

A Reference Model for Reputation Systems

Sokratis Vavilis^{a,*}, Milan Petković^{a,b}, Nicola Zannone^a

^a*Eindhoven University of Technology, Den Dolech 2, Eindhoven 5612AZ, Netherlands*

^b*Philips Research Eindhoven, High Tech Campus 34, Eindhoven 5656AE, Netherlands*

Abstract

Recent advances in ICT have led to a vast and expeditious development of e-services and technology. Trust is a fundamental aspect for the acceptance and adoption of these new services. Reputation is commonly employed as the measure of the trustworthiness of users in on-line communities. However, to facilitate their acceptance, reputation systems should be able to deal with the trust challenges and needs of those services.

The aim of this survey is to propose a framework for the analysis of reputation systems. We elicit the requirements for reputations metrics along with the features necessary to achieve such requirements. The identified requirements and features form a reference framework which allows an objective evaluation and comparison of reputation systems. We demonstrate its applicability by analyzing and classifying a number of existing reputation systems. Our framework can serve as a reference model for the analysis of reputation systems. It is also helpful for the design of new reputation systems as it provides an analysis of the implications of design choices.

Keywords: Reputation Systems, Reference Model, Trust Requirements

1. Introduction

Advances of ICT have led to overall digitization of processes in data lifecycle management and consequently resulted in improved efficiency and cost savings. Each of us is constantly exposed to emerging digital technologies, either at home or at work, with an increasing number of business transactions daily carried out over the Internet. However, to fully exploit

*Corresponding author

Email addresses: s.vavilis@tue.nl (Sokratis Vavilis),
milan.petkovic@philips.com (Milan Petković), n.zannone@tue.nl (Nicola Zannone)

the potentials of e-services (e.g., e-commerce, e-business, e-health) and facilitate their adoption, it is important to establish and manage trust amongst the parties involved in the transactions [1].

Reputation systems play an important role in the process of trust establishment and management. When a user needs to make a trust decision whether to engage or not in an interaction with an e-service, he takes very much into account the reputation of the service. The user's past experience as well as the experience of the other users with the service influences his decision whether to repeat this interaction in the future. Therefore, a reputation system, which helps in managing reputations in digital world (for example by collecting, distributing and aggregating feedback about entity's behavior), becomes a fundamental component of the trust and security architecture of any ICT system or service.

However, the application and adoption of reputation systems in e-services rely on their ability to address the trust challenges that such services have to deal with. Therefore, the design of reputation systems requires identifying the trust needs of e-services and of the application domain in which such services are deployed. In addition, when selecting a reputation system to be applied to an e-service, it is important to verify whether the selected reputation system meets the trust requirements for such a service.

In this paper we address these issues by presenting a framework for the analysis of reputation systems. In particular, we:

- elicit the requirements for reputation systems from a literature study;
- identify the features necessary to fulfill these requirements;
- present a reference framework for the analysis and evaluation of reputation metrics;
- demonstrate the applicability of the framework by comparing and classifying several well-known reputation systems.

The analysis presented in this work aims to serve both researchers that develop reputation systems and practitioners that intend to employ reputation systems for their services. On the one hand, it provides researchers an analysis of the implication of design decision. On the other hand, it provides practitioners a reference framework that can assist them in the selection of a reputation system that meets their needs.

Several surveys on reputation systems can be found in the literature [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13]. Similarly to most of these surveys, our work identifies the main features to be supported by reputation systems

and evaluate existing systems against the defined features. In contrast to them, this work mainly focuses on the trust information used to assess the reputation and aggregation method, offering a more fine-grained analysis of reputation metrics. Moreover, existing surveys present features as abstract concepts. In contrast, we identify features from the requirements for reputation systems; features employed in this work are, thus, closer to the real needs of reputation systems.

The remainder of the paper is organized as follows. Next section introduces the basic concepts of reputation systems. Section 3 discusses the requirements for reputation systems along with the features needed for their fulfillment. Section 4 analyzes existing reputation systems with respect to the identified features and requirements. Section 5 provides guidelines on the application of the framework. Finally, Section 6 discusses related work, and Section 7 concludes the paper.

2. Overview of Reputation Systems

Reputation has been proposed as a measurement of a user's trustworthiness based on his past behavior [14] and it is used to predict his future behavior [15]. Typically, users rate other users on the basis of their interactions. In particular, a *rating* is the judgment that a user (*origin*) gives to another user (*target*) about a certain interaction that occurred between them (*scope*). Reputation systems assess the reputation of a user by aggregating the ratings that other users have given to that user.

Reputation systems can be analyzed from three dimensions [9], namely *formulation*, *calculation* and *dissemination*. The formulation dimension describes the mathematical model and input for the assessment of reputation values. It includes two main aspects: the reputation *measure* and the mathematical model (*metric*) used to aggregate ratings. Reputation can be measured using discrete or continuous values. Metrics can be based on simple summation or average of ratings [16], fuzzy logic [17, 18], flow-based models [19, 20, 21, 22], probabilistic models such as Bayesian systems [23, 24, 25], beta probability density [26, 27] and subjective logic [28, 29].

The calculation dimension addresses the practical design and implementation of the algorithm for assessing reputation, whereas the dissemination dimension focuses on the mechanisms for the distribution and storage of ratings and reputation values among entities within the system. The main feature of the calculation and dissemination dimensions is the *structure* of the reputation system, which can be either *centralized* or *decentralized*. Centralized systems like auction or expert sites are characterized by the presence

of a central authority which is responsible for the collection and storage of user’s ratings, and for the calculation of reputation values and their dissemination. On the other hand, decentralized systems like Peer-to-Peer (P2P) networks and Multi-Agent Systems (MAS) do not have neither a central authority nor a fixed network topology that can be used to control the entities within the system; rather each entity is responsible for controlling its data and resources. In these systems, the storage of ratings and calculation of reputation are distributed among the entities within the system.

3. Requirements and Features for Reputation Systems

This section presents the main requirements for reputation systems along with the features needed to achieve such requirements. The elicited requirements and features form a reference model for the analysis and comparison of reputation systems.

3.1. Requirements

For the correct and secure functioning of reputation systems several desirable requirements should be satisfied. In this work, we are mainly interested in requirements which ensure that assessed reputation values reflect the actual trustworthiness of users. Based on a literature study, we have identified three groups of requirements (Table 1). Requirements in the first group focus on the formulation dimension. The other two groups contain requirements about the fair treatment of newcomers and the integrity of reputation values (w.r.t. the calculation and dissemination dimensions).

The first group of requirements (R1 to R8) focuses on the information and aggregation method used for the assessment of reputation values. R1 and R2 require ratings and reputation values to accurately discriminate user behavior. R3 and R4 focus on the “quality” of information used to assess reputation values. A user can (un)intentionally provide incorrect ratings about an interaction he had. For instance, a malicious user may give negative ratings to a user with the purpose of decreasing its reputation [9, 22]. Moreover, he can subvert the reputation system by first creating a large number of pseudonymous entities, and then using them to influence the reputation of a target user [30]. In particular, if users are able to rate themselves, they can provide a series of self-promoting ratings, leading to an unfair increasing of their own reputation [9]. To prevent such attacks, reputation systems should be able to discriminate “incorrect” ratings (R3). R4 refines R3 by explicitly forbidding self-rating.

ID	Requirement	Source
R1	Ratings should discriminate user behavior	[5],[9]
R2	Reputation should discriminate user behavior.	[5],[9],[12],[13],[22]
R3	The reputation system should be able to discriminate “incorrect” ratings.	[5],[9],[11],[12],[13],[28]
R4	An entity should not be able to provide rating for itself.	[9],[11],[20]
R5	Aggregation of ratings should be meaningful.	[13],[27],[34]
R6	Reputation should be assessed using a sufficient amount of information.	[12],[13],[34]
R7	The reputation system should differentiate reputation information by the interaction it represents.	[11]
R8	Reputation should capture the evolution of user behavior.	[11],[12],[13],[34]
R9	Users should not gain advantage of their new status.	[9],[12],[22]
R10	New users should not be penalized for their status.	[20],[22]
R11	Users should not be able to directly modify ratings.	[9],[11],[20]
R12	Users should not be able to directly modify reputation values.	[9],[11],[20]
R13	Users should not be responsible to directly calculate their own reputation.	[9],[11],[20]

Table 1: Requirements for reputation systems

Requirements R5 to R8 deal with the type and amount of information used to assess reputation values. Interactions between entities can differ significantly in their nature, making it difficult to draw conclusions about the reputation of entities. For instance, aggregating the ratings referring to different types of interactions would result in reputation values that may not accurately reflect the trustworthiness of entities. Therefore, reputation values should be assessed using comparable trust information (R5). However, the reliability of reputation values depends on the amount of information used to calculate them [31]. Due to the restrictions imposed by R5, reputation may only rely on a small amount of information. R6 relaxes R5 by allowing the use of a larger amount of information in the assessment of reputation, while ensuring that the obtained reputation values remain meaningful. Interactions between entities may also differ in their “cost” (e.g., economic transactions) [32]. For example, in an online auction site, a user can build his reputation through transactions involving small amount of money, and then take advantage of his gained reputation to cheat other users in a transaction involving substantial amount of money. R7 requires reputation systems to differentiate ratings with respect to the cost of transactions. Finally, R8 focuses on temporal aspects. Reputation is build upon the knowledge of past interactions. The behavior of a user can change over time. For instance, a malicious entity might be fair in his interactions for a period in order to build positive reputation and be able to successfully

deploy his attack. Therefore, user behavior evolution should be captured in the assessment of reputation to reflect the actual trustworthiness of users.

The second group of requirements (R9 and R10) addresses the fair treatment of new users. When new users join the system, their behavior is unknown. Typically, reputation systems assign a default reputation value to new users. Such a value, however, should not penalize them for their status. If newcomers are treated as users with bad reputation, they may be never selected by other users and thus they cannot build their reputation [20]. At the same time, reputation systems should prevent users to gain advantage of their new status. Indeed, to avoid the consequences of their actions, users with bad reputation may change their identity by re-joining as a new user (the so called white-washing attack [33]). Requirements R9 and R10 define the boundaries in the selection of the reputation value for new users.

The last group of requirements (R11 to R13) addresses calculation and dissemination issues regarding the integrity of reputation values and ratings. R11 and R12 impose reputation systems to protect ratings and reputation values from unauthorized manipulation during transmission and storage. This, however, may not be sufficient to guarantee their integrity. If users are involved in the calculation of their own reputation, they may influence the obtained value. R13 aims to prevent such a malicious behavior.

3.2. Features

The requirements in Table 1 constitute the basic and desirable characteristics that a reputation system should satisfy. Their fulfillment can be achieved by *features* or *technical solutions*. Features identify the types of trust information that should be considered when assessing reputation as well as the properties of reputation values and aggregation method, whereas technical solutions are mechanisms that can be employed to satisfy the requirements in Table 1. In this article we mainly focus on features related to the formulation dimension of reputation systems. At the end of the section we also discuss some technical solutions that can be used to satisfy the elicited requirements.

Reputation should accurately characterize user behavior. As users may have different behavior, being either good, neutral or bad, ratings and reputation values should be able to accurately represent such a range of behaviors. We represent the ability of reputation metrics to express the entire range of user behaviors using the *trust and distrust* feature. In addition, reputation value should be expressed in an absolute way (*absolute reputation values*) rather than in relation with the other users (e.g., ranking) [22]. Indeed, a ranking may provide misleading perception of users' behavior; in a

reputation system in which reputation is only able to rank users (in contrast to determine their actual trustworthiness), users who are not trustworthy can have a high position in the ranking.

Several aspects of user interactions should be considered in the assessment of reputation. Interactions may differ in their topic. For instance, users might interact for medical advices on different diseases. Users, for example medical professionals, might not have the ability to provide the same quality of medical advices for every disease as it might be out of their area of expertise. We use *interaction scope* to indicate whether reputation systems are able to discriminate ratings and reputation values with respect to the type of interaction. However, different interaction scopes might have a close semantic relation. For instance, a medical advice on infectious diseases has a high degree of similarity with a medical advice on the influenza virus. *Scope similarity* is used to denote whether the degree of similarity between scopes is used to assess reputation. *Interaction context* denotes the ability of reputation metrics to discriminate interactions on the basis of their cost. Finally, we use *timestamp* to represent whether the (exact) time in which an interaction occurred is considered in the assessment of reputation.

In real life situations, users do not know all the other users within the system. *Trust transitivity* has been introduced to derive new trust relations from existing trust relations [2, 34]. Intuitively, trust transitivity determines the indirect trust between two users due to their trust relation with other users. However, the ability of a user to provide a service may differ from his ability to provide *recommendations* about other users. For instance, an entity might be honest in its interactions but provide dishonest recommendations to decrease the reputation of other entities (e.g., competitors) for personal benefit. In this respect, it is necessary to distinguish *functional trust*, i.e., the trust on a user's ability to provide a service, and *referral trust*, i.e., the trust on a user's ability to provide recommendations [5, 28]. In addition, users typically prefer to relay on the opinions of users they trust [35]. We use *(un)certainty* to represent the ability of reputation metrics to specify the level of confidence on trust information.

The satisfaction of the requirements in Table 1 may also require employing technical solutions. For instance, transaction proofs [36] can be employed to ensure that a transaction took place and thus to ensure that ratings are not artificial. Moreover, technical solutions such as the application of cryptographic mechanisms [37] and Public Key Infrastructure (PKI) [38] provide the necessary infrastructure for integrity and authentication.

Features \ Requirements	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10
F1 Trust/Distrust	✓	✓	✓	✗	✗	✗	✗	✗	✓	✗
F2 Absolute Rep. Values	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗
F3 Origin/Target	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗
F4 (un)Certainty	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓
F5 Interaction Scope	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗
F6 Scope Similarity	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
F7 Trust Transitivity	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
F8 Functional vs Referral	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
F9 Interaction Context	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗
F10 Timestamp	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗

Table 2: Requirements/Features

3.3. Mapping between Features and Requirements

In this section we discuss which features are needed to satisfy the elicited requirements. A summary is presented in Table 2. Notice that the table presents necessary but not sufficient conditions for the fulfillment of requirements. Indeed, additional technical solutions like cryptographic mechanisms or transaction proofs may be necessary to achieve the requirements.

To precisely characterize the actual behavior of an entity (R1 and R2), reputation systems should be able to capture the range of user behavior (trust and distrust). Moreover, they should be able to capture entities’ confidence on trust information ((un)certainty). Finally, reputation values should represent absolute measurements (absolute reputation values).

The protection against “incorrect” ratings (R3) requires identifying the origin of ratings. This makes it possible to discriminate the ratings on the basis of the reputation of the origin: if the origin is not trustworthy, the ratings given by such a user might have a low influence in the assessment of reputation values. In addition, reputation systems should accurately characterize user behavior (trust and distrust). The combination of these features allows a reputation system to reduce the effect of incorrect ratings. In addition, the separation between functional and referral reputation is needed because honest behavior during interactions does not imply that entities also provide honest recommendations. Finally, a reputation system should specify the level of uncertainty characterizing trust information to prevent that reputation is assessed on the basis of unreliable information. To reduce the risk of self-promoting attacks (R4), the origin of the ratings should be identified and be different from the target. This feature alone, however, is not sufficient to fulfill R4. It should be coupled with technical solutions, for instance identification and authentication mechanisms, to prevent that an entity can create sybils or re-enter the system as a new user.

For the aggregation of ratings to be meaningful (R5), only comparable information should be aggregated. Therefore, aggregated ratings should have the same interaction scope. Functional and referral reputation should also be distinguished and properly aggregated. Nonetheless additional trust information may be used to assess reputation (R6) by inferring it using ratings with similar interaction scope (similarity scope) and trust transitivity. To ensure that reputation reflects the actual behavior of users, interaction context is used to differentiate trust information by the interaction it represent (R7), and timestamp to capture users' behavior evolution (R8).

The fair treatment of new users (R9 and R10) can be accomplished through features and by employing technical solutions. On the one hand, a default value for those users should be carefully chosen in such a way that they can be distinguished by existing users. To this end, reputation values should be able to express the entire range of entities' behavior (trust and distrust). For instance, problems may arise in systems in which only negative ratings are used. Here, new users automatically get high reputation. In contrast, assigning them a bad/neutral reputation may lead to the inactivity of new users. Thus, uncertainty can be used to characterize new users. On the other hand, technical solutions for authentication and identification can be employed to forbid malicious users re-entering the system as new users. In addition, the creation of a new identity and re-entering the system should be costly for users. For instance, users might have to pay an entry fee which may include computational or other costs [39, 40]. In addition, technical solutions can be employed to allow a degree of randomness in the user selection [20], giving thus new users the possibility to prove them selves.

To protect ratings and reputation values against unauthorized manipulation (R11 to R13), information integrity solutions such as the application of cryptographic mechanisms, can be applied. Message authenticity can be assured using mechanisms such as PKI. The structure of the reputation system can also provide a solution itself to the problem of integrity. In particular, in a centralized system users cannot interfere with the calculation and dissemination of the reputation value, as long as the system is secure.

4. Survey of Reputation Systems

This section analyzes existing reputation systems with respect to the requirements and features introduced in Section 3.

4.1. Research Design and Reputation System Selection

A large variety of reputation systems have been proposed in the literature. Each system has been designed to provide a solution for a specific environment and for a specific purpose. For instance, reputation systems are used in P2P and MAS to isolate malicious users; in auction sites like eBay they provide information about the credibility of sellers and buyers. However, reputation systems share common concepts in their design regardless of the particular application domain.

In this work, we use the reputation metric as the main criterion for the analysis of reputation systems. In particular, we use the features presented in Section 3.2 as a baseline for the analysis. For each reputation metric, we verify whether it provides support for the identified features. Conclusions on requirement satisfaction are drawn by analyzing the supported features against the mapping between requirements and features (Section 3.3) as well as by assessing the technical solutions employed by reputation systems.

Several types of metrics are used for assessing the reputation. In this section, we focus on summation and average of ratings (referred as counting), flow-based, probabilistic and fuzzy logic metrics. We have selected the most representative reputation systems for each type of metric.

The literature on reputation systems also presents several frameworks like SuperTrust [41] and TrustGuard [36] that provide guidelines for the application of reputation systems along with additional technical solutions (i.e., usage of cryptography and transaction proofs). The analysis of these proposals, however, is out of the scope of this work as they do not focus on the reputation metric, but rather provide a general framework in which existing metrics can be accommodated.

4.2. Analysis with respect to Features

In this section, we analyze existing reputation systems with respect to features presented in Section 3.2. The results of the analysis are summarized in Table 3. The analyzed reputation systems are grouped with respect to the type of reputation metric which is employed. Feature support is expressed in a scale of values, which ranges from full support (✓), partial support (*), limited support (†), and no support (✗). In particular, partial support represents that the feature is partially supported, while limited support that only very basic characteristics are supported.

Summation and average of ratings are the simplest ways of assessing reputation. In the first case, reputation is the sum of positive and negative ratings. In the second case, an average is calculated. The main advantage

of this approach is that the model is straightforward. Moreover, due to the nature of the calculation, the resulting reputation values are absolute (i.e., not expressed in relation to other users). The eBay online auction community [16] is a representative example of reputation system using summation. When an object is sold, the buyer and seller have the possibility to rate each other. Ratings can be positive, neutral or negative and thus cover the entire range of user behavior (trust and distrust). Moreover, eBay partially supports the specification of the interaction scope: a user can provide ratings for a number of pre-defined scopes such as “item as described” and “shipping”. In addition, eBay records the time in which a transaction occurred.

REGRET [42] is a reputation system based on weighted average. In REGRET, ratings specify the rating value, origin and target of an interaction. They also contain other information about the interaction such as outcome of the interaction, scope and time in which the interaction occurred. REGRET supports trust transitivity. In addition, reliability metrics based on the interaction context are employed to represent the origin’s confidence on the provided rating. Finally, REGRET supports an ontological dimension of interactions which provides a metric to assess the similarity between scopes. FIRE [43] is an extension of REGRET that uses the same model for evaluating the reputation of an entity. On top of REGRET, FIRE implements a witness reputation system that provides a form of referral reputation.

Another reputation system based on weighted average is CORE [44]. CORE has been designed to assess reputation for different scopes such as packet forwarding and routing tasks in MANETs. The main characteristic of CORE is that only positive ratings are propagated, resulting in partial support for trust and distrust. Moreover, CORE partially supports uncertainty by providing a confidence metric based on the number of ratings used for the calculation and their variance. This reputation system also uses timestamps to record the time in which an interaction occurred.

The next class of reputation systems we have analyzed consists of flow-based reputation systems. Such reputation systems use Markov chains as the mathematical foundation. One of their main features is trust transitivity: reputation is assessed based on iteration through arbitrary chains, and ratings are weighted with respect to the reputation of each rating’s origin. However, they often assume that users who are trustworthy to provide a service are also trustworthy to provide recommendation. Accordingly, they do not distinguish between functional and referral trust. A reputation system representative of this class is EigenTrust [20]; other examples are PageRank [19] and SALSA [21]. In EigenTrust the local trust value of a user is calculated as the difference between satisfactory and unsatisfactory ratings

	eBay [16]				REGRET [42]				FIRE [43]				CORE [44]				EigenTrust [20]				Absolute EigenTrust [22]				Peertrust [45]				Beta [26]				Travos [46]				Dirichlet [47]				Subjective logic [28, 29]				CertainTrust [48]				Hedaquin [27, 49]				Fuzzy Trust [51]				FuzzyTrust/eTrust [18]			
Metric	Counting				Flow				Probabilistic				Fuzzy																																															
Structure	C	D	D	D	D	C	D	C	D	C	C	D	C	D	D																																													
Measure	D	C	C	C	C	C	C	C	C	C	C	C	C	C	C																																													
F1	✓	✓	✓	*	*	✓	*	✓	✓	✓	✓	✓	✓	✓	✓																																													
F2	✓	✓	✓	✓	✗	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓																																													
F3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓																																													
F4	✗	*	*	✗	✗	✗	✗	*	*	✗	✓	✓	✓	✗	✗																																													
F5	*	✓	✓	✓	†	†	†	†	†	†	✓	✓	✓	✗	✗																																													
F6	✗	*	*	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗																																													
F7	✗	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	✓	✓	✓																																													
F8	✗	✗	✓	✗	✗	✗	*	✓	✓	✗	✓	✓	✓	✓	✗																																													
F9	✗	✓	✓	*	✗	✗	✓	✗	✗	✗	✗	✗	✗	✓	✓																																													
F10	✓	✓	✓	✓	✗	✗	✓	✓	✗	✓	✓	✗	✓	✓	✓																																													

Legenda

Structure: C: Centralized D: Decentralized
Measure: C: Continuous D: Discrete
Features: ✓: Full support †: Limited support * : Partial support ✗: No support

Table 3: Survey of Reputation Systems and Classification w.r.t. Features

received by the user. Then, local trust values are normalized by dividing local trust value by the sum of the local trust value of all users. During the normalization process, negative local trust values are replaced by zero in order to obtain a Markov chain. This normalization procedure has two major drawbacks. First, normalized trust values do not distinguish users with negative reputation from users with neutral reputation. Moreover, trust values form a ranking rather than providing an absolute measure of trust. Finally, EigenTrust has limited support for scope specification as it assumes that reputation is solely for file exchange in the network. Absolute EigenTrust [22] extends EigenTrust by assessing absolute reputation values and aggregating both positive and negative ratings.

PeerTrust [45] is another flow-based reputation system for P2P. Although PeerTrust shares many similarities with EigenTrust, more factors are taken into account when assessing user trustworthiness. In PeerTrust the repu-

tation of a peer that had no interaction remains undefined. In addition, recommendations from other users are weighted with respect to their credibility. Thus, PeerTrust is marked to have partial support for referral trust. PeerTrust provides support for the interaction context which specifies information such as the number of transactions, value of the transaction and number of shared files. Finally, PeerTrust records the time when a transaction occurred, although it is part of the transaction context rather than a separate feature.

Other reputation metrics use probability density functions (PDF) to represent the expected outcomes of an interaction. When the rated events are of binary nature (e.g., satisfactory and unsatisfactory), the beta PDF can be used to compute reputation. Reputation can be represented in the form of beta PDF parameters which represent the number of positive (denoted by r) and negative interactions (denoted by s), or in the form of the probability expectation value of the beta PDF. Reputation can be also accompanied by a confidence parameter. Uncertainty is implicitly represented by the number of ratings: a high value of $r + s$ indicates high certainty, whereas a low value indicates low certainty.

The Beta reputation system [26] is an example of reputation system based on beta PDF. This system uses two operators for trust transitivity and aggregation, namely discounting and consensus. Discounting is used to build trust chains, whereas consensus is used to aggregate opinions of different users with the same scope. These operators, however, act on fixed trust networks, hence discarding (uncertain) information from the computation of reputation. The Beta reputation system also supports the aging of ratings. In particular, this mechanism provides a longevity factor which, with the usage of timestamps, gives more weight to recent ratings.

Another reputation system based on beta PDF is Travos [46]. Differently from the Beta reputation system, Travos provides a mechanism for detecting misleading and malicious recommendations. This is accomplished based on the past recommendation a user has given and the outcome of the corresponding interactions. This mechanism resembles the distinction between functional and referral reputation because it provides a distinction between the ability to provide a service and the ability to provide recommendations.

A disadvantage of reputation systems based on beta PDF is that they only support two event states and ratings (i.e., satisfactory and unsatisfactory) [47]. This excludes the possibility of fine-grained ratings (e.g., bad, mediocre, good, excellent). The Dirichlet reputation system [47] extends the Beta reputation system by adopting the Dirichlet distribution, the multivariate generalization of the beta distribution, as PDF. However, the Dirichlet

reputation system does not support discounting and, thus, trust transitivity.

Subjective logic [28, 29] extends probabilistic logic with an explicit notion of uncertainty. In subjective logic ratings are formed as opinions about a certain scope. An opinion is a tuple (b, d, u) where b , d and u represent belief, disbelief and uncertainty respectively such that $b + d + u = 1$. Belief represents the probability of a statement to be true, disbelief the probability to be false, and uncertainty to be unknown. Moreover, subjective logic support trust transitivity and a rating aging factor similarly to the Beta reputation system. CertainTrust [48] is a reputation system similar to subjective logic; it has, however, some notable differences. First, uncertainty is independent from belief and disbelief. In particular, it is defined on the basis of the amount of the available information rather than through belief and disbelief. Moreover, CertainTrust does not support temporal aspects.

Hedaquin [27, 49] is a reputation system for ensuring reliability of information in healthcare that builds on the Beta reputation system and subjective logic. Ratings in Hedaquin fall into two classes, namely functional and referral ratings, which result to functional and referral reputation respectively. Ratings are similar to the ones in Beta reputation system, but they are enriched with a certainty factor which is related to the amount and quality of the available information. Ratings also have a timestamp allowing the systems to give more importance to recent interactions. The assessment of reputation is as in the Beta systems; however, each rating is weighted according to its certainty. In addition, Hedaquin allows the specification of arbitrary interaction scopes and uses scope similarity metrics for weighting ratings with respect to their scope similarity.

Some reputation systems use fuzzy concepts to describe an entity's behavior while offering the ability to handle imprecision of information. In particular, membership functions are used to define the degree to which a fuzzy variable is a member of a set (e.g., very trustworthy, trustworthy, partially trustworthy, untrustworthy, very untrustworthy or unknown) [50]. Therefore, reputation based on fuzzy logic can represent both trust and distrust in an absolute way. A reputation system based on this type of metrics is Fuzzy Trust [51]. Each entity records local ratings that include the rating value and the time in which the interaction occurred. The assessment of reputation uses trust transitivity. Recommendations in Fuzzy Trust contain the rating value, detailed time information about the interaction and information about the context. The received ratings are weighed according to the time and credibility of the rating, which is a form of referral trust. Another reputation system based on fuzzy logic is presented in [18]. This reputation system is similar to Fuzzy Trust; however, in [18] the received ratings are

weighted with respect to the general reputation of the recommender and the interaction context. Therefore, it does not support referral trust.

4.3. Analysis with respect to Requirements Satisfaction

In this section we analyze which requirements in Table 1 are satisfied by existing reputation systems. The results of the analysis are presented in Table 4. They are obtained by combining Tables 2 and 3. In this section we mainly discuss the cases where the results are not straightforward.

Requirements R1 and R2 are only fulfilled by approaches based on subjective logic. Approaches based on counting and Beta PDF partially fulfill these requirements due to limited or no support of uncertainty. In addition, CORE allows the propagation of only positive values and, thus it does not provide the information necessary to discriminate users' behavior. Flow-based systems except Absolute EigenTrust, are not able to accurately measure users' trustworthiness due to the normalization process.

Reputation systems that support all the required features are able to discriminate incorrect ratings (R3). Systems that support the features but provide limited or no support for uncertainty fulfill this requirement only partially. Systems that support at least partially the majority of features provide limited satisfaction. Finally, systems that do not support neither uncertainty nor functional and referral recommendations do not fulfill R3. An interesting case is eBay: although it does not support all the necessary features, it allows rating only when an interaction took place, which is similar to the use of interaction proofs.

All analyzed systems are able to identify the origin of ratings (R4). However, reputation systems should be coupled with identification and authentication mechanisms in order to verify the identity of entities. Nevertheless, the problem of identification and authentication is often not addressed within reputation systems. The reputation systems studied in this work mainly focus on the mathematical model of reputation systems rather than on the deployment of an application to be used in practice. Accordingly, identification issues are often left outside of their scope. Only EigenTrust, PeerTrust and Hedaquin address these issues. EigenTrust provides a simple solution based on the IP address to identify users. Peertrust uses PKI to uniquely identify users. Finally, Hedaquin assumes that users and their medical devices are registered within a healthcare system. However, this does not guarantee that measurements are actually of the registered patient.

The primary feature needed for the fulfillment of R5 is the interaction scope. In Table 4, reputation systems that have only limited support for the

	eBay [16]	REGRET [42]	FIRE [43]	CORE [44]	EigenTrust [20]	Absolute EigenTrust [22]	Peertrust [45]	Beta [26]	Travos [46]	Dirichlet [47]	Subjective logic [28, 29]	CertainTrust [48]	Hedaquin [27, 49]	Fuzzy Trust [51]	FuzzyTrust/eTrust [18]
R1	*	*	*	*	✗	*	✗	*	*	*	✓	✓	✓	*	*
R2	*	*	*	✗	✗	*	✗	*	*	*	✓	✓	✓	*	*
R3	*	†	*	✗	✗	*	†	*	*	✗	✓	✓	✓	*	✗
R4	*	*	*	*	✓	*	✓	*	*	*	*	*	*	*	*
R5	*	*	✓	*	✗	✗	*	*	*	✗	✓	✓	✓	✗	✗
R6	✗	✓	✓	*	*	*	†	*	*	✗	*	*	✓	*	*
R7	✗	✓	✓	*	✗	✗	✓	✗	✗	✗	✗	✗	✗	✓	✓
R8	✓	*	*	*	✗	✗	*	*	✗	*	*	*	*	*	*
R9	†	†	†	†	*	†	*	†	†	†	†	†	✓	†	†
R10	✗	†	†	✗	*	†	✗	†	†	✗	†	†	†	✗	✗
R11	✓	N/A	N/A	N/A	†	✓	*	✓	N/A	✓	✓	N/A	✓	N/A	N/A
R12	✓	N/A	N/A	N/A	†	✓	*	✓	N/A	✓	✓	N/A	✓	N/A	N/A
R13	✓	N/A	N/A	N/A	†	✓	*	✓	N/A	✓	✓	N/A	✓	N/A	N/A

Legenda

✓: Full satisfaction * : Partial satisfaction †: Limited satisfaction ✗: No satisfaction

Table 4: Survey of Reputations Systems w.r.t. Requirements

interaction scope do not fulfill the requirement. If the distinction of functional and referral ratings is fully supported and the support for interaction scope is limited, we consider R5 partially fulfilled. Finally, reputation systems such as subjective logic, that support interaction scope and distinguish functional and referral trust, satisfy the requirement.

FIRE, REGRET and Hedaquin incorporate all the needed features to fulfill R6. The other reputation systems lack some fundamental features like trust transitivity or scope similarity. Among the reputation systems we have analyzed, only REGRET, FIRE, PeerTrust and Fuzzy Trust provide support for the interaction context and, therefore, fulfill R7. CORE fulfills R7 only partially due to partial support of the interaction context.

Evolution of user behavior (R8) is captured by taking temporal aspects into account. Although most of the analyzed systems make use of timestamps to record the time in which interaction occurred, they exploit it only partially. Indeed, most reputation systems use timestamps to order ratings

and use such an order (instead of the actual time) to weight them. In particular, reputation systems usually give more weight to recent ratings based on the intuition that an entity is most likely to maintain its recent behavior. An exception is CORE which gives more weight to old ratings based on the assumption that a user may temporally behave better to increase its reputation and then behave again in an inappropriate way. eBay is the only reputation system in our analysis which is able to capture user behavior in a specific time period. In particular, eBay displays reputation values for the last year and last six months. This, in addition with the overall reputation, provides a view of user behavior over time.

Avoiding that users take advantage of the new user status (R9) requires both the capability of discriminating user behavior and identification mechanisms. In Table 4, reputation systems that only support identification mechanisms partially satisfy R9, whereas reputation systems which only support trust and distrust feature are marked with limited support.

Although a few reputation systems support all the feature needed for the fulfillment of R10, EigenTrust is the only one that actually addresses the problem of the bootstrap of new users. It uses Distributed Hash Tables (DHTs) [52, 53] and a randomness factor in the peer selection. However, due to the lack of the uncertainty feature it does not fully fulfill R10.

Centralized reputation systems satisfy requirements R11-R13 as long as the central reputation service is secure. For decentralized systems cryptographic solutions should be employed. Most systems, however, do not discuss these requirements in their implementation and, therefore, the satisfaction of these requirements cannot be determined (marked as N/A in Table 4). EigenTrust and PeerTrust are the only proposals that address them. In particular, EigenTrust uses a secure structure based on DHTs and superpeers that are randomly selected to compute the reputation of a user. The major flaw of EigenTrust approach is that superpeers might be compromised. In addition, cryptographic solutions are not used. PeerTrust adopts a similar approach for dealing with this issue. In particular, it employs DHTs enhanced with PKI technologies.

5. Discussion

The analysis presented in the previous section shows that none of the studied systems satisfy all elicited requirements. In particular, most systems lacks support for scope similarity and interaction context (see Table 4). Nevertheless, proposals based on subjective logic and, in particular, Hedaquin satisfy the majority of them.

However, it is worth noting that the reference model presented in this work (i.e., the requirements and features in Section 3) is intended to be general. When used to select a reputation system for a given application domain, the model has to be tuned with respect to the trust requirements of the application domain. Indeed, not all requirements may have the same importance in every application domain. For instance, interaction context (R7) is not relevant for the assessment of reputation in the healthcare domain, whereas it has a crucial role in other domains like auction sites.

Therefore, the practitioners should analyze the trust requirements of the application domain in which the reputation system should be applied, and thus identify which requirements should be included in the analysis. Table 2 should be then used to identify the features that should be supported by the reputation system. The candidate reputation systems should be compared against those features. Conclusions on the more appropriate reputation system can be drawn by mapping back the features to requirements along with an analysis of the technical solutions employed by the reputation systems.

6. Related Work

Several surveys on reputation systems can be found in the literature; they can be classified into three categories. The first category of surveys (e.g., [2, 3]) discusses trust management issues in general, where reputation systems are usually presented as a solution. The second category of surveys (e.g., [4, 5, 6, 7, 8, 13, 12]) focuses specifically on the functionality of reputation systems, whereas surveys in the third category (e.g., [9, 10, 11, 54]) consists of studies on vulnerabilities and attacks against reputation systems. For instance, Friedman et al. [54] present an approach for the analysis of three threats to reputation systems, namely whitewashing, incorrectly reported feedback and sybil attacks based on game theory.

These surveys usually have a broader scope than our work. Hoffman et al. [9] analyze reputation systems with respect to three dimensions, namely formulation, calculation and dissemination. Liu and Munro [13] represent the structure of reputation systems using five components, namely input, processing, output, feedback loop, storage. For each component, they identify a number of characteristics to compare reputation systems. In this article, we use the dimensions proposed in [9] as they provide a clearer conceptual model for reasoning on reputation systems. With respect to these dimensions, our work mainly focuses on the formulation dimension.

Similarly to our work, most surveys define the main features to be supported by reputation systems and evaluate existing systems against the de-

	Jesang et al. [5]	Koutrouli et al. [6]	Rouhoma et al. [7]	Sabater et al. [8]	Liu and Munro [13]	Govindan et al. [12]	Hoffman et al. [9]	Yao et al. [11]	Current Survey
Metric	✓	✓	✓	✗	✓	✓	✓	✓	✓
Structure	✓	✗	✓	✗	✳	✓	✓	✓	✓
Measure	✗	✓	✓	✓	✗	✓	✓	✓	✓
F1: Trust/Distrust	✗	✓	✗	✗	✓	✓	✓	✓	✓
F2: Absolute Rep. Values	✗	✓	✓	✗	✗	✗	✗	✗	✓
F3: Origin/Target	✗	✗	✗	✗	✓	✓	✗	✗	✓
F4: (un)Certainty	✗	✗	✓	✓	✗	✓	✗	✓	✓
F5: Interaction Scope	✓	✓	✗	✓	✓	✓	✗	✓	✓
F6: Scope Similarity	✗	✗	✗	✗	✗	✗	✗	✗	✓
F7: Trust Transitivity	✓	✓	✗	✗	✗	✓	✗	✗	✓
F8: Functional vs. Referral	✓	✗	✗	✓	✓	✗	✗	✗	✓
F9: Interaction Context	✗	✓	✗	✓	✗	✗	✗	✓	✓
F10: Timestamp	✓	✓	✓	✗	✓	✓	✓	✓	✓
F11: Recommender Credibility	✗	✓	✓	✓	✗	✓	✗	✓	✳
F12: Value Range	✗	✓	✓	✓	✗	✓	✓	✗	✗
F13: Subjective/Objective	✓	✗	✗	✗	✗	✗	✗	✗	✗
F14: Specific/General	✓	✗	✗	✗	✓	✗	✗	✗	✗

Legenda

✓: Considered ✗: Not considered ✳: Inferred

Table 5: Analysis coverage

finer features. In Table 5, we present a comparison of the coverage of the analysis provided by the considered surveys with respect to the formulation dimension. As shown in the table, our work provides a more fine-grained and comprehensive analysis of reputation metrics. Nonetheless, other surveys consider features that have not been considered in our study (bottom part of Table 5). For instance, several surveys [7, 8, 11] analyze reputation systems with respect to the ability of determining the credibility of the recommender (F11). Although this feature is not included explicitly in our analysis framework, it can be expressed as a combination of F1, F3 and F8. A number of surveys [9, 6, 7, 8] also consider the range in which ratings and reputation values are expressed (F12). We have not considered such a feature as it is irrelevant for the satisfaction of the elicited requirements

(Table 1). Moreover, values lying in a certain interval (and even qualitative values) can easily be mapped into the desired range of values [22]. We have not considered features F13 and F14 for similar reasons. It is worth noting that most surveys define features as abstract concepts that reputation systems should provide. In contrast, we have first elicited the requirements for reputation systems and, based on those requirements, we have identified the features needed for their fulfillment. Consequently, the features considered in our analysis are closer to the real needs for reputation systems.

7. Conclusions

Nowadays, advances of ICT have led to the spread of several e-services. The adoption of these services, however, depends on the trust that end-users have in them. Reputation systems are becoming the “solution” to assess the trustworthiness of users and reliability of information in online communities. However, in order to be used and, therefore, facilitate the adoption of such services, reputation systems should be able to deal with the trust challenges and needs of the application domain in which the reputation system is deployed. This requires, on the one hand, understanding the capabilities and limits of existing reputation systems as well as the requirements of the application domain in which the reputation system will be eventually deployed. On the other hand, rigorous methods to compare existing reputation systems are necessary for practitioners who build systems and applications that rely on reputation systems.

To address these issues, we have presented a framework for evaluating and comparing reputation systems. We have elicited the requirements that reputation systems should satisfy by reviewing the literature. To objectively evaluate to what extent a reputation system meets the elicited requirements, we have identified for each requirement the features that reputation metrics should support in order to fulfill the requirements. We have demonstrated the applicability of the proposed framework by reviewing and comparing several state-of-the-art reputation systems.

Acknowledgments. This work has been funded by the Dutch national program COMMIT under the THeCS project.

- [1] S. Vavilis, M. Petkovic, N. Zannone, Impact of ICT on home healthcare, in: ICT Critical Infrastructures and Society, IFIP AICT 386, Springer, 2012, pp. 111–122.

- [2] D. Artz, Y. Gil, A survey of trust in computer science and the Semantic Web, *Web Semant.* 5 (2007) 58–71.
- [3] S. Ruohomaa, L. Kutvonen, Trust management survey, in: *Trust Management*, LNCS 3477, Springer, 2005, pp. 77–92.
- [4] A. Jøsang, Online reputation systems for the health sector, *eJHI* 3 (2008) e8.
- [5] A. Jøsang, R. Ismail, C. Boyd, A survey of trust and reputation systems for online service provision, *Decis. Support Syst.* 43 (2007) 618–644.
- [6] E. Koutrouli, A. Tsalgaidou, Reputation-based trust systems for P2P applications: design issues and comparison framework, in: *Trust and Privacy in Digital Business*, LNCS 4083, Springer, 2006, pp. 152–161.
- [7] S. Ruohomaa, L. Kutvonen, E. Koutrouli, Reputation management survey, in: *Proc. of Int. Conf. on Availability, Reliability and Security*, IEEE, 2007, pp. 103–111.
- [8] J. Sabater, C. Sierra, Review on computational trust and reputation models, *Artificial Intelligence Review* 24 (2005) 33–60.
- [9] K. Hoffman, D. Zage, C. Nita-Rotaru, A survey of attack and defense techniques for reputation systems, *ACM Comput. Surv.* 42 (2009) 1–31.
- [10] S. Marti, H. Garcia-Molina, Taxonomy of trust: Categorizing P2P reputation systems, *Computer Networks* 50 (2006) 472–484.
- [11] Y. Yao, S. Ruohomaa, F. Xu, Addressing common vulnerabilities of reputation systems for electronic commerce, *JTAER* 7 (2012) 1–15.
- [12] K. Govindan, P. Mohapatra, Trust computations and trust dynamics in mobile adhoc networks: a survey, *IEEE Commun. Surv. Tutor.* 14 (2012) 279–298.
- [13] L. Liu, M. Munro, Systematic analysis of centralized online reputation systems, *Decis. Support Syst.* 52 (2012) 438–449.
- [14] M. Witkowski, A. Artikis, J. Pitt, Experiments in building experiential trust in a society of objective-trust based agents, in: *Trust in Cyber-societies*, LNCS 2246, Springer, 2001, pp. 111–132.
- [15] L. R. Smeltzer, The Meaning and Origin of Trust in Buyer-Supplier Relationships, *Journal of Supply Chain Management* 33 (1997) 40–48.

- [16] P. Resnick, R. Zeckhauser, Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system, *Advances in Applied Microeconomics* 11 (2002) 127–157.
- [17] K. Bharadwaj, M. Al-Shamri, Fuzzy computational models for trust and reputation systems, *Electron. Commer. Res. Appl.* 8 (2009) 37–47.
- [18] S. Song, K. Hwang, R. Zhou, Y. Kwok, Trusted P2P transactions with fuzzy reputation aggregation, *Internet Computing* 9 (2005) 24–34.
- [19] S. Brin, L. Page, The anatomy of a large-scale hypertextual web search engine, *Computer networks and ISDN systems* 30 (1998) 107–117.
- [20] S. Kamvar, M. Schlosser, H. Garcia-Molina, The EigenTrust algorithm for reputation management in P2P networks, in: *Proc. of Int. Conf. on World Wide Web*, ACM, 2003, pp. 640–651.
- [21] R. Lempel, S. Moran, The stochastic approach for link-structure analysis (SALSA) and the TKC effect, *Computer Networks* 33 (2000) 387–401.
- [22] A. Simone, B. Skoric, N. Zannone, Flow-based reputation: more than just ranking, *Int. J. Inf. Technol. Decis. Mak.* 11 (2012) 551–578.
- [23] W. T. L. Teacy, M. Luck, A. Rogers, N. R. Jennings, An efficient and versatile approach to trust and reputation using hierarchical bayesian modelling, *Artificial Intelligence* 193 (2012) 149–185.
- [24] M. Tavakolifard, S. Knapskog, A probabilistic reputation algorithm for decentralized multi-agent environments, *Electron. Notes Theor. Comput. Sci.* 244 (2009) 139–149.
- [25] A. Whitby, A. Jøsang, J. Indulska, Filtering Out Unfair Ratings in Bayesian Reputation Systems, *J. of Management Research* 4 (2005).
- [26] T. Muller, P. Schweitzer, On beta models with trust chains, in: *Trust Management VII*, IFIP AICT 401, Springer, 2013, pp. 49–65.
- [27] T. van Deursen, P. Koster, M. Petković, Hedaquin: A reputation-based health data quality indicator, *Electron. Notes Theor. Comput. Sci.* 197 (2008) 159–167.
- [28] A. Jøsang, R. Hayward, S. Pope, Trust network analysis with subjective logic, in: *Proc. of Australasian Computer Science Conf.*, Australian Computer Society, Inc., 2006, pp. 85–94.

- [29] A. Jøsang, Subjective Logic, Tech. Rep., University of Oslo, 2013.
- [30] J. R. Douceur, The Sybil Attack, in: Peer-To-Peer Systems, LNCS 2429, Springer, 2002, pp. 251–260.
- [31] G. E. Bolton, E. Katok, A. Ockenfels, Cooperation among strangers with limited information about reputation, *J. Public Econ.* 89 (2005) 1457–146.
- [32] P. Nurmi, A bayesian framework for online reputation systems, in: Proc. of the Advanced Int. Conf. on Telecommunications and Int. Conf. on Internet and Web Applications and Services, IEEE, 2006, pp. 121–126.
- [33] S. Abbas, M. Merabti, D. Llewellyn-Jones, Deterring whitewashing attacks in reputation based schemes for mobile ad hoc networks, in: Proc. of IFIP Wireless Days Conf., IEEE, 2010, pp. 1–6.
- [34] A. Josang, T. Azderska, S. Marsh, Trust transitivity and conditional belief reasoning, in: Trust Management VI, IFIP AICT 374, Springer, 2012, pp. 68–83.
- [35] P. Avesani, P. Massa, R. Tiella, A trust-enhanced recommender system application: Moleskiing, in: Proc. of Symposium on Applied Computing, ACM, 2005, pp. 1589–1593.
- [36] M. Srivatsa, L. Xiong, L. Liu, TrustGuard: countering vulnerabilities in reputation management for decentralized overlay networks, in: Proc. of Int. Conf. on World Wide Web, ACM, 2005, pp. 422–431.
- [37] H. C. A. van Tilborg, S. Jajodia (Eds.), *Encyclopedia of Cryptography and Security*, 2nd Ed, Springer, 2011.
- [38] C. Koscielny, M. Kurkowski, M. Srebrny, Public key infrastructure, in: *Modern Cryptography Primer*, Springer, 2013, pp. 175–191.
- [39] A. Mohaisen, J. Kim, The Sybil Attacks and Defenses: A Survey, *Smart Computing Review* 3 (2013) 480–489.
- [40] A. Nandi, T.-W. J. Ngan, A. Singh, P. Druschel, D. S. Wallach, Scrivener: providing incentives in cooperative content distribution systems, in: *Middleware*, LNCS 3790, Springer, 2005, pp. 270–291.

- [41] T. Dimitriou, G. Karame, I. T. Christou, SuperTrust - A Secure and Efficient Framework for Handling Trust in Super Peer Networks, in: Distributed Computing and Networking, LNCS 4904, Springer, 2008, pp. 350–362.
- [42] J. Sabater, C. Sierra, Reputation and social network analysis in multi-agent systems, in: Proc. of Int. Joint Conf. on Autonomous Agents and Multiagent Systems, ACM, 2002, pp. 475–482.
- [43] T. D. Huynh, N. R. Jennings, N. R. Shadbolt, An integrated trust and reputation model for open multi-agent systems, AAMAS 13 (2006) 119–154.
- [44] P. Michiardi, R. Molva, Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks, in: Advanced Communications and Multimedia Security, IFIP AICT 100, Kluwer Academic Publishers, 2002, pp. 107–121.
- [45] L. Xiong, L. Liu, Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities, TKDE 16 (2004) 843–857.
- [46] W. Teacy, J. Patel, N. Jennings, M. Luck, Travos: Trust and reputation in the context of inaccurate information sources, AAMAS 12 (2006) 183–198.
- [47] A. Jøsang, J. Haller, Dirichlet reputation systems, in: Proc. of Int. Conf. on Availability, Reliability and Security, IEEE, 2007, pp. 112–119.
- [48] S. Ries, S. M. Habib, M. Mühlhäuser, V. Varadharajan, Certainlogic: A logic for modeling trust and uncertainty, in: Trust and Trustworthy Computing, Springer, 2011, pp. 254–261.
- [49] T. van Deursen, P. Koster, M. Petkovic, Reliable personal health records, in: Proc. of Int. Congress of the European Federation for Medical Informatics, Studies in health technology and informatics 136, IOS Press, 2008, pp. 484–489.
- [50] N. Griffiths, K.-M. Chao, M. Younas, Fuzzy Trust for Peer-to-Peer Systems, in: Proc. of Int. Conf. on Distributed Computing Systems Workshops, IEEE, 2006, pp. 73–78.

- [51] S. Schmidt, R. Steele, T. S. Dillon, E. Chang, Fuzzy trust evaluation and credibility development in multi-agent systems, *Appl. Soft. Comput.* 7 (2007) 492–505.
- [52] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Shenker, A scalable content-addressable network, *Comput. Commun. Rev.* 31 (2001) 161–172.
- [53] I. Stoica, R. Morris, D. Liben-Nowell, D. Karger, M. Kaashoek, F. Dabek, H. Balakrishnan, Chord: a scalable peer-to-peer lookup protocol for internet applications, *IEEE/ACM Trans. on Networking* 11 (2003) 17–32.
- [54] E. Friedman, P. Resnick, R. Sami, Manipulation-resistant reputation systems, *Algorithmic Game Theory* (2007) 677–697.