. Correnti, A. Longheu, M. Malgeri, and G. Mangioni, "Exploiting Trust into E-Learning: Adding Reliability to Learning Paths," *Int'l J. Technology Enhanced Learning*, vol. 1, no. 4, pp. 253-265, 2009.
[28] D. Artz and Y. Gil, "A Survey of Trust in Computer Science and the Semantic Web," *J. Web Semantics*, vol. 5, no. 2, pp. 58-71, 2007.
[29] J.L. Steel, "Interpersonal Correlates of Trust and Self-Disclosure," *Psychological Reports*, vol. 68, pp. 1319-1320, 1991.
[30] B. Friedman, P.H. Kahn Jr., and D.C. Howe, "Trust Online," *Comm. ACM*, vol. 43, no. 12, pp. 34-40, 2000.
[31] W. Nejdl, D. Olmedilla, and M. Winslett, "Peertrust: Automated Trust Negotiation for Peers on the Semantic Web," *Proc. 30th Workshop Secure Data Management in a Connected World (SDM '04)*, pp. 118-132, 2004.
[32] T. Yu and M. Winslett, "Policy Migration for Sensitive Credentials in Trust Negotiation," *Proc. ACM Workshop Privacy in the Electronic Soc. (WPES '03)*, pp. 9-20, 2003.
[33] D.G. Johnson and K. Miller, "Anonymity, Pseudonymity, or Inescapable Identity on the Net (Abstract)," *SIGCAS Computer Soc.*, vol. 28, no. 2, pp. 37-38, 1998.
[34] S. Marsh, "Formalising Trust as a Computational Concept," PhD dissertation, Univ. of Stirling, 1994.
[35] J. Golbeck and J. Hendler, "Accuracy of Metrics for Inferring Trust and Reputation in Semantic Web-Based Social Networks," *Proc. Eng. Knowledge in the Age of the SemanticWeb*, 2004.
[36] L.A. Adamic, "The Small World Web," *Proc. Third European Conf. Research and Advanced Technology for Digital Libraries*, Jan. 1999.
[37] C. Ridings and M. Shishigin, "Pagerank Uncovered," technical report, 2002.
[38] I. Agudo, M.C.F. Gago, and J. Lopez, "A Multidimensional Reputation Scheme for Identity Federations," *Proc. Sixth European Conf. Public Key Infrastructures, Services and Applications (EuroPKI)*, pp. 225-238, 2009.
[39] F. Pingel and S. Steinbrecher, "Multilateral Secure Cross-Community Reputation Systems for Internet Communities," *Proc. Fifth Int'l Conf. Trust, Privacy and Security in Digital Business (TrustBus)*, pp. 69-78, 2008.
[40] L. Lilien and B.K. Bhargava, "A Scheme for Privacy-preserving Data Dissemination," *IEEE Trans. Systems, Man, and Cybernetics*, vol. 36, no. 3, pp. 503-506, May 2006.

Mohd Anwar received the PhD degree from the University of Saskatchewan. Currently, he is working as a visiting research assistant professor in the School of Information Sciences of the University of Pittsburgh.

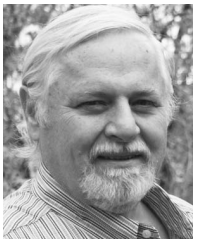

Jim Greer received the PhD degree from the University of Texas at Austin and has been a faculty member at the University of Saskatchewan for more than 20 years. Currently, he is working as a professor of computer science and also serves as a director of the University Learning Center.