Asymptotically Idempotent Aggregation Operators for Trust Management in Multi-agent Systems

Roberto Ghiselli Ricci
Università degli Studi di Sassari
DEIR – Via Torre Tonda, 34
07100 Sassari (SS) – Italia
ghiselli@unimo.it

Marco Viviani
Università degli Studi di Milano
DTI – Via Bramante, 65
26013 Crema (CR) – Italia
viviani@dti.unimi.it

Abstract

The study of trust management in multi-agent system, especially distributed, has grown over the last years. Trust is a complex subject that has no general consensus in literature, but has emerged the importance of reasoning about it computationally. Reputation systems takes into consideration the history of an entity's actions/behavior in order to compute trust, collecting and aggregating ratings from members in a community. In this scenario the aggregation problem becomes fundamental, in particular depending on the environment. In this paper we describe a technique based on a class of asymptotically idempotent aggregation operators, suitable particulary for distributed anonymous environments.

Keywords: Trust management, multi-agent systems, P2P, aggregation operators, asymptotical idempotency.

1 Introduction

Trust management is still nowadays a complex subject, since there is no general consensus in literature on what trust is; however many researchers have recognized the value of modeling trust and reasoning about it computationally.

According to the definition given by McKnight and Chervany [14], trust is "the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible". The definition given in Mui *et al.* [15] add to the previous one a reference to past encounters, introducing the concept of trust based on *reputation*, described as "a subjective expectation a party has about another's future behavior based on the history of their encounters".

Aspects connected to trust and reputation has been investigated in the last years particulary in multi-agent systems [8]. These environments borrow the concept of social network from sociology: they are composed of autonomous agents related each other by means of interconnection roles or communication links. In the resulting network, individual components interact to achieve some overall objective, transmitting information, services and data to each other. The basic problem is that information about transactions performed between components is dispersed throughout the network so that every party can only build an approximation of the global view. Furthermore components storing and processing trust related data cannot be considered as unconditionally trustworthy themselves and their possible malicious behavior must be taken into account [1].

In these environments *trust management* is therefore concerned with collecting the information required to make a decision based on trust, evaluating the criteria related to

trust, and monitoring and re-evaluating existing trust values. The history of an entity's actions/behavior is used to compute trust, and in the absence of first-hand knowledge, referral-based trust (information from others) can be used. In such a scenario, reputation can be considered as a collective measure of trustworthiness (in the sense of reliability) based on referrals or ratings from members in a community; it is a social notion generally built by combining trust assessments given by a group or agents to obtain a single value representing an estimate of reputation.

In this paper we will focus our attention on the aspect of aggregating ratings of members of a multi-agent system about a generic "entity", that could be either a single agent or a resource. In particular we will describe aspects connected to the use of numerical approaches and we will illustrate a technique based on asymptotically idempotent aggregation operators [9][10] studied to provide meaningful reputation-based trust values (reputation score) after a process of collection and aggregation of ratings.

2 Reputation Systems

A *reputation system* is built on top of a multi-agent system; it collects and aggregates all the information about the past behavior of a group of entities. According to Resnick *et al.* [17] ratings about past interactions must guide decisions about current interactions.

The network architecture determines how ratings and reputation scores (derived from the aggregation of ratings) are communicated between participants in a reputation systems. The two main types are centralized and distributed architectures.

In a centralized system, a central authority (reputation center) collecting all the ratings typically derives a reputation score for every participant and makes all scores publicly available.

In a distributed system there is no central location for submitting ratings or obtaining reputation scores of others. Instead, there can be distributed stores where ratings can be submitted, or each participant simply records the opinion about each experience with other parties and provides this information on request from relying parties. Peer-to-Peer (P2P) networks represent nowadays an environment well suited for distributed reputation management. In P2P networks, every node plays the role of both client and server, and is therefore sometimes called a servent. This allows the users to overcome their passive role typical of web navigation, and to engage in an active role by providing their own resources.

2.1 Computing Reputation

In [12] are described various principles for computing reputation: belief models, discrete models, fuzzy models, flow models. In this paper we will concentrate on numerical/treshold-based models. They are the most commonly used when agents represent the trustworthiness of others in numerical intervals, typically $[-1, +1]$ or $[0, 1]$. The lower bound corresponds to complete distrust and the upper bound to blind trust. The numerical value representing the reputation of an agent is updated by some function after each interaction, and usually these methods tend to use a threshold to define different levels of trustworthiness. Probabilistic methods are a subset of numerical approaches in which trust is represented in the interval $[0, 1]$ and this number represents a probability with a clear semantics associated with it. Among many probabilistic approaches there are those based on Bayesian probability distributions.

2.2 Related Work

Depending on the reputation system architecture, several methods of aggregating ratings have been described in literature.

In *centralized reputation systems*, numerical methods have been developed in particular in auction sites. The eBay reputation system [18] sums the number of positive ratings and negative ratings separately, and keep the total score as the positive score minus the negative one. Such a simplistic aggregation of ratings

can be unreliable, particulary when some buyers do not return ratings. A positive sum of ratings could be due to the fact that there are people that do not report bad ratings; furthermore the lack of ratings is not considered in the aggregation process, it is simply discarded.

Amazon.com, Epinions.com and other auction sites feature reputation systems like eBay's, with variations such as a rating scale from $1 - 5$, use slightly more advanced schemes, or using several measures (friendliness, prompt response, quality product, etc.), or averaging rather than totaling feedback scores. Bizrate.com rates registered retailers by asking consumers to complete a survey after each purchase.

So-called "expert sites" (Allexperts.com, Askme.com) provide Q&A forums in which experts provide answers for questions posted by clients in exchange for reputation points and comments.

Product review sites (Epinions.com) offer rating services for product reviewersthe better the review, the more points the reviewer receives. iExchange.com tallies and displays reputations for stock market analysts based on the performance of their picks.

In *distributed reputation systems*, the problem of the aggregation of referrals has been handled in several and different ways depending on the trust model adopted.

In some of them, ratings may have not to be aggregated at all. *Game-theoretic reputation models* for example take a different approach: if the reputation system is designed properly, trustworthy behavior emerges as the most convenient one. Several game-theoretical approaches to trust management have been proposed by economic systems research [16]. Unfortunately, game theoretical approaches need a relatively high number of transactions to reach equilibrium, making them less suitable than direct aggregation for many applications.

An evolution of these techniques is *network-based* reputation aggregation. This class of approaches normally implies the aggregation of all reputation information available on a (previously established) trust graph. This process requires checking all paths on the trust graph from the computation initiator to the candidate partner and aggregating reputation values along them; finally, path reputations are merged into a network-wide value. Network-based aggregation of reputation is at the basis of several proposals, including the EigenTrust system by Hector Garcia-Molina *et al.* [13]. It computes agent trust scores in P2P networks through repeated and iterative multiplication and aggregation of trust scores along transitive chains until the trust scores for all agent members of the P2P community converge to stable values. However, its complexity is high and its overhead in terms of messages is not negligible.

Probabilistic approaches [19] use *Bayesian networks*, taking binary ratings as inputs and computing trust scores by statistically updating probability density functions (PDF). In [2], Karl Aberer and Zoran Despotovic show that a simple probabilistic technique, called *maximum likelihood estimation*, can substantially reduce overhead when employed as the feedback aggregation strategy.

The same line of reasoning in favor of straightforward aggregations can be applied to numerical approaches based on aggregation operators [3][6].

P2PRep [5] is a reputation-based protocol where reputation and trust are represented as single values in the interval $[0, 1]$. If an agent has some information to offer and wants to express an opinion on another one, it generates a vote based on its experience. An entity receives a set of votes from agents having expressed their opinions on the same agent, and evaluates the votes to collapse any set of votes that may belong to a clique and explicitly selects a random set of votes for verifying their trustworthiness [7]. Later, an aggregated community-wide reputation value is computed from the set of ratings collected before. Based on this value, it is possible to take a decision on whether using data, information or services provided by a certain agent.

In the Yu and Singh model [21] the absence of information is taken into account, using the Dempster-Shafer theory of evidence to model information received [20]. An agent may receive a good or bad ratings ($+1$ or -1) and the lack of belief in their model is considered as a state of uncertainty. Beliefs obtained from various sources are combined by Dempster's rules.

3 Choice of the Aggregation Operator

From literature and experience emerge that choosing the right operator for aggregating ratings is by no means straightforward. In numerical approaches in particular it is necessary choosing an aggregation operator that has a non-trivial behavior in managing ratings and computing a final reputation score.

We know that a classic aggregation operator has to satisfy the *identity property*, the *boundary conditions* and has to be *monotonic* in each aggregation function [4].

Definition 3.1. A mapping

$$A_n : [a, b]^n \to [a, b], \quad n \in \mathbb{N}$$

is called an n-ary *aggregation function* (AF) acting on $[a, b]$ if it is non-decreasing monotone in its components, that is

$$A_n(x_1, ..., x_n) \le A_n(y_1, ..., y_n)$$

whenever $a \le x_i \le y_i \le b$ for all $i \in \{1, ..., n\}$.

An aggregation operator can be introduced by means of a family $\mathbf{A} = \{A_n\}_n$ of n-ary aggregation functions, where, in general, A_n and A_m for different n and m need not be related.

Definition 3.2. A sequence $\mathbf{A} = \{A_n\}_n$ of n-ary AFs acting on $[a, b]$ is called an *aggregation operator* (AO) on $[a, b]$ if the following conditions hold:

$$A_1(x) = x \ , \tag{1}$$

$$A_n(a, \dots, a) = a \ ,$$
$$A_n(b, \dots, b) = b \ , \tag{2}$$
$$\text{for all } n \in \mathbb{N}.$$

When numerical values have to be aggregated, it can be useful that other two additional properties are satisfied.

1. *Unanimity*: it occurs when, if all the partial scores are equal to a certain value, also the global one is equal to that value. It is represented by the *idempotency* of the aggregation operator.

   Definition 3.3. An AO $\mathbf{A} = \{A_n\}_n$ defined on $[a, b]$ is *idempotent* if, for all $x \in [a, b]$

   $$A_n(x, ..., x) = x \ .$$

2. *Anonymity*: it occurs when the knowledge of the order of the input values is irrelevant. It is represented by the *commutativity* of the aggregation operator.

   Definition 3.4. An AO $\mathbf{A} = \{A_n\}_n$ defined on $[a, b]$ is *commutative* if, for all $n \in \mathbb{N}$, for all $x_1, \dots, x_n \in [a, b]$

   $$A_n(x_1, \dots, x_n) = A_n(x_{\alpha(1)}, \dots, x_{\alpha(n)}) \ ,$$

   for all permutations $\alpha = (\alpha(1), \dots, \alpha(n))$ of $\{1, \dots, n\}$.

In the reputation computation field, the first property is fundamental: if all the agents regard another agent/resource as trustworthy, the global score assigned to that agent/resource has to reflect their positive degree of trustworthiness. The second property assume a different importance depending on the kind of environment we are dealing with.

- *Anonymous systems*: No identity is assigned to agents in the system so the order of evaluation is irrelevant.

- *Non-anonymous systems*: The relevance of the rating may depends on the sources.

Distributed systems, where anonymity is a major feature, need both the above properties. In order to correctly choose a suitable aggregation operator for aggregating ratings in distributed reputation systems, we make some consideration reasoning about data.

We suppose that input values are numerical votes re-scaled on the interval $[-1, 1]$, in order to take into consideration the concept of *distrust* [11]. The more the vote is near to 1, the more the resource on which the vote is expressed is "positive" (trusted) (the same symmetrically for -1, untrusted). A vote equal to 0 can represent a never expressed rating or insufficient information to express a judgement.

We also suppose that for two different resources, the following votes are collected:

RESOURCE 1: n votes equal to 0, m votes equal to 1,
RESOURCE 2: p votes equal to 0, q votes equal to 1.

It is possible that the number of the inputs is very large, and that the number of votes is different for the two different resources:

$$\Rightarrow n + m \neq p + q \ ,$$
$$\Rightarrow n + m, p + q \gg 1 \ .$$

A classical idempotent aggregation operator $\mathbf{A} = \{A_n\}_n$, not considering the n and p votes equal to 0 for the first and the second resource respectively, will give:

RESOURCE 1: $A_m(1, ..., 1) = 1$,
RESOURCE 2: $A_q(1, ..., 1) = 1$.

In general we have $m \neq q$. Supposing $m \gg q$, using an aggregation operator idempotent in a "classic way", the final judgement concerning two blocks of votes having the same values, expressing for example the maximum positive evaluation, is the same, despite the higher cardinality of the first block in respect to the second. It is now clear that

1. the AO we need has to take into consideration all data inputs, without a preselection of valid votes. The presence of non-significant votes does not have to influence the output value;

2. the AO has to discriminate the weight of a consistent block of votes of the same value with respect to a less consistent one with the same values.

The first request is satisfied by the introduction of a neutral element.

Definition 3.5. Let $\mathbf{G} = \{G_n\}_n$ be a commutative AO on $[a, b]$. Then an element $e \in [a, b]$ is called a *neutral element* (NE) for $\mathbf{G}$ if

$$G_{n+1}(x_1, ..., x_n, e) = G_n(x_1, ..., x_n)$$

for all $x_1, ..., x_n \in [a, b]$.

In our case we set $e = 0$. From a theoretical point of view, we distinguish two classes of commutative AOs with neutral element:

$\mathcal{CA}[a, b]$, where e is a or b ,
$\mathcal{CB}[a, b]$, where e is a value $\in \ ]a, b[$.

Without loss of generality, and disregarding isomorphisms, the first class is mapped on $[0, 1]$, with $e = 0$, while the second on $[-1, 1]$ and again $e = 0$.

The second request is satisfied by the introduction of the concept of asymptotical idempotency.

4 Asymptotically Idempotent Aggregation Operators

Definition 4.1. An aggregation operator $\mathbf{G} = \{G_n\}_n$ on $[a, b]$ is called *asymptotically idempotent* (AI) if

$$\lim_{n \to \infty} G_n(x, ..., x) = x \quad \text{for all } x \in [a, b] \ . \quad (3)$$

This way we obtain the *unanimity* in a *weak form*, that is the final value tends to unanimity asymptotically, when the number of votes having the same values grows over the time.

Unfortunately, there exist AI AOs which do not meet (1) and (2); on the other hand, there exist AOs, under the classical definition, which do not satisfy (3). All these cases are detailed in [10].

A reasonable way to bypass this "cul-de-sac" is to weaken the definition of AO, omitting conditions (1) and (2).

Definition 4.2. ("Extended" AO) An aggregation operator (AO) on $[a, b]$ is given by any sequence $\mathbf{G} = \{G_n\}_n$ of aggregation functions on $[a, b]$.

This way, the focal point characterizing any aggregation operator is the monotonicity each single aggregation function preserves.

4.1 Building an AI AO

We now illustrate a general procedure for building a *symmetrical* AI AO (with respect to the neutral element) $\mathbf{G} = \{G_n\} \in \mathcal{CB}$, that is such that

$$G_n(x_1, ..., x_n) = -G_n(|x_1|, ..., |x_n|)$$

for all $x_1, ..., x_n \in [-1, 0]$.

4.1.1 First Step on $[0, 1]$

We explicitly build $\mathbf{G} = \{G_n\} \in \mathcal{CA}$. We define

$$G_n(x_1, ..., x_n) := \max_{i=1,...,n} \{x_i\} \cdot \psi(h_n(x_i)) \quad (4)$$

where (i) $\psi : [0, \infty] \to [0, 1]$ is a non-decreasing function such that

$$\psi(0) = 0, \quad \psi(\infty) = 1 ;$$

and (ii) $\mathbf{H} = \{h_n\}$ is a symmetrical AO on $[0, \infty]$, admitting 0 as neutral element, such that

$$h_1(0) = 0, \quad \sup_{n \to \infty} h_n(x, ..., x) = \infty .$$

(i) The class of the applications ψ is very rich.

(ii) It is possible to explicitly build $\mathbf{H} = \{h_n\}$ as follows:

Definition 4.3. Let $\mu : [0, \infty[\to [0, \infty[$ be a not decreasing function such that $\mu(0) = 0$ and $\mu(t) > 0$ for $t > 0$

$$h_n(x_1, ..., x_n) = \sum_{i=1}^{n} \mu(x_i) .$$

$\mathbf{H} = \{h_n\}$ so defined satisfies all the requested properties.

Example 4.1. Let $\psi(t) = \frac{\sqrt{t}}{1+\sqrt{t}}$ and $\mu(t) = t^2$. We have

$$G_n(x_1, ..., x_n) = \max_{i=1,...,n} \{x_i\} \cdot \frac{\sqrt{\sum_{i=1}^{n} x_i^2}}{1 + \sqrt{\sum_{i=1}^{n} x_i^2}} ,$$

for all $x_1, ..., x_n \in [0, 1]$.

4.1.2 Second Step on $[-1, 0]^n$

We extend every G_n on $[-1, 0]^n$, necessarily

$$G_n(x_1, ..., x_n) := -G_n(|x_1|, ..., |x_n|)$$

because of the symmetry with respect to the neutral element, for all $x_1, ..., x_n \in [-1, 0]$.

This way we have fixed the aggregation operator behavior in conjunction with a block of votes all positive or all negative, in a perfectly symmetric way.

4.1.3 Last Step on $[-1, 1]^n$

We extend every G_n on $[-1, 1]^n \setminus I_n$, where $I_n := [0, 1]^n \cup [-1, 0]^n$.

Now we have to determine the aggregation behavior in the case of a block of votes either positive or negative. We denote $(x_1^*, ..., x_n^*)$ any permutation of an arbitrary data n-ple $x_1, ..., x_n \in [-1, 1]^n \setminus I_n$ such that

$$x_1^*, ..., x_k^* \geq 0 \quad \text{and} \quad x_{k+1}^*, ..., x_n^* < 0,$$

for some $k \in \{1, ..., n-1\}$.

Definition 4.4. Let $f : [-1, 1] \to [-1, 1]$ be a strictly increasing bijection such that $f(0) = 0$. We set

$$g_n(x_1, ..., x_n) := f(G_n(x_1, ..., x_n))$$

for all $x_1, ..., x_n \in I_n$. We obtain

$$G_n(x_1^*, ..., x_n^*) := f^{-1}(g_k(x_1^*, ..., x_k^*) + g_{n-k}(x_{k+1}^*, ..., x_n^*)) . \quad (5)$$

$\mathbf{G} = \{G_n\} \in \mathcal{CB}$ and it does not present obvious compensation effects.

4.2 A Worked-out Example

Using the the method described above, we illustrate in this Section two different approaches in assigning relevance to the data values, and showing the lack of obvious compensation effects[1]. Choosing $\psi(t) = \frac{\sqrt{t}}{1+\sqrt{t}}$ and

---

[1]For the sake of simplicity henceforth in the examples we will avoid the notation x_i^*, simply writing x_i.

$\mu(t) = t^2$, from (4) and (5) we have

$$G_n(x_1, ..., x_n) =$$

$$f^{-1}\left(f\left(\max_{i=1,...,k}\{x_i\} \cdot \frac{\sqrt{\sum_{i=1}^{k} x_i^2}}{1 + \sqrt{\sum_{i=1}^{k} x_i^2}} \right) + \right.$$

$$\left. f\left(\min_{j=k+1,...,n}\{x_j\} \cdot \frac{\sqrt{\sum_{j=k+1}^{n} x_j^2}}{1 + \sqrt{\sum_{j=k+1}^{n} x_j^2}} \right) \right) .$$

$$(6)$$

for all $x_1, \ldots, x_n \in [-1, 1]^n \setminus I_n$.

Let us consider a function of the kind

$$f(t) = \begin{cases} t^\alpha & t \geq 0 \\ t^\beta & t < 0 \end{cases}$$

where, for $\alpha < \beta$, $\alpha = 2n$, $\beta = 2n + 1$, and for $\alpha > \beta$, $\alpha = 2n$, $\beta = 2n - 1$, for $n \in \mathbb{N}$.

In the sequel we will show how it is possible to obtain a "confident" or a "diffident" behavior of the aggregation operator by choosing the "right" values for α and β. In a confident approach, in order to enhance positive values, we set $\alpha < \beta$; conversely, in a diffident one, we set $\alpha > \beta$.

From (6) we obtain

$$G_n(x_1, ..., x_n) = \begin{cases} \sqrt[\alpha]{(\cdots)^\alpha + (\cdots)^\beta} & (i) \\ \sqrt[\beta]{(\cdots)^\alpha + (\cdots)^\beta} & (ii) \end{cases}$$

$$(7)$$

(i) if $(\cdots)^\alpha + (\cdots)^\beta \geq 0$,

(ii) if $(\cdots)^\alpha + (\cdots)^\beta < 0$.

In the following examples, in order to show in the most immediate way the behavior or the operator in the two approaches, we take as input two groups of votes with the same cardinality, with the same values, but with opposite sign.

4.2.1 The "confident" approach

In this example of confident approach, we choose $n = 1$ obtaining

$$f(t) = \begin{cases} t^2 & t \geq 0 \\ t^3 & t < 0 \end{cases}$$

Supposing

$$n = 6,$$
$$\{x_1, ..., x_k\} = \{x_1, x_2, x_3\} = \{0.8, 0.5, 0.3\}$$
$$\{x_{k+1}, ..., x_n\} = \{x_4, x_5, x_6\} =$$
$$\{-0.8, -0.5, -0.3\}$$

from (6) and (7) we obtain:

$$G_6(0.8, 0.5, 0.3, -0.8, -0.5, -0.3) = 0.309 .$$

4.2.2 The "diffident" approach

In this example of diffident approach we choose $n = 2$. The resulting function f is therefore defined as follows

$$f(t) = \begin{cases} t^4 & t \geq 0 \\ t^3 & t < 0 \end{cases} .$$

With the same group of data used for the confident approach, from (6) and (7) in the "diffident" one we will obtain

$$G_6(0.8, 0.5, 0.3, -0.8, -0.5, -0.3) = -0.336 .$$

4.2.3 Behavior of the AI AO

Using both the "confident" and the "diffident" approaches, we will summarize in the following different aspects connected to the use of an asymptotically idempotent aggregation operator provided by negative and neutral values in the computation of a final reputation score for a resource. We suppose that different users $\mathbf{U1}, ..., \mathbf{U8}$ vote on a resource assigning seven possible ratings, three values of *trustworthiness* $(0.8, 0.5, 0.2)$, the *neutral element* (0), and three values of *untrustworthiness* $(-0.2, -0.5, -0.8)$.

Depending on the chosen aggregation approach and the number and the values of ratings assigned to a certain resource, in Table 1 we show different cases where emerge interesting behaviors of the aggregation operator.

(a) The final aggregated value is calculated using the "diffident" approach.

(b) The final aggregated value is calculated using the "confident" approach.

Table 1: Examples of an AI AO behavior.

	U1	U2	U3	U4	U5	U6	U7	U8	T
(a)	0.5	0.5	0.5	0.5	-0.5	-0.5	-0.5	-0.5	-0.23
(b)	0.5	0.5	0.5	0.5	-0.5	-0.5	-0.5	-0.5	0.22
(c)	0.8	0.8	0	0	0	0	0	0	0.42
	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.55
(d)	-0.8	-0.8	0	0	0	0	0	0	-0.42
	-0.8	-0.8	-0.8	-0.8	-0.8	-0.8	-0.8	-0.8	-0.55

(c) Positive values: the aggregation function is sensitive to the number of the inputs, the final value asymptotically tend to their positive values.

(d) Negative values: the aggregation function is sensitive to the number of the inputs, the final value asymptotically tend to their negative values.

5 Conclusions and Further Work

Assessing the trustworthiness of remote entities offering services or data on the global net is fundamental in the development of multi-agent systems, especially distributed like P2P systems. In this paper we have described a numerical technique conceived for taking into consideration the properties of the environments in order to provide a meaningful measure of trustworthiness. Furthermore the growing interest for explicit context representation in large scale P2P and network systems encourages developing a theory of parametric, context-sensitive aggregation operators, semantically handling diverse contexts. We believe asymptotically idempotency will play a crucial role in this research.

References

[1] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *CIKM '01: Proceedings of the tenth international conference on Information and knowledge management*, pages 310–317, New York, NY, USA, 2001. ACM Press.

[2] K. Aberer and Z. Despotovic. Maximum likelihood estimation of peers' performance in p2p networks. In *The Second Workshop on the Economics of Peer-to-Peer Systems*, 2004.

[3] R. Aringhieri. Assessing efficiency of trust management in peer-to-peer systems. In *WETICE '05: Proceedings of the 14th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprise*, pages 368–374, Washington, DC, USA, 2005. IEEE Computer Society.

[4] T. Calvo, A. Kolesárová, M. Komorníková, and R. Mesiar. Aggregation operators: properties, classes and construction methods. In *Aggregation operators: new trends and applications*, pages 3–104. Physica-Verlag GmbH, Heidelberg, Germany, 2002.

[5] F. Cornelli, E. Damiani, S. De Capitani di Vimercati, S. Paraboschi, and P. Samarati. Choosing reputable servents in a p2p network. In *WWW '02: Proceedings of the 11th international conference on World Wide Web*, pages 376–386, New York, NY, USA, 2002. ACM Press.

[6] E. Damiani, P. Ceravolo, and M. Viviani. Bottom-up extraction and trust-based refinement of ontology metadata. *IEEE Transactions on Knowledge and Data Engineering*, 19(2):149–163, 2007.

[7] E. Damiani, S. De Capitani di Vimercati, S. Paraboschi, and P. Samarati. Managing and sharing servents' reputations in p2p systems. *IEEE Trans. Knowl. Data Eng.*, 15(4):840–854, 2003.

[8] J. Ferber. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.

[9] R. Ghiselli Ricci. Finitely and Absolutely Non Idempotent Aggregation Operators. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 12(2):201–218, 2004.

[10] R. Ghiselli Ricci. Asymptotically Idempotency. In *FUZZ-IEEE: Proceedings of the 2007 IEEE International Conference on Fuzzy Systems, July 24-26, 2007, London, England*, 2007.

[11] T. Grandison and M. Sloman. A Survey of Trust in Internet Applications. *IEEE Communications Surveys and Tutorials*, 3(4), September 2000.

[12] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decis. Support Syst.*, 43(2):618–644, 2007.

[13] 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.

[14] D. H. McKnight and N. L. Chervany. The meanings of trust. Technical Report WP 96-04, University of Minnesota, Carlson School of Management, 1996.

[15] L. Mui, M. Mohtashemi, and A. Halberstadt. Notions of reputation in multi-agents systems: a review. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 280–287, New York, NY, USA, 2002. ACM Press.

[16] M. Rabin. Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5):1281–1302, 1993.

[17] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation Systems. *Commun. ACM*, 43(12):45–48, December 2000.

[18] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. In M. R. Baye, editor, *The Economics of the Internet and E-Commerce*, volume 11 of *Advances in Applied Microeconomics*, pages 127–157, Amsterdam, The Netherlands, 2002. Elsevier Science.

[19] Y. Wang and J. Vassileva. Bayesian network-based trust model. In *WI '03: Proceedings of the IEEE/WIC International Conference on Web Intelligence*, page 372, Washington, DC, USA, 2003. IEEE Computer Society.

[20] R. R. Yager, J. Kacprzyk, and M. Fedrizzi, editors. *Advances in the Dempster-Shafer theory of evidence*. John Wiley & Sons, Inc., New York, NY, USA, 1994.

[21] B. Yu and M. P. Singh. An evidential model of distributed reputation management. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 294–301, New York, NY, USA, 2002. ACM Press.