A Taxonomy to Express Open Challenges in Trust and Reputation Systems

Mozhgan Tavakolifard
Centre for Quantifiable Quality of Service in Communication Systems
Norwegian University of Science and Technology
Email: mozhgan@q2s.ntnu.no

Kevin C. Almeroth
Department of Computer Science
University of California Santa Barbara, CA 93106-5110
Email: almeroth@cs.ucsb.edu

Abstract— **During the past decade, online trust and reputation systems have provided cogent answers to emerging challenges in the global computing infrastructures relating to computer and network security, electronic commerce, virtual enterprises, social networks and cloud computing. The goal of these systems in such global computing infrastructures is to allow entities to reason about the trustworthiness of other entities and to make autonomous decisions on the basis of trust. This requires the development of computational trust models that enable entities to reason about trust and to verify the properties of a particular interaction. The robustness of these mechanisms is one of the critical factors required for the success of this technology. In this paper, we briefly present characteristics of existing online trust and reputation models and systems through a multidimensional framework that can serve as a basis to understand the current state of the art in the area. The critical open challenges that limit the effectiveness of today's trust and reputation systems are discussed by providing a comprehensive literature review. Furthermore, we present a set of our contributions as a way to address some of these challenges.**

Index Terms— **Challenges, Context, Reputation systems, Taxonomy, Trust.**

I. Introduction

The deployment of a global computing infrastructure raises new and difficult security and privacy issues. Traditional security mechanisms are of questionable effectiveness in the new global computing era. Part of the reason is that no common infrastructure can be assumed to enforce any notion of correct behavior, in part because even defining a common and acceptable standard is impossible. No single authority can define and enforce rules, and therefore, online interactions cannot be governed by common rules as before. Trust-based security mechanisms have emerged as a solution, that expand the scope of traditional security models. Trust enables humans to accept risks and deal with uncertainty. These new mechanisms provide weaker security guaranties, but serve greater application areas. However, the online environments such as the web, search engines, peer-to-peer networks, and new applications built on highly complex social networks introduce several challenges in the interpretation and use of online trust and reputation systems. For example, some of these challenges have their roots in the subjective nature of feedback and some of them are related to the ease with which online identities can be attacked. Before online reputation systems will be accepted as legitimate trust solutions, a better understanding is needed of how such systems can be compromised and how these problems can be solved.

Despite the promise of online trust and reputation systems, there remain significant challenges requiring further research and commercial development. The primary objective of this paper is to describe the critical open challenges that limit the effectiveness of trust and reputation systems and have prevented their integration into large-scale distributed applications. Integrating reliable reputation solutions will contribute tremendously towards increasing user cooperation, thereby improving the performance of these applications [1], [2]. Our goal is to identify challenges that weaken trust and reputation systems, and to survey prominent strategies to overcome these challenges. In addition, we summarize our solutions to some of these problems and we propose a future research agenda.

The remainder of this paper is organized as follows. Section II provides background on trust and reputation systems and a multidimensional framework for categorization and comparison of them. We describe the main problems and solutions in Section III. An overview of the surveys that can be found in related work is given in Section IV. We briefly describe some of our relevant work in Section V and present future directions in Section VI. Finally, Section VII offers concluding remarks.

II. Background on Trust/Reputation Systems

Trust and reputation systems represent a significant evolution in support for Internet services, especially in helping users decide among a growing number of choices, from which movies to rent to which data sources to trust. In this section, we first describe the concept and nature of trust by indicating what is not trust [3]. Trust is not simply "confidence" because

*"Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence" appointed by The Research Council of Norway, funded by the Research Council, NTNU, and UNINETT. http://www.q2s.ntnu.no/

doi:10.4304/jcm.7.7.538-551

trust is about what the perception of what someone is willing to do, while confidence is about what another is capable of doing. Moreover, trust is not "reliability" since a person may not have a choice on whom she relies. Trust also differs from "hope" in terms of available choices. When a potentially risky action has to be taken, a person hopes that it will result in a satisfactory outcome.

A universally accepted definition of trust is still lacking despite extensive studies from philosophers, sociologists, and psychologists. One of the most commonly accepted definitions is from the sociologist Diego Gambetta [4]: "... trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent will perform a particular action, both before [we] can monitor such action (or independently of his capacity of ever be able to monitor it) and in a context in which it affects [our] own action". As stated in this definition, some of the characteristics of trust are: subjectivity, context-dependency, and dynamicity. It is easier to determine the properties of trust than to define exactly what trust itself is. The reason for this difficult is that trust involves a combination of interrelated cognitive and non-cognitive constructs, some of which may or may not be called on depending on the entities and situations involved.

A *trust relationship* exists between two agents when one agent has an opinion about the other agent's trustworthiness and a *recommendation* is an opinion about the trustworthiness from a third party agent. If the referrer is not known by the recommendation requester, the requester can obtain recommendations about the unknown referrer as well. There is also the scenario where a recommendation requester may carry out a network search for a particular party and the received recommendation may be the result of the request being forwarded through a number of intermediary referrers. In both scenarios, when a referrer recommends another referrer, the result is a recommendation chain.

Reputation is defined as an "expectation about an agent's behavior based on information about or observations of his past actions." Therefore, reputation can be considered a collective measure of trustworthiness (in the sense of reliability) based on the referrals or ratings from members in a community. An individual's subjective trust can be derived from a combination of received referrals and personal experience.

The basic idea in existing online trust and reputation systems is to let parties generate feedback about each other after completion of a transaction, and aggregating the feedback to derive a reputation score. The reputation score is used to assist others in deciding whether or not to trust that party in the future. Resnick et al. identifies three phases as being fundamental to any reputation system: (i) feedback generation, (ii) feedback distribution, and (iii) feedback aggregation [5].

In addition to reputation systems, some applications can use *collaborative filtering*. Collaborative filtering techniques calculate a personalized rating estimation of an item for a user as the weighted average of previous ratings given to that item by other users. The weights are proportional to the similarity between a current user and the previous users. The user similarity can be calculated using the users' profiles or as a function of the correlation between users' ratings assigned to a common set of items. For example, if User A likes Items X and Y, and User B likes Item Y, it is likely that User B will like Item X too.

There are similarities between collaborative filtering and reputation systems. Both types of systems collect ratings from members in a community/social network. The usefulness of the former arises when the emphasis is on the content, and the latter can be used when the source of information is a more important factor. Therefore, they are complimentary decision mechanisms for use in decision systems [6].

In the following, we present several dimensions for classification of the current state of the art in trust and reputation systems. This classification can serve as a basis to understand the current state of the art in trust and reputation systems, to give an overview of research areas, and to help distinguish the areas that require more work.

Information type: Trust and reputation systems can use explicit or implicit information for decision making. Examples of implicit trust information can be found in social networks such as Facebook or LinkedIn. Entities within a social network can extract some degree of trust for information gathered through friends of friends. Although neither Facebook nor LinkedIn directly implement a reputation system, members of both systems are able to utilize reputable connections through friends within the environment. Another implicit form of trust information is the use of topological analysis in online social networks to determine reputation [7]. In the Google search engine, reputation is determined by the number of links that point at a page, and from where the links originate. A link originating at a page with a high reputation is likely to mean that the target page has some value.

Trust value representation: Degrees of trust are represented as either discrete or continuous levels of trust. Humans are often better able to rate performance in the form of discrete verbal statements, than they are continuous measures. This limitation is also valid for determining trust measures. The discrete levels differ from one model to the next, with some using a bounded range [8] and others allowing the value to extend to infinity [9]. Moreover, discrete values can be binary or multinomial. A binary trust representation allows complete trust in another agent or no trust at all. Binary trust representations are simple constructs and allow unambiguous implementations. The concept is, nevertheless, rather restrictive because users are forced to choose between trusting another agent completely or not at all. The ability to handle degrees of trust in multinomial form [6], [8] allows users to proceed in situations where the amount of trust in another agent is not complete, but sufficient for the situation concerned. The disadvantage of discrete measures is that they do not easily lend themselves to sound computational principles. Instead, heuristic mechanisms like look-up tables must be used. On the other hand, continuous values [10], represented as real numbers, are modeled as either objective or subjective probability values. The objective probabilities represent purely syntactic forms of trust values, *e.g.*, the beliefs of the agent does not influence the value, and subjective probabilities are intuitive "likelihood" measurements given by the agent depending on its current beliefs. All in all, it is

better to maintain reputation values as multiple component scores. Applying different functions to the scores allows a rating best suited for the given situation to be calculated. Many proposed systems suggest maintaining multiple statistics about each user. For example, keeping separate ratings on a user's likelihood to defect on a transaction (its "trustworthiness") or user's likelihood to recommend malicious users (its "reliability" as a referral) [11].

Some proposals divide the span of trust into strata and assign qualitative labels to them [3]. For example, the stratification is given as the set of Very Trustworthy, Trustworthy, Untrustworthy, and Very Untrustworthy [3]. The use of strata with qualitative labels may initially be considered a good solution to the problem of subjectivity because it seems to provide a clear semantics and avoids the ambiguity associated with numerical values. Nevertheless, in order for it to have the claimed effect, a qualitative label such as "trustworthy" should hold the same meaning for one person as it does for another. This assumption is not necessarily the case because persons with different personality cultures may associate the same experience with different strata. For example, based on her own perception of trust, what is viewed by someone as "very trustworthy" may be judged as only "trustworthy" by another person. Previous work either only considered the positive values of trust or ignorance (absence of trust or no opinion about the trustworthiness) or considered distrust as well [10].

Network architecture: The network architecture determines how feedback and reputation scores are communicated between participants in a reputation system. The two main types are *centralized* (or hierarchical) and *distributed* (or peer-based) architectures. In a centralized reputation system, a central authority (reputation center) collects feedback about a given participant from other members in the community who have had direct experience with that participant. The central authority derives a publicly available reputation score. Centralized structures work well within closed networks or where decentralized approaches are not suitable for management and control purposes. On the other hand, in a distributed reputation system [8], [12], each participant simply records an opinion about each experience with other parties and provides the information on-demand. Any user can compute a reputation score based on the received feedback from others and his/her own direct experiences.

Algorithm: A reputation system uses a specific method (*e.g.*, averaging, probabilistic-based or belief-based) to compute reputation values based on the collection of feedback from others. Some of the various methods for computing reputation and trust measures include.

1) Rank ordering: This method has no explicit reputation score and acts as an implicit indicator of reputation. For instance, in Slashdot, an online discussion board, readers rate posted comments and postings are prioritized or filtered according to the ratings they receive from readers.

2) Simple summation or average of ratings: This method is the simplest form of computing reputation scores. The score is the sum of the number of positive ratings and negative ratings, for example, positive scores minus negative scores (*e.g.*, eBay) or the average (*e.g.*, Epinions and Amazon).

3) Probabilistic models: The reputation score is computed by updating Probability Density Functions (PDFs). The updated reputation score is computed as a combination of the previous reputation score and the new rating.

4) Fuzzy models: These methods represent trust and reputation as linguistically fuzzy concepts, where membership functions describe to what degree an agent can be described as trustworthy or not. Fuzzy logic provides rules for reasoning with fuzzy measures of this type.

5) Flow models: A participant's reputation increases as a function of incoming flow, and decreases as a function of outgoing flow (e.g., Google's PageRank and Advogato). In the case of Google, many hyperlinks to a web page contribute to increased PageRank whereas many hyperlinks from a web page contributes to a decreased PageRank for that web page.

6) Game theoretical models: Problematic social situations can be described as trust games with two players and two periods of play. A Trust Game is a one-sided Prisoners Dilemma Game. The restrictiveness of the social conditions under which problematic social situations have to be solved can be reduced by adding the notion of reputation (the possibility of obtaining or spreading information about a trustee's trustworthiness) and third parties. This can be explained by the fact that the principal effect of information from third parties is to reduce uncertainty about the behavior of the trustee.

7) Stochastic models: Events are modeled by Markov decision processes and reputation is aggregated using stochastic system theory [13].

8) Belief models: Dempster-Shafer theory of evidence is an extension to probability theory with the advantage of being able to model uncertainty. It is a widely used model which provides the means for approximate reasoning under uncertainty. According to it, there is no direct relationship between a hypothesis and its negation and as a result the summation of probabilities of atomic elements may not necessarily result in a value of one. In this case, the remaining probability is interpreted as a state of uncertainty [6].

TABLE I.
A COMPARISON OF EXISTING ONLINE TRUST AND REPUTATION SYSTEMS ACROSS SEVERAL DIMENSIONS.

Model	Info type	Value representation	Architecture	Algorithm	Info source	Context awareness	Parameters
FIRE [14]	Explicit	Continuous	Decentralized	Custom-designed	Direct experiences, Referrals recommendation, Certifications	Single	Role-based trust, Time, Reliability, Credibility
Confidant [15]	Explicit	Real number in the range [0, 1]	Decentralized	Custom-designed	Direct experiences, Referrals recommendation	Single	Time
Yu and Singh [16]	Explicit	Real number in the range [0, 1]	Decentralized	Belief-based approach	Direct experiences, Referrals recommendation	Single	Reilability, Time
TRAVOS [17]	Explicit	Binary ratings and reputation value as a real number in [0, 1]	Decentralized	Probabilistic approach (Bayesian)	Direct experiences, Referrals recommendation	Single	Time, Reliability of referrals
Conner et al. [18]	Explicit	Continuous	Decentralized	Custom-designed	Direct experiences	Multiple	
EigenTrust [19]	Explicit	Simple summation	Decentralized	Flow model	Direct experiences, Referrals recommendation	Single	
Gupta et al. [20]	Explicit	Continuous	Centralized	Custom-designed	Direct experiences	Multiple	Time
PeerTrust [21]	Explicit	Continuous	Decentralized	Custom-desinged (normalization)	Direct experiences, Referrals recommendation	Single	Context compatibility, reliability of referrals, Time
H-Trust [22]	Explicit	Bounded continuous	Decentralized	Custom-designed (inspired by H-index)		Single	Credibility of referrals,
REGRET [23]	Explicit Ł	[0, 1]	Decentralized	Fuzzy approach	Direct experiences, Referrals recommendation	Multiple	Criteria compatibility, Reliability of referrals, Time
ŁŁŁŁ **HISTOS [24]**	Explicit	Continuous	Decentralized	?	Direct experiences, Referrals recommendation	Single	Reliability of referrals, time
RATEWeb [25]	Explicit	Continuous	Decentralized	Custom-designed (weighted-avarage)	Direct experiences, Referrals recommendation	Multiple	Criteria compatibility, context compatibility, credibility of referrals, time
P-Grid [26]	Explicit	Binary	Decentralized	Custom-designed	Direct experiences, referral's recommendation	Single	
PowerTrust [27]	Explicit	Binary	Decentralized	Probabilistic-based approach (Bayesian method)	Direct experiences	Single	Time
PRIDE	Explicit	Discrete and unbounded	Decentralized	Custom-designed (certification)	Direct experiences	Single	
TrustMe [28]	Explicit	Continuous	Decentralized	Custom-designed	Direct experiences, referral's recommendation	Single	Time
XRep [29]	Explicit	Binary, Discrete/Continuous	Decentralized	Custom-designed	Referral's recommendation	Single	
Amazon	Explicit	1-5 stars	Centralized	Average of ratings	Direct experiences	Single	
eBay	Explicit	-1,0,1	Centralized	Summation	Direct experience	Single	
SlashDot	Implicit	Strata & continuous	Centralized	Rank ordering	Direct experiences	Single	Credibility of feedback
ePinions	Implicit	Strata	Centralized	Average	Direct experiences	Single	Credibility of feedback

9) Semantic web and ontologies: This is a logical approach for trust formalization and mainly focuses on trust's semantic structure and its logical conditions and effects. As opposed to other approaches that focus on the uncertainty of trust, trust quantification, trust dynamics, and trust computings models and algorithms. The semantics of trust relationships are modeled using ontologies [30].

10) Spread activation networks: This is an example of a cognitive science approach. Spread activation models simulate human comprehension through semantic memory, and are commonly described as "models of retrieval from long term memory in which activation subdivides among paths emanating from an activated mental representation" [31].

11) Social network measures: These approaches attempt to find an answer to the question: in which way does a truster's level of trust in a trustee depend on his "local" network position and on the global network structure? In other words, they evaluate the effects of density, outdegree centrality, and centralization on the level of trust a trustor can have in a trustee [32].

12) Custom-desinged models: In these models, trust values are calculated from handcrafted formula to yield the desired results. The flexibility of these approaches enables trust and reputation systems to define a composite trust metric to aggregate the essential parameters and factors they have been considered in their models. Basically, they include credibility of witnesses and time or a recency factor as the main variables. However, some advanced models may use other important variables such as the transitivity rate and context and criteria similarity into account as well.

Information source: The majority of trust models consider two types of knowledge in estimating the trustworthiness of a trustee in an interaction: direct experiences and referral's recommendations (or witness observations). Personal experience typically carries more weight than second hand recommendations or reputation, but in the absence of personal experience, trust often has to be based on recommendations from others. Furthermore, some trust and reputation systems designed particular information components for the situations when neither of these information sources is available. For the FIRE model [14] uses *certifications* by target members. The other information source is called *role-based trust* [14], [33], in which agents trust each other based on the predefined roles and relationships that exist among them.

Context awareness: A single-context trust and reputation model is designed to associate a single trust value per partner without taking into account the context. These systems entail information being collected from a single method and being interpreted in a predefined way. This means all of the information collected about an entity is related to one explicit aspect of that entity's actions A multi-context model has the mechanisms to deal with several contexts at a time maintaining different trust values associated to these contexts for a single partner [18]. Multi-context reputation systems can take advantage of numerous sources to gather information, or collect the information such that it can be used from different perspectives.

Parameters: This dimension addresses crucial parameters which may increase the accuracy of the expected reputation value.

1) Time: The time parameter is an essential parameter in any reputation calculation and indicates the recency and freshness of information. Older information should have less influence in the calculation.

2) Credibility of referrals: In a recommendation chain, recommendations from known referrals who already have had interaction with the requested party should have more weight as first-hand recommendations than those who are known but have not had any previous interactions with the requested party or those who are unknown.

3) Reliability of referrals: It is also important to consider the reliability and honesty of the referrals, even for well known referrals by the requesting party, when their recommendations are used to be able to calculate the confidence level of the generated recommendations and to alleviate the effect of dishonest information providers and spurious ratings. This can be measured by assessing the trend line and behavior of the referrer in the time interval [33].

4) Context compatibility: If the information used for calculation has not been generated in exactly the same context as the decision making, then information should be weighted based on the similarity between the current context and the context of the information in order to determine to what extent the received information should be taken into account.

5) Criteria compatibility: This factor determines the similarity between the criteria used to evaluate the outcome of an interaction.

Table I provides a comparison of some of the exiting trust and reputation systems against the aforementioned dimensions. Table I highlights the differences between the selected trust and reputation systems and compares them across the dimensions of the framework. Table I shows that some of the systems address a wider range of dimensions than others. This fact may not necessarily imply better quality and applicability of such systems. Instead, one should consider the context in which these systems are employed and evaluate how well they accomplish the goals and requirements of that particular environment.

III. Challenges

To better explain what future work is possible, we first describe some of the basic problems and proposed solutions for online reputation systems. We consider each phase of operation for such systems, namely: feedback generation, feedback distribution, and feedback aggregation. Each of these components needs safeguarding against a variety of adversarial threats. As a case in point, reliability in terms of reputation accuracy is a critical requirement for the aggregation component. This section, therefore, studies the extent to which existing research efforts counter these threats.

A. Feedback Generation

One of the most important tasks in a reputation system is generating accurate and representative feedback. Not only

must a qualitative, opinion-based process be reduced to quantitative facts, but also users will sometimes try to game the system. We have identified the following challenges.

Low incentive for providing feedback: There are two main reasons for this problem [34]. First, feedback constitutes a public good and once available, everyone can benefit, yet the provider benefits very little. Second, providing feedback presupposes that the provider will assume the risks of the transaction, risks that are typically higher for new products or users. To solve this problem, some models propose payments and financial rewards for honest feedback [35], [36]. For example, Epinions provides incentives for reviewers, whereby they can earn money based on general use of reviews by consumers. Bizrate, a customer certified merchant, gives discounts as an incentive to fill out surveys. An alternative approach is to build incentives into the feedback aggregation equation. This goal can be accomplished by providing a small increase in reputation whenever a user provides reputation feedback to others [37]. For instance, Amazon gives some members status as a top reviewer. Another approach is to use implicit feedback, where users' actions are recorded and the feedback is inferred from the recorded data [34]. For example, an assumption in Google's reputation score is that if enough people consider a page to be important enough to place links to it, and if the pointing pages are "reputable" themselves, then the information contained on the target page is likely to be valuable.

Bias toward positive feedback: It can be difficult to elicit negative feedback because of reciprocity. For example, the observed ratings on eBay are surprisingly positive. Of all ratings provided, less than 1% are negative, less than 0.5% are neutral and about 99% are positive [38]. Providing anonymity may help to avoid this problem. It was also found that there is a high correlation between buyer and seller ratings, suggesting that there is a degree of reciprocation of positive ratings and retaliation for negative ratings. A possible remedy could be to not let sellers rate buyers.

Initialization and cold-start problem: Bootstrapping a reputation mechanism is not trivial. In many systems, users start with a neutral reputation. Newcomers are offered only a limited number of resources and so struggle initially to build their reputations. As other users in the system tend to interact with high reputable users, the chance of a new user being selected for interaction is generally rare (*e.g.*, in eBay, many users will not deal with individuals with a low reputation score [37]). Hence, it is hard for a new user to raise her reputation score. This challenge may be a barrier to entry into the marketplace or community. Solutions include taking into consideration the interconnections among reputation systems and social networks. For example, the location of a given member of a community within a social network can be used to infer some properties about her degree of expertise, *i.e.*, her reputation [39].

Subjectivity: Feedback information is strongly influenced by subjectivity factors such as the feedback provider's taste and cultural background. One solution based on collaborative filtering, is to personalize the feedback by weighting it in inverse proportion to "taste distance" between the provider and the receiver of the feedback. Therefore, it will be easier for the receiver of the feedback to interpret it because it consists of opinions from like-minded people [34].

False feedback: When users incorrectly report their feedback, it creates errors in the system. We categorize the different forms of false feedback as follows:

1) Dishonest and unfair reports: This problem happens because of the low cost of submitting online feedback and the relative anonymity of the raters. Unfair ratings can be excluded using their statistical properties [40], [41] or by using the rater's reputation [42], [43]. Slashdot addresses this issue by using the judgment of longstanding users as *a priori* trusted agents.

2) Collusion: Collusion occurs when two of more peers collectively boost each others reputations or conspire against one or more peers in the network. Dellarocas identifies three types of collusion misbehavior [34]:

a) *Ballot stuffing*: Parties engage in many fake transactions to artificially inflate their reputations and ratings. This problem is solved in eBay by only allowing participants to rate each other after the completion of a transaction, and charging a fee for each transaction. In the Sporas model [24], when a user rates another more than once, only the most recent rating is considered.

b) *Bad-mouthing*: This problem occurs when a malicious collective conspires against one or more users in the community and hurt their reputation by assigning unfairly low ratings to them.

c) *Positive and negative discrimination*: Discriminatory behavior can occur both when providing services and when providing feedback. A seller can, for example, provide good quality to all buyers except one in particular. Feedback about that particular seller will indicate that she is trustworthy except for the feedback from the victim buyer. Filtering techniques will give false positives, *i.e.*, judge the buyer victim unfairly in such situations. Only systems that are able to recognize the victim buyer as trustworthy would be able to handle this situation.

Cheap pseudonyms: In online environments where new identities may be created with minimal cost, these multiple identities create several problems, including the following:

1) Sybil-based collusion: Malicious entities may acquire multiple identities for the sole purpose of creating phantom feedback in the system. Proposed solutions to deal with Sybil attacks fall into centralized and decentralized approaches. In a centralized approach, a central authority issues and verifies credentials unique to each entity. To increase the cost of obtaining multiple identities, the central authority may require monetary or computational payment for each identity. In decentralized approaches, some proposed solutions include binding a unique identifier, such as IP addresses, to public keys or using network coordinates to detect nodes with multiple identities (*e.g.*, Kuro5hin allows only one rating from any single IP address). Other solutions take advantage of social knowledge to propagate reputations originating from trusted sources along the edges of a "web of trust". Thus, the effect of the attackers will be limited based on the expense of requiring social interactions [44].

2) Re-entry problem or churn attacks: In online communities, it is usually easy for members to disappear and re-register under a completely different online identity with zero or very low cost (*e.g.*, eBay). Models that treat unknown users and disreputable ones differently [10], [45], [46], [46] are vulnerable to this problem, however, models that penalize newcomers are resistant [24]. There are two classes of approaches to this issue [34]: either making it more difficult to change online identities (*e.g.*, by using cryptographic authentication techniques), or making it unprofitable to exit and re-enter with a new identity (e.g., by imposing an upfront cost to each new entrant such as a fee or an implicit cost of having to go through an initial reputation-building phase with low or negative profits).

B. Feedback Distribution

Assuming reputation information can be collected and processed correctly and without malicious influence, the next challenge is to get the feedback to those who need it to make their decisions. Some of the challenges in this part of the process include:

Reputation lag problem: There is usually a time lag between an instance of a transaction and the corresponding effect on the reputation score (*e.g.*, in eBay, the buyer pays before the seller ships the item). A user has the opportunity to make use of this time lag to provide a large number of low quality services over a short period before the reputation score suffers any significant degradation [47]. Further, the re-entry problem can be combined with this problem in a way that a seller may re-enter the market each time a buyer learns of a dishonest seller. In this way, a seller can repeatedly take advantage of reputation lag.

Lack of portability between systems: The limited distribution of feedback limits its effectiveness. As a solution, Amazon allowed users to import their ratings from eBay [5]. Obviously only users with good reputations will take advantage of this feature, thereby diluting the value of the scores.

Inability to filter or search: Online communities run into several information overload problems due to the sheer size of many of these sites. The ability to filter and search by reputation would greatly improve their usability [37].

Categorization: A reputation score is too general in most systems (*e.g.*, eBay) and there is little ability to use reputation scores in different categories. Reputation categories could enhance systems by providing better granularity. For example, a user might have a good reputation in one area (*e.g.*, quality of products) and a bad reputation in another area (*e.g.*, on-time delivery). This concept could work in conjunction with a search and filtering feature [37].

C. Feedback Aggregation

Assuming reputation information can be collected and processed correctly and then delivered to a user, there is still the challenge in aggregating and displaying feedback so that it is truly useful in influencing future decisions about whom to trust. Some of the challenges in this part of the process include:

Inaccurate equations: Simple reputation schemes such as eBay's reputation score (*i.e.*, the sum of positive ratings minus the sum of negative ratings) can be misleading. For example, a user in eBay with 100 positive and 10 negative ratings would have the same total reputation score as a user with 90 positive and no negative ratings; however, the former should appear less reputable. This problem results in a vulnerability caused by "increased trust by increased volume." That is, a user can increase his/her trust value by increasing his/her transaction volume, thereby hiding the fact that she frequently misbehaves.

Value imbalance problem: In many reputation models (*e.g.*, eBay), all feedback is weighted equally regardless of the transaction value. This problem encourages Sybil attacks and collusion. A user can take advantage of this property to build a good reputation by honestly executing a number of small-value trades, and then using the accumulated reputation to cheat in a very high-value transaction [48].

Spread of false rumors: This problem occurs when the reputation of the feedback providers is not considered. One approach to this problem is to rely on pre-trusted identities. Another approach is to employ statistical methods to build robust formulations (*e.g*, a Bayesian framework) that can be reasoned about in a precise fashion [44].

Unlimited memory: Most reputation calculation algorithms use all transactions when calculating the overall score, thus, a new user might not understand how a site functions [37]. Besides, a user can perform short duration malicious attacks with little risk of negative consequences because a lengthy previous history can heavily outweigh current actions. This problem can have a large impact on the system as the malicious users will continue to have a high reputation for a substantial period of time during which the system is slow to identify the malicious behavior and unable to sufficiently lower the user's reputation [44]. Therefore, the memory should be de-emphasized in some way, though this is not easy in practice. For example, a simple cut-off function handicaps the user by providing only the most recent information. Further, new users will likely require some time to become familiar with the mores of a site and they should not be penalized for initial bad behavior if the behavior is unintentional. One solution is to give less weight to negative feedback for new users and more weight for old users. Instead of a strict cut-off, this approach leads to a gradual change in the importance of more recent feedback [37].

Dependence on profit margins: Reputation effects can induce users to accept short-term losses in order to realize larger long-term gains provided that the latter exceeds the former. In other words, the remaining horizon must be long enough and the profit per transaction must exceed a threshold. This result can have at least two potential interpretations. First, reputation mechanisms are not effective in highly competitive markets. Second, prices tend to be higher in markets where trust is based on reputation than in markets with perfect information [34].

Time sensitivity of reputation: treating old positive behavior equal to new negative behavior may result in attackers abusing the system by using previous altruism to hide current malicious behavior. Techniques have been proposed that use

more aggressive short-term history and give more weight to recent negative behavior [44].

Denial of service attacks: Attackers may seek to subvert the mechanisms underlying the reputation system in centralized architectures, causing a denial of service. For instance, attackers can attempt to cause the central entity to become overloaded by attacking its network or computational resources. Attackers are then able to perform malicious actions without their negative reputation being known or without being punished for their negative behavior. Distributed architectures with enough redundancy are often less vulnerable to this attack. Techniques to cope with denial of service attacks are similar with the ones used by many routing protocols and include: use of acknowledgments, multi-path dissemination, gossip mechanisms, and forward error correction codes [44].

Playbooks: A playbook is a sequence of actions that maximizes profit of a participant according to certain criteria. A typical example is to act honestly and provide quality services over a period to gain a high reputation score, and then to subsequently milk the high reputation score by providing low quality services at a low production cost [47].

Exit problem: Since there is no incentive for a party leaving a system to maintain a good reputation, the entire accumulated reputation can be used for cheating (*e.g.*, in eBay). One solution to this problem is to introduce community membership rules that elicit good behavior. For example, online communities can levy a sufficiently high entrance fee that is refundable subject to maintaining a good reputation upon exit or reputation scores can be viewed as assets that can be bought and sold in a market [34].

While an innumerable variety of attacks can be devised by malicious peers, our above discussion identifies attack strategies most commonly observed in reputation systems. We now discuss some additional related work that surveys solutions to some of these attacks.

IV. Related Work

There have been many papers on reputation systems dealing with all manner of perceived challenges, from how to calculate the most accurate reputation score to how to protect the reputation system from some of the attacks just described. There has now been enough work in the field that there are now also many surveys that have been published. Below is a "survey of surveys" that addresses some of the problems and challenges in online reputation systems.

Jøsang et al. describe and analyze the state of the art in trust and reputation systems [6]. They give an overview of existing and proposed systems as a way of analyzing current trends and developments and proposing a research agenda. In addition, they describe the main problems in reputation systems and provide an overview of literature that proposes solutions to these problems.

Kerr and Cohen identify the theoretical possibility of a number of vulnerabilities that permit effective and profitable cheating despite the use of reputation systems [47], [47]. They propose a testbed formulation designed to support systematic experimentation and evaluation of reputation systems.

Dellarocas surveys the progress in understanding the new possibilities and challenges that online reputation mechanisms represent and describes the most important issues related to designing, evaluating, and using these mechanisms [34]. The paper provides an overview of relevant work in game theory and economics on the topic or reputation.

The work by Hoffman et al. focuses on attacks and defense mechanisms in reputation systems [44]. They present an analysis framework that allows for a general decomposition of existing reputation systems and classifies attacks against reputation systems by identifying which system components and design choices are the target of attacks. Furthermore, the paper provides a survey of defense mechanisms employed by existing reputation systems.

Resnick et al. present an introduction to the area of reputation systems [5]. Their work describes three challenges for a successful reputation system: first, entities must be long-lived to account for accurate reputations; second, feedback must be captured and distributed; and third, reputations should help distinguish between trustworthy and untrustworthy partners.

Malaga examines current approaches in reputation management systems, outlines six main problems with them, and suggests how these problems can be solved [37].

Swamynathan et al. identify a thorough taxonomy on reputation management, and use it as a framework to classify adversarial threats that affect reliable operation of reputation systems [1]. In addition, they present solutions to address the user collusion problem and short-lived online identities.

While many of the systems described in these surveys have made useful contributions, there is still much work that needs to be done to develop these techniques such that they are mature enough to be used in modern systems.

V. A New Philosophy

We outline our proposed solutions for some of the problems described in the previous sections. Our work improves the utility and accuracy of trust management systems by proposing methods on how to use *contextual information*. More specifically, we address the following problems in trust and reputation systems and propose solutions based on the use of contextual information:

- Problems related to explicit feedback such as initialization, cold-start, and subjectivity
- Lack of categorization and the ability to filter or search
- Lack of portability between systems

Based on our discussion about the shortcomings in trust and reputation systems in the previous sections, we recognize that there is a need for providing methods to initialize trust properly, to bundle trust information with its category (subject or scope) in order to be able to search and filter the information.

In our approach, we are mainly motivated by what humans do in traditional trust and reputation systems such as analogical, abductive, and inductive reasoning. The main idea is to consider contextual information, as a special kind of implicit feedback, in trust computations and the goal is to bring additional knowledge to the reasoning process by use of available auxiliary data or Meta-data (contextual data). Context

qualifies a trust opinion, describing what the truster's belief in another's trustworthiness is really about. The introduction of context as an explicit notion may improve problem solving efficiency by better grounding what knowledge is used in decision making in the real world. With more information and better context, trust and reputation systems can make more well-informed decisions.

The advantages and importance of using contextual information is also recognized by other researchers. For example, Neisse et al. attempt to reduce the complexity in management of trust relationships [49]. Neisse et al. [50] and Gray et al. [51] focus on the improvement of the trust recommendation process. Holtmanns and Yan investigate how to infer trust information in context hierarchies [52]. Rehak et al. improves the performance of trust management systems [53]. They also provide protection against changes of identity and first time offenders in trust management systems. Bagheri and Ghorbani [54], Bagheri et al. [55], and Gray et al. [51] provide methods that correlate trust information among various contexts.

Contextual information has been represented in several different forms such as Context-aware domains [50], Intentional Programming [56], Multi-dimensional goals [57], Clustering [58], and Ontologies [59].

In our work, we distinguish between *external* and *internal* context. External context is related either to the properties of the trustee or the object to be acted on (*e.g.*, information to be exchanged or something to be bought). These are the facts that exist independent of the reasoner. They are independent in the sense that they are there before and after the reasoner notices them. Internal context (*i.e.*, subjective/cognitive context), on the other hand, characterizes the mental and emotional state of the reasoner, the truster. A trust evaluation process is complicated by context in two ways: (1) trust is situation-specific (the effect of external context); a typical example is that a person may trust her financial adviser about investment analysis but does not trust the same adviser related to health-care issues, and (2) trust is person-specific (the effect of internal context); judgments of two persons on the same matter or event are often quite different.

We describe a holistic trust management approach that deals both with the situation-sensitivity of trust and the subjectivity problem. The impact of internal context (subjectivity) and four types of external context: time, similarity, situation, and stereotypes are modeled and assessed. Our conception of different trust evaluation scenarios and the reasoning methods appropriate in each of them is illustrated in Figure 1. The two fundamental types of knowledge are acquired for trust evaluation through one's own experiences and from recommendations by third parties. If the truster previously has had interactions with the same trustee in the same situation, she can immediately use her past experiences to predict the outcome of the new interaction and make a decision on that basis. On the other hand, if the truster has had interactions with the trustee but in different situations, she can still use her past experiences, but should map the old and new situations and make necessary adaptations in order to draw a conclusion. We have used Case-Based Reasoning (CBR) to handle such situation-specificity of trust. "Situation" in Figure 1 refers to external context. The notion of internal context comes into play when recommendations from a third party is used to evaluate trust.

We now describe some of our work that attempts to solve the five problems identified in Section III.

Unlimited memory and time sensitivity of reputations: In an attempt to model the effect of time as a type of external context, we have proposed a formula for a dynamic longevity factor, $\lambda \in [0, 1]$, that is able to discount past ratings correctly [60]. The longevity factor, λ, controls the rate at which past ratings are aged and discounted as a function of time. With $\lambda = 0$, past ratings are completely forgotten after a given time period. With $\lambda = 1$, past ratings are never forgotten. We propose to adjust λ after each interaction based on the similarity between the estimated and real outcome of the interaction. The higher the similarity, the larger the increase in the value of λ, and the larger the memory size (*i.e.*, time window). If the real outcome is not similar to what is expected, we are facing a change in behavior and a change in the value of λ. As a result, the size of memory for remembering the past rating/behavior will be decreased. The amounts of increase and decrease are decided based on the application. For example, in risky applications, after a change in the behavior of the trustee, the value of λ should be decreased sharply. The initial value of λ should be zero. Only after a number of successful interactions is λ allowed to increase.

The sparsity and cold start problem: We consider the patterns in how individuals and groups trust as another external context factor. Based on the characteristics of how and what individuals and groups trust, we have proposed that the like-mindedness of individuals and groups can be utilized to identify other trust relationships. For instance, if one knows that, with respect to a specific property, two parties are trusted by a large number of different trusters, one can assume that the two parties have similar trust characteristics. Thus, if one has a certain degree of trust in the first party, one can safely assume a similar trustworthiness for the other party. In an attempt to provide high quality recommendations and proper initial trust values, even when no complete trust path or user profile exists, we propose `TILLIT`, a model based on a combination of trust inferences and user similarities. Similarity is derived from the structure of a trust graph and users' trust behavior as opposed to other collaborative-filtering-based approaches that use ratings of items or users' profiles. We describe an algorithm realizing the approach based on a combination of trust inferences and user similarities, and validate the algorithm using a large, real-world data set [61].

Categorization: It is possible to draw information from feedback that is generated in a variety of situations, but for the feedback to be useful in other situations. For example, in online auctions, there are common factors between buying and selling activities that affect trust formation. Therefore, the feedback about a user as a buyer might be useful to calculate a reputation of the same user as a seller. We present a knowledge intensive and model-based, case-based reasoning framework that supports a system that can infer such information. The suggested method augments other work in environments where

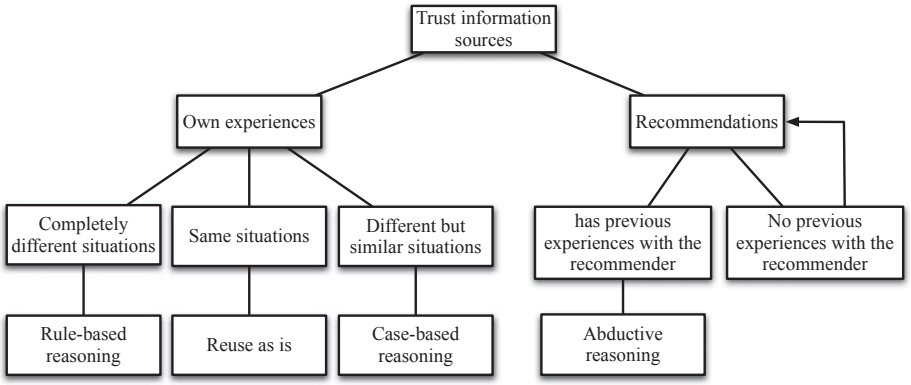

Figure 1. Different trust evaluation scenarios and their reasoning methods.

information is typically sparse (*e.g.*, there are many buyers and sellers, and it is unlikely that there is a previous transaction on which to calculate an accurate trust value). A trust rating can be calculated by inferring the lack of relationship information using other situational conditions. Such a solution allows better support for situation-aware trust and reputation management.

The CBR technique is particularly useful for tasks that are experience-intensive, involve plausible reasoning and have incomplete rules to apply. The fundamental principle of the CBR technique is similar to that of the human analogical reasoning process which employs solutions of past problems to solve current ones. The reasoning process is generally composed of three stages: remembering, reusing, and learning. Remembering is the case-retrieval process, which recalls relevant and useful past cases. In the reusing step, the CBR system uses the recalled cases to find an effective solution to the current problem. Learning is the process of case base enhancement. At the end of each problem-solving session the new case and the problem-solving experiences are incorporated into the casebase [62]. In our approach, the role of context is to generate candidate cases. This hypothesis-generation activity of the reasoner can be thought of as an instance of "cued recall" in cognitive psychology terminology. Context has been shown to have a major influence on remembering cases and its inclusion in case-based problem solving empowers the case-based approach. The strong dependence between the context and a powerful memory-retrieval arise most probably from the role context plays in similarity assessment of two cases (*i.e.*, the new and a past case). We proposed a rule-based reasoning model (far left in Figure 1) for decision making when the truster does not have own similar past experiences or available recommendations about the trustee either (this we have addressed in a paper under preparation). The trust judgment then resorts to a set of domain-specific association rules.

Our framework can be coupled with existing models to make them situation-aware. Our model uses the underlying model of trust and reputation management to transfer information between situations and can also be used to transfer information from one system into another to provide more potability. We validate the proposed framework for the Subjective Logic Model [6] and evaluate it by conducting experiments on a large, real-world data set [63].

Our second motivation for this work is trust transitivity. Trust is not always transitive in real life. For example, the fact that A trusts B to fix her car and B trusts C to look after his child does not imply that A trusts C to fix a car, or for child care. However, under certain semantic constraints, trust can be transitive and a trust referral system can be used to derive transitive trust. The semantic constraint is that the subject of trust should be the same along the entire path, for example all trust subjects should be "a good car mechanic" or "looking after a child". However, trust relations with the same subject are not always available. This constraint is relaxed in our work by introducing the notion of situation. We suggest that trust situations along a transitive trust path can be different but similar to each other. For instance, trust situations can be "to be a good car mechanic" or "to be a good motor mechanic". In this way, we are able to use trust information from available similar situations.

Initialization and low incentive for providing feedback: When a user first comes into a system, there is little information available to use to build a trust recommendation. Further, gathering such information is difficult when there is little incentive to provide feedback. We categorize the decision making process with respect to these two factors based on the familiarity of the truster with the situation and the trustee. Different combinations of incomplete knowledge are:

1) Unfamiliar situation, familiar trustee: If the truster has had previous interactions with the trustee or similar other trustees, but in different situations, she can still use her past experiences. But in these new situations, the truster needs to map the old and new situations and make the necessary adaptations in order to draw a conclusion. As we mentioned earlier, case-based reasoning is used to handle such situation-specificity of trust.

2) Familiar situation, unfamiliar trustee: If the truster has had previous interactions in the same situations with other trustees (*i.e.*, a stereotype of the trustee), the trust judgment then resorts to a set of domain-specific association rules. We propose a rule-based reasoning algorithm to handle this situation. Past trustees are grouped based on a common attribute

with the current trustee and the general trustworthiness of the group can be summarized. Then, an opinion about those trustees as a group is formed and the current trustee is included in that group. In this way, the opinion about the group is effectively transformed into an opinion about the prospect.

3) Unfamiliar situation, unfamiliar trustee: If there is no situational or trustee information, the trust model uses a default trust value since there is no information to be used for the initialization of trust.

Subjective and unfair ratings: Our approach to modeling the impact of subjectivity is based on the idea that a feedback provider's judgment method can be inferred and the target entity can be (re-)evaluated according to the value system of the receiver of the feedback. The judgment method is a function that maps an attribute value (*e.g.*, delivery time = late) to the value the feedback provider attached to that attribute (*e.g.*, unsatisfied). Thus, the receiver of the feedback will be able to translate (*i.e.*, eliminating subjectivity) for subsequent feedback from this particular user based on what she has learned. This method of extracting judgment information involves abductive reasoning.

Abductive inference is typically relied upon in imperfect domains, *i.e.*, in the face of incomplete or inconsistent information as well as in cases where the domain does not provide a strong theory. Abductive reasoning, in general, is reasoning from consequences to antecedents and describes the process of discovering hypotheses (*i.e.*, antecedent), and assesses the likelihood that a specific hypothesis entails a given a conclusion (*i.e.*, consequence). Inference of "it must have rained", upon seeing the grass wet is based on experiences: "when it rains, the grass gets wet. The grass is wet, then it must have rained." However, if there is a person near her car and a hose is on the grass, the inference would tend towards "when a person washes her car, the grass gets wet." This person may have washed her car "since she is near her car and there is a hose on the grass" [64].

We envisage that the truster can infer the *judgment method* of the recommender by observing the recommender's ratings and corresponding trustee's properties. For example, in cases where a recommender is known to consistently bias its ratings (*e.g.*, always exaggerating positively or negatively, or always reporting the opposite of what she thinks), it is in fact possible to "re-interpret" the ratings. This can be done by extracting the conditional relation between the trustee's properties as antecedents and the recommenders' ratings as consequences from the history of interactions[1]. Based on this information, the truster will be able to translate in future a new rating provided by the recommender into the actual properties of the trustee by employing abductive reasoning. Figure 2(a) shows the trust value computation by the truster without considering the subjective difference. The recommender sends a rating about the trustee to the truster based on his own observations of the trustee's properties and the truster simply uses this rating as is in her own trust model (decision making model) as if she has generated this rating herself. Figure 2(b) shows the same process, however the truster considers the subjective differences this time and re-interprets the rating from the recommender by way of inferring the judgment method of the recommender from the historical data.

This proposal is implemented using subjective logic [6]. This approach has been quantitatively compared with two other methods. The experiments show that an extended version of the "Beta trust model" [6], a trust model without the elimination of subjectivity, with our method, in which subjectivity is eliminated, outperforms the original model. Although our method is not aimed at addressing the deception problem, it is able to cope with deception when a majority of feedback providers give deceptive, yet consistent ratings. In addition, our suggested method for trust and reputation systems may also be applied to other systems that include a rating mechanism such as recommender systems.

VI. Future Research Directions

There are several unexplored areas for trust and reputation systems that present fertile opportunities for future research. The following list contains what we consider to be the most important open areas of research:

1) Most of the current trust and reputation mechanisms are centralized, resource-based, and personalized, which leaves space to research the suitability of many other types of system attributes. In addition, effective solutions need to be developed for the problems identified in Section III such as sufficient participation, easy identity changes, and strategic manipulation of online feedback.

2) Proposals from the academic community are not always deployable and are usually designed from scratch. Only in a very few cases do authors build on proposals from others. Hence, there is a need for a set of sound, accepted principles for building trust and reputation systems. The design space and limitations of mediated trust and reputation mechanisms should be explored and a set of design parameters that work best in different settings should be understood. Formal models of those systems in both monopolistic and competitive settings should be developed.

3) Universal testbeds and evaluation metrics for comparison of the relative efficiency of trust and reputation mechanisms compared to that of more established systems are needed and theory-driven guidelines should be developed to decide which set of mechanisms to use.

4) A comprehensive set of robustness evaluation methods and criteria and a standardized set of attack types should be defined. Trust and reputation system robustness can be evaluated by implementing them in a real environment or from a theoretical perspective by third parties.

5) New domains where reputation mechanisms can be usefully applied need to be defined.

6) A calculated trust value should be presented to users in ways that they can rely on this value. For example, the trust value should be accompanied with an explanation of the estimation grounds and an uncertainty value, which shows how much data has been used for this estimation. The importance

[1]The history contains two kinds of information for each interaction: the rating that the truster received from the recommender regarding the trustee before an interaction and the truster's own observation of the trustee's properties after the completion of the interaction.

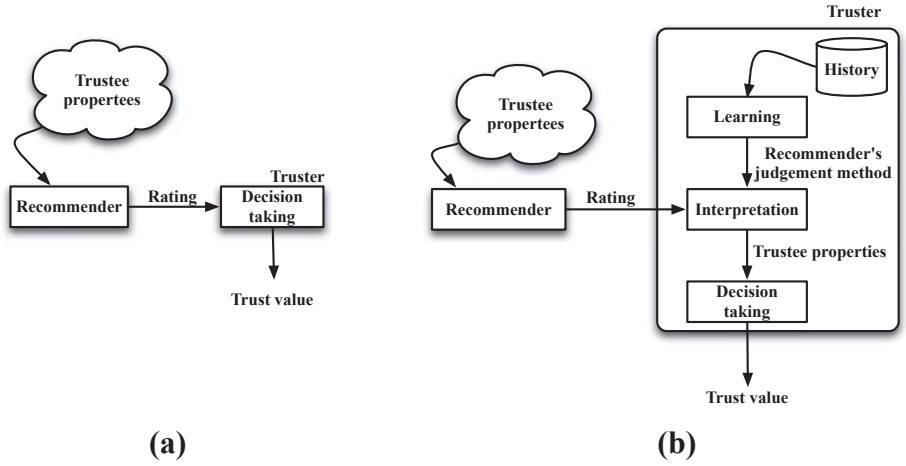

Figure 2. Trust value computation (a) without and (b) with subjectivity consideration.

of explanation interfaces in providing system transparency and thus increasing user acceptance has been well recognized in a number of fields.

7) A decision to trust is a decision tied with risk. Even when the expectations are well grounded, there is an element of risk in trust, a chance that those who are trusted will not act as expected. The risk should be justified in order to confirm the current trust and to strengthen it, otherwise if the other party defects, trust decreases dramatically. The estimation of this risk remains a problematic area. Game theory is a powerful tool for this purpose.

8) There are fundamental differences between traditional and online environments. Therefore, adequate online substitutes for the traditional cues to trust and reputation in the physical world should be found, and new information elements, specific to a particular online application, which are suitable for deriving measures of trust and reputation, should be identified.

9) Social acceptance of trust and reputation systems is another critical factor. For the more widespread and general usage of these systems, social acceptance by all parties is an issue that needs to be considered.

VII. Conclusions

Reputation systems confront many complex challenges, many of which yield no easy solutions. Efforts are underway to address these problems using a variety of approaches. This paper examines current techniques used in reputation management systems and outlines a set of problems and proposed solutions. Furthermore, we present a summary of solutions to address some of these problems and we propose a research agenda for trust and reputation systems.

References

[1] G. Swamynathan, K. Almeroth, and B. Zhao, "The design of a reliable reputation system," *Electronic Commerce Research*, vol. 10, no. 3-4, pp. 239–270, December 2010.

[2] G. Swamynathan, C. Wilson, B. Boe, B. Zhao, and K. Almeroth, "Do social networks improve e-commerce: A study on social marketplaces," in *Proceedings of the ACM Sigcomm Workshop on Online Social Networks (WOSN)*, August 2008.

[3] A. Abdul-Rahman, "A framework for decentralised trust reasoning," Ph.D. dissertation, University College London, 2004.

[4] D. Gambetta, "Can we trust trust," *Trust: Making and breaking cooperative relations*, pp. 213–237, 2000.

[5] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman, "Reputation systems," *Communications of the ACM*, vol. 43, no. 12, pp. 45–48, December 2000.

[6] A. Jøsang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Decision Support Systems*, vol. 43, no. 2, pp. 618–644, March 2007.

[7] J. Pujol, R. Sangüesa, and J. Delgado, "Extracting reputation in multi agent systems by means of social network topology," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2002, pp. 467–474.

[8] R. Chen and W. Yeager, "Poblano: A distributed trust model for peer-to-peer networks. Sun Microsystems, inc. White Paper," 2001.

[9] U. Maurer, "Modelling a public-key infrastructure," in *Proceedings of the European Symposium on Research in Computer Security (ESORICS)*, 1996, pp. 325–350.

[10] S. Marsh, "Formalising trust as a computational concept," Ph.D. dissertation, Dept. of Computing Science and Mathematics, University of Stirling, April 1994.

[11] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins, "Propagation of trust and distrust," in *Proceedings of the International Conference on the World Wide Web*, 2004, pp. 403–412.

[12] M. Blaze, J. Feigenbaum, and J. Lacy, "Decentralized trust management," in *Proceedings of the IEEE Symposium on Security and Privacy*, 1996, pp. 164–173.

[13] K. Regan, R. Cohen, and P. Poupart, "The advisor-pomdp: A principled approach to trust through reputation in electronic markets," in *Proceedings of the Conference on Privacy Security and Trust*, 2005.

[14] T. Huynh, N. Jennings, and N. Shadbolt, "An integrated trust and reputation model for open multi-agent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 13, no. 2, pp. 119–154, 2006.

[15] Y. Rebahi, V. Mujica-V, and D. Sisalem, "A reputation-based trust mechanism for ad hoc networks," in *Proceedings of the IEEE Symposium on Computers and Communications (ISCC)*, 2005, pp. 37–42.

[16] B. Yu and M. Singh, "An evidential model of distributed reputation management," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2002, pp. 294–301.
[17] J. Patel, W. Teacy, N. Jennings, and M. Luck, "A probabilistic trust model for handling inaccurate reputation sources," *Trust Management*, pp. 413–419, 2005.
[18] W. Conner, A. Iyengar, T. Mikalsen, I. Rouvellou, and K. Nahrstedt, "A trust management framework for service-oriented environments," in *Proceedings of the International Conference on the World Wide Web*, 2009, pp. 891–900.
[19] S. Kamvar, M. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks," in *Proceedings of the International Conference on the World Wide Web*, 2003, pp. 640–651.
[20] M. Gupta, P. Judge, and M. Ammar, "A reputation system for peer-to-peer networks," in *Proceedings of the International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, 2003, pp. 144–152.
[21] L. Xiong and L. Liu, "Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 7, pp. 843–857, 2004.
[22] H. Zhao and X. Li, "H-trust: A group trust management system for peer-to-peer desktop grid," *Journal of Computer Science and Technology*, vol. 24, no. 5, pp. 833–843, 2009.
[23] J. Sabater and C. Sierra, "Social regret, a reputation model based on social relations," *ACM SIGecom Exchanges*, vol. 3, no. 1, pp. 44–56, 2001.
[24] G. Zacharia, A. Moukas, and P. Maes, "Collaborative reputation mechanisms in electronic marketplaces," in *Proceedings of the Annual Hawaii International Conference on System Sciences*, January 1999.
[25] Z. Malik and A. Bouguettaya, "Rateweb: Reputation assessment for trust establishment among web services," *The VLDB Journal*, vol. 18, no. 4, pp. 885–911, 2009.
[26] K. Aberer, P. Cudré-Mauroux, A. Datta, Z. Despotovic, M. Hauswirth, M. Punceva, and R. Schmidt, "P-grid: a self-organizing structured p2p system," *ACM SIGMOD Record*, vol. 32, no. 3, pp. 29–33, 2003.
[27] A. Rahbar and O. Yang, "Powertrust: A robust and scalable reputation system for trusted peer-to-peer computing," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 18, no. 4, pp. 460–473, 2007.
[28] A. Singh and L. Liu, "Trustme: anonymous management of trust relationships in decentralized p2p systems," in *Proceedings of the International Conference on Peer-to-Peer Computing*, 2003, pp. 142–149.
[29] E. Damiani, D. di Vimercati, S. Paraboschi, P. Samarati, and F. Violante, "A reputation-based approach for choosing reliable resources in peer-to-peer networks," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2002, pp. 207–216.
[30] J. Golbeck, B. Parsia, and J. Hendler, "Trust networks on the semantic web," *Cooperative Information Agents VII*, pp. 238–249, 2003.
[31] C. Ziegler and G. Lausen, "Propagation models for trust and distrust in social networks," *Information Systems Frontiers*, vol. 7, no. 4, pp. 337–358, 2005.
[32] V. Buskens, "The social structure of trust," *Social Networks*, vol. 20, no. 3, pp. 265–289, 1998.
[33] J. Sabater and C. Sierra, "Review on computational trust and reputation models," *Artificial Intelligence Review*, vol. 24, no. 1, pp. 33–60, 2005.
[34] C. Dellarocas, "The digitization of word-of-mouth: Promise and challenges of online reputation systems," *Management Science*, vol. 49, no. 10, pp. 1407–1424, October 2003.
[35] R. Jurca and B. Faltings, "An incentive compatible reputation mechanism," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, July 2003, pp. 1026–1027.
[36] N. Miller, P. Resnick, and R. Zeckhauser, "Eliciting Honest Feedback in Electronic Markets," *Working Paper Series*, August 2002.
[37] R. Malaga, "Web-based reputation management systems: Problems and suggested solutions," *Electronic Commerce Research*, vol. 1, no. 4, pp. 403–417, 2001.
[38] R. Zeckhauser and P. Resnick, "Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system," *The Economics of the Internet and E-commerce*, pp. 127–157, 2002.
[39] J. Golbeck and J. Hendler, "Accuracy of metrics for inferring trust and reputation in semantic web-based social networks," in *Proceedings of the International Conference on Knowledge Engineering and Knowledge Management*, October 2004.
[40] M. Chen and J. Singh, "Computing and using reputations for internet ratings," in *Proceedings of the ACM Conference on Electronic Commerce*, October 2001, pp. 154–162.
[41] C. Dellarocas, "Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior," in *Proceedings of the ACM Conference on Electronic Commerce*, October 2000, pp. 157–164.
[42] S. Buchegger and J. Le Boudec, "A robust reputation system for mobile ad-hoc networks," in *Proceedings of the Workshop on the Economics of Peer-to-Peer Systems*, June 2004.
[43] F. Cornelli, E. Damiani, S. di Vimercati, S. Paraboschi, and P. Samarati, "Choosing reputable servents in a P2P network," in *Proceedings of the International Conference on the World Wide Web*, May 2002, pp. 376–386.
[44] K. Hoffman, D. Zage, and C. Nita-Rotaru, "A survey of attack and defense techniques for reputation systems," *ACM Computing Surveys*, vol. 42, no. 1, pp. 1–31, December 2009.
[45] N. Griffiths, "Task delegation using experience-based multi-dimensional trust," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2005, pp. 489–496.
[46] T. Tran and R. Cohen, "Improving user satisfaction in agent-based electronic marketplaces by reputation modelling and adjustable product quality," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, July 2004, pp. 828–835.
[47] R. Kerr and R. Cohen, "Modeling trust using transactional, numerical units," in *Proceedings of the International Conference on Privacy, Security and Trust: Bridge the Gap Between PST Technologies and Business Services*, October/November 2006, pp. 21:1–21:11.
[48] C. Dellarocas, "Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality," in *Workshop on Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, July 2002, pp. 93–112.
[49] R. Neisse, M. Wegdam, and M. van Sinderen, "Context-Aware Trust Domains," in *Proceedings of the European Conference on Smart Sensing and Context*, 2006.
[50] R. Neisse, M. Wegdam, M. van Sinderen, and G. Lenzini, "Trust Management Model and Architecture for Context-Aware Service Platforms," *Lecture Notes in Computer Science*, vol. 4804, p. 1803, 2007.
[51] E. Gray, Y. Chen, and C. Jensen, "Initial Investigation into Cross-context Trust and Risk Assessment," in *Proceedings of the IASTED International Conference on Communication, Network, and Information Security*, 2003, pp. 56–61.
[52] S. Holtmanns and Z. Yan, "Context-Aware Adaptive Trust," in *Proceedings of the Ambient Intelligence Developments Conference*, 2006.
[53] M. Rehak, M. Gregor, M. Pechoucek, and J. Bradshaw, "Representing Context for Multiagent Trust Modeling," in *Proceedings*

of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT), 2006, pp. 737–746.
[54] E. Bagheri and A. A. Ghorbani, "Behavior analysis through reputation propagation in a multi-context environment," in *Proceedings of the International Conference on Privacy, Security and Trust: Bridge the Gap Between PST Technologies and Business Services*, 2006, pp. 40:1–40:7.
[55] E. Bagheri, M. Barouni-Ebrahimi, R. Zafarani, and A. Ghorbani, "A Belief-Theoretic Reputation Estimation Model for Multi-context Communities," *Lecture Notes in Computer Science*, vol. 5032, p. 48, 2008.
[56] K. Wan and V. Alagar, "An intensional functional model of trust," *Trust Management II*, pp. 69–85, 2008.
[57] N. Gujral, D. DeAngelis, K. Fullam, and K. Barber, "Modeling multi-dimensional trust," in *the Proceedings of the Workshop on Trust in Agent Societies*, 2006, pp. 8–12.
[58] M. Rehak and M. Pechoucek, "Trust modeling with context representation and generalized identities," in *Proceedings of the International Workshop on Cooperative Information Agents XI (CIA)*, 2007, pp. 298–312.
[59] S. Toivonen, G. Lenzini, and I. Uusitalo, "Context-aware trust evaluation functions for dynamic reconfigurable systems," in *Proceedings of the Models of Trust for the Web Workshop*, May 2006.
[60] M. Tavakolifard and S. Knapskog, "A probabilistic reputation algorithm for decentralized multi-agent environments," *Electronic Notes in Theoretical Computer Science*, vol. 244, pp. 139 – 149, August 2009.
[61] M. Tavakolifard, P. Herrmann, and S. Knapskog, "Inferring trust based on similarity with TILLIT," *Trust Management III*, pp. 133–148, 2009.
[62] C. Jung, I. Han, and B. Suh, "Risk Analysis for Electronic Commerce Using Case-Based Reasoning," *Int. J. Intell. Sys. Acc. Fin. Mgmt*, vol. 8, pp. 61–73, 1999.
[63] M. Tavakolifard, P. Herrmann, and P. Öztürk, "Analogical trust reasoning," *Trust Management III*, pp. 149–163, 2009.
[64] G. Harman, "Enumerative induction as inference to the best explanation," *The Journal of Philosophy*, vol. 65, no. 18, pp. 529–533, 1968.

Mozhgan Tavakolifard is a PhD candidate at Norwegian University of Science and Technology. Her research interests are in social networks, trust and reputation systems, recommender systems and social computing in general. She has a MSc in Information Technology from Amirkabir University (Tehran Polytechnic), and a BSc from Sharif University of Technology.

Kevin C. Almeroth is currently a Professor in the Department of Computer Science at the University of California in Santa Barbara where his main research interests include computer networks and protocols, wireless networking, multicast communication, large-scale multimedia systems, and mobile applications. At UCSB, Dr. Almeroth is the Associate Director of the Center for Information Technology and Society (CITS), a founding faculty member of the Media Arts and Technology (MAT) Program, Technology Management Program (TMP), and the Computer Engineering (CE) Program. In the research community, Dr. Almeroth has authored nearly 200 refereed papers. He is a Member of the ACM and a Senior Member of the IEEE.