Using Trust and Possibilistic Reasoning to Deal with Untrustworthy Communication in VANETs

Andrew Koster¹

Andrea Tettamanzi²

Ana L. C. Bazzan¹

Célia da Costa Pereira²

Abstract—VANETs allow for unprecedented amounts of information to be sent between participants in traffic. Unfortunately, without countermeasures, they also allow selfish agents to take advantage of communication to improve their own utility. In this paper we present a novel framework for dealing with potentially untrustworthy information. The framework consists primarily of two components: a computational trust model for estimating the amount of uncertainty in received information and a possibilistic beliefs-desires-intentions agent system for reasoning about this uncertain information in order to achieve the driver's goals. We demonstrate the framework's effectiveness in an easy to understand but realistic scenario of a freeway system in which we also show that deceit may have a larger impact on traffic flow than previously thought.

I. INTRODUCTION

A recent article reports that, by 2015, 80% of the cars that the Ford Motor Company sells on the US Market will have wireless communication technology built in [1]. The same article cites other suppliers of car to car (C2C) communication technology; all of them expect a large growth of sales in the near future. This will enable the upcoming generation of cars to communicate directly with each other using an ad hoc network. Some initial applications of the technology are already being experimented with, such as the Cooperative Forward Collision Warning, which assists the driver in avoiding rear-end collisions with other vehicles [2]. Nevertheless, much of the potential for this technology is still unexplored. In particular, the use of communication in order to assist the driver for non-safety related tasks, such as signalling other drivers about upcoming congestion, is still very much an open issue.

One major problem in inter-vehicular communication is that selfish agents may be better served communicating false information than the truth. Agents may not just try to choose actions that optimize their utility, but actively attempt to deceive others in order to improve their own utility. Such deceitful agents are incompatible with approaches to cooperative driving, such as the one described by Bejan and Lawrence [3], where agents' truthfulness is a necessary prerequisite to obtain accurate information about the state of the road. Such truthfulness, however, cannot be expected in many traffic situations. For instance, if a deceitful driver knows that the highway he wants to take is congested, he is best suited convincing other drivers that it is clear, in order to avoid heavy traffic on the alternative road. Kraus et al. [4] analyzed

a traffic scenario in which deceitful agents can exploit C2C communication over a VANET, in order to optimize their own travel time. Their research indicates that the effect of deception on travel times is low, however the simulation is for a downtown environment with all cars moving in different directions. We expect that misinformation is more serious on major arteries, such as those used by commuter traffic, where the origin and destination of the drivers are more similar.

Kraus et al. also present some approaches to obtain truthful information in environments with deceiving agents. Firstly, they propose to take historical information about the road into account when evaluating received information. This is a way of evaluating the trustworthiness of received information, by comparing it with another information source. We consider historical data about the road as an alternative source of information, but using it to evaluate communicated information *a priori* defeats the purpose of communication in the first place: we require communication specifically when other sources of information are insufficient. In a similar way, they consider communications from trusted sources, such as ambulances or police cars. This is similar to considering a centralized source of information, such as a GPS path finding service, or the news broadcast by the radio. Such sources have no incentive to lie, but are often unavailable, or provide outdated information.

Instead of relying only on trusted channels, which may be expensive or unavailable, we propose a new framework for reasoning about communicated information. We emphasize that while the individual techniques used are not new, their combination and application to evaluating the truthfulness of communication in VANETs is. We use a possibilistic beliefs-desires-intentions (BDI) framework, as first presented by da Costa Pereira and Tettamanzi [5], to reason about the uncertainties stemming from the various sources of information. The use of an intelligent agent model has various advantages. The first is that it is essential for the automated transmission of messages: deciding what message to send and when to send it is a complex decision, especially as bandwidth in a VANET is limited. Second, an intelligent agent can help with the interaction with the driver. As more and more information is available to the user, it becomes increasingly important to reduce the information overload. Intelligent agents are eminently suited to this task [6]. Moreover, a BDI framework allows for reasoning about the communication and the source; thereby providing a context in which the trustworthiness of his message can be assessed.

The BDI framework we are using here has two important features: it allows to consider information from sources

¹A. Koster and A. Bazzan are with the Department of Informatics at the Federal University of Rio Grande do Sul, Porto Alegre, Brazil

²A. Tettamanzi and C. da Costa Pereira are with the Laboratoire I3S at the Université de Nice Sophia Antipolis, Sophia Antipolis, France

which can be partially trusted and updates the agent’s beliefs with respect to both new information and the associated trust degree and it proposes a reasoning model to generate the agent’s goals under the new situation. Because our approach is possibilistic, unlike Kraus’ approach, which necessarily needs historical data, ours can work even when the available data is incomplete and a qualitative ordering of the trustworthiness of the sources is the only thing that is available.

To deal with deception in C2C communication we use a computational trust model. Many computational trust models have been proposed for a variety of domains [7]. As Zhang [8] points out, however, C2C communication faces problems that are not addressed in such models. Most conventional trust models rely on repeat interactions with an individual agent to build up a trust relationship. In the absence of such repeated interactions, they turn to their peers, who may have repeated interactions, or a centralized authority that holds reputational information. Unfortunately, in communication in a VANET, none of these methods are available. The massive, decentralized nature of a VANET makes repeat interactions with a single car unlikely, and precludes the presence of a centralized source of reputation information. We present a trust model in Section II-A that considers the VANET as a whole as a single source, and also deals with aggregating information from other sources such as a GPS path finding service or institutional sources.

In Section III we discuss the experimental setup, in which we demonstrate the effect deceitful agents can have on traffic flow, and how our model improves the situation. We conclude this paper in Section IV.

II. OUR FRAMEWORK

Our method for dealing with uncertainty in C2C communication uses an autonomous agent framework to reason about all received information, and to use this intelligently to achieve the user’s goals; regardless of whether this is communication about congested roads, as we discuss in this paper, or communication about free parking spots or hazardous situations on the road. The main building blocks of this system are (1) a computational trust model that is able to assess the trustworthiness of various information sources, in particular C2C communication, and (2) a possibilistic BDI agent, to reason about uncertain information and make decisions. We describe both these systems in detail in this section.

A. Trusting communication

Cars equipped with modern communication technology will be able to receive information from a variety of sources. We consider C2C communication devices in particular, but other information sources are taken into account as well. For instance, GPS-based path planning services, such as Google Maps or TomTom already provide information about traffic conditions to their clients, and government authorities could use wireless communication with cars, in addition to the already present methods — such as digital information

boards on freeways or counters with free parking spaces in the inner city — for communicating to cars on the road.

We do not consider any of these sources to be inherently better than any other, but rather use a computational trust model to evaluate information from each of them. Additionally, we consider different types of messages separately, allowing the context to be taken into account when considering each source. For example, consider that government authorities may be very trustworthy when supplying information about traffic congestion, but their information about available parking spots is found to be out of date. In this case, we should treat information from the government authority about congestion differently from that about parking spaces: the trustworthiness of the source is dependent on the context of the message. After analyzing the trustworthiness of each source individually, the information from various sources is combined, using a consensus operator [9] and passed on to the possibilistic BDI agent, which reasons about how the communication affects its plans.

1) *Trust and context*: Before we discuss the trust model and consensus operator, we need to define the messages that are sent. We assume that the content of the message uses a shared communication language $\mathcal{L}_{Comm}$. An example of such a language, which might be used in a full-scale implementation, is the one specified by the Ontology of Transportation Systems [10]. In this work, however, we restrict the communication to the set of literals of a First-Order Language. We introduce this restriction for the sake of simplicity. Firstly, bandwidth and computational power is at a premium in VANETs: by restricting messages to literals, we restrict their size, as well as the computation required for processing them. Secondly, it allows us to define message contexts as predicate symbols. In a richer language, message contexts can also be defined (using, for instance, clustering with a distance measure over the language), but this is outside the scope of this paper. The context of a message allows us to generalize over multiple messages with similar content: these other messages allow for the trustworthiness to be evaluated.

To obtain the context of a message, we take advantage of the structure of First-Order Logic and we define each predicate (also called a property) as a separate context. Messages, being literals, communicate a property of a constant (which represents an object in the world). For any literal p in the communication language, we denote its context using the function $context(p)$. Because communication is restricted to literals, a message will always belong to a single context.

The computational trust model assigns a trust evaluation to each message, dependent on the source’s trustworthiness with regards to the message’s context. For most sources, we can rely on existing trust models, because these sources are persistent. We can use the truthfulness of past messages to evaluate the source’s trustworthiness. We therefore do not consider this in much detail, but assume that a trust evaluation is a numerical evaluation in the range $[0, 1]$, which we can interpret as the likelihood that the source will communicate truthfully. We refer to Pinyol et al. [7] for a recent survey of trust models for multi-agent systems.

They also show that most contemporary models compute an evaluation that can be interpreted in the manner described.

Information from C2C communication, however, must be treated differently. A VANET allows for direct communication with hundreds, or even thousands, of other cars and the chance of a repeat interaction with any individual car is small. Furthermore, it would require a large amount of storage to maintain a database of messages previously received from each car. We therefore consider all C2C communication together as a single source of information. In other words, the VANET, rather than individual cars, is considered as a source of information, whose trustworthiness must be evaluated. Additionally, when doing this, we need to take into consideration when a message was sent: which messages do we aggregate together and consider as a single piece of information that is provided by the C2C communication source.

2) *Trust in a VANET*: In order to combine information received over a VANET with information from other sources, we must aggregate it into a similar structure: a single predicate in $\mathcal{L}_{Comm}$ and an associated trustworthiness. To assess this, we use Maximum Likelihood Estimation.

We make a slight adjustment in that we only consider recently received information. We define a threshold $P_{context}$, a maximum period of time that may have passed for the message to still be relevant. This threshold is dependent on the context, because some messages describe events that are only true for a very short period of time — for instance, a bridge being open — whereas others, such as congestion, take longer to become invalidated.

When communicating, a source states that either the atom p is true or false. We thus consider each such message as evidence for the statement p : all messages stating p support it, and all messages stating $\neg p$ conflict with it. Over the VANET, the agent may receive many messages both supporting and conflicting an atom. We therefore consider that C2C communication gives a likelihood for the message content. The frequency of any atom p , as given by C2C communication, is $FC_{2C}(p) = \frac{|M(p)|}{|M(p)| + |M(\neg p)|}$, where $M(p)$ is the set of messages received at most $P_{context}$ time ago, that contain p .

However, this is only the likelihood that the new information is truthful, not taking into account prior communications. To obtain the actual trustworthiness in the information, we must further take the trustworthiness of the source into account. The trust we can place in the communication of atom p , as communicated through a VANET is as in Eq. (1), where $Trust(C2C, p)$ is the trustworthiness of the source C2C in $context(p)$.

$$T_{C2C}(p) = Trust(C2C, p) \cdot FC_{2C}(p) + (1 - Trust(C2C, p)) \cdot FC_{2C}(\neg p) \quad (1)$$

To calculate $Trust(C2C, p)$ we use BRS [11], a simple, but widely used statistical trust model. In particular, we use one of the extensions proposed by Jøsang and Ismail, which allows for discounting older information. This is crucial in such a dynamic system as a traffic network, where the number of trustworthy cars may change over time. Moreover,

we can only consider those messages for which we know the true state of the world. This information could be obtained by evaluating sensor data, such as the odometry of the car, or by obtaining information afterwards; for instance, from a trusted, centralized, database when the car is in its home garage. This provides the agent with a *knowledge base*, and we evaluate the truth of messages in comparison to this knowledge base. The trustworthiness of a source is the posterior likelihood of communicating a truthful message.

3) *Aggregating over different sources*: When evaluating how trustworthy the information concerning a specific atom p is, we must aggregate information about p from various different sources, each with their own trustworthiness. Similar to the case of C2C communication, we must take into account that the sources may disagree. This is a similar problem to that encountered in data fusion for sensor networks, but most data fusion methods rely on the measurements over time and are not applicable to the problem we consider. Jøsang proposes to use a consensus operator [9], which is specifically designed to take advantage of the trustworthiness of different sources.

We wish to consider all sources equally, weighted by their trustworthiness. We thus use the same insight that we used in computing the trustworthiness of a message in C2C communication (see Eq. (1)): we consider events as binary and if a source communicates p with trustworthiness t , it can also be seen as communicating $\neg p$ with trustworthiness $1 - t$. We use this property in the fusion of information from different sources and compute the trustworthiness of p as the mean of the trustworthiness of all the sources' communication of p .

We use the result in reasoning about the communicated information. In particular, we have to reason about its trustworthiness in comparison to other sources of information, such as the car's odometry data or a GPS signal. For this we need to interpret the trustworthiness of the information in terms of beliefs. We thus update the agent's belief base with the new information and its associated trustworthiness. The possibilistic BDI framework takes care of the rest, as briefly discussed in the next section.

B. Reasoning about uncertain information in traffic

The communicated information must be combined with other available information and a decision must be made based on this in order to best achieve the user's goals. Such goals can be, for example, to get home as fast as possible, or park the car near a supermarket. The information an agent has available is uncertain and possibly incomplete. We thus need a goal-oriented reasoning system that can deal with uncertain, incomplete information. For this we use the possibilistic BDI model that was first proposed by da Costa Pereira and Tettamanzi [5]. The model has two features in particular that make it uniquely equipped for decision-making in traffic. The first is that it performs automated belief revision. The agent has multiple sources of information. So far we have focused on communication, but this is only part of the picture: it must use this information

together with information from other sources, such as a GPS signal, odometry data or lidar. Such information is more or less accurate and must be integrated correctly in a belief base. The proposed model provides a way of doing so in accordance with the AGM postulates of belief revision [12].

Even more important is that the model takes the uncertainty of its beliefs into account when selecting goals. The best set of goals to be pursued is not only dependent on the utility that can be obtained, but also on the feasibility of actually achieving the goal. The process of selecting goals must thus deal with the uncertainty of information. For instance, if the GPS signal fails, then it is unfeasible to give the user navigation instructions and instead of doing so erratically, it may be better to display other information, such as a roadmap with the route and last known location. Similarly, if we received trustworthy communication that there is no street parking available downtown, it may be better to direct the user straight to a parking garage, rather than waste time going around the block, whereas going around the block may be the preferred action if there is more uncertainty about parking availability.

In this paper we explain the use of the possibilistic BDI framework using an example, which we will use for evaluating the framework in Section III. For technical details of the framework we refer to [5].

III. EMPIRICAL EVALUATION

In the previous section we presented a framework for dealing with uncertain communication. We empirically demonstrate its applicability in this section. The scenario represents a freeway with a single alternative route. For example, consider the map in Figure 1a. The thick red line represents the major freeway between Utrecht and Amsterdam, two large cities in The Netherlands. This freeway is often congested with commuter traffic between the two cities and the surrounding towns. However, all the alternatives are slow roads that twist around the countryside (we have drawn one of them in blue). Usually, as it is the case here, there are a number of alternatives, but we focus on just one (the shortest, as calculated by Google Maps) to keep the example scenario clear. The framework as presented can equally well deal with multiple different routes: it is simply more information that must be taken into account.

A. Scenario

The road network is represented in the graph of Figure 1b. All agents have the aim of traveling from A to C. In node B they can switch from the freeway onto the side road, or vice versa. We consider the cost of each edge i as the time it takes to travel along it from start to finish, which is dependent on the speed with which drivers can move along it. This, in turn, is dependent on the amount of traffic T on the edge. We base this relationship on the fundamental diagram of traffic [13] and some basic assumptions about the road. The maximum speed v_{max} along the freeway is 120 distance/time, whereas the side road has a speed limit of 60 distance/time. An edge is never completely blocked, and

Fig. 1: Experimental scenario

there is a minimum speed v_{min} of 10 distance/time along the freeway and 5 distance/time along the side road. We use the sigmoid function $v(T) = v_{max} \frac{v_{max} - v_{min}}{1 + e^{5 - T/1000}}$ to represent the fundamental diagram in a continuous function, resulting in the speeds dependent on the amount of traffic T . The speed is normalized around a maximum occupancy of both roads at 10,000 cars, an easy number for use in the simulation. The cost is then simply the distance over the speed. In addition to allowing a higher speed, the freeway is shorter. The two edges along the freeway are 120 distance units each, whereas the side road is 1.5 times as long, with each edge being 180 units long. Note that all values are arbitrarily chosen, although it is normal for side roads to be both longer and slower than the main route.

With 10,000 cars traveling along the network, there is a unique mixed Nash equilibrium, in which each car chooses to go along the freeway with probability 0.6358 and otherwise along the side road, resulting in a cost of 3.69 each. However, we will assume cars with no prior information do not know about the amount of traffic and will always choose to travel along the freeway. The only situation in which a driver considers the side road is if he receives some information indicating that the freeway is congested. In such cases the side road may be a better alternative for reaching his destination.

A driver considers the freeway congested if he is forced to drive slower than half the maximum speed, which occurs at an occupancy of 5,183 cars. For the sake of simplicity, if the freeway is congested at the start of the simulation, there are 6,000 cars on it. We run the simulation with a further 4,000 cars that have the ability to communicate and reason about the communication. We further assume that 20% of the reasoning agents have access to some form of prior information, informing them whether an edge is congested or not. They may then choose to communicate truthfully or not. We further assume that any car that does not have prior information does not attempt to create chaos and fabricate messages. We choose 20% of the cars having prior information in order to ensure there are enough cars that may communicate.

B. Agents

The use of trust in interpreting communication is a straightforward application of the equations in Section II-A. To reason about the communication and decide whether to take the freeway or the alternative road, we program a possibilistic BDI agent as follows.

Firstly, there are only two beliefs to consider: (1) $\text{in}(n)$, the belief that the agent is in node n of the network, and (2) $\text{isCongested}(e)$, the belief that edge e is congested. Secondly, the agent must decide which road to take based on its beliefs. For this there are three high-level actions: $\text{freeway}(xy)$ (take the freeway to go from x to y), $\text{sideroad}(xy)$ (take the side road to go from x to y) and stop.

Reasoning about which action to take is performed using desire-generation rules as follows:

```
if  $\mathcal{B}(\text{in}(A))$  and 0.01 then  $\text{freeway}(AB)$ ,  
if  $\mathcal{B}(\text{in}(A) \wedge \text{isCongested}(AB))$  then  $\text{sideroad}(AB)$ ,  
if  $\mathcal{B}(\text{in}(A) \wedge \neg \text{isCongested}(AB))$  then  $\text{freeway}(AB)$ ,  
if  $\mathcal{B}(\text{in}(B))$  and 0.01 then  $\text{freeway}(BC)$ ,  
if  $\mathcal{B}(\text{in}(B) \wedge \text{isCongested}(BC))$  then  $\text{sideroad}(BC)$ ,  
if  $\mathcal{B}(\text{in}(B) \wedge \neg \text{isCongested}(BC))$  then  $\text{freeway}(BC)$ ,  
if  $\mathcal{B}(\text{in}(C))$  then stop.
```

Here, $\mathcal{B}(\phi)$ represents the degree to which the agent believes formula ϕ , and the 0.01 constants appearing in the antecedents are a generalization of the logical constants $\top$ and $\perp$, i.e., 0.01 is an atom whose trustworthiness is 0.01. The trustworthiness of beliefs gives a priority ordering over the rules. In this case, if $\text{congested}(AB)$ is believed with a trustworthiness greater than 0.01, then the sideroad will be chosen instead of the freeway. In practice this means that any amount of communication will lead to either the second or the third rule to be prioritized over the first. However, if we raise the trustworthiness of the default rule, then it may be prioritized if there is high uncertainty about the beliefs (ergo, the trustworthiness is low). This allows an agent to reason about the trustworthiness of its available information.

This default rule is not chosen arbitrarily. In a real setting, this default choice to take the freeway comes from some prior information, for instance, historical data such as that used by Kraus et al. [4]. The trustworthiness associated with this rule corresponds to the trust the agent puts in this data. In this case, while defaulting to taking the freeway, the agent has no trust in this action and any other information may change it. However, if an agent were to have more trust in the default choice, then it could raise this value: in this case the communication must be more trustworthy in order to override the default rule.

C. Communication and deception

As mentioned in the introduction, deceitful drivers who have some form of *a priori* information about the state of the freeway can improve their travel time by lying about this information. If there is congestion, then they want as few people as possible on the side road, so will try to deceive others into believing the freeway is clear. In the reverse situation, if there is no congestion, then the fewer drivers on the freeway, the better, however minimal the gain.

Either situation is severely detrimental to a naive driver who believes the communication.

Nevertheless, not all drivers act in their own best interest all the time. Because communication in cars is a new technology, no research has been done yet on whether drivers send trustworthy communication (or their intelligent agents do that on their behalf). Instead we rely on research done on altruism in traffic in a broader sense. Mujcic et al. [14] found that 40% of drivers cede way at intersections when they do not have to. Other researchers find similar numbers, or lower of altruistic drivers in traffic scenarios. We will run experiments with different numbers of altruistic agents, and adopt a baseline of 40% for the results in Table I.

There are two opportunities for communication. The first round is before the simulation starts, and the second is before the second choice point in node B of Figure 1b. In the second round of communication, all agents have additionally received information about the true state of the first stretch of freeway, allowing them to assess the trustworthiness of the first round of communication. For the sake of simplicity, we assume the population is homogenous and the percentage of altruistic agents stays at the same throughout the simulation.

D. Results

The experiment serves firstly to demonstrate that deceitful behaviour has a detrimental effect upon traffic flow in the network, and secondly to show that our method for reasoning about the trustworthiness of information allows agents to improve their performance.

1) *The use of trust:* We run the simulation scenario with four different settings. The first is that none of the agents have any prior information about the state of the roads. The second is to provide 20% of the agents with truthful prior knowledge about the state of the freeway. In the third setting, we add the ability to communicate: 60 % of the agents are deceitful. Without trust, we use a simple majority vote to interpret the C2C communication. Finally we allow agents to use trust and reasoning to interpret C2C communication. In this final setting, the trust model needs to be initialized with a default value for trust. We use a naive setting: before being able to evaluate, C2C communication is a fully trusted information source. The results can be found in Table I for two different traffic scenarios: one with, and one without congestion on the freeway.

The first thing to note is that if there is no congestion, then deceitful agents gain little from lying to others, however the agents who rely entirely on C2C communication for their information about the state of the freeway are tricked into taking the sideroad, resulting in their travel time doubling. So while it may not be worthwhile for the deceitful agents, the victims of the deception are significantly affected. In the case of a congested freeway the reverse is true: the agents without information were going to get stuck in the traffic jam on the freeway in any case. Similarly, the deceitful agents with prior knowledge of the congestion were going to take the side road in any case, so in this case their deceitful communication alone does not have any affect on the performance.

	No congestion			Both freeway edges congested		
	Average	Deceitful	Others	Average	Deceitful	Others
No knowledge	2.65			22.35		
Knowledge, no communication	2.65	2.65 ^a		17.75	6.08 ^a	
Knowledge, communication	5.94	2.03	6.90	17.75	6.08	20.64
Knowledge, communication and trust	4.28	2.34	4.77	12.85	7.02	14.31

^a There is no communication, so these are all of the agents (20%) who have knowledge about congestion.

TABLE I: Travel time through graph of Figure 1b using different settings

When using trust and possibilistic reasoning to interpret the communication and act accordingly, we see the situation changes significantly. The cars trust the information (erroneously) in the first round of C2C communication, but in the second round, they are able to adjust and obtain the underlying truth. This mitigates the negative impact of deceitful communication in the case where there is no congestion, and significantly improves agents' performance when there is congestion, if they rely only on C2C communication.

2) *Balancing information sources*: The way our trust model works, the smaller the minority group, regardless of whether they are truthful or lying agents, the higher the trust we can place in VANET communication. The 40% truthful agents is thus not a favourable example for us, as the agents have very low trustworthiness in communication. If 50% of the agents are truthful, then we can say nothing at all about the trustworthiness of the VANET. The further from 50%, the better we can estimate the trustworthiness of the VANET.

In order to decide whether to use the communicated information, the trustworthiness in the default choice is given by prior reasoning about some statistical data regarding the expected state of the freeway. Because the environment is dynamic, and the trust model is capable of adapting over time, it may be the case that in some situations the ad-hoc communication over the VANET overrides the default choice, while at others, trust in the VANET is low and the agent advises the car to follow the default choice. This is precisely the kind of adaptive behaviour that we wish to obtain by using a possibilistic BDI agent.

IV. DISCUSSION

The experiment in the previous section is a proof-of-concept demonstration of the framework for reasoning about C2C communication. We show that under some basic assumptions about the behaviour of traffic, selfish, deceitful agents have a significant impact on traffic flow. Furthermore, our framework is able to learn how to interpret deceitful information and improve the agent's functioning. Nevertheless, we acknowledge that further experimentation is necessary. Particularly in a microsimulation, where it is possible to generate numerous different behaviours, which can change over time. In addition, we intend to gather data about how human participants in traffic use the ability to deceive when being able to communicate easily between each other, in order to calibrate the experimentation scenario better. Furthermore, C2C communication should be considered within the overall problem of congestion management: it is not clear what cars should communicate and when, even if the information is trustworthy.

The framework we proposed performs well in a homogeneous population, but it may be possible to consider C2C communication as not simply a single source of information, but make a more fine-grained distinction. While we cannot evaluate each car's trustworthiness individually, there may be some characteristics that can be associated with trustworthy behaviour: for instance, cars traveling in the opposite direction on the freeway have less incentive to lie about congestion than cars traveling in the same direction. This could be achieved by a richer modeling of other agents in the system, including their possible goals.

C2C communication through VANETs is a valuable tool for providing drivers with more, and more accurate information. However, to process this information it is necessary to take into account that it may be false and to reason about what to present to the user. We presented a comprehensive framework for doing so.

Acknowledgements: Andrew Koster is supported by CAPES (PNPD). Both Andrew Koster and Ana Bazzan are supported in their work by CNPq.

REFERENCES

- [1] R. Lever, "Wi-fi cars hitting the information superhighway," *Agence France-Presse*, March 26 2011.
- [2] C2C-CC, "Car 2 car communication consortium manifesto," 2007.
- [3] A. Bejan and R. Lawrence, "Peer-to-peer cooperative driving," in *Proceedings of ISCIS*, Orlando, USA, 2002, pp. 259–264.
- [4] S. Kraus, R. Lin, and Y. Shavitt, "On self-interested agents in vehicular networks with car-to-car gossiping," *IEEE Transactions on Vehicular Technology*, vol. 57, no. 6, pp. 3319–3332, 2008.
- [5] C. da Costa Pereira and A. G. B. Tettamanzi, "An integrated possibilistic framework for goal generation in cognitive agents," in *AAMAS'10*, Toronto, Canada, 2010, pp. 1239–1246.
- [6] P. Maes, "Agents that reduce work and information overload," *Comm. of the ACM*, vol. 37, no. 7, pp. 30–40, 1994.
- [7] I. Pinyol and J. Sabater-Mir, "Computational trust and reputation models for open multi-agent systems: a review," *Artificial Intelligence Review*, In Press.
- [8] J. Zhang, "A survey on trust management for VANETs," in *Proceedings of IEEE AINA'11*, Biopolis, Singapore, 2011, pp. 105–112.
- [9] A. Jøsang, "A logic for uncertain probabilities," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 9, no. 3, pp. 279–312, 2001.
- [10] B. Lorenz, H. J. Ohlbach, and L. Yang, "Ontology of transportation networks," University of Munich, REVERSE Project, Tech. Rep. A1-D4, 2005.
- [11] A. Jøsang and R. Ismail, "The beta reputation system," in *Proceedings of the Fifteenth Bled Electronic Commerce Conference e-Reality: Constructing the e-Economy*, Bled, Slovenia, 2002.
- [12] C. E. Alchourrón, P. Gärdenfors, and D. Makinson, "On the logic of theory change: Partial meet contraction and revision functions," *Journal of Symbolic Logic*, vol. 50, no. 2, pp. 510–530, 1985.
- [13] B. D. Greenshields, "A study of traffic capacity," in *Proceedings of the 14th Annual Meeting of the Highway Research Board*, 1935, pp. 448–481.
- [14] R. Mujcic and P. Frijters, "Altruism in society: Evidence from a natural experiment involving commuters," *IZA Discussion Papers*, Tech. Rep. 5648, 2011. [Online]. Available: <http://ftp.iza.org/dp5648.pdf>