Fuzzy Trust for Peer-to-Peer Systems

Nathan Griffiths
Dept. of Computer Science
University of Warwick
Coventry, CV4 7AL, UK
nathan@dcs.warwick.ac.uk

Kuo-Ming Chao
School of MIS
Coventry University,
Coventry, CV1 5FB, UK
k.chao@coventry.ac.uk

Muhammad Younas
Dept. of Computing
Oxford Brookes University
Oxford, OX33 1HX, UK
m.younas@brookes.ac.uk

Abstract

Peer-to-peer (P2P) systems are based upon the cooperative interactions of member peers. Typically, peers are both autonomous and self-interested, meaning that there is no hierarchy of control or power, and that individuals seek to maximise their own goal achievement, rather than acting in a benevolent or socially-oriented manner. Consequently, interaction outcomes are uncertain, since peers can break their commitments or provide sub-standard contributions or services. Thus, when a peer cooperates it is entering into an uncertain interaction, that has an associated risk of failure or reduced performance. For peers to be effective they need some mechanism for managing this risk of failure. In this paper we show how peers can use trust to manage this risk. Our model of trust uses fuzzy logic to allow peers to represent and reason with uncertain and imprecise information regarding others' trustworthiness.

1. Introduction

Over the last decade peer-to-peer (P2P) systems have become increasing popular, and have been used to provide solutions in areas as diverse as distributed computation, voice over IP, and knowledge sharing [15]. P2P systems rely on the collaboration of two or more peers using appropriate information and communication systems, without the necessity for central coordination [12]. There is no explicit hierarchy of control or power; peers are equal and autonomous. A peer has individual capabilities, knowledge and resources that are made available to others, potentially in return for imposing some cost. Individuals have private preferences regarding the nature of their interactions, in particular with respect to balancing the cost, quality, timeliness etc. of an interaction. Using their knowledge of others' capabilities and previous performance a peer can, according to its preferences, select appropriate peers for cooperation.

Peers have varying degrees of reliability, quality and honesty, and interactions may fail, produce substandard results, or cost more and take longer than expected. Interactions are inherently uncertain, since high quality, timely, and on budget outcomes are not guaranteed. Since peers are autonomous they determine for themselves when to cooperate, when to cease cooperating, and how to conduct themselves. For example, a peer may choose to delay the provision of information, and reduce its quality. To function effectively peers must manage the risk of interactions failing or having reduced performance. In this paper we show how trust can be used to manage this risk. Our approach uses fuzzy logic to represent and reason with uncertain and imprecise information regarding peers' trustworthiness. We provide a mechanism for a peer to select appropriate interaction partners. Additionally, we show how peers can incorporate the related notions of distrust, untrust and undistrust into their reasoning.

2. Background

2.1 Trust and reputation

Trust and reputation are related, but distinct, concepts. The former represents a peer's *individual* assessment of the reliability, honesty etc. of another, while the latter is a *social* notion corresponding to a group assessment of such issues. Reputation is often built from a combination of individual trust assessments, and the process of combining individual assessments into a group notion, requires peers to make their private assessments of others publicly available. In some situations this can be undesirable from an individual's perspective, since it involves revealing private information that may reduce future effectiveness. For example, suppose that peer α frequently cooperates with β, who provides a reliable high quality and timely service. If α were to make information regarding β's reliability and quality (i.e. its trustworthiness) public, then β may become overloaded and unreliable for α's future interactions.

In providing trust information to establish reputation, a peer might reduce the effectiveness of its own future interactions. For a peer to provide such information, there must be some intrinsic motivation for information sharing. In the absence of such a motivation, there will be insufficient information to assess reputation. There are also general issues with reputation concerning the subjectivity and context-specific nature of feedback [4]. Although in many situations the benefits of reputation might outweigh the individual cost of trust information sharing, it is useful in general to consider trust and reputation as separate, enabling peers to use trust without considering reputation.

Many of the existing applications of trust to P2P systems use reputation, in which individual trust is globally aggregated into a reputation assessment [13, 14, 16]. Such research tends not to address how trust itself can be used in an individual's decision making. In this paper we focus specifically on trust for P2P systems, and do not consider reputation further. Our approach is complimentary to reputation-based models, and we view trust and reputation as both playing an important role in a complete system.

Existing approaches to trust can be categorised according to how trust is used: for security or for quality of service (QoS) [3]. In this paper we focus on the QoS perspective to enable peers to maximise the "quality" of their interactions according to their current preferences. Previous work on trust in P2P systems, tends to concentrate on addressing the security aspects, with little use of trust for QoS [1, 15].

2.2 Fuzzy logic for trust

Fuzzy logic offers the ability to handle uncertainty and imprecision effectively, and is therefore ideally suited to reasoning about trust. Fuzzy inference copes with imprecise inputs, such as assessments of quality or timeliness, and allows inference rules to be specified using imprecise linguistic terms, such as "very high quality" or "slightly late". Existing models have successfully used fuzzy logic to represent trust [7, 10] including in P2P systems [13]. However, each of these approaches uses trust as a means of establishing reputation, rather than focusing on individual trust. Moreover, the notions of distrust, untrust and undistrust are not considered. In this paper we describe a method that uses fuzzy logic to make assessments about various aspects of trust, and allows peers to make decisions based on trust. Before describing our trust model, however, we introduce some basic fuzzy concepts.

2.3 Basic fuzzy concepts

In classical set theory the membership of an object in a set is clearly defined: it is either a member or it is not. For example, a person of age 10 might be a member of the set $young$, and not of the set old. Such sets are required to have well-defined boundaries. However, the *concept* of young does not have a clear boundary, and in some contexts age 30 might be considered to be young, and not in others. *Fuzzy sets* are based on the notion of a *membership function*, $\mu(x)$, which defines the degree to which a fuzzy variable x is a member of a set. Full membership is represented by 1, and no membership by 0. The membership function $\mu(x)$ maps x into the interval $[0, 1]$. For example, age 35 might have a membership of 0.8 in a fuzzy set $\widetilde{y}$, representing young ages, and a 0.1 membership in the set $\widetilde{o}$ representing old ages. We use a tilde accent, $\tilde{x}$, to indicate that a set x is a fuzzy set. The *universe of discourse* (UoD) of a fuzzy set corresponds to the range of values that are considered, such as $[0, 130]$ for age. Fuzzy sets are used to define *terms* with respect to a *variable*. For example, the sets $\widetilde{y}$ and $\widetilde{o}$ define the terms $\widetilde{young}$ and $\widetilde{old}$ respectively, on the variable age. Terms can be subjected to *modifiers* (also called *linguistic hedges*), such as very or slightly, which serve to modify, or hedge, the membership function from its original definition. The former example *concentrates* the membership function, while the latter *dilates* it. The mathematical definition of such modifiers is beyond the scope of this paper, but we adopt Zadeh's definitions which follow the intuitive linguistic meaning [17]. *Fuzzy inference rules* of the form:

if x_1 is m_1 $\widetilde{t_1}$ **and** x_2 is m_2 $\widetilde{t_2}$ **then** y is m' $\widetilde{t'}$

define the relationship between antecedents (inputs) x_i and consequent (output) y, described by terms t_i and t' and optional modifiers m_i and m'. For example, we might have rules such as the following.

(R1) **if** age is $\widetilde{young}$ **and** $income$ is very $\widetilde{high}$
then $customerPotential$ is $\widetilde{high}$
(R2) **if** age is $\widetilde{old}$ **and** $income$ is $\widetilde{low}$
then $customerPotential$ is $\widetilde{medium}$

Rules are applied in parallel, and the conclusion membership degrees are aggregated by superimposing the resultant membership curves. We adopt a Mamdani min-max approach to inference, such that the membership degree of rule conclusions is clipped at a level determined by the minimum of the maximum membership values of the intersections of the fuzzy value antecedent and input pairs [6]. This ensures that the degree of membership in the antecedents is reflected in the output. We give an example of fuzzy inference for trust in Section 5. A crisp value can be obtained from the result of inference by *defuzzifying* the aggregated consequents. There are many methods for defuzzification, but for simplicity we take the centre of the area bounded by the membership curve. (Further discussion of the concepts introduced in this section can be found in [11].)

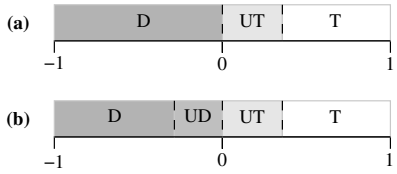

Figure 1. Marsh's notions of trust (a) and our addition of undistrust (b), where D, UD, UT, and T correspond to distrust, undistrust, untrust, and trust respectively.

3. Trust

Our proposed mechanism builds on existing work using service-oriented trust in agent-based systems. Trust is generally taken to be the belief that an agent will act in the best interests of another (i.e. will cooperate), even if given the opportunity to do otherwise (i.e. to defect) [2]. Most previous work on trust has concentrated on this positive view of trust, and has largely ignored the notion of *distrust*. Distrust is not simply the negation of trust [5], but rather it is a belief that an agent will act against the best interests of another [8]. Alternatively, *untrust* corresponds to the space between distrust and trust, in which an agent is positively trusted, but not sufficiently to cooperate with. This view of trust, proposed by Marsh, is illustrated in Figure 1(a). Marsh argues that distrust is an important concept, that can play an important role in an agent's reasoning, complimenting trust itself [8]. We concur with this view, and in this paper we provide a mechanism for agents to use distrust in their decision making. In addition to distrust, untrust and trust, however, we propose a fourth notion of *undistrust*. Untrust is defined as positive trust, but insufficient to support cooperation. For distrust to play a useful role in an agent's reasoning, we argue that a similar region of undistrust is needed, namely negative trust, but insufficient to make definite conclusions in the reasoning process. Figure 1(b) illustrates our definition of the notions of trust, distrust, untrust and undistrust.

4. Interaction histories

Trust is based on an individual peer's experiences, and so peers need to track their interaction histories, and record whether their expectations have been met. We take a multi-dimensional approach and view trust as comprising the combination of the different dimensions of an interaction, such as the quality of a task or the cost imposed for executing it. Cooperative interactions are typically more than simple succeed or fail tasks. For example, peers cooperate with an expectation of successful performance to a given quality for some anticipated cost. In addition to possible failure, a task may succeed but be of lower than expected quality or at a higher cost. Peers can model such characteristics as *dimensions of trust*, which taken together give an assessment of a peer's trustworthiness. For illustrative purposes, in this paper we consider the dimensions of success, cost, quality and timeliness.

In order to assess trust a peer must evaluate its experiences in each of the trust dimensions. For each interaction, and in each dimension, a peer's expectations will have been either met or not met. Peers maintain a history of the interactions that they have had with each other peer, and track the number of successful and unsuccessful interactions for each dimension, in terms of whether their expectations were met. Thus, for each dimension, d, and peer that has been cooperated with, α, a peer maintains a value I_α^{d+} which corresponds to the number of interactions in which its expectations were met, and a value I_α^{d-} in which they were not met. From these values, the *experience*, E_α^d, in each dimension d, for each peer α, can be calculated as:

$$E_\alpha^d = \frac{I_\alpha^{d+} - I_\alpha^{d-}}{I_\alpha^{d+} + I_\alpha^{d-}}$$

These experiences are crisp values in the interval $[-1, 1]$ and must be translated into fuzzy values in order to reason about trust. Experience values are based directly on a peer's interaction histories, and so they are not uncertain in themselves. Therefore, they can each be fuzzified by translating them into a fuzzy value defined by the singleton fuzzy set whose membership function is 0 at all points except for E_α^d which has a membership of 1. Thus, the fuzzified experience is given by $\widetilde{E}_\alpha^d = fuzzySingleton(E_\alpha^d)$.

Peers keep track of the outcomes of their interactions by using a window of experiences that is maintained for each peer. This window is bounded, such that there is an upper limit on the number of interactions that are recorded for a peer. The interaction window acts as a first-in first-out queue, and when full it is the earliest experiences that are removed to be replaced by new ones. Over time, however, the information stored may become outdated if the peer's environment has changed and its previous experiences are no longer relevant. Peers may change, and a peer that was reliable previously may no longer be so. To address this problem a peer purges outdated experiences from its interaction windows after a certain predefined period. Thus, even if an interaction window is not full, the record of experiences will be removed over time. The delay between the occurrence of an interaction and the removal of its record from the interaction window is called the *purge lag*, and has a direct influence on how quickly a peer's trust assessments respond to changes in its environment. A small purge lag, meaning that interaction records do not persist for long, means that the effect of previous experiences decays quickly and trust assessments respond quickly to changes. However, a small

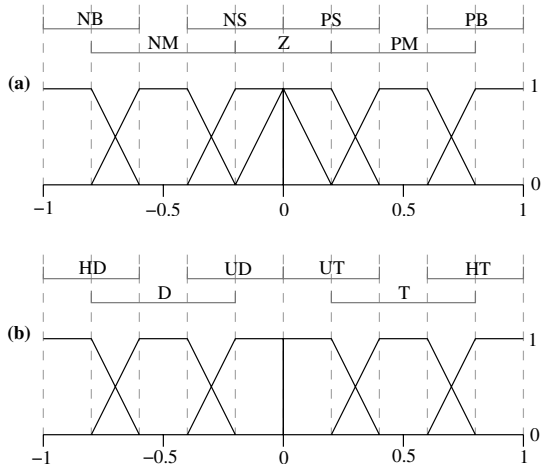

Figure 2. Definition of terms for: (a) experience, where NB, NM, NS, Z, PS, PM, and PB correspond to negative big, negative medium, negative small, zero, positive small, positive medium and positive big respectively, and (b) trust, where HD, D, UD, UT, T, and HT correspond to high distrust, distrust, undistrust, untrust, trust, and high trust respectively.

purge lag also reduces the extent of the experiences that can be used to determine trust. If the purge lag is too small there will be insufficient experiences on which to base trust, and small perturbations will have a significant effect on trust.

In determining trust it is important that an peer has sufficient experience on which to calculate trust. We define the confidence level in the experience for a particular dimension as the total number of interactions on which it is based.

$$confidence_d = I_\alpha^{d+} + I_\alpha^{d-}$$

If this confidence level is below a predefined threshold then the peer will use a default value for trust, rather than basing its assessment on a small number of experiences.

5 Fuzzy trust

We define fuzzy terms for experience in each of the dimensions in which peers record their interactions, in our case success, cost, quality, and timeliness. Fuzzy terms are defined in reference to fuzzy variables, and for experience we define fuzzy variables for each trust dimension. Thus, for our chosen dimensions we introduce E_α^s, E_α^c, E_α^q and E_α^t corresponding to success, cost, quality, and timeliness for peer α respectively. The UoD of these fuzzy variables is $[-1, 1]$. For each of these variables we define the terms:

(R1) **if** E_α^d is $\widetilde{negativeBig}$ **then** T_α is $\widetilde{highDistrust}$
(R2) **if** E_α^d is $\widetilde{negativeMedium}$ **then** T_α is very $\widetilde{distrust}$ or $\widetilde{undistrust}$
(R3) **if** E_α^d is $\widetilde{negativeSmall}$ **then** T_α is $\widetilde{undistrust}$
(R4) **if** E_α^d is $\widetilde{zero}$ **then** T_α is $\widetilde{undistrust}$ or $\widetilde{untrust}$
(R5) **if** E_α^d is $\widetilde{positiveSmall}$ **then** T_α is $\widetilde{untrust}$
(R6) **if** E_α^d is $\widetilde{positiveMedium}$ **then** T_α is very $\widetilde{trust}$ or $\widetilde{untrust}$
(R7) **if** E_α^d is $\widetilde{positiveBig}$ **then** T_α is $\widetilde{highTrust}$
...
(Rn) **if** T_α is $\widetilde{highTrust}$ and F_α^c is $\widetilde{medium}$ and F_α^q is very $\widetilde{high}$ **then** R_α is $\widetilde{high}$
(Rm) **if** T_α is $\widetilde{distrust}$ and F_α^c is $\widetilde{medium}$ and F_α^q is $\widetilde{high}$ **then** R_α is $\widetilde{reject}$

Figure 3. Example fuzzy inference rules.

negative big, negative medium, negative small, zero, positive small, positive medium, and positive big. Other terms are possible, but these are sufficient for our purposes. The fuzzy sets that describe these terms are illustrated in Figure 2(a).

In order to use fuzzy inference to determine trust, T_α, in a peer α we must also define trust as a fuzzy variable, with an associated set of fuzzy terms. The UoD for trust is also $[-1, 1]$, i.e. complete distrust to complete trust, and we define the terms: high distrust, distrust, undistrust, untrust, trust, and high trust. These terms allow us to represent the attitudes of untrust and undistrust, in addition to two degrees of trust and distrust. Again, other definitions are possible, but these are sufficient for our application. The fuzzy sets that describe these terms are illustrated in Figure 2(b).

For each dimension we define a set of fuzzy inference rules that take the fuzzified experiences as antecedents and make conclusions regarding trust. The definition of these rules is the responsibility of the system developer, and we do not prescribe a particular rule set. In the simulation described in Section 6 we use the rules R1–R7 given in Figure 3. Other rules are, of course, possible and can be easily incorporated into the system.

To determine the trustworthiness of the potential interaction partners we must consider the inference rules for each of the trust dimensions. First, the peer checks whether its has sufficient confidence for each dimension, i.e. that $confidence_d > minConfidence$. Provided that there is sufficient confidence then fuzzy inference begins. Each rule is considered in turn, and if there is a match between the input (i.e. E_α^d) and the fuzzy set defined by the antecedent of the rule, then the rule is fired. For example, if there is an overlap between the input E_α^d and the area defined by the

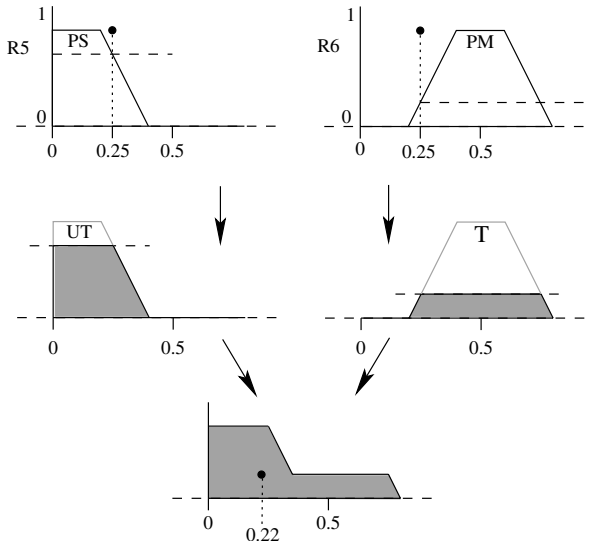

Figure 4. A simple inference example showing the firing of rules R5 and R6.

term $negativeBig$ then rule R1 is fired. If there is insufficient confidence in a particular dimension, the peer uses a default "experience" value for that dimension, $default_d$. This value is determined by the peer's trusting disposition, with optimists using higher values than pessimists.

By way of example, suppose that for the success dimension we have determined that $E^s_\alpha = 0.25$ for peer α based on the experiences recorded in the interaction window. This crisp value is fuzzified as described above, and the fuzzy rules are then applied. In this case the input set matches with the antecedents of rules R5 and R6, i.e. $fuzzySingleton(0.25)$ overlaps with the sets defined by the terms $positiveSmall$ and $positiveMedium$. Using Mamdani min-max inference the membership of the conclusion fuzzy set is clipped by the degree of membership of the antecedent. The outputs of the rules are then aggregated by taking the fuzzy union. This is shown graphically in Figure 4. The process is then continued for the other dimensions, with the outputs from any matching rules being combined with the existing output by taking the fuzzy union. Once rules R1–R7 have be applied for all dimensions we have determined a fuzzy value for trust T_α. A crisp value can be determined by defuzzifing as shown in Figure 4, in this case resulting in a trust of 0.22.

5.1 Additional decision factors

Peers might simply use trust to select which peer to cooperate with, by selecting the most trusted. However, typically there is additional information with which to make a decision. For example, each of the alternative peers may advertise a cost and quality for the interaction. In this case, the selecting peer can incorporate such information into its decision making. Since these advertised values represent uncertain information (i.e. the actual cost and quality are unknown at the point of making a decision), they lend themselves to fuzzy inference. Thus, we introduce fuzzy rules that combine trust with each of the other decision factors and determine a rating for each alternative peer. Each of these factors F_i is a crisp value, which can be fuzzified as a singleton set. We define a set of inference rules that have fuzzy trust and the fuzzy decision factors as antecedents and the *rating* for a peer as conclusions. These factors are domain specific. In our example, we use advertised cost and quality from peer α, denoted F^c_α and F^q_α respectively. Suppose that we have defined the fuzzy terms low, $medium$ and $high$ for these factors, according to the UoD defined by the range of potential advertised cost and quality values. Similarly, suppose that we have terms low, $medium$, $high$ and $reject$ defined for ratings, which has a UoD of $[0, 1]$. We then define a set of rules of form illustrated by Rn and Rm in Figure 3. Rule Rn states that if a peer is trusted, has a medium advertised cost and a high advertised quality, then it has a high rating. Similarly, rule Rm states that if a peer is highly distrusted, has a medium advertised cost and high advertised quality, then it should be rejected. Rule Rm is an example of how distrust is used in reasoning. As for calculating trust, each of these rules is applied in parallel using Mamdani min-max inference, and a crisp rating value for peer α is obtained by defuzzifying the fuzzy rating. To balance the importance of the various decision factors (including trust), peers can scale the inputs before performing inference. For example, if timeliness is not currently important then the input E^t_α would be multiplied by some reduction factor, r, where $0 < r < 1$.

In order to select a peer to cooperate with, the selecting peer calculates the rating value for each alternative, and selects the one with the highest rating. After the interaction, the interaction window is updated according to whether the interaction was successful, and whether the expected (as determined by advertised value) cost and quality were met.

5.2 Bootstrapping

Initially peers have insufficient experience for reasoning. Therefore, each peer goes through a bootstrapping phase in which partner peers are chosen randomly by way of exploration. If a particular peer is found to be highly distrusted then it is excluded from the remainder of the bootstrapping phase. However, during bootstrapping undistrusted, untrusted, and trusted peers have an equal chance of being selected.

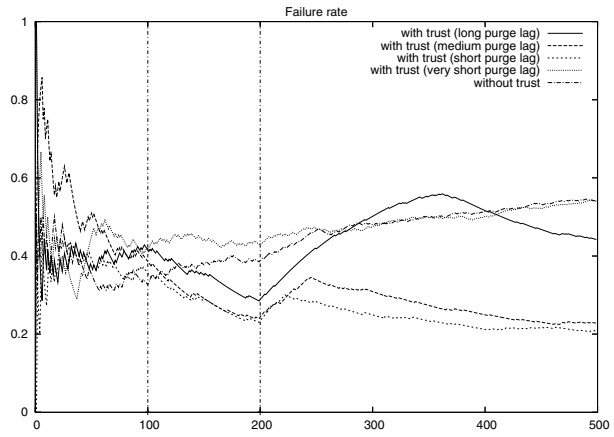

Figure 5. Failure rate with and without trust.

6. Example results

Our approach has been validated experimentally. The decision mechanism has been implemented using the NRC FuzzyJ Toolkit [9], and we have experimented with peers situated in a simulated environment. We constructed a stub that represents all interactions with other peers, and we are able to simulate interactions with a predefined number of other peers over some specified duration. Figure 5 shows an example of our initial results, giving failure rate versus interactions for fuzzy trust versus using advertised quality and cost only. These results are for a single peer situated in an environment of 50 others that are randomly initialised. The vertical line after 100 interactions signifies the end of the bootstrapping phase, and at 200 interactions we randomly changed the reliability of some of the simulated peers. The lowest failure rate is obtained using trust with a short purge lag. A very short purge lag does not give sufficient information to reason with, and the failure rate is similar to without trust. Longer purge lags eventually give similar failure rates, but take significantly longer to stabilise (both after bootstrapping and after changes in environment). We do not discuss our results further, due to space constraints.

7. Conclusions

In this paper we have shown how fuzzy logic can be used to represent trust, and select appropriate peers for cooperation. We have proposed a new notion of *undistrust* and incorporated this into the reasoning process. Our system is flexible; the fuzzy rules are specifiable by a system designer, and peers are able to scale inputs according to their current preferences. Initial experimental results indicate that the approach provides a significant reduction in failure rate. There are many areas of ongoing work, including additional experimentation, and integration with existing models of reputation.

References

[1] M. Bursell. Security and trust in P2P systems. In R. Subramanian and B. D. Goodman, editors, *Peer-to-Peer Computing: The Evolution of a Disruptive Technology*, pp. 145–165. Idea Group Publishing, 2005.

[2] D. Gambetta. Can we trust trust? In D. Gambetta, editor, *Trust: Making and Breaking Cooperative Relations*, pp. 213–237. Basil Blackwell, 1988.

[3] N. Griffiths. Trust: Challenges and opportunities. *AgentLink News*, 19:9–11, 2005.

[4] N. Griffiths and K.-M. Chao. Experience-based trust: Enabling effective resource selection in a grid environment. In *Proc. of Int. Conf. on Trust Management*, pp. 240–255. 2005.

[5] N. Luhmann. Familiarity, confidence, trust: Problems and alternatives. In D. Gambetta, editor, *Trust: Making and Breaking Cooperative Relations*, pp. 94–107. Basil Blackwell, 1988.

[6] E. H. Mamdani and S. Assilian. An experiment in linguistic synthesis with a fuzzy logic controller. *Int. J. of Man-Machine Studies*, 7(1):1–13, 1975.

[7] D. W. Manchala. E-commerce trust metrics and models. *IEEE Internet Computing*, 4(2):36–44, 2000.

[8] S. Marsh and M. R. Dibben. Trust, untrust, distrust and mistrust — an exploration of the dark(er) side. In *Proc. of Int. Conf. on Trust Management*, pp. 17–33. 2005.

[9] NRC Institute for Information Technology. The FuzzyJ toolkit. www.iit.nrc.ca/IR_public/fuzzy/fuzzyJToolkit.html.

[10] S. D. Ramchurn, C. Sierra, L. Godo, and N. R. Jennings. Devising a trust model for multi-agent interactions using confidence and reputation. *Artificial Intelligence*, 18(9–10):833–852, 2004.

[11] T. J. Ross. *Fuzzy Logic With Engineering Applications*. Wiley, 2nd ed., 2004.

[12] D. Schoder and K. Fischbach. Peer-to-peer prospects. *Communications of the ACM*, 46(2):27–29, 2003.

[13] S. Song, K. Hwang, R. Zhou, and Y.-K. Kwok. Trusted P2P transactions with fuzzy reputation aggregation. *IEEE Internet Computing*, 9(6):24–34, 2005.

[14] N. Stakhanova, S. Basu, J. Wong, and O. Stakhanov. Trust framework for P2P networks using peer-profile based anomaly technique. In *Proc. Int. Work. on Security in Distributed Computing Systems*, pp. 203–209, 2005.

[15] M. Waldman, L. F. Cranor, and A. Rubin. Trust. In A. Oram, editor, *Peer-To-Peer: Harnessing the Power of Disruptive Technologies*, pp. 242–270. O'Reilly, 2001.

[16] L. Xiong and L. Liu. PeerTrust: Supporting reputation-based trust in peer-to-peer communities. *IEEE Trans. on Knowledge and Data Engineering*, 16(7):843–857, 2004.

[17] L. A. Zadeh. A fuzzy-set-theoretic interpretation of linguistic hedges. *J. of Cybernetics*, 2(3):4–34, 1972.