2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology

Intelligent Agents in Mobile Vehicular Ad-Hoc Networks: Leveraging Trust Modeling Based on Direct Experience with Incentives for Honesty

Umar Farooq Minhas Jie Zhang† Thomas Tran‡ Robin Cohen
David R. Cheriton School of Computer Science, University of Waterloo, Canada
† School of Computer Engineering, Nanyang Technological University, Singapore
‡School of Information Technology and Engineering, University of Ottawa, Canada

***Abstract*—In this paper we introduce a multi-faceted trust model of use for the application of ad-hoc vehicular networks (VANETs) – scenarios where agents representing drivers exchange information with other drivers regarding road and traffic conditions. We argue that there is a need to model trust in various dimensions and that combining these elements effectively can assist agents in making transportation decisions. We then introduce two new elements to our proposed model: i) distinguishing direct and indirect reports that are shared ii) employing a penalty for misleading reports, to promote honesty. We demonstrate how these two elements together serve to increase the value of the trust model, through a series of experiments of simulated traffic. In brief, we present a framework to facilitate the effective sharing of information in VANET environments between agents representing the vehicles.**

I. Introduction

An application area of increasing interest to the multiagent systems community is that of traffic and transportation management [1]. In this context, intelligent agents represent the drivers of the vehicles on a road. Agents enter into communication with other agents in order to obtain timely information of use in proposing actions for the drivers to take, thus providing decision-making support.

Mobile ad-hoc vehicular networks (VANETs) arise when the agents continuously determine, in real-time, which other agents to contact and which advice to consider and then initiate communication with the other vehicles in the environment [2]. In order for effective decisions to be made on the basis of the information that has been received, it then becomes important for the agents to be modeling the trustworthiness of the other agents in the environment: the challenge of malicious reports has been examined by several researchers, acknowledging the presence of self-interested agents taking part in the communication [2].

The topic of modeling trust in multiagent systems has been examined by a variety of researchers, promoting the use of learning (reinforcement learning in [3] or probabilistic reasoning in [4], [5], [6]) in order to effectively model agent trustworthiness, in advance of communication with fellow agents. When connections between the agents have been somewhat sparse, various researchers have then promoted the use of social networks of advisors (other agents that may have had experience with the agents whose trustworthiness needs to be assessed). With the introduction of advisors, however, it becomes important to also be modeling just how trustworthy these advisors are, as well [4], [5], [6].

In this paper we introduce a multi-faceted trust model of particular use in settings such as VANETs, where the environment is changing dynamically, there may be special roles for agents to play, there may be a well-established social network, connections between agents may be infrequent, over time and the advice that is offered may be more or less valued depending on its time and location. We then discuss in more detail the importance of distinguishing whether the report that is received is on the basis of direct or indirect evidence, in order to determine its potential value.

We also explore the importance of encouraging honest reporting in the VANET environment, through the introduction of a penalty system. With this mechanism in place, we demonstrate how agents are able to reason more effectively with the information that is received from other agents, in order to make decisions about transportation.

In particular, we offer experimental results in a simulation where vehicles are travelling through roads with various traffic conditions. We are able to demonstrate the value of our proposed framework in comparison with models that offer less complete trust modeling or that fail to invoke a penalty mechanism. As such, we argue that our proposed framework is of value in the design of VANETs that operate on the basis of intelligent agent communication and trust modeling.

II. The Core Model

Our proposed model for reasoning about the trustworthiness of advice provided by other agents in VANETs builds upon a core model that promotes a multi-faceted approach to trust modeling.

We first acknowledge that certain vehicles in the environment may play a particular **role** and, on this basis, merit greater estimates of trustworthiness. For example, there may be vehicles representing the police and other traffic authorities or ones representing radio stations dedicated to determining accurate traffic reports by maintaining vehicles in the vicinity of the central routes. Our proposal for considering roles is

978-0-7695-4191-4/10 $26.00 © 2010 IEEE
DOI 10.1109/WI-IAT.2010.66

to group agents together according to their designated role, but then to still be able to order each collection of role-based agents from the most to the least trustworthy, on the basis of past experience with the agents.

Consideration of any past personal **experiences** with agents allows the model to include any learning about particular agents due to previous encounters, specifically modeling trustworthiness each time and adjusting the level of trust to be higher or lower, based on the outcome of the advice that is offered. Experience-based trust is particularly useful when there is a set of agents with common experience, for example a group of agents that are all regular commuters on a particular route.

We propose that the agent requesting information first of all sets a fixed number of agents, n, whose advice will be taken into consideration in order to reach a decision about the action to take, based on the road conditions that are reported by these agents. We are primarily considering scenarios where one agent may ask questions such as "Is there any significant problem with the traffic on Road X?" to which agents may provide "Yes" or "No" replies. From the set of n agents that are consulted, the algorithm attempts to determine whether there is a clear **majority opinion** response as shown in the following computations:

Step 1: Each agent maintains, in its internal database, an ordered list of agents to ask for advice. The list will be partitioned into groups as follows:

$$\begin{bmatrix} G_1: & a_{11}, & a_{12}, & a_{13}, & ..., & a_{1k} \\ G_2: & a_{21}, & a_{22}, & a_{23}, & ..., & a_{2k} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ G_j: & a_{j1}, & a_{j2}, & a_{j3}, & ..., & a_{jk} \end{bmatrix}$$

This priority list is ordered from higher roles to lower roles, for example, G_1 being the highest role. Within each group of agents of similar roles, the group is ordered from higher personal experience-based trust ratings to lower ratings. Thus, a_{ij} represents the agent in role class i that is at the j^{th} level of experience-based trust, relative to other agents at that level[1]. Hence, role-based trust and experience-based trust are combined into this priority-based approach. These two trust metrics will be further discussed later in this section.

Step 2: Depending on the task at hand, set a value n = number of agents whose advice will be considered. This incorporates task-based trust. For example, if you need a very quick reply, you may limit n to be relatively small, say $n \leq 10$; if you are planning ahead and have time to process responses, n could potentially be larger.

Step 3: When an agent requires advice, the procedure is to ask the first n agents from its ordered list of agents the question, receive the responses and then perform some majority-based trust measurement.

Step 3A: The processing of the responses is as follows: if there is a majority consensus on the response, up to some tolerance that is set by the asker (e.g. I want at most 30% of the responders to disagree), then this response is taken as the advice and is followed. We will formalize this majority-based trust in Section II-A below.

Step 3B: Once this advice is followed, the agent evaluates whether this advice was reliable and if so, personal experience-based trust values of those agents are increased; if not, personal experience-based trust values of those agents are decreased. Detailed formalization of this process will be given in Section II-A.

Step 3C: If a majority consensus cannot be reached, then requiring majority consensus for advice is abandoned. Instead, the agent relies on role-based trust and experience-based trust (e.g., taking the advice from the agent with highest role and highest experience trust value)[2].

Step 4: In order to eventually admit new agents into consideration, when advice is sought, the agent will ask a certain number of agents beyond agent a_n in the list. The responses here will not count towards the final decision, but will be scrutinized in order to update personal experience-based trust values and some of these agents may make it into the top n agents, in this way. As we will see below, this is particularly useful to incentivize agents to be honest, if they have previously had their trustworthiness reduced or have been relegated to a position far down in a priority list.

A. Detailed Calculations

Experience and majority opinion, in turn, break down into more specific calculations. If we define the range of all personal experience-based trust values to be the interval $(-1, 1)$, where 1 represents absolute trust, -1 represents absolute distrust and 0 represents a neutral trust value (initially given to a new agent), then we can use the following scheme to update an agent's personal experience trust value, as suggested by [3]:

Let $T_A(B) \in (-1, 1)$ be the trust value indicating the extent to which agent A trusts (or distrusts) agent B according to A's personal experience in interacting with B. After A follows an advice of B, if the advice is evaluated as reliable, then the trust value $T_A(B)$ is increased by

$$T_A(B) \leftarrow \begin{cases} T_A(B) + \alpha(1 - T_A(B)) & \text{if } T_A(B) \geq 0 \\ T_A(B) + \alpha(1 + T_A(B)) & \text{if } T_A(B) < 0 \end{cases} \quad (1)$$

where $0 < \alpha < 1$ is a positive increment factor.

Otherwise, if B's advice is evaluated as unreliable, then $T_A(B)$ is decreased by

$$T_A(B) \leftarrow \begin{cases} T_A(B) + \beta(1 - T_A(B)) & \text{if } T_A(B) \geq 0 \\ T_A(B) + \beta(1 + T_A(B)) & \text{if } T_A(B) < 0 \end{cases} \quad (2)$$

where $-1 < \beta < 0$ is a negative decrement factor.

The absolute values of α and β are dependent on several factors because of the dynamics of the environment, such as

[1]There is no need for each group to have the same number of elements. We provide here only a simplified example.

[2]Note that an additional motive for modeling the trustworthiness of a variety of agents is to be able to learn about these agents for future interactions, for example in the calculations of experience-based trust and majority-opinion trust.

the data sparsity situation and the event/task specific property. For example, when interaction data is sparse, these values should be set to be larger, giving more weights to the available data. For life-critical events (i.e. collision avoidance), $|\alpha|$ and $|\beta|$ should be larger, in order to increase or decrease trust values of reporting agents more rapidly. Also note that we may set $|\beta| > |\alpha|$ by having $|\beta| = \mu|\alpha|$ and $\mu > 1$ to implement the common assumption that trust should be difficult to build up, but easy to tear down.

As for the majority opinion calculation, suppose agent A in VANET receives a set of m reports $\mathcal{R} = \{R_1, R_2, ..., R_m\}$ from a set of n other agents $\mathcal{B} = \{B_1, B_2, ..., B_n\}$ regarding an event. Agent A will consider more heavily the reports sent by agents that have higher level roles and larger experience-based trust values. When performing majority-based process, we also take into account the location closeness between the reporting agent and the reported event, and the closeness between the time when the event has taken place and that of receiving the report. We define C_t (time closeness), C_l (location closeness), T_e (experience-based trust) and T_r (role-based trust). Note that all these parameters belong to the interval $(0, 1)$ except that T_e needs to be scaled to fit within this interval.

For each agent B_i $(1 \leq i \leq n)$ belonging to a subset of agents $\mathcal{B}(R_j) \subseteq \mathcal{B}$ that report the same report $R_j \in \mathcal{R}$ $(1 \leq j \leq m)$, we aggregate the effect of its report according to the above factors. The aggregated effect $E(R_j)$ from reports sent by agents in $\mathcal{B}(R_j)$ can be formulated as follows:

$$E(R_j) = \sum_{B_i \in \mathcal{B}(R_j)} \frac{T_e(B_i)T_r(B_i)}{C_t(R_j)C_l(B_i)} \tag{3}$$

In this equation, experience-based trust and role-based trust are discounted based on the two factors of time closeness and location closeness. The summation is used to provide the aggregated effect of the reporting of the agents.

Note that location closeness $C_l(B_i)$ depends only on the location of agent B_i while time closeness $C_t(R_j)$ depends on the time of receiving the report R_j. $C_t(R_j)$ can also be written as $C_t(B_i)$ because we can assume that each report is sent by an unique agent in possibly different time.

To consider the effect of all the different reports, the majority opinion is then

$$M(R_j) = \max_{R_j \in \mathcal{R}} E(R_j) \tag{4}$$

which implies the report that has the maximum effect, among all reports.

A majority consensus can be reached if

$$\frac{M(R_j)}{\sum_{R_j \in \mathcal{R}} E(R_j)} \geq 1 - \varepsilon \tag{5}$$

where $\varepsilon \in (0, 1)$ is set by agent A to represent the maximum error rate that A can accept. A majority consensus can be reached if the percentage of the majority opinion (the maximum effect among different reports) over all possible opinions is above the threshold set by agent A.

If the majority consensus is reached, the majority opinion is associated with a confidence measure. This measure takes into account the number of interactions taken for modeling experience-based trust values of reporting agents and the maximum accepted error rate ε. We define $N(R_j)$ as the average of the discounted number of interactions used to estimate experience-based trust values of the agents sending the majority report R_j. Based on the Chernoff Bound theorem [7], the confidence of the majority opinion can be calculated as:

$$\gamma(R_j) = 1 - 2e^{-2N(R_j)\varepsilon^2} \tag{6}$$

III. Expanding the Core Model

A. Distinguishing Direct Experience

The first new element that we introduce is to distinguish direct and indirect experience, when information is provided by an agent to another agent. In this context, we require each agent to self declare whether its information has been derived from firsthand experience or not. We initially assume that this declaration is truthful, and determine which action to take through a weighting of the advice that has been provided. If an agent is not a direct witness but claims to be one, then it will run the risk of having its trust value reduced more severely (as detailed below) when its advice is verified to be unreliable. The main idea is that an agent that asks for information from other agents will value advice from the direct witnesses more than that from the indirect ones. The algorithm and computation steps followed by each asking agent described in Section II are now modified as follows:

Step 1 and **Step 2** are the same as those in Section II.

Step 3: When an agent needs advice, it will ask the first n agents from its ordered list of agents. Suppose that the agent receives reports (i.e. responses) from these n agents and m of them declare that the information is from their direct experience. The processing of the reports is as follows:

Step 3A: The asking agent determines whether there are sufficient direct witnesses such that it can make a decision based solely on their reports. This can be done by comparing m with a minimum threshold N_{min} representing the minimum number of direct witnesses from which the agent has confidence to draw majority opinion from their reports. N_{min} can be calculated using a variation of Equation (6):

$$N_{min} = -\frac{1}{2\varepsilon^2} ln \frac{1-\gamma}{2} \tag{7}$$

where γ is the confidence level set by the asking agent and ε is the maximum error rate that the agent can accept, as explained in Equation (5).

Step 3B: If $m \geq N_{min}$, then the asking agent will only consider the reports from the direct witnesses and follow the majority opinion of these reports. The process to determine and formalize majority consensus is now similar to that described in Section II-A using Equations (3), (4), and (5).

Step 3C: If $m < N_{min}$, then there are insufficient direct witnesses. In this case, the asking agent will consider reports from both direct and indirect witnesses, assigning different

weight factors to them. This can be done by using a variation of Equation (3) to calculate the aggregated effect of a report R_j taking into account whether or not agent B_i that sent this report is a direct witness:

$$E(R_j) = \sum_{B_i \in \mathcal{B}(R_j)} \frac{T_e(B_i)T_r(B_i)}{C_t(R_j)C_l(B_i)W(B_i)} \quad (8)$$

where the meanings of parameters are the same as in Equation (3) except that the new weight factor $W(B_i)$ is set to 1 if agent B_i that sent report R_j is an indirect witness, and $W(B_i)$ is set to a value in $(0, 1)$ if agent B_i is a direct witness[3]. From this point, the asking agent will obtain and follow the majority opinion with the procedure for majority consensus being similar to that in Section II-A, i.e. using Equations (4), and (5).

Step 3D: Once the actual road conditions are verified, the asking agent adjusts the trust values of the reporting agents based on personal experience. It penalizes more strongly those agents that reported incorrect information in the direct evidence case and less severely those agents with incorrect advice from indirect experience. Similarly, it rewards those agents that reported correct information in the direct evidence case more than in the indirect evidence one. This process of adjusting trust values is performed using Equations (1) and (2) in which the positive increment factor and the negative decrement factor (α and β) are replaced by α_D and β_D for the case of direct evidence, and are replaced by α_I and β_I for the case of indirect evidence, with $\alpha_D > \alpha_I$ and $\beta_D < \beta_I$ (to implement the above awarding and penalization policies).

Step 4: Is the same as that in Section II.

B. Incentives to Honesty

We now integrate an element designed to encourage truthful reporting from agents.

An agent B is considered dishonest or deceitful by an agent A if the personal experience trust value that A has on B falls below some value that A can accept. In other words, B is regarded as dishonest by A if $T_A(B) < \theta_A$ where $-1 < \theta_A < 0$. The value θ_A is agent A's specific constant[4]. Each agent that seeks advice from other agents maintains a set of dishonest agents to whom it will not respond when asked, as a penalty to these dishonest agents.

We envision that with this penalty system in place dishonest agents are penalized and consequently honesty is promoted. Consider a simple scenario where there are 10 agents, $A_1, ..., A_{10}$, with agent A_{10} being the dishonest one. As mentioned in Section II-A, initially agents $A_1, ..., A_9$ give one another and agent A_{10} a neutral personal experience trust value of zero. However, after several rounds of interactions with agent A_{10}, their experience trust values on A_{10} are more and more decreased due to the untruthful reports of this agent, up to a point where all the trust values fall below the θ_{A_i} values ($1 \leq i \leq 9$). At that time, agents $A_1, ..., A_9$ stop responding to A_{10}, resulting in A_{10} not being able to benefit from the network at all.

The effect of the proposed penalty mechanism is further demonstrated via our experimental results.

[3]For example, setting $W(B_i) = 1/2$ for the case of direct witnesses indicates that the asking agent values direct evidence twice more than indirect evidence.

[4]If A sets θ_A too low, dishonest or deceitful agents will not be identified as they should be. However, if A sets θ_A too high, some novice agents may be mistakenly considered as deceitful agents. Considering these, we suggest that θ_A take values in $[-0.9, -0.7]$.

IV. Experimental Results

In order to examine the value of the core model and its extended version, we conducted experiments that simulated a traffic environment in a city, using SWANS (Scalable Wireless Ad-hoc Network Simulator, jist.ece.cornell.edu) with STRAW (STreet RAndom Waypoint) mobility model [8]. SWANS is entirely implemented in Java and can simulate networks with potentially thousands of nodes while using incredibly small amount of memory and processing power. STRAW allows to simulate real world traffic by using real maps with vehicular nodes that follow rules such as speed limits, traffic signals, stop signs etc. For our experiments we fix the total number of vehicles to 100 and run the simulation for a total duration of 900 seconds of simulation framework time.

We present experimental results to clearly show the improved performance by distinguishing direct experience, the value of having incentives for truthfully revealing directness of evidence, and social implications of lying on an agent.

1) Improved Performance by Distinguishing Direct Experience: One of the applications of V2V communication is to be able to route traffic effectively through the VANET and to avoid congestion or hot spots. Malicious agents in the network may send untruthful traffic information, to mislead other agents and cause traffic congestion. We measure the performance of our proposed trust model by observing to what extent it can cope with deceptive information sent by malicious agents. According to [8], we can measure congestion based on the average speed of vehicles. Lower average speed implies more traffic congestion. The performance of our model can then be measured as the increase in average speed of all agents by incorporating our model under the environment where malicious agents exist.

In this experiment we combine role-based and experience-based trust and measure average speed. As we can see in Figure 1(a), combining the two dimensions produces greater speed than using any one dimension alone (and that combining facets of trust is valuable). We can also see that distinguishing direct and indirect experience further improves the overall performance of our trust model. It is demonstrated to be advantageous to distinguish direct and indirect experience shared by other agents.

2) Incentive for Truthfully Revealing the Directness of Evidence: This experiment demonstrates that in our model agents will be motivated to truthfully reveal direct or indirect nature of their evidence. Our model will more aggressively penalize an agent that presents an indirect evidence as direct. The initial trust value for all agents is set to 0.5. We then simulate reports

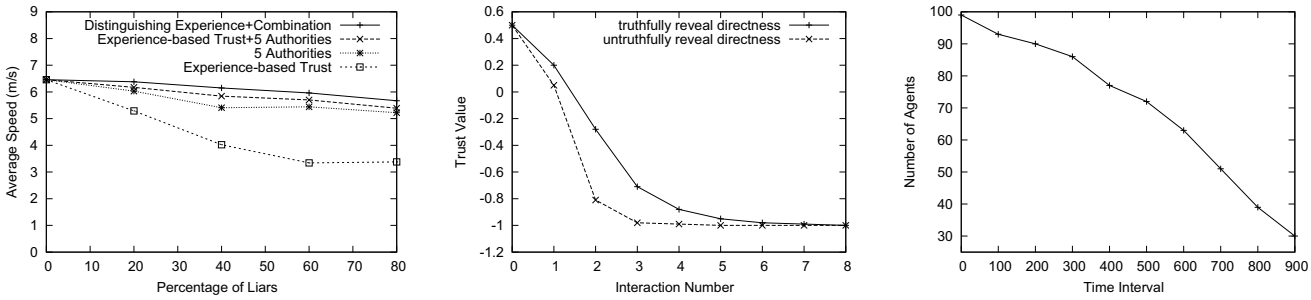

Fig. 1. (a) Average Speed of All Cars with Role-based and Experience-based Trust; (b) Incentive for Truthfully Revealing the Directness of Evidence; (c) Social Implications of Lying

from malicious (lying) agents and measure how the trust values of these agents decrease with every interaction when truthfully or untruthfully revealing the directness of evidence. We present the results in Figure 1(b). On the x-axis, we have the number of interactions between an observer agent B and a malicious agent A. There are a total of 8 interactions between the two agents during the entire simulation. On the y-axis, we report the trust value of the observer agent B in the malicious agent A. Agent B receives reports from agent A declared as either direct or indirect. Agent B then updates the trust value in agent A depending on whether A has truthfully or untruthfully reported the directness of reports.

Figure 1(b) presents two cases: (a) the agent truthfully reveals the directness of its reports, labeled as *truthfully reveal directness* and (b) the agent untruthfully reveals the directness of its reports, labeled as *untruthfully reveal directness*. As we can see, our model aggressively penalizes the malicious agent A in case (b) where agent B's trust value in agent A drops rapidly with every interaction. On the other hand, agent A is less aggressively penalized in case (a). This shows that in our model, it is in an agent's best interest to truthfully report the evidence as direct or indirect.

3) Social Implication: In this experiment, we quantify the social implications of lying in our model. We observe how a particular malicious agent A is singled out by other agents over the duration of the simulation. Figure 1(c) presents a graph with time interval from 0 to 900 seconds on the x-axis, and the number of agents who trust the malicious agent A on the y-axis. Initially, all the agents will trust agent A. But as the time passes, more agents interact with and detect the lying behavior of agent A. Thus the total number of agents who trust agent A drops gradually as shown in Figure 1(c). This shows that if the lying behavior of agent A continues, it will ultimately be distrusted by all agents and thus will neither be consulted nor provided reports by any other agent – practically making agent A a social outcast.

V. Discussion

In this paper, we presented a multi-faceted trust model and argued for its value when judging the trustworthiness of agents in VANET environments. We then introduced two important extensions to this model, leading to an adjusted algorithm for decision making by agents in VANET environments, on the basis of information received by other agents. Experimental evidence demonstrated the specific value of the extended model for evaluating trustworthiness and confirmed the effectiveness of the proposed penalty system, for promoting honesty.

For future work, it would be valuable to explore in greater detail the potential value of indirect reports from agents for future interactions and advice. In particular, agents may be asked to not only indicate whether their report is an indirect one but also to declare their internal trustworthiness of the agent from whom they have received their report. This would then allow the agent who is seeking advice to consider a subset of the indirect reports as more valuable, weighted more heavily when determining a decision. Research on the use of advisor networks such as that of [9] would be a useful staring point to explore the ongoing use of those providing advice indirectly.

References

[1] U. F. Minhas, J. Zhang, T. Tran, and R. Cohen, "Promoting effective exchanges between vehicular agents in traffic through transportation-oriented trust modeling," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) Workshop on Agents in Traffic and Transportation*, 2010.

[2] R. Lin, S. Kraus, and Y. Shavitt, "On the benefit of cheating by self-interested agents in vehicular networks," in *Proceedings of International Autonomous Agents and Multi Agent Systems (AAMAS)*, 2007.

[3] T. Tran, "A reliability modelling based strategy to avoid infinite harm from dishonest sellers in electronic marketplaces," *Journal of Business and Technology (JBT), Special Issue on Business Agents and the Semantic Web*, vol. 1, no. 1, pp. 69–76, 2005.

[4] A. Jøsang and R. Ismail, "The beta reputation system," in *Proceedings of the 15th Bled Electronic Commerce Conference*, 2002.

[5] W. Teacy, J. Patel, N. R. Jennings, and M. Luck, "Travos: Trust and reputation in the context of inaccurate information sources," *Auton Agent Multi-Agent Sys*, vol. 12, pp. 183–198, 2006.

[6] J. Zhang and R. Cohen, "Trusting advice from other buyers in e-marketplaces the problem of unfair ratings," in *Proceedings of the Eighth International Conference on Electronic Commerce*, 2006.

[7] L. Mui, M. Mohtashemi, and A. Halberstadt, "A computational model of trust and reputation," in *Proceedings of the 35th Hawaii International Conference on System Science (HICSS)*, 2002.

[8] D. R. Choffnes and F. E. Bustamante, "An integrated mobility and traffic model for vehicular wireless networks," in *Proceedings of VANET*, 2005.

[9] B. Yu and M. P. Singh, "A social mechanism of reputation management in electronic communities," in *Proceedings of the 4th International Workshop on Cooperative Information Agents*, 2000, pp. 154–165.