Collaborative Reputation Mechanisms in Electronic Marketplaces

Giorgos Zacharia, Alexandros Moukas and Pattie Maes
MIT Media Laboratory
20 Ames Street, Room 305,
Cambridge MA 02139 USA
lysi,moux,pattie@media.mit.edu

Abstract

The members of electronic communities are often unrelated to each other, they may have never met and have no information on each other's reputation. This kind of information is vital in Electronic Commerce interactions, where the potential counterpart's reputation can be a significant factor in the negotiation strategy. This paper proposes two complementary reputation mechanisms that rely on collaborative rating and personalized evaluation of the various ratings assigned to each user. While these reputation mechanisms are developed in the context of electronic commerce, we believe that they may have applicability in other types of electronic communities such as chatrooms, newsgroups, mailing lists etc.

1. Introduction

Consumer to consumer electronic transaction systems like Kasbah [1], eBay [2] and "ONSALE Exchange Auction Classifieds" [3] create online market places that bring together users unknown to each other. Kasbah is an ongoing research project to help realize a fundamental transformation in the way people transact goods -- from requiring constant monitoring and effort, to a system where software agents do much of the bidding and negotiating on a user's behalf. A user wanting to buy or sell a good creates an agent, gives it some strategic direction, and sends it off into the agent marketplace. Kasbah agents pro-actively seek out potential buyers or sellers and negotiate with them on their creator's behalf. Each agent's goal is to make the "best deal" possible, subject to a set of user-specified constraints, such as a desired price, a highest (or lowest) acceptable price, and a date to complete the transaction [1]. OnSale Exchange is an online auction specializing in computer products, consumer electronics, sporting goods, and auction classifieds, where sellers can list their items for sale and buyers compete in the auction-like bidding system to buy the posted items.

These kinds of online marketplaces introduce two major issues of trust among the users of the system:

1. The potential buyer has no physical access to the product of interest while he/she bids or negotiates. Therefore the seller could misrepresent the condition or the quality of his/her product in order to get more money.
2. The seller or buyer may decide not to abide by the agreement reached at the electronic marketplace asking at some later time to renegotiate the price, or even refusing to commit the transaction.

An approach, which could solve the above mentioned problems, would be to incorporate in the system a reputation brokering mechanism, so that each user can actually customize his/her pricing strategies according to the risk implied by the reputation values of the potential counterparts.

Reputation is usually defined as the amount of trust inspired by the particular person in a specific setting or domain of interest [4]. In "Trust in a Cryptographic Economy" [5] reputation is regarded as asset creation and it is evaluated according to its expected economic returns.

Reputation is conceived as a multidimensional value. An individual may enjoy a very high reputation in one domain, while he/she has a low reputation in another. For example, a Unix guru will naturally have a high rank regarding Linux questions, while he may not enjoy that high a reputation for questions regarding Microsoft's operating systems. These individual reputation standings are developed through social interaction among a loosely connected group that shares the same interest.

We are developing methods through which we can automate the social mechanisms of reputation for the purposes of an electronic marketplace. These reputation mechanisms are implemented and tested in the Kasbah electronic marketplace. In Kasbah, the reputation values of the individuals trying to buy or sell books or CDs are a major parameter of the behavior of the buying, selling or finding agents of the system.

In this paper we describe two reputation mechanisms:

0-7695-0001-3/99 $10.00 (c) 1999 IEEE

a) Sporas is a simple reputation mechanism which can be implemented irrespectively of the number of rated interactions, and
b) Histos is a more complex reputation mechanism that assumes that the system has been somehow bootstrapped (probably by using Sporas) so that there is an abundance of rated interactions to create a dense web of pairwise ratings.

In the first section of the paper we outline the problem we are trying to solve and the problems we faced during the initial implementation of the system. The second section describes related work and the third section outlines specific problems inherent to online marketplaces like Kasbah. In the fourth and the fifth section we describe the solution adopted in the case of Kasbah, and in the sixth section we present results from simulations. Finally the last section outlines future goals for research in reputation mechanisms for online communities.

2. Related Work

The study of reputation brokering for online communities has generally focused on content filtering. We can divide the related work into two major categories: the non-computational ones like the local Better Business Bureaus [6] and the computational ones. The computational methods cover a broad domain of applications, from rating of newsgroup postings and webpages, to rating people and their expertise in specific areas. In this section we focus on the related computational methods and compare their major features in Table 1.

One approach to building a reputation mechanism is to have a central agency that keeps records of the recent activity of the users on the system, very much like the scoring systems of credit history agencies [7]. This central agency also keeps records of complaints by users in textual formats and even publishes warnings against possibly malicious users, pretty much like the local Better Business Bureaus in the US [6]. Adopting such a solution requires a lot of overhead on the behalf of the providers of the online community.

Other proposed approaches are more distributed. For example, approaches such as Yenta [8], Weaving a web of Trust [9], or the Platform for Internet Content Selection (PICS) [10], (like the Recreational Software Advisory Council [11]) would require that users give a rating for themselves and either have a central agency or other trusted users verify their trustworthiness. We can make the reasonable assumption that no user would ever label him/herself as a non-trustworthy person. Thus all new members would have to await the validation by other trustworthy users of the system. A user would end up evaluating their counterparts' reputation as a consequent of the number and the trustworthiness of the recommendations for each user.

Friend of a Friend Finder (FFF) [12], Yenta and Weaving Web of Trust introduce computational methods for creating personal recommendation systems, the former two for people and the latter for webpages. FFF and the Weaving a Web of Trust rely on the existence of a connected path between two users, while Yenta clusters people with shared interests according to the recommendations of users that know each other and can verify the assertions they make about themselves. All three systems require the a priori existence of social relationships among the users of their online community, while in the online marketplaces, deals are brokered among people who probably have never met each other.

Collaborative filtering is a technique used to detect patterns among the opinions of different users and to make recommendations to people, based on others who have shown similar taste. It essentially automates the process of "word of mouth" to produce an advanced, personalized marketing scheme. Examples of collaborative filtering systems are HOMR [13], Firefly [13] and GroupLens [14]. GroupLens is a collaborative filtering solution for rating the content of Usenet articles and presenting them to the user in a personalized manner. Users are clustered together according to the ratings they give to the same articles. The user sees the articles with a value equal to the average of the ratings given to the article by users in the same cluster.

The most relevant computational methods to our knowledge, are the reputation mechanism of the OnSale Exchange [3] and the eBay [2]. OnSale allows its users to rate and submit textual comments about sellers and overall reputation value of a seller is the average of his/her ratings through his usage of the OnSale system. In eBay, sellers receive +1, 0 or –1 as feedback for their reliability in each auction and their reputation value is calculated as the sum of those ratings over the last six months. In OnSale, the newcomers have no reputation until someone eventually rates them, while on eBay they start with zero feedback points. However, bidders in the OnSale Exchange auction system are not rated at all. OnSale tries to ensure the bidders' integrity through a rather psychological measure: bidders are required to register with the system by submitting a credit card. OnSale believes that this requirement helps to ensure that all bids placed are legitimate, which protects the interests of all bidders and sellers. In both sites the reputation value of a seller is available with any textual comments that may exist to the potential bidders.

Table 1 Comparison of reputation systems

System	Computational	Pair-wise rating	Personalized	Textual comments
GroupLens	Yes	rating of articles	Yes	
Elo & Glicko	Yes	result of game		
OnSale	Yes	buyers rate sellers		
FairIsaac	Yes		Yes	
Local BBB's				Yes
Web of Trust	Yes	Self rating of cost	Yes	
Kasbah	Yes	Yes	Yes	
Firefly	Yes	Rating of recommendations	Yes	
Ebay	Yes	buyers rate sellers		Yes

The latest release of Kasbah [1] features a Better Business Bureau service that implements the reputation mechanisms we describe below.

3. Desiderata for online reputation systems

While the above discussed reputation mechanisms have some interesting qualities, we believe they are not perfect for maintaining reputations in online communities and especially in online marketplaces. This section describes some of the problems of online communities and their implications for reputation mechanisms.

In online communities, it is relatively easy to adopt a new or change one's identity. Thus, if a user ends up having a reputation value lower than the reputation of a beginner, he/she would have an incentive to discard his/her initial identity and start from the beginning. Hence, it is desirable that while a user's reputation value may decrease after a transaction, it will never fall below a beginner's value. We therefore decided for the reputation mechanisms described in the following section that a beginner cannot start with an average reputation.

We also want to make sure that even if a user starts receiving very low reputation ratings, he/she can improve his/her status later at almost the same rate as a beginner. If the reputation value is evaluated as the arithmetic average of the ratings received since the user joined the system, users who perform relatively poorly in the beginning, have an incentive to adopt a new identity so that they get rid of their bad reputation history.

Another problem is that the overhead of performing fake transactions in both Kasbah and OnSale Exchange is relatively low (OnSale does not charge any commission on its Exchange service yet). Therefore two friends might decide to perform some dozens of fake transactions, rating each other with perfect scores so as to both increase their reputation value.

Even if we allow each user to rate another only once, another way to falsely increase one's reputation would be to create fake identities and have each one of those rate the user's real identity with perfect scores. A good reputation system would avoid both these problems.

We have to ensure that those ratings given by users with an established high reputation in the system are weighted more than the ratings given by beginners or users with low reputations.

In addition the reputation values of the users should not be allowed to increase at infinitum like the case of eBay, where a seller may cheat 20% of the time but he/she can still maintain a monotonically increasing reputation value.

Finally we have to consider the effect of the memory of our system [4]. The larger the number of ratings used in the evaluation of reputation values the highest the predictability of the mechanism it gets. However, since the reputation values are associated with human individuals and humans change their behavior over time it is desirable to disregard very old ratings. Thus we ensure that we the predicted reputation values are closer to the current behavior of the individuals rather their overall performance.

4. Sporas: A reputation mechanism for loosely connected online communities

Sporas provides a reputation service based on the following principles:

1. New users start with a minimum reputation value, and they build up reputation throughout their activity on the system.

2. The reputation value of a user should not fall below the reputation of a new user no matter how unreliable the user is.

3. After each rating the reputation value of the user is updated based on the feedback provided by the other party to reflect his/her trustworthiness in the latest transaction.

4. Two users may rate each other only once. If two users happen to interact more than once, the system keeps the most recently submitted rating.

5. Users with very high reputation values experience much smaller rating changes after each update. This approach is similar to the method used in the Elo [15] and the Glicko [16] system for pairwise ratings.

Each user has one reputation value, which is updated as follows:

$$R_{t+1} = \frac{1}{\theta}\sum_{1}^{t}\Phi(R_i)\bullet R_{i+1}^{other}\bullet\left(W_{i+1} - E(R_{i+1})\right)$$

$$\Phi(R) = 1 - \frac{1}{1+e^{\frac{-(R-D)}{\sigma}}}$$

$$E(R_{t+1}) = R_t / D$$

Equation 1 Sporas formulae

Where,

- t is the number of ratings the user has received so far,
- θ is a constant integer greater than 1,
- W_i represents the rating given by the user i,
- R^{other} is the reputation value of the user giving the rating
- D is the range of the reputation values,
- σ is the acceleration factor of the dumping function Φ.

The smaller the value of σ, the steeper the dumping factor $\Phi(R)$.

New users start with reputation equal to 0 and can advance up the maximum of 3000. The reputation ratings vary from 0.1 for terrible to 1 for perfect. Since the reputation of a user in the community is the weighted average of non-negative values, it is guaranteed that no user can ever have a negative reputation value, thus no user can ever have lower than that of a beginner. Also the weighed average schema guarantees that no user exceeds the maximum reputation value of 3000. If a user has a persistent real reputation value, the iteration of Equation 1 over a large number of ratings will give as an estimate very close to that value [Figure 1].

As we can see from Equation 1, the change in the reputation value of the user receiving a rating of W_i from user R_i^{other}, is proportional to the reputation value R_i^{other} of the rater himself. The expected rating of a user is his/her current reputation value over the maximum reputation value allowed in the system. Thus if the submitted rating is less than the expected one the rated user loses some of his reputation value.

The value of θ determines how fast the reputation value of the user changes after each rating. The larger the value of θ, the longer the memory of the system. Thus, just like credit card history schemes [7], even if a user enters the system being really unreliable in the beginning, if he/she improves later, his/her reputation value will not suffer forever from the early poor behavior.

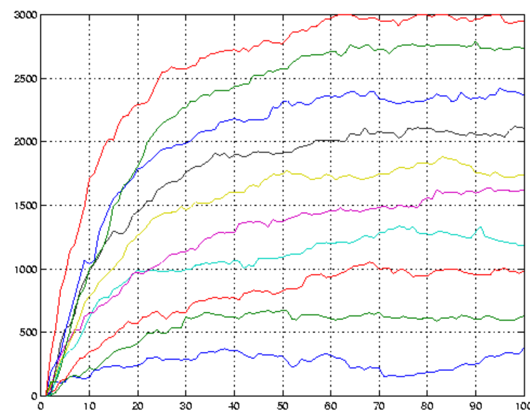

Figure 1 Change of reputation for 10 different users over 100 ratings with θ=10

5. Histos: A reputation mechanism for highly connected online communities

"Although an application designer's first instinct is to reduce a noble human being to a mere account number for the computer's convenience, at the root of that account number is always a human identity." Weaving a Web of Trust [9].

The reputation mechanism described in the previous section provides a global reputation value for each member of the online community, which is associated with them as part of their identity. Besides the online agent mediated interaction, our users will eventually have to meet each other physically in order to commit the agreed transaction, or they may even know each other through other social relationships. The existing social relationships as well as the actual physical transaction process create personalized biases on the trust relationships between those users. FFF [12] and the PGP web of Trust [17] use the idea that as social beings we tend to trust a friend of a friend more than a total stranger.

Following a similar approach, we decided to build a more personalized system. In Weaving a Web of Trust [9], what matters is that there is a connected path of PGP signed webpages between two users. In our case we have to take into consideration the different reputation ratings connecting the users of our system.

We can represent the pairwise ratings in the system as a directed graph, where nodes represent users and weighted edges represent the most recent reputation rating given by one user to another, with direction pointing towards the rated user. If there exists a connected path between two users, say from A to A_L, we can compute a more personalized reputation value for A_L.

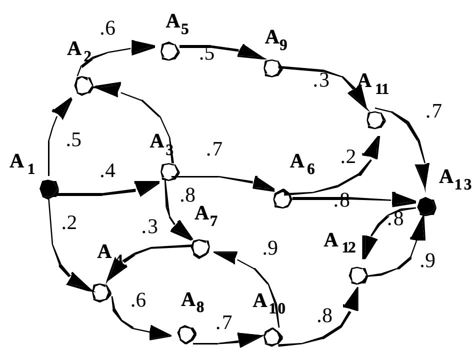

Figure 2 A directed graph representing the rating paths between user A_1 and A_{13}

When the user A submits a query for the Histos reputation value of a user A_L we perform the following computation:

a) The system uses a Breadth First Search algorithm to find all directed paths connecting A to A_L that are of length less than or equal to N. As described above we only care about the chronologically θ most recent ratings given to each user. Therefore, if we find more than θ connected paths taking us to user A_L, we are interested only in the most recent θ paths with respect to the chronological order of the rating events represented by the last edge of the path.

b) We can evaluate the personalized reputation value of A_L if we know all the personalized reputation ratings of the users at the last node of the path before A_L. Thus, we create a recursive step with at most θ paths with length at most N-1.

c) If the length of the path is only 1, it means that the particular user, say C, was rated by A directly. The direct rating given to user C is used as the personalized reputation value for user A. Thus, the recursion terminates at the base case of length θ and has an order of growth bounded by:

$$O(q \bullet N)$$

Equation 2 Order of growth of Histos

Note that for any length θ user A may have even been among the last θ users that have rated A_L directly. However, user A has the option of getting other peoples' opinions about A_L by evaluating his personalized value for A_L in a more collaborative fashion. Also for the purpose of calculating the personalized reputation values, we use a slightly modified version of the reputation function described above. For each user A_L, with m connected paths coming towards A_L from A, we calculate the reputation of A_L as follows:

$$R_{t+1} = \frac{1}{q'} \sum_{t-q'}^{t} \Phi(R_{i+1}) \bullet \left(R_{i+1}^{other} \bullet W_{i+1}\right) \Big/ \sum_{t-q'}^{t} R_{i+1}$$
$$q' = \min(q, m)$$
$$m = \deg(A_L)$$

Equation 3 Histos formulae

Where deg (A_L) is the number of connected paths from A to A_L with length less than or equal to the current value of L. In the base case where L=1, since we have a connected path it means that A has rated A_1 him/herself, the personalized value for A_1 is naturally the rating given by A.

In order to be able to apply the Histos mechanism we need a highly connected graph. If there does not exist a path from A to A_L with length less than or equal to N, we fall back to the simplified Sporas reputation mechanism.

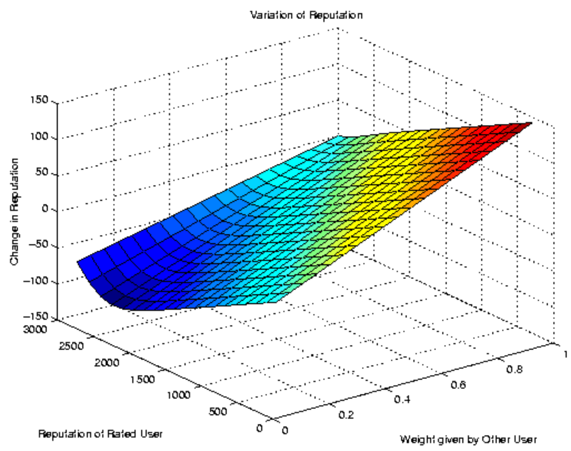

Figure 3 Change of Reputation with respect to value of the other user and the weight received.

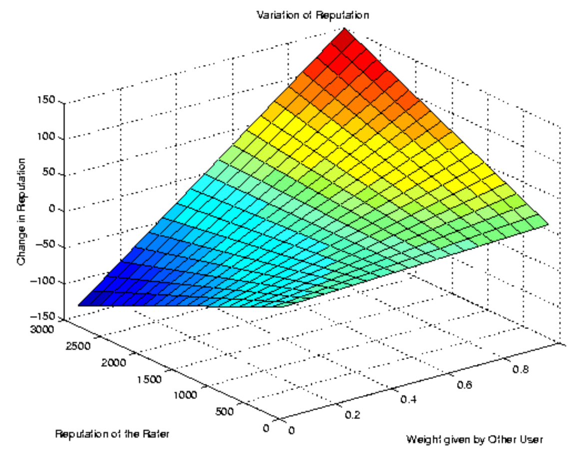

Figure 4 Change of Reputation with respect to the value of the two users.

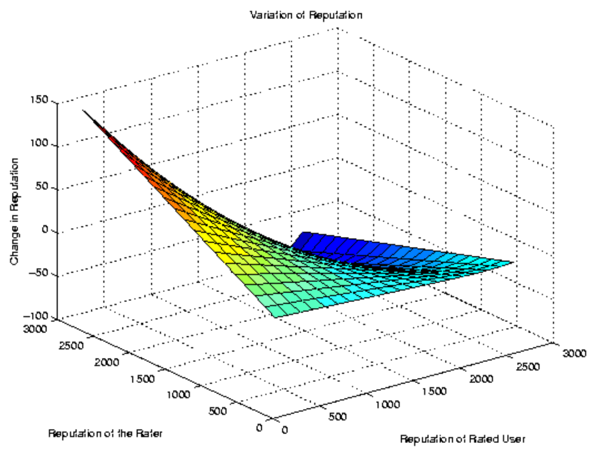

Figure 5 Change of Reputation with respect to the value of the user rated and the weight received.

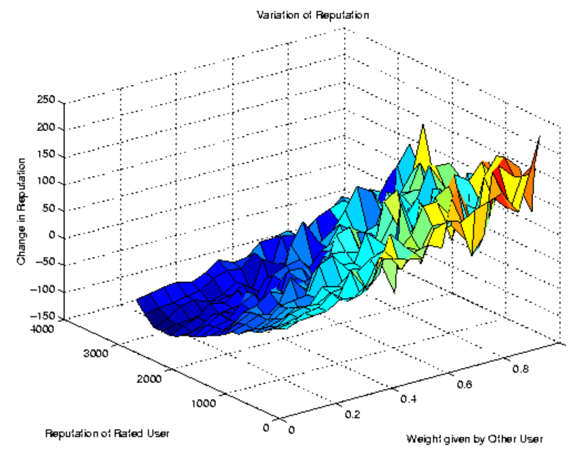

Figure 6 Change of Reputation with respect to the value of the other user and the weight received.

6. Results

While we are still gathering real data from our experiment with Kasbah, we ran some simulations in order to test our system. The four figures above represent the results from some preliminary simulations we have run in order to evaluate our proposed solution for Sporas. Figure 6 shows the results of an older version of Sporas, where the reputation value is calculated in the same way as in Histos, where only the last θ ratings count towards the total current reputation value of he user. This is how we actually calculate the reputation values in each step in Histos, without showing the effect of transitive perspectives.

In the first three graphs we calculated the reputation R_t as follows:

$$R_{t+1} = \frac{1}{q}\Phi(R_i)R_{i+1}^{other}\left(W_{i+1} - {R_t}/{D}\right) + R_t$$

Equation 4 Simulation formula

- Figure 3 and Figure 6 give the change of the Reputation of a user A, with average reputation (1500), rated by 20 users with reputations varying from 150 to 3000. The graph shows how much the reputation of the user would change if he/she received any rating between 0.1 and 1.
- Figure 4 gives the change of the reputation of a user A, who receives an average rating with respect to his/her own reputation and the user B who rates A.

The graph shows how the change in the reputation of A varies if the reputations of users A and B vary from 150 to 3000.

- Figure 5 gives the change in the reputation of a user A, if A is rated by a user B with an average reputation (1500), with respect to the previous reputation of user A and the rating which user B gives to user A. Like the two previous cases the ranking of user B varies from 150 to 3000, and the weight B gives to A varies from 0.1 and 1.

These graphs demonstrate the desired behavior by satisfying all the desiderata in a pairwise reputation system for an online community. As we can see from Figure 4 and Figure 5, even if the user giving the feedback has a very high reputation, he/she cannot affect significantly the reputation of a user with an already very high reputation. However if the user rated has a low reputation rating, he/she occurs much more significant updates whenever he/she receives a new feedback. In Figure 4 we can also see that when the user giving the feedback has a very low reputation, the effect of the rating is very small unless the user being rated has a very low reputation value him/herself. In this case the effect is actually negative for the user being rated. In Figure 5 we observe exactly the same phenomena with respect to the weight given as feedback.

7. Conclusion

Collaborative filtering methods have been around for some years now, but they have focused on content rating and selection. We have developed two collaborative reputation mechanisms that establish reputation ratings for the users themselves. Incorporating reputation mechanisms in online communities may induce social changes in the way users participate in the community.

We have discussed desiderata for reputation mechanisms for online communities and presented 2 systems that were implemented in Kasbah, an electronic marketplace.

However, further testing is required in order to evaluate the effects of our mechanisms in both the domains of Electronic Commerce and online communities. We are currently developing a web-based gateway that will act as the reputation server of a listserv list.

12. References

[1] Anthony Chavez and Pattie Maes. "Kasbah: An Agent Marketplace for Buying and Selling Goods". Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96). London, UK, April 1996

[2] eBay http://www.ebay.com

[3] OnSale http://www.onsale.com/exchange.htm

[4] Stephen Paul Marsh, Formalising Trust as a Computational Concept, PhD Thesis, University of Stirling, April 1994.

[5] Joseph M. Reagle Jr., Trust in a Cryptographic Economy and Digital Security Deposits: Protocols and Policies, Master Thesis, Massachusetts Institute of Technology, May 1996.

[6] Better Business Bureau, http://www.bbb.org

[7] Fair Isaak Co, http://www.fairisaac.com

[8] Lenny Foner, “Yenta: A Multi-Agent, Referral Based Matchmaking System”, The First International Conference on Autonomous Agents (Agents '97), Marina del Rey, California, February 1997.

[9] Rohit Khare and Adam Rifkin, "Weaving a Web of Trust", summer 1997 issue of the World Wide Web Journal (Volume 2, Number 3, Pages 77-112).

[10] Paul Resnick and Jim Miller, PICS: Internet Access Controls Without Censorship, Communications of the ACM, 1996, vol. 39(10), pp. 87-93.

[11] Recreational Software Advisory Council http://www.rsac.org

[12] Nelson Minar and Alexandros Moukas, Friend of a Friend Finder, Software Agents Group, MIT Media Laboratory, Working paper 1998

[Firefly] http://www.firefly.net

[13] Upendra Shardanand and Pattie Maes, "Social Information Filtering: Algorithms for Automating 'Word of Mouth’, " Proceedings of the CHI-95 Conference, Denver, CO, ACM Press, 5/1995.

[14] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, John Riedl, GroupLens: An Open Architecture for Collaborative Filtering of Netnews, from Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work, Chapel Hill, NC: Pages 175-186

[15] Elo, A. E., "The Rating of Chessplayers, Past and Present", Arco Publishing, Inc. (New York) 1978.

[16] Mark E. Glickman, Paired Comparison Models with Time-Varying Parameters, PhD Thesis, Harvard University, May 1993.

[17] Simson Garfinkel. PGP: Pretty Good Privacy, O'Reilly and Associates, 1994.