A Reputation System for Electronic Negotiations

Omid Tafreschi[1⋆], Dominique Mähler[1], Janina Fengel[2⋆⋆], Michael Rebstock[2⋆⋆] and Claudia Eckert[1]

[1] Technische Univertiät Darmstadt
Department of Computer Science

[2] Darmstadt University of Applied Sciences
Faculty of Economics and Business Administration

Abstract. In this paper we present a reputation system for electronic negotiations. The proposed system facilitates trust building among business partners who interact in an ad-hoc manner with each other. The system enables market participants to rate the business performance of their partners as well as the quality of offered goods. These ratings are the basis for evaluating the trustworthiness of market participants and the quality of their goods. The ratings are aggregated using the concept of Web of Trust. This approach leads to robustness of the proposed system against malicious behavior aiming at manipulating the reputation of market participants.

1 Introduction

Markets provide the basis for all types of business. According to [3], they have three main functions, i.e., matchmaking between buyers and sellers, providing an institutional infrastructure and the facilitation of transactions. A typical market transaction consists of 5 phases [14], which are depicted in Figure 1:

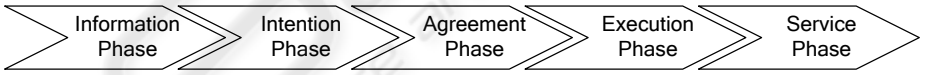

Fig. 1. Transaction Phases on a Market.

In the information phase market participants gather relevant information concerning products, market partners, etc. During the intention phase market participants submit offers concerning supply and demand. Then, the terms and conditions of the transaction are specified and the contract is closed in the agreement phase. The agreed-upon contract is operationally executed in the execution phase. During the service phase support, maintenance and customer services are delivered.

⋆ The author is supported by the German Federal Ministry of Education and Research under grant 01AK706C, project Premium.

⋆⋆ These authors are supported by the German Federal Ministry of Education and Reserach under grant 1716X04, project ORBI.

Tafreschil O., Mählerl D., Fengel J., Rebstock M. and Eckertl C. (2007).
A Reputation System for Electronic Negotiations.
In *Proceedings of the 5th International Workshop on Security in Information Systems*, pages 53-62
DOI: 10.5220/0002415300530062

A negotiation can be defined as a decentralized decision-making process by at least two parties [5]. It is performed until an agreement is reached or the process is terminated without reaching an agreement by one or all the partners involved. The objective of this process is to establish a legally binding contract between the parties concerned, which defines all agreed-upon terms and conditions next to regulations in case of failure of their fulfillment [14]. Beam and Segev define electronic negotiations "as the process by which two or more parties multilaterally bargain resources for mutual intended gain, using the tools and techniques of electronic commerce." [4] Electronic negotiations enable market participants to interact with each other in an ad-hoc manner over large geographical distances. However, these opportunities lead to undesirable side-effects. Over open and anonymous networks, market participants have to cope with much higher mount of uncertainty about the quality of products and the trustworthiness of other participants [18]. Ackerlof showed that knowledge about the trustworthiness of a seller is vital for the functioning of a market [1]. Therefore building trust in virtual environments is of high importance. Information about the reputation can be used to build up trust among strangers, e.g., new business partners. Learning from experiences from other users by seeing their assessment of a potential partners' business performance as well as the quality of the delivered goods can reduce this uncertainty.

In this paper we propose a reputation system for electronic negotiations, which is robust against unfair behavior. The remainder of this paper is organized as follows: In Section 2, we present a business web supporting electronic negotiations, explain the notion of reputation, and introduce reputation systems and their requirements. In Section 3, we propose a reputation system and discuss its robustness against opportunistic behavior. Section 4 presents related work. We conclude and discuss future work in Section 5.

2 Background

2.1 A Business Web Forming an Electronic Market

The MultiNeg Project [8] has developed components for forming a business web which shapes an electronic market. The market participants can offer their products, negotiate with each other and transfer transaction results into their ERP systems. The business web consists of four applications, i.e., marketplace, MultiNeg, Authentication and Authorization Server (AAS) and ERP. Figure 2 shows the overall architecture of the business web.

The *marketplace* allows sellers to publish catalogs containing their products or services, which can be subject to negotiations.

MultiNeg is an electronic negotiation support system for bilateral, multi-attributive negotiations. It facilitates the negotiation of multiple items with arbitrary variable attributes. The key objectives for the development have been the design of an architecture suitable for different industries, company sizes and products and a communication interface design that allows the integration of inter-organizational with intra-organizational applications. The functionalities are conceptualized for usage in a decentralized deployment and are based on open standards, e.g., HTTP, XML, SOAP, WSDL. Thus they allow the seamless electronic integration of internal and external business processes.

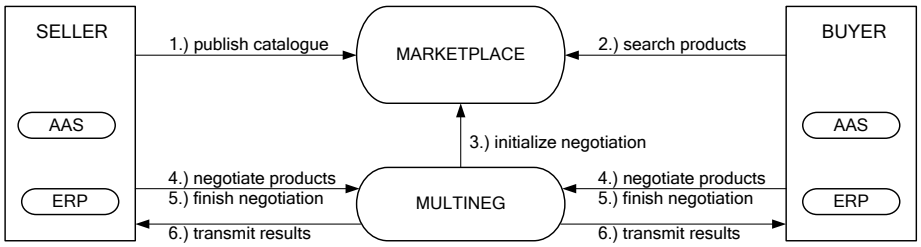

Fig. 2. Architecture of Business Web.

Each company has an *Authentication and Authorization Server (AAS)*. This allows market participants to setup security policies for authentication and authorization. Authentication policies define how business partners of a company can authenticate themselves in order to interact with the owner of the AAS. Authorization policies support the implementation of access control to the resources of a company, e.g., catalog data.

As mentioned in the introduction, trust building is crucial for electronic negotiations. Information about reputation can be used for that purpose. In the following we describe the notion of reputation.

2.2 The Notion of Reputation

According to [17], reputation is a concept that arises in repeated game settings when there is uncertainty about some property of one or more players in the mind of other players. If "uninformed" players have access to the history of past stage game outcomes, reputation effects then often allow informed players to improve their long-term payoffs by gradually convincing uninformed players that they belong to the type that best suits their interests.

Reputation information is useful in two respects. First, it supports a market participant with regard to decision support. This is necessary when a buyer has to choose between several sellers offering the same product with comparable conditions. Second, it makes opportunistic behavior unprofitable, since this would lead to a poor rating, which in turn may deter potential business partners from considering the owner of the poor rating.

2.3 Reputation Systems and their Requirements

A reputation system collects, distributes and aggregates feedback about participants' past behavior [15]. Feedback is usually collected at the end of each transaction. The authority responsible for reputation management asks the participants to submit a rating. Aggregation of feedback means that the available information about an entity's transactions is evaluated and condensed into a few values that allow users to easily make a decision. Feedback is distributed finally to the participants, i.e., it has to be made available to everyone who wants to use it for making decisions on whom to trust.

Figure 3 shows the components and actors of a reputation system [16]. The target of a rating is called ratee. The collector gathers ratings from agents called raters. This

information is processed and aggregated by the processor. The emitter makes the results available to other requesting agents.

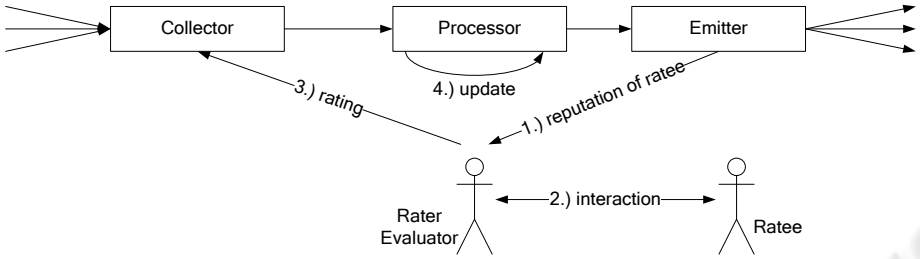

Fig. 3. Components and Actors of Reputation Systems [16].

The main objective of a reputation system is to prevent opportunistic behavior in virtual environments by making it worthless in the long term. However, since the same communication facilities are used for implementing reputation systems, these systems face themselves some problems caused by opportunistic behavior of market participants using reputation systems.

In [9], Dellarocas presents two scenarios in which market participants try to manipulate their own reputation or the reputation of others, i.e., ballot stuffing and badmouthing. Ballot stuffing is the result of unfairly high ratings, which are given by a colluded group of buyers to a single seller in order to improve his reputation. Badmouthing describes a situation in which colluded buyers give unfairly negative ratings to a seller in order to drive him out of the market.

In [10], Friedman and Resnick describe the dilemma of cheap pseudonyms. This occurs when a user of an online service can change his virtual identity by obtaining a new pseudonym at little cost. As a consequence, market participants can easily get rid of negative reputations. Such market participants are referred to as whitewashers: users who leave a system and rejoin with new identities to avoid reputational penalties. Cheap pseudonyms allow another adversarial strategy, i.e., sybil attacks [7]. Sybils are different identities of one user. This user improves his reputation by using his sybils.

In addition to the issues mentioned above, reputation systems supporting mutual ratings should prevent that one party can learn the rating he is supposed to receive before he has issued the rating for his business partner. This should ensure that a rating reflects its issuer's experience in the context of the corresponding transaction and that a rating is independent from the rating of the opposite side.

3 A Reputation System for Electronic Negotiations

In this section we present a reputation system for the business web mentioned above. It enables market participants, i.e., buyers and sellers, to rate each other after a concluded transaction. They can assess the quality of the traded goods as well as the partner's performance. This mechanism enables participants to make better decisions considering several aspects. For example a seller is offering products of excellent quality, but tends

to deliver late. This information is valuable for potential buyers who have to plan their production in advance. If they decide to buy some goods from the seller in question, they have to schedule reserve time for maybe late delivery.

Our reputation system encompasses two phases, i.e., rating phase and rating aggregation phase. In the rating phase two business partners evaluate each other after a concluded transaction by giving a rating to the opposite side. This phase affects the last three transaction phases on a market during which two business partners interact with each other. During these three phases including the negotiation they form an opinion about each other, which is later reflected by the corresponding ratings. These ratings are used in the rating aggregation phase to get information about the reputation of a market participant. The rating aggregation phase takes place during the first three transaction phases. In the information phase and the intention phase a buyer wants to learn more about his potential negotiation partners, i.e., sellers. For this purpose he has to evaluate the rating of sellers. During the actual negotiation both buyer and seller can use ratings of their opposite to map out their negotiation strategy. Figure 4 depicts the rating phase and the rating aggregation phase during a market transaction.

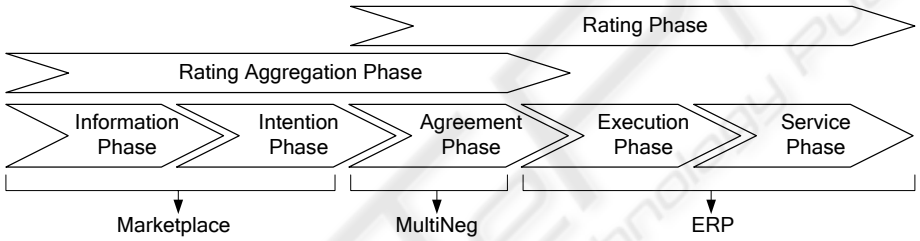

Fig. 4. Rating Phase and Rating Aggregation Phase During a Market Transaction.

3.1 Rating Phase

Figure 5 shows the process of a negotiation and the subsequent rating phase. We describe the process in the following:

Steps 1-2 After concluding a negotiation the involved parties evaluate each other by issuing corresponding ratings. A rating ranges from 0, i.e., poor rating, to 100, i.e., excellent rating. We support two different types of rating, i.e., buyer rating and seller rating. A buyer rating is given by a seller to a buyer, whereas a seller rating is given by a buyer to a seller. Each kind of rating has a set of specific categories, which can be evaluated separately. For example, a seller can rate the payment practice of a buyer and a buyer can rate the quality of products delivered by a seller. A weighting can be assigned to each category to reflect the importance of this category to the evaluator. The sum of all weightings of a single rating is 100. The default value of a weighting is 100 divided by the number of the categories. Additionally, a category can be annotated with a comment. Since the data of a negotiation are stored on MultiNeg, negotiation partners use MultiNeg to create ratings for each other.

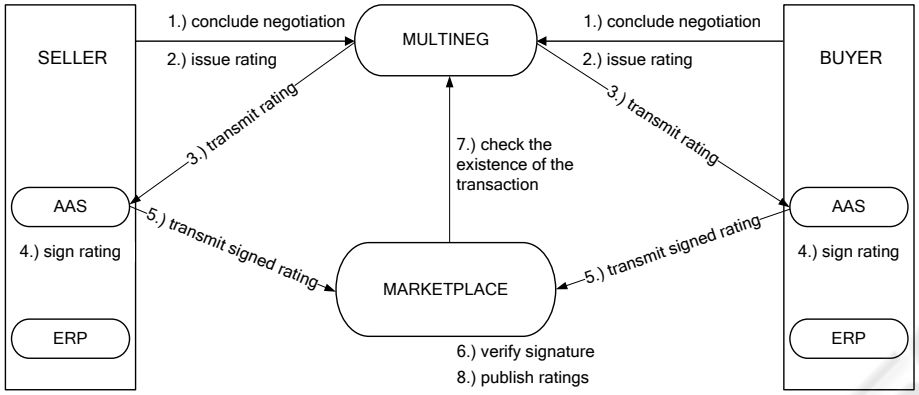

Fig. 5. Rating Phase.

Steps 3-5 After the receipt of the rating data MultiNeg forwards it to the AAS of the corresponding party. The AAS displays the rating to its issuer and allows him to digitally sign it. The signed rating is transferred to the marketplace to be published.

Steps 6-8 The signature of a signed rating is verified at the marketplace. In case of a positive result the marketplace checks whether the rating belongs to an existing negotiation. This step ensures that only authentic ratings are published. If MultiNeg's answer is positive, the marketplace publishes the rating if both ratings of a negotiation have been issued. Publishing two ratings simultaneously prevents that one negotiation partner changes his mind (rating) after learning the rating given to him by his opposite. However, this could protect one party from receiving a poor rating. In this case the party does not give a rating for his opposite. To overcome this shortcoming the marketplace publishes a single rating without the second corresponding rating after a defined period of time has passed. This rule is known to all market participants.

3.2 Rating Aggregation Phase

Existing ratings of market participants are stored and displayed at the marketplace where buyers search the catalogs of sellers and choose partners for a negotiation. The ratings of market participants can be used to support that decision. Our reputation system provide a set of different views to the ratings. An evaluator can look at details of one rating, e.g., its categories and the corresponding comments. In order to simplify the assessment of a certain market participant based on his ratings, our system allows calculating the arithmetic average of all ratings, i.e., $\overline{ratings}$, of one market participant mp_i with the help of the following formulas. We assume that the market participant has n ratings and each rating consists of m categories.

$$\overline{rating} = \frac{\sum_{j=1}^{m} category_j \cdot weighting_j}{100} \tag{1}$$

$$\overline{ratings}_{mp_i} = \frac{\sum_{l=1}^{n} \overline{rating}_l}{n} \quad (2)$$

Formula 1 calculates the average of one rating. We apply this to all ratings of mp_i and calculate the average of them with Formula 2. Although this approach is easy to implement and to understand, it is not robust against ballot stuffing and bad-mouthing, since it considers only the number of ratings independently from their issuers. In the following we show how to overcome this shortcoming.

If we assume that trust is transitive, we can apply the concept of Web of Trust for the calculation of trust values based on rating lists. Pretty Good Privacy (PGP) was the first popular system which used the concept of Web of Trust to establish the authenticity of the binding between a public key and a user in a decentralized and open environment, e.g., the Internet [19]. We use the concept of Web of Trust to calculate trust values in our reputation system. For this purpose, the evaluator first has to build up a mesh between himself and the peer for which he wants to calculate the trust value. The mesh should consist of one trust path at least. The start vertex of a trust path is the evaluator and the end vertex is the market participant to be evaluated. In addition, a trust path contains other vertices, i.e., market participants which have been rated by the evaluator or by other vertices in the mesh except the end vertex. In order to keep the calculation base meaningful, we limit the length of a trust path, i.e., the number of its vertices [11]. The average of the rating which was given to a vertex is calculated with Formula 1. If there are several ratings, i.e., weightings to be assigned to one edge, we aggregate them with Formula 2. The result is assigned as weighting to the edge directing to the vertex in question. It is not always possible to find a trust path, e.g., when no in between vertices exist. Our reputation system indicate such cases to the evaluator.

We explain the algorithm for the aggregation of ratings based on Web of Trust with the help of an example. One buyer (Company A) enters the marketplace and wants to purchase products which are offered by the seller (Company Z). These two market participants have never negotiated with each other. But there are two other companies which have negotiated with both, seller and buyer. The corresponding ratings are shown in Figure 6. We use the existing ratings to build up a mesh between buyer and seller. If

Rater	Ratee	Rating
Company A (buyer)	Company B	95%
Company A (buyer)	Company C	70%
Company A (buyer)	Company C	90%
Company B	Company Z (seller)	99%
Company C	Company Z (seller)	95%

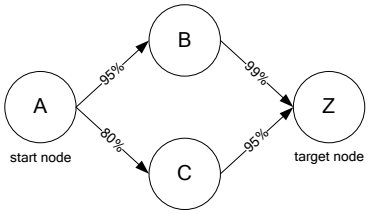

Fig. 6. Example for Web of Trust.

there are several ratings between two parties we aggregate them into one value, which is the arithmetic average. Figure 6 shows the corresponding mesh of our example. The mesh is the input for the following algorithm proposed by Caronni [6].

Algorithm 1: calculateTrust.

Input: mesh W, start node A, target node Z
Output: TrustLevel

```
1  trust in A = 1;
2  TrustLevel = 0;
3  for all direct neighbors of Z do
4      direct neighbor is X;
5      if X is the start node then
6          TrustLevel = 1 − (1 − (trust in X) · (edge_{X,Z}) · (1 − TrustLevel));
7      else
8          create a copy C of W;
9          drop target node Z in C;
10         calculateTrust(C,A,X);
11         TrustLevel = 1 − (1 − return · edge_{X,Z}) · (1 − TrustLevel);
12 return TrustLevel;
```

The algorithm is recursive and has three runs for our example mesh. The direct neighbors from Z are node B and node C. The intermediate steps are depicted in Figure 7. The algorithm starts with node B, which is not the start node. As a consequence, in the left mesh the company Z (seller) has been removed and company B is the new target node (see Figure 7(a)). We find out the direct neighbors of node B, which is node A (start node). Line 6 of the algorithm contains the formula for calculating the trust level of node A with regard to node B ($TrustLevel_{A,B} = 1-(1-1\cdot 0,95)\cdot(1-0) = 0,95$). This value is returned from the recursive call. It is the input in the formula in line 11 of the algorithm, where we calculate the TrustLevel of node A with regard to node Z ($TrustLevel_{A,B,Z} = 1-(1-0,95\cdot 0,99)\cdot(1-0) = 0,9405$).

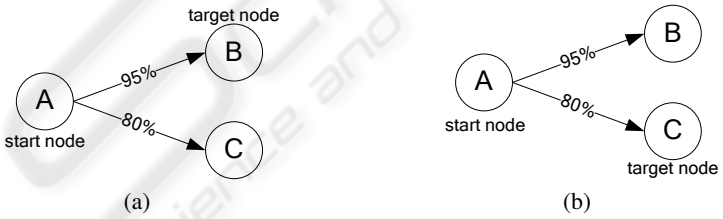

Fig. 7. Iteration Steps for Calculating a Trust Level.

We repeat the sketched process again for node C as next neighbor of node Z (see Figure 7(b)). The resulting value is $TrustLevel_{A,C} = 1-(1-1\cdot 0,8)\cdot(1-0) = 0,8$. In the final step $TrustLevel_{A,B,Z}$ and $TrustLevel_{A,C}$ are the inputs for the formula in line 11 where the return value of the algorithm is calculated $TrustLevel_{A,Z} = 1-(1-0,8\cdot 0,95)\cdot(1-0,9405) = 0,98572$. This value is presented to Company A in addition the arithmetic average of Company B's ratings. This additional information provides an alternative for decision-making considering the perspective of the evaluator.

The aggregation of ratings based on the concept of Web of Trust does not totally prevent ballot stuffing and bad-mouthing, but these kinds of malicious behavior lose some of their impact. Recall that ballot stuffing requires that a (small) set of raters give a large number of unfair good ratings to a ratee in order to improve his reputation. Since the approach based of Web of Trust considers the average of ratings given to a market participant, it does defuse the impact of several unfair ratings. The same reasoning applies to bad-mouthing.

4 Related Work

Au et al. present a framework for enabling crossorganisational trust establishment [2]. They introduce trust tokens which are issued by trust servers. These tokens are required when a user wants to access a protected resource. Au et al. also use the concept of Web of Trust to model the trust relationship between the trust servers. However, in contrast to their work, we build up a Web of Trust based on ratings and can calculate the trust level in an ad-hoc manner.

Kamvar et al. propose the EigenTrust algorithm for reputation management in P2P networks [12]. This work is close to our system, since both are based on the concept of transitive trust. But they have different design goals. EigenTrust minimizes the impact of malicious peers on the performance of a P2P system.

In [13], Maurer proposes an approach to model a user's view of a public key infrastructure (PKI). From this view, a user draws conclusions about the authenticity of other entities' public keys and possibly about the trustworthiness of other entities. A user's view consists of the user's statements about the authenticity of public keys and the trustworthiness of their owners, as well as a collection of certificates and recommendations obtained or retrieved from the PKI. This approach is similar to our work. However, we use other information, i.e., transaction ratings, to calculate the trustworthiness of market participants.

5 Conclusion and Future Work

In this paper, we proposed a reputation system for electronic negotiations to enable market participants to evaluate each other and their offered goods. For this purpose, we developed a system architecture which allows business partners to evaluate each other with the help of ratings after a concluded transaction. These ratings are aggregated to support a market participant to find a trustworthy business partner. This aggregation is based on the concept of Web of Trust to make the reputation systems robust against malicious behavior aiming at manipulating the reputation of some market participants.

In addition to the aggregated values, our approach offers differentiated views for assessing a business partners' capabilities. Instead of an overall information, a potential partner can get insight about the business behavior and the quality of the offered goods separately, since those do not necessarily depend on each other.

We have implemented the proposed reputation system and plan to do an empirical study of the usability of our system and to analyze its efficiency.

References

1. G. A. Akerlof. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quartely Journal of Economics*, 84(3):488–500, 1970.
2. R. Au, M. Looi, and P. Ashley. Automated Cross-Organisational Trust Establishment on Extranets. In *ITVE '01: Proceedings of the Workshop on Information Technology for Virtual Enterprises*, pages 3–11, Washington, DC, USA, 2001. IEEE Computer Society.
3. Y. Bakos. The Emerging Role of Electronic Marketplaces on the Internet. *Communincations of the ACM*, 41(8):35–42, 1998.
4. C. Beam and A. Segev. Automated Negotiations: A Survey of the State of the Art. *Wirtschaftsinformatik*, 39(3):263–268, 1997.
5. M. Bichler. A Roadmap to Auction-based Negotiation Protocols for Electronic Commerce. In *HICSS '00: Proceedings of the 33rd Hawaii International Conference on System Sciences*, volume 6, Washington, DC, USA, 2000. IEEE Computer Society.
6. G. Caronni. Walking the Web of Trust. In *WETICE '00: Proceedings of the 9th IEEE International Workshops on Enabling Technologies*, pages 153–158, Washington, DC, USA, 2000. IEEE Computer Society.
7. A. Cheng and E. Friedman. Sybilproof Reputation Mechanisms. In *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems*, pages 128–132, New York, NY, USA, 2005. ACM Press.
8. The MultiNeg Consortium. http:www.fbw.h-da.de/multineg, 2007.
9. C. Dellarocas. Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior. In *EC '00: Proceedings of the 2nd ACM Conference on Electronic commerce*, pages 150–157, New York, NY, USA, 2000. ACM Press.
10. E. Friedman and P. Resnickl. The Social Cost of Cheap Pseudonyms. *Journal of Economics & Management Strategy*, 10(2):173–199, 2001.
11. A. Jøsang, R. Hayward, and S. Pope. Trust Network Analysis with Subjective Logic. In *ACSC '06: Proceedings of the 29th Australasian Computer Science Conference*, pages 85–94, Darlinghurst, Australia, Australia, 2006. Australian Computer Society, Inc.
12. S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The Eigentrust Algorithm for Reputation Management in P2P Networks. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pages 640–651, New York, NY, USA, 2003. ACM Press.
13. U. M. Maurer. Modelling a Public-Key Infrastructure. In E. Bertino, H. Kurth, G. Martella, and E. Montolivo, editors, *Computer Security - ESORICS 96, 4th European Symposium on Research in Computer Security*, volume 1146 of *Lecture Notes in Computer Science*, pages 325–350. Springer, 1996.
14. M. Rebstock and M. Lipp. Webservices zur Integration interaktiver elektronischer Verhandlungen in elektronische Marktplätze. *Wirtschaftsinformatik*, 45(3):293–306, 2003.
15. P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.
16. A. Schlosser, M. Voss, and L. Brückner. On the Simulation of Global Reputation Systems. *Journal of Artificial Societies and Social Simulation*, 9(1), 2005.
17. R. Wilson. Reputations in Games an Markets. *Game Theoretic Models of Bargaining*, 1985.
18. W. Zhang, L. Liu, and Y. Zhu. A Computational Trust Model for C2C Auctions. In *Services Systems and Services Management, Proceedings of ICSSSM '05. 2005 International Conference*, 2005.
19. P. R. Zimmermann. *The official PGP User's Guide*. MIT Press, Cambridge, MA, USA, 1995.