On Self-Interested Agents in Vehicular Networks With Car-to-Car Gossiping

Sarit Kraus, Raz Lin, *Member, IEEE*, and Yuval Shavitt, *Senior Member, IEEE*

***Abstract*—As more and more cars are equipped with Global Positioning System (GPS) and Wi-Fi transmitters, it becomes easier to design systems that will allow cars to autonomously interact with each other, e.g., regarding traffic on the roads. Indeed, car manufacturers are already equipping their cars with such devices. Although, currently, these systems are a proprietary, we envision a natural evolution where agent applications will be developed for vehicular systems, e.g., to improve car routing in dense urban areas. Nonetheless, this new technology and agent applications may lead to the emergence of self-interested car owners, who will care more about their own welfare than the social welfare of their peers. These car owners will try to manipulate their agents such that they transmit false data to their peers. Using a simulation environment, which models a real transportation network in a large city, we demonstrate the benefits that are achieved by self-interested agents if no countermeasures are implemented. We then proceed to describe the mechanisms for minimizing the effect of the malicious agents on other agents in the network.**

***Index Terms*—Agent-based deployed applications, intelligent agents, peer-to-peer networks, transportation networks.**

I. Introduction

As TECHNOLOGY advances, more and more cars are being equipped with devices that enable them to act as autonomous agents. An important advancement in this respect is the introduction of ad hoc communication networks (such as Wi-Fi), which enable the exchange of information between cars, e.g., for locating road congestion [1] and optimal routes [2] or improving traffic safety [3].

Agent technology, which allows cars to interact autonomously, is becoming recognized by car manufactures as an important aspect in the deployment of future intelligent cars [4], [5]. For example, General Motors [4] develops vehicles with a "sixth sense" that, through vehicular-to-vehicular (V2V) communication, allows vehicles to detect the movement of other vehicles and use this technology to provide more safety for the driver. The U.S. Department of Transportation is also promoting the Vehicle Infrastructure Integration initiative [6] with the vision wherein every car manufactured in the United States will be equipped with a communication device and a GPS unit so that data can be exchanged via a nationwide instrumented roadway system. In addition, "vehicles could serve as data collectors and anonymously transmit traffic and road-condition information from every major road within the transportation network" [6].

We build on the notion of *Gossip Networks*, which were introduced by Shavitt and Shay [2], in which the agents can obtain road congestion information by gossiping with peer agents using ad hoc communication. We first investigate the attraction of being a selfish agent in vehicular networks. That is, we investigate the benefits that are achieved by car owners who tamper with on-board devices and incorporate their own self-interested agents in them, which act in their benefit by exchanging false data with other agents.

We recognize two typical behaviors on which the self-interested agents could embark in the context of vehicular networks. In the first behavior, described in Section IV, the objective of the self-interested agents is to maximize their own utility, expressed by the duration of their average journey. This situation can be modeled in real life by car owners whose aim is to reach their destination as fast as possible and who would like to have their route free of other cars. To this end, the self-interested agents would let their agents cheat the other agents by injecting false information into the network. This is achieved by reporting heavy traffic values for the roads on their route to other agents in the network in the hope of making the other agents believe that the route is jammed, causing them to choose a different route.

The second type of behavior, which is described in Section V, is modeled by the self-interested agents' objective to cause disorder in the network more than they are interested in maximizing their own utility. This kind of behavior could be generated, for example, by vandals or terrorists, who aim to cause as much mayhem in the network as possible.

We note that the introduction of self-interested agents into the network would most probably motivate other agents to try and detect these agents to minimize their effect. This is similar, although in a different context, to the problem introduced by Lamport *et al.* [7] as the *Byzantine Generals Problem*. However, the mechanism that is introduced in [7] and a long line of consequent works that deal with self-interested agents are costly and time consuming. In this paper, we mainly focus on the attractiveness of the selfish behavior by these agents, although we also provide some insights into the possibility of detecting self-interested agents and minimizing their effect on other agents in the network.

Manuscript received December 29, 2007; revised June 4, 2008 and June 12, 2008. First published June 27, 2008; current version published November 12, 2008. This work was supported in part by the Israeli Science Foundation under Grant 1685 and in part by the National Science Foundation under Grant 0705587. The review of this paper was coordinated by Dr. T. Zhang.

S. Kraus is with the Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel, and also with the Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 USA (e-mail: sarit@cs.biu.ac.il).

R. Lin is with the Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel (e-mail: linraz@cs.biu.ac.il).

Y. Shavitt is with the School of Electrical Engineering, Tel-Aviv University, Ramat-Aviv 69978, Israel (e-mail: shavitt@eng.tau.ac.il).

Digital Object Identifier 10.1109/TVT.2008.927986

0018-9545/$25.00 © 2008 IEEE

Because of the complexity of mathematically analyzing dynamic networks, to demonstrate the benefits that are achieved by self-interested agents, we have used a simulation environment that models the transportation network in a central part of a large real city. To this end, we extended the microsimulation tool (see [8]–[10] for other microsimulators) that is proposed in [11], which supports the use of gossiping between individual cars to support the different behaviors of agents (see Sections IV–VI for the description of the different behaviors). The simulation environment is further described in Section III. Our simulations provide insights into the benefits of self-interested agents that cheat. Our findings can motivate future research in this field to minimize the effect of selfish agents. Finally, in Section VI, we describe mechanisms for minimizing the effect of the malicious agents on other agents in the network. In Section VII, we show the results of malicious agents forming coalitions to overcome the protection mechanisms that are implemented by gossip agents.

We begin by reviewing related work in the field of self-interested agents and V2V communication.

II. Related Work

In their seminal paper, Lamport *et al.* [7] describe the *Byzantine Generals Problem*, in which processors need to handle malfunctioning components that give conflicting information to different parts of the system. They also present a model in which not all agents are connected, and thus, an agent is not able to send a message to all the other agents. Dolev *et al.* [12] have built on this problem and have analyzed the number of faulty agents that can be tolerated to eventually reach the right conclusion about the true data. Similar work is presented by Minsky and Schneider [13], who discuss techniques for constructing gossip protocols that are resilient to up to t malicious host failures. As opposed to the above works, this paper focuses on vehicular networks, in which agents constantly roam the network and exchange data. Also, the domain of transportation networks introduces dynamic data, as the load of the roads is subject to change. In addition, transportation network systems include a feedback mechanism since the load of the roads depends on the reports and the movement of the agents themselves.

Malkhi *et al.* [14] present a gossip algorithm for propagating information in a network of processes in the presence of malicious parties. Their algorithm prevents the spread of spurious gossip and diffuses genuine data. This is done in time, which is logarithmic in the number of processes and linear in the number of corrupt parties. Nevertheless, their work assumes that the network and the agents are static (they discuss a network of processes). This is not true for transportation networks. Transportation networks, by nature, are dynamic. The agents constantly move, and the data change over time. For example, in our model, agents might gossip about a heavy traffic load on a specific road, which is currently jammed. Nonetheless, this information might be false several minutes later, leaving the agents to speculate whether the spreading agents are indeed malicious. In addition, as the agents are constantly moving, each agent cannot choose with whom it interacts and exchanges data.

In the context of analyzing the data and its correctness, researchers have focused on distributed reputation systems or on mechanisms to decide whether to share data. Yu and Singh [15] built a social network of agents' reputation. In their network, every agent keeps a list of its neighbors, which can be changed over time, and computes the trustworthiness of other agents by updating the current values of testimonies that are obtained from reliable referral chains. After a bad experience with another agent, every agent decreases the rating of the "bad" agent and propagates this bad experience throughout the network so that other agents can accordingly update their ratings. This approach could be implemented in our domain to allow the agents, by gossiping with their peer agents, to identify self-interested agents and, thus, minimize their effect. However, the implementation of such a mechanism is an expensive addition to the infrastructure of autonomous agents in transportation networks, mainly due to the dynamic nature of the list of neighbors in such networks.

Leckie and Kotagiri [16] study when to share information between the agents in the network. Their domain involves monitoring the distributed sensors. Each agent monitors a subset of the sensors and evaluates a hypothesis based on the local measurements of its sensors. If the agent believes that a hypothesis is likely, he/she exchanges this information with the other agents. In their domain, the goal of all the agents is to reach a global consensus about the likelihood of the hypothesis. In our domain, however, as the agents constantly move, they have many samples, which they exchange with each other. Also, the data are dynamic (e.g., a road might be reported as jammed, but a few minutes later, it could be free), thus making it harder to decide whether to trust the agent who sent the data. Moreover, the agent might lie only about a subset of its samples, thus making it even harder to detect his/her cheating.

Some work has been done in the context of gossip networks or transportation networks regarding the spreading of data and their dissemination. Datta *et al.* [17] focus on information dissemination in mobile ad hoc networks. They propose an autonomous gossiping algorithm for an infrastructure-less mobile ad hoc networking environment. Their autonomous gossiping algorithm uses a greedy mechanism to spread data items in the network. The data items are spread to immediate neighbors that are interested in the information and avoid ones that are not interested. The decision of which node is interested in the information is made by the data item itself using heuristics. However, its work concentrates on the movement of the data themselves and not on the agents who propagate the data. This is different from our scenario in which each agent maintains the data it has gathered, while it roams the road and is responsible (and has the capabilities) for spreading the data to other agents in the network.

Das *et al.* [18] propose a cooperative strategy for content delivery in vehicular networks. In their domain, peers download a file from a mesh and exchange parts of the file among themselves. We, on the other hand, are interested in vehicular networks in which there is no rule forcing the agents to cooperate.

Shibata *et al.* [19] propose a method for cars to cooperatively and autonomously collect traffic jam statistics to estimate the arrival time to destinations of each car. The communication is based on IEEE 802.11 without having to utilize a fixed infrastructure on the ground. Although we use the same domain, we focus on a different problem. Shibata *et al.* [19] mainly focus on efficiently broadcasting the data between agents (e.g., avoid duplicates and communication overhead), whereas we focus on the case where agents are not cooperative by nature and on how selfish agents affect other agents and the network load.

Kraus *et al.* [11] describe a simulation tool that supports the use of gossiping between individual cars to support the different behavior of each car. In their model, they assume that drivers learn the expected congestion on the roads, and some of them have a gossiping system that helps them learn about congestion on distant roads. They study the information propagation speed in an urban network and quantify its advantage to drivers on the road. Although we use the same simulation tool, we focus on a different problem and investigate the effects of self-interested and malicious agents on the other drivers in the network.

Wang *et al.* [20] also assert that individual agents are likely to selfishly act in the context of wireless networks. They design a protocol for communication in networks in which *all* agents are selfish. Their protocol motivates every agent to maximize its profit only when it behaves truthfully (an *incentive compatibility* mechanism). However, the domain of wireless networks is quite different from the domain of transportation networks. In the wireless network, a wireless terminal is required to contribute its local resources to transmit data. Thus, Wang *et al.* [20] use a payment mechanism, which attaches costs to terminals when transmitting data and, thus, enables them to maximize their utility when transmitting data instead of acting selfishly. Disparately, in the context of transportation networks, constructing such a mechanism is not quite a straightforward task, as self-interested agents and regular gossip agents might incur the same cost when transmitting data. The difference between the two types of agents only exists with regard to the credibility of the data that they exchange.

In Section III, we will describe our transportation network model and gossiping between the agents. We will also describe the types of agents in our system.

III. Model and Simulations

In our simulations, we wanted to model a scenario in which drivers roam the city, with the objective of traveling from one point to another, and observe what happens when self-interested drivers are also present. To this end, we devised different scenarios and settings. We first describe our transportation network model, and then, we depict the simulations' designs.

A. Transportation Network Model

Following Shavitt and Shay [2], Kraus *et al.* [11], and Parshani [21], the transportation network is represented by a directed graph $G(V, E)$, where V is the set of vertices representing junctions, and E is the set of edges representing roads. An edge $e \in E$ is associated with a weight $w > 0$, which specifies the time it takes to traverse the road that is associated with that edge. The roads' weights vary in time according to the network (traffic) load. Each car, which is associated with an autonomous agent, is given a pair of origin and destination vertices. A *journey* is defined as the (not necessarily simple) path that is taken by an agent between the origin vertex and the destination vertex. We assume that there is always a path between a source and a destination. A *journey length* is defined as the sum of all weights of the edges constituting it. Every agent aims to minimize its journey length.

At a given time, an agent may have inaccurate information about the weights and no information on how the weights may change over time. We assume that an agent, which travels from vertex $v_1 \in V$ to $v_2 \in V$, will search for the shortest path between these two vertices, based on its current available information, and will move accordingly. Once its information about the network has been updated, it will randomly decide whether to recalculate the shortest path or to keep on moving and follow its current route. If there is more than one path that is associated with the shortest distance, one of them will be chosen randomly.

Initially, agents are ignorant about the state of the roads. *Regular agents* are only capable of gathering information about the roads as they traverse them. However, we assume that some agents have means of intervehicle communication (e.g., IEEE 802.11) with a given communication range, which enables them to communicate with other agents with the same device. Those agents are referred to as *gossip agents*. Since the communication range is limited, the exchange of information using gossiping is done in one of two ways: 1) between gossip agents passing one another or 2) between gossip agents that are located at the same junction. We assume that each agent stores the most recent information that it has received or gathered around the edges in the network. Note that we assume a limited communication range. This assumption can be extended to allow a broader communication range. However, such an extension would also raise other issues such as complexity (e.g., maintaining a larger set of information) and applicability (e.g., how much would the data gathered at time t on a given junction be relevant for another agent that would not arrive at the said junction within the near future). In addition, this could create a similar effect as the results of increasing the percentage of gossiping agents. However, as we discuss in Section III-B, our simulations were conducted when the percentage of gossip agents was shown to be most efficient. Thus, in this paper, we only investigate the limited communication model.

A subset of the gossip agent are agents that are self-interested and manipulate the devices for their own benefit. We will refer to these agents as *self-interested agents*. A detailed description of their behavior is given in Sections IV and V.

B. Simulation Design

Building on [11] and [21], the network in our simulations[1] replicates a large city center and consists of 50 junctions and 150 main roads. Each simulation consists of six iterations. The basic time unit of the iteration is a step, which is equivalent to about 30 s. Each iteration simulates 6 h of movement. The average number of cars passing through the network during

[1] See http://www.cs.biu.ac.il/~linraz/vehicularAgents.htm for the simulation tool.

the iteration is about 70 000, and the average number of cars in the network at a specific time unit is about 3500 cars. In each iteration, the same agents are used, such that each agent has the same origin and destination points in the different iterations, whereas the data collected in earlier iterations are preserved for future iterations (which is referred to as the history of the agent). This allows us to simulate a somewhat daily routine in the transportation network (e.g., a working week).

Each of the experiments that we describe below was run with five different traffic scenarios. Each of these traffic scenarios differed from one another by the initial load of the roads and the designated routes of the agents (cars) in the network. Five simulations were run for each scenario, thereby creating a total of 25 simulations for each experiment.

Kraus *et al.* [11] and Parshani [21] showed that the information propagation in the network is very efficient when the percentage of gossiping agents is 10% or more. Yet, due to congestion that is caused by too many cars rushing to what was reported as the less-congested part of the network, 20%–30% of gossiping agents led to the most efficient routing results in their experiments. Consequently, in our study, we focus only on the simulations in which the percentage of gossip agents is 20%.

The simulations were done with different percentages of self-interested agents. Each simulation was run with changes in the set of gossip agents and the set of self-interested agents.

To attain a similar ordinal scale, the results were normalized. The normalized values were calculated by comparing each agent's results to its results when the same scenario was run with no self-interested agents. This was done for all iterations. Using the normalized values enabled us to observe how worse (or better) each agent would perform compared to the basic setting. For example, if the journey length of a certain agent in iteration 1 with no self-interested agent was 50, and the length was 60 in the same scenario and iteration in which self-interested agents were involved, then the normalized value for that agent would have been $60/50 = 1.2$. We refer to a change of $\pm 3\%$ in the normalized value as a small effect on the agent, whereas higher changes are considered to have large effects.

The simulations were all done at the system level. In particular, we did not model the medium-access-control performance and signal propagation. The simulator with documentation is available at http://www.cs.biu.ac.il/~linraz/vehicularAgents.htm.

Further details of the simulations are presented in Sections IV and V.

IV. Spreading Lies; Maximizing Utility

In the first set of experiments, we investigated the benefits that are achieved by self-interested agents, whose aim was to minimize their own journey length. The self-interested agents adopted a cheating approach, whereby they sent false data to their peers.

In this section, we first describe the simulations with the self-interested agents. Then, we model the scenario as a game with two types of agents and prove that the equilibrium result can only be achieved when there is no efficient exchange of gossiping information in the network.

A. Modeling the Self-Interested Agents' Behavior

Whereas the gossip agents gather data and send them to other agents, the self-interested agents' behavior is modeled as follows.

1) Calculate the shortest path from the origin to the destination.
2) Communicate the following data to other agents.
   a) If the road is *not* on the agent's route—send the true data about it (e.g., data about roads that it has received from other agents).
   b) For all the roads on the agent's route, which the agent has not yet traversed, send a random high weight.

Basically, the self-interested agent acts in the same manner as the gossip agent. It collects data regarding the weight of the roads (either by traversing the road or by obtaining the data from other agents) and sends the data it has collected to other agents. However, the self-interested agent acts differently when the road is on its route. Since the agent's goal is to reach its destination as fast as possible, the agent will falsely report that all the roads on his/her route are heavily congested, which frees the path for itself by making other agents recalculate their paths, this time without including the roads on the self-interested agent's route. To this end, for all the roads in its route, which the agent has not yet passed, the agent generates a random weight, which is above the average weight of the roads in the network. It then associates these new weights with the roads on its route and sends them to the other agents.

Although an agent can also divert cars from his/her route by falsely reporting congested roads that are parallel to his/her route as free, this behavior is not very likely since other agents attempting to use the roads will find the mistake within a short time and spread the true situation of the road. On the other hand, if an agent manages to persuade other agents not to use a road, it will be harder for them to detect that the said roads are not congested.

In addition, to avoid being influenced by their own lies and other lies that are spread throughout the network, all self-interested agents will ignore the data that are received about the roads with heavy traffic (note that the data about the roads that are not congested will not be ignored).[2]

In Section IV-B, we describe the simulation results involving the self-interested agents.

B. Simulation Results

We ran several experiments to test the benefits of self-interested agents that are cheating. In the first set of experiments, we created a scenario in which a small group of self-interested agents spread lies about the same route and tested their effect on the journey length of all the agents in the network. To try and maximize the effect of the lies on agents that are traveling that route, we selected several cars that had

[2]In other simulations that we have run, in which there were several occurrences of real congestion in the network, we, indeed, observed that even when the roads were jammed, the self-interested agents were less affected if they ignored all reported heavy traffic since, consequently, they also discarded all disinformation roaming the network.

TABLE I
NORMALIZED JOURNEY LENGTH VALUES BY ITERATION WHEN SIX SELF-INTERESTED AGENTS, WITH THE SAME ORIGIN AND DESTINATION, SPREAD LIES ABOUT THEIR ROUTE; ONE ROAD, ON THE ROUTE OF THESE AGENTS, WAS PARTIALLY BLOCKED

Iteration Number	Self-Interested Agents	Gossip - Same	Gossip - Others	Regular Agents
1	1.38	1.27	1.06	1.06
2	0.95	1.56	1.18	1.14
3	1.00	1.86	1.28	1.17
4	1.06	1.93	1.35	1.17
5	1.13	2.00	1.40	1.17
6	1.08	2.02	1.43	1.18

TABLE II
NORMALIZED JOURNEY LENGTH VALUES BY ITERATION WHEN ONE SELF-INTERESTED AGENT HAS THE OBJECTIVE OF HELPING ANOTHER BENEFICIARY AGENT WITH THE SAME ORIGIN AS ITS OWN

Iteration Number	Beneficiary Agent	Gossip - Same	Gossip - Others	Regular Agents
1	1.10	1.05	0.94	1.11
2	1.09	1.14	0.99	1.14
3	1.04	1.19	1.02	1.14
4	1.03	1.26	1.03	1.13
5	1.05	1.32	1.05	1.12
6	0.92	1.39	1.06	1.11

the same origin and destination to serve as the self-interested agents. In this simulation, we selected only six agents to be part of the self-interested agent group to investigate the effect that is achieved by only a small number of agents.

In this experiment, six different agents were randomly chosen in each simulation to be part of the self-interested agent group, as described above. In addition, one road, on the route of these agents, was randomly selected to be partially blocked, letting only one car go through at each time step. About 8000 agents were randomly selected as regular gossip agents, and the other 32 000 agents were designated as regular agents.

We analyzed the average journey length of the self-interested agents compared with the average journey length of other regular gossip agents traveling along the same route. Table I summarizes the normalized results for the self-interested agents, the gossip agents (those having the same origin and destination as the self-interested agents, denoted *Gossip—Same*, and all other gossip agents, denoted *Gossip—Others*), and the regular agents as a function of the iteration number.

The results that are presented in Table I reveal that the first time (iteration 1) self-interested agents travel the route while spreading false data about the roads does not help them (using the paired t-test, we show that the agents have significantly lower journey lengths in the scenario in which they do not spread any lies, with p-value < 0.01). This is mainly due to the fact that the lies do not advance ahead of the self-interested agent and reach other cars that are ahead of the self-interested car on the same route. Thus, spreading the lies in the first iteration does not help the self-interested agent to free the route it is about to travel during the first iteration.

Only when the self-interested agents repeat their journey in the next iteration (iteration 2) will the disinformation significantly help them (p-value $= 0.04$). The reason for this is that other gossip agents have received these data and have used them to recalculate their shortest path, thus avoiding the roads that are the subject of the disinformation. It is also interesting to note the large value that is attained by the self-interested agents in the first iteration. This is mainly due to several self-interested agents that enter the jammed road. This situation occurs since the self-interested agents ignore all heavy traffic data and, thus, ignore the fact that the road is jammed. As they begin to spread lies about this road, more cars shift from this route, thus making the road free for future iterations.

However, we also recall that the self-interested agents ignore all information about roads with heavy traffic. Thus, when the network becomes congested, more self-interested cars are affected since they might enter jammed roads, which they would otherwise not have entered. This can be seen, for example, in iterations 4–6, in which the normalized value of the self-interested agents increases above 1.00. Using the paired t-test to compare these values with the values that are achieved by these agents when no lies are used, we reveal that there is no significant difference between the two scenarios.

As opposed to the gossip agents, we observe how little effect the self-interested agents have on the regular agents. In comparison to the gossip agents on the same route that travel as much as 93% more when self-interested agents are introduced, the average journey length for the regular agents only increases by about 15%. This result is even lower than the effect on other gossip agents in the entire network.

Since we noticed that self-interested agents do not benefit by cheating in the first iteration, we devised another set of experiments. In the second set of experiments, the self-interested agents have an objective of helping another agent that is supposed to enter the network some time after the self-interested agent has entered. We refer to the latter agent as the beneficiary agent. Similar to a self-interested agent, the beneficiary agent also ignores all data regarding the heavy traffic. In real life, this can be modeled, for example, by a husband who would like to help his wife find a faster route to her destination. Table II summarizes the normalized values for the different agents. As in the first set of experiments, five simulations were run for each scenario, with a total of 25 simulations. In each of these simulations, one agent was randomly selected as a self-interested agent, and then, another agent with the same origin as the self-interested agent was randomly selected as the beneficiary agent. The other 8000 and 32 000 agents were designated as regular gossip agents and regular agents, respectively.

We can see that the higher the number of iterations, the lower the normalized value for the beneficiary agent. In this scenario, as in the previous one, in the first iterations, not only does the beneficiary agent not avoid the jammed roads, since it ignores all heavy traffic, but it also does not benefit from the lies that are spread by the self-interested agent. This is due to the fact that the disinformation has not yet been incorporated by other gossip agents. Thus, if we compare the average journey length

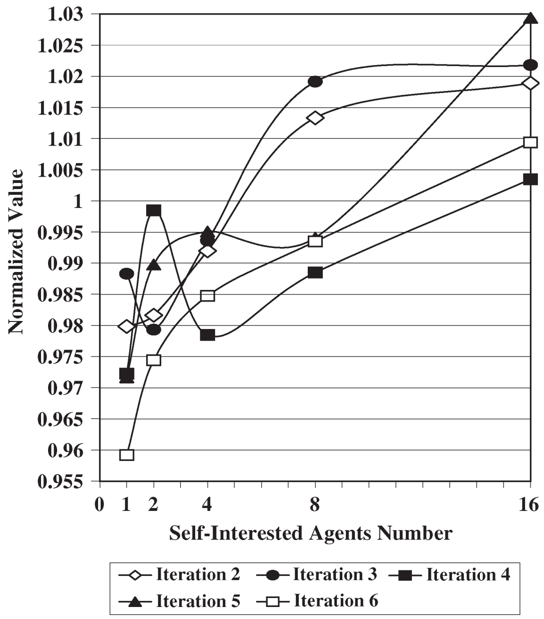

Fig. 1. Self-interested agent normalized values as a function of the number of self-interested agents. Self-interested agents have the objective of minimizing their average journey length.

in the first iteration when lies are spread and when there are no lies, the average is significantly lower for the latter case (p-value < 0.03). On the other hand, if we compare the average journey length in all the iterations, there is no significant difference between the two settings. Nonetheless, in most of the iterations, the average journey length of the beneficiary agent is longer than in the case when no lies are spread.

We can also see the impact on the other agents in the system. While the gossip agents, which are not on the route of the beneficiary agent, are virtually not affected by the self-interested agent, those on the route and the regular agents are affected and have higher normalized values. That is, even with only one self-interested car, we can see that both the gossip agents that begin the same route (i.e., the same origin and destination points) as the self-interested agents spreading the lies and other regular agents significantly increase their journey length (p-value < 0.015 for the gossip agents and p-value < 0.002 for the regular agents) by more than 17%, on average.

In our third set of experiments, we examined a setting whereby there is an increasing number of self-interested agents, which do not necessarily have the same origin and destination points. To model this, we randomly selected self-interested agents, whose objective was to minimize their average journey length, assuming that the cars repeat their journeys (that is, more than one iteration was performed). As opposed to the first set of experiments, in this set, the self-interested agents were selected randomly, and we did not enforce the constraint of having the same origin and destination points.

As in the previous sets of experiments, we ran five different simulations per scenario. In each simulation, 11 runs were made, each run with a different number of self-interested agents: 0 (no self-interested agents), one, two, four, eight, and 16. Each agent adopted the behavior that is modeled in Section IV-A. Fig. 1 shows the normalized value that is achieved by the self-interested agents as a function of their number. The figure shows these values for iterations 2–6. The first iteration is not shown intentionally, as we assume repeated journeys.

Using these simulations, we examined the possible threshold of the number of randomly selected self-interested agents, which will allow them to benefit from their selfish behavior. We can see that with up to eight self-interested agents, the average normalized value is below 1. That is, the self-interested agents benefit from their malicious behavior. In the case of one self-interested agent, a significant difference is revealed between the average journey length of when misinformation is spread by the agent and when no lies are spread (p-value < 0.001). However, when there are two, four, eight, and 16 self-interested agents, there is no significant difference. Yet, as the number of self-interested agents increases, the normalized value also increases. In such cases, the normalized value is larger than 1, and the self-interested agents' journey length becomes significantly higher than their journey length in cases where there are no self-interested agents in the system.

In Section IV-C, we analyze the scenario as a game and show that, when in equilibrium, the exchange of gossiping between the agents becomes inefficient.

C. When Gossiping Is Inefficient

We continued by modeling our scenario as a game to find the equilibrium.

In our game, there are two possible types of agents: 1) regular gossip agents and 2) self-interested agents. Each of these agents is a representative of its group, and thus, all agents in the same group have similar behavior.

We note that the advantage of using gossiping in transportation networks is to allow the agents to detect anomalies in the network (e.g., traffic jams) and to quickly adapt by recalculating their routes [11]. We also assume that the objective of the self-interested agents is to minimize their own journey length; thus, they spread lies on their routes, as described in Section IV-A. Furthermore, we assume that sophisticated methods for identifying the self-interested agents or managing reputation are not used. This is mainly due to the complexity of incorporating and maintaining such mechanisms, as well as due to the dynamics of the network, in which the interactions between different agents are frequent; agents may leave the network, and data about the road might change as time progresses (e.g., a road might be reported by a regular gossip agent as free at a given time, and yet, currently, it may be jammed due to heavy traffic on the road). Nevertheless, we discuss the mechanisms for overcoming malicious agents in Section VI.

We should also note the fact that the Nash solution does not necessarily mean the optimal solution, but rather a stable solution. Also, research has shown that humans (and we assume that the self-interested agents model human drivers) do not necessarily follow an equilibrium strategy (e.g., see [22] and [23]). Even as such, we should note the different assumptions that were the basis of this analysis and were not part of the simulations with which we experiment.

- We assume that there are two groups of agents—self-interested agents and regular gossip agents. We give a similar weight to both groups, although in our simulation, there are far fewer self-interested agents than gossip agents (as we assume is the case in real life).

- The dynamics of the network, the propagation of information, and whether the data are an update or not are not taken into consideration in the game modeling.
- We assume that self-interested agents and gossip agents have information regarding the average time that it takes to traverse each edge (this, however, can be assumed in real life as well).

We proceed by analyzing the game's equilibrium. Let $T_{\rm avg}$ be the average time it takes to traverse an edge in the transportation network (that is, the average load of an edge). Let $T_{\max}$ be the maximum time it can take to traverse an edge. We will investigate the game in which the self-interested and regular gossip agents can choose the following actions. The self-interested agents can choose how much to lie, that is, they can choose to spread the information about how long (not necessarily the true duration) it takes to traverse certain roads. Since the objective of the self-interested agents is to spread messages as though some roads are jammed, the traversal time they report is obviously larger than the average time. We denote the time that the self-interested agents spread as T_s, such that $T_{\rm avg} \leq T_s \leq T_{\max}$. Motivated by the results of the simulations we have described above, we observed that the agents are less affected if they discard the heavy traffic values. Thus, the regular gossip cars, attempting to mitigate the effect of the liars, can choose a strategy to ignore abnormal congestion values that are above a certain threshold T_g. Obviously, $T_{\rm avg} \leq T_g \leq T_{\max}$. To prevent the gossip agents from detecting the lies and simply discarding the values, the self-interested agents send lies within the given range ($[T_{\rm avg}, T_{\max}]$) with an inverse geometric distribution, that is, the higher the T value, the higher its frequency.

Now, we construct the utility functions for each type of agent, which is defined by the values of T_s and T_g. If the self-interested agents spread traversal times that are higher than or equal to the regular gossip cars' threshold, they will not benefit from their lies. Thus, the utility value of the self-interested agents in this case is zero. On the other hand, if the self-interested agents spread misinformation stating traversal times that are lower than the threshold, they will gain a positive utility value (to ensure that the utility value will always be larger than 0, we added 1 in the calculations). From the regular gossip agents' point of view, if they accept messages from the self-interested agents, then they incorporate the lies into their calculation, thereby losing utility points. On the other hand, if they discard the false values that the self-interested agents send, i.e., they do not incorporate the lies, they will gain utility values. Formally, we use u^s to denote the utility of the self-interested agents and u^g to denote the utility of the regular gossip agents. We also denote the strategy profile in the game as $\{T_s, T_g\}$. The utility functions are defined as

$$u^s = \begin{cases} 0, & \text{if } T_s \geq T_g \\ T_s - T_{\rm avg} + 1, & \text{if } T_s < T_g \end{cases} \tag{1}$$

$$u^g = \begin{cases} T_g - T_{\rm avg}, & \text{if } T_s \geq T_g \\ T_s - T_g, & \text{if } T_s < T_g. \end{cases} \tag{2}$$

We are interested in finding the *Nash equilibrium*. Recall from Osborne and Rubinstein [24, ch. 2] that the Nash equilibrium is a strategy profile where no player has anything to gain by deviating from his/her strategy, given that the other agent follows his/her strategy profile. Formally, let (S, u) denote the game, where S is the set of strategy profiles, and u is the set of utility functions. When each agent $i \in \{\text{regular gossip}, \text{self-interested}\}$ chooses a strategy T_i, resulting in a strategy profile $T = (T_s, T_g)$, then agent i obtains a utility of $u^i(T)$. A strategy profile $T^* \in S$ is a Nash equilibrium if no deviation in the strategy by any single agent is profitable, i.e., if for all i, $u^i(T^*) \geq u^i(T_i, T^*_{-i})$. In other words, (T_s, T_g) is a Nash equilibrium if the self-interested agents have no other value T'_s such that $u^s(T'_s, T_g) > u^s(T_s, T_g)$ and similarly for the gossip agents.

We now present the following theorem.

Theorem 1: $(T_{\rm avg}, T_{\rm avg})$ is the only Nash equilibrium.

Proof: First, we will show that $(T_{\rm avg}, T_{\rm avg})$ is a Nash equilibrium. Assume, by contradiction, that the gossip agents choose another value $T_{g'} > T_{\rm avg}$. Thus, $u^g(T_{\rm avg}, T_{g'}) = T_{\rm avg} - T_{g'} < 0$. On the other hand, $u^g(T_{\rm avg}, T_{\rm avg}) = 0$. Thus, the regular gossip agents have no incentive to deviate from this strategy. The self-interested agents also have no incentive to deviate from this strategy. By contradiction, again, assume that the self-interested agents choose another value $T_{s'} > T_{\rm avg}$. Thus, $u^s(T_{s'}, T_{\rm avg}) = 0$, whereas $u^s(T_{\rm avg}, T_{\rm avg}) = 0$.

We will now prove that the above solution is unique. We will demonstrate that any other pair (T_s, T_g), such that $T_{\rm avg} < T_g \leq T_{\max}$ and $T_{\rm avg} < T_s \leq T_{\max}$, is not a Nash equilibrium.

We have three cases. In the first case, $T_{\rm avg} < T_g < T_s \leq T_{\max}$. Thus, $u^s(T_s, T_g) = 0$, and $u^g(T_s, T_g) = T_g - T_{\rm avg}$. In this case, the regular gossip agents have an incentive to deviate and choose another strategy $T_g + 1$ since, by doing so, they increase their own utility: $u^g(T_s, T_g + 1) = T_g + 1 - T_{\rm avg}$.

In the second case, $T_{\rm avg} < T_s < T_g \leq T_{\max}$. Thus, $u^g(T_s, T_g) = T_s - T_g < 0$. Also, the regular gossip agents have an incentive to deviate and choose another strategy $T_g - 1$ in which their utility value is higher: $u^g(T_s, T_g - 1) = T_s - T_g + 1$.

In the last case, $T_{\rm avg} < T_s = T_g \leq T_{\max}$. Thus, $u^s(T_s, T_g) = T_s - T_g = 0$. In this case, the self-interested agents have an incentive to deviate and choose another strategy $T_g - 1$, in which their utility value is higher: $u^s(T_g - 1, T_g) = T_g - 1 - T_{\rm avg} + 1 = T_g - T_{\rm avg} > 0$. ■

The above theorem proves that the equilibrium point is reached only when the self-interested agents send the time to traverse certain edges equal to the average time, and on the other hand, the regular gossip agents discard all data regarding roads that are associated with an average time or higher. Thus, for this equilibrium point, the exchange of gossiping information between agents is inefficient, as the gossip agents are unable to detect congestion and heavy traffic in the network.

Although, above, we prove that the equilibrium states that gossiping is inefficient under the assumptions we have laid, this theoretical result is relevant to these extreme cases. Moreover, this proof provides a guideline on how to ensure that the gossip will remain effective, i.e., preventing the assumption of the theoretical model from coming true.

TABLE III
NORMALIZED JOURNEY LENGTH VALUES FOR THE FIRST ITERATION. INCREASING THE NUMBER OF SELF-INTERESTED AGENTS WITH THE OBJECTIVE OF MINIMIZING THE AVERAGE JOURNEY LENGTH

Number of Self-Interested Agents	Self-Interested Agents	Gossip Agents	Regular Agents
1	0.98	1.01	1.05
2	1.09	1.02	1.05
4	1.07	1.02	1.05
8	1.06	1.04	1.05
16	1.03	1.08	1.06
32	1.07	1.17	1.08
50	1.12	1.28	1.10
64	1.14	1.39	1.13
80	1.15	1.50	1.14
100	1.17	1.63	1.16

TABLE IV
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS. INCREASING THE NUMBER OF SELF-INTERESTED AGENTS WITH THE OBJECTIVE OF MINIMIZING THEIR AVERAGE JOURNEY LENGTH

Number of Self-Interested Agents	Self-Interested Agents	Gossip Agents	Regular Agents
1	0.98	1.01	1.06
2	1.00	1.02	1.07
4	1.00	1.04	1.07
8	1.01	1.18	1.11
16	1.02	1.53	1.17
32	1.06	2.13	1.25
50	1.13	2.21	1.29
64	1.21	2.21	1.32
80	1.21	2.12	1.27
100	1.26	2.10	1.27

In Section V, we describe another scenario for the self-interested agents, in which they are not concerned with their own utility but, rather, are interested in maximizing the average journey length of other gossip agents.

V. SPREADING LIES; CAUSING CHAOS

Another possible behavior that can be adopted by self-interested agents is characterized by their goal to cause disorder in the network. This can be achieved, for example, by maximizing the average journey length of all agents, even at the cost of maximizing their own journey length.

To understand the vulnerability of the gossip-based transportation support system, we ran five different simulations for each scenario. In each simulation, different agents were randomly chosen (using a uniform distribution) to act as gossip agents from which self-interested agents were chosen. Each self-interested agent behaved in the same manner as described in Section IV-A.

Every simulation consisted of 11 runs, with each run comprising different numbers of self-interested agents: 0 (no self-interested agents), one, two, four, eight, 16, 32, 50, 64, 80, and 100. Also, in each run, the number of self-interested agents was increased incrementally. For example, the run with 50 self-interested agents consisted of all the self-interested agents that were used in the run with 32 self-interested agents but with an additional 18 self-interested agents. Also, recall that, in each run, the average number of cars passing through the network during an iteration was about 70 000.

Tables III and IV summarize the normalized journey length for the self-interested agents, the regular gossip agents, and the regular (nongossip) agents for the first iteration and for the average of all iterations, respectively. Fig. 2 demonstrates the changes in the normalized values for the regular gossip agents and the regular agents as a function of the iteration number. Similar to the results in our first set of experiments, described in Section IV-B, we can see that randomly selected self-interested agents that follow different randomly selected routes do not benefit from their malicious behavior (that is, their average journey length does not decrease). However, when only one self-interested agent is involved, it does benefit from the malicious behavior, even in the first iteration. The results also indicate that the regular gossip agents are more sensitive to malicious behavior than regular agents—the average journey length for the gossip agents increases significantly (e.g., with 32 self-interested agents, the average journey length for the gossip agents was 113% higher, which is significantly higher with p-value < 0.01, than in the setting with no self-interested agents, as opposed to an increase of only 25% for the regular agents). In addition, these results also indicate the effects of the self-interested agents' behavior on the network load. It is also interesting to see that the highest normalized value for the gossip agents is achieved when there are 50 malicious agents. When the number of malicious agents increases, the normalized

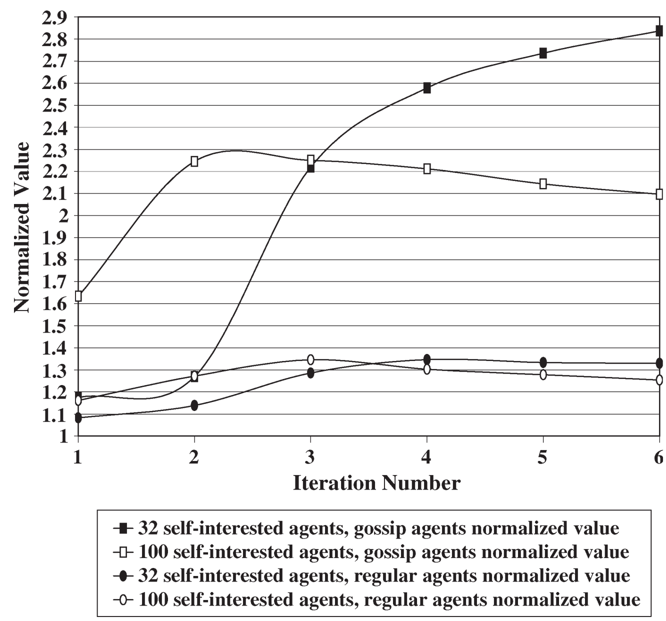

Fig. 2. Gossip and regular agent normalized values as a function of the iteration. Thirty-two and 100 self-interested agents with the objective of minimizing their average journey length.

TABLE V
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS. INCREASING THE NUMBER OF SELF-INTERESTED AGENTS WITH THE OBJECTIVE OF MINIMIZING THEIR AVERAGE JOURNEY LENGTH. THIRTEEN MAIN ROADS ARE JAMMED

Number of Self-Interested Agents	Self-Interested Agents	Gossip Agents	Regular Agents
1	1.07	1.02	1.22
2	1.09	1.04	1.23
4	1.06	1.06	1.23
8	1.09	1.15	1.26
16	1.11	1.55	1.39
32	1.14	2.25	1.56
50	1.30	2.25	1.60
64	1.35	2.47	1.63
80	1.51	2.41	1.64
100	1.68	2.61	1.75

value begins to decrease. This can be explained by the fact that the malicious agents were chosen randomly, and thus, they spread lies that more routes are highly congested. This, in turn, virtually makes different routes have the same (high and inaccurate) weights and allows the regular gossip agents to choose routes that eventually turn out to be uncongested.

Since the goal of the self-interested agents in this case is to cause disorder in the network rather than use the lies for their own benefit, the question arises as to why the behavior of the self-interested agents would be to send lies about their routes only. Furthermore, we hypothesize that if they all send lies about the same major roads, the damage they might inflict on the entire network would be larger than had each of them sent lies about its own route. To examine this hypothesis, we designed another set of experiments. In this set of experiments, all the self-interested agents spread lies about the same 13 main roads in the network. However, the results show quite a smaller impact on other gossip and regular agents in the network. The average normalized value for the gossip agents in these simulations was only about 1.07 as opposed to 1.7 in the original scenario. When analyzing the results, we revealed that although the false data were spread, they did not cause other gossip cars to change their route. The main reason was that the lies were spread on roads that were not on the route of the self-interested agents. Thus, it took the data longer to reach agents on the main roads, and when the agents reached the relevant roads, these data were "too old" to be incorporated in the other agents' calculations.

We also examined the impact of sending lies to cause chaos when there is already congestion in the network. To this end, we simulated a network in which 13 main roads are jammed. The behavior of the self-interested agents is the same as described in Section IV-A, and the self-interested agents spread lies about their own route. The simulation results, which are detailed in Table V, show that there is a greater incentive for the self-interested agents to cheat when the network is already congested, as their cheating causes more damage to other agents in the network. For example, whereas the average journey length of the regular agents increased only by about 18% in the original scenario with an uncongested network (see Table IV), in this scenario, the average journey length of the agents significantly increased: by about 60% (p-value < 0.03).

VI. MECHANISMS FOR OVERCOMING MALICIOUS AGENTS IN VEHICULAR NETWORKS

In the previous section, we demonstrated the effects of the malicious agents on other agents, i.e., mostly gossip agents, in the network. Although the effect is relatively low, it still increases the average journey length that is incurred by the other gossip agents. Therefore, we proceeded to implement two mechanisms to show how they can significantly reduce the influence of the malicious or self-interested agents in the network. Unlike mechanisms of distributed reputation, our proposed mechanisms are neither costly nor time consuming. The first mechanism we propose is mainly incorporated into the agents themselves. A history of the roads is maintained and used to update the belief regarding each road. The second mechanism is implemented in the network with the introduction of trusted agents in the network. For example, ambulances or police cars (agents) are flagged, and their data are always assumed to be true. Thus, each agent can use these data as a reference to evaluate the data on each road. We elaborate on these mechanisms below.

When implementing mechanisms to overcome the effects of malicious agents, we should take into consideration the special dynamics and characteristics of transportation networks. Since the communication range is limited, there is a bound on the amount of information that two cars can exchange. A complex mechanism would turn out to be costly, as well as inefficient, since it would significantly reduce the data that are exchanged on road conditions. Even if we attempt to incorporate only a simple mechanism of distributed reputation, the tradeoff between communicating reputation and data exists.

To this end, we began by implementing two mechanisms, and, using these simulations, we show their efficacy in significantly decreasing the effects of malicious agents on other agents in the system. For both mechanisms, we characterize the data about a given road as having a *true value* (e.g., an agent that is gathering data about a road as he/she traverses it will characterize these data as being true for his/her local evaluations) or as having an *unknown value* (e.g., the data received from other agents, even if they are characterized as true in their local evaluations).

A. Maintaining a History

In this mechanism, a history is maintained for each road. Each agent maintains a constant size array of values per road (history) and uses these values to update its belief regarding the road load. We continue with a description of how the history is updated and how a belief about the road load is updated.

1) Description of the Mechanism: When receiving new data, the agent can distinguish between two cases. First, when the history array is not yet full, the data are simply added to the array of the given road. In the second case, when the history array is already full, the agent needs to decide whether the newly received data should override any existing data. Basically, the agent gives a higher priority to data that are known to be true over other data. A major difference between this algorithm and our initial version, which is described in Section III, is that the agent distinguishes between the data items it collects itself while traversing the road and the data items that are received from other agents. Since the data that are collected by the agent are characterized as having a true value, whereas the other data are not, its own data receive a higher priority, even if newer data about the same roads are received. This allows the agents to be more selective when updating their history. Let t_{recv} be the time of the newly received data. Specifically, the agent needs to distinguish between two possibilities. t_{recv} can override the data in the history only if they are either *more* recent than any data in the history or within a given time threshold from the oldest data. If this is the case, the new data will override the existing data either if the new data are characterized as having a *true value* or if the data in the agent's history are not characterized as having a *true value*. In addition, the history is maintained per road, and there cannot be more than one data item per road's history that was generated from the same agent. This is to protect against malicious agents that are aware of the fact that the gossip agents maintain a history and try to manipulate it to their advantage by bombarding them with misinformation regarding the same road.

Another important decision when using the history mechanism is which of the data in the history should be used—both for gossiping purposes and for local calculations. If any of the data of the history are characterized as having a true value, then these data are used (if there are several items in the history of the road having a true value, then the most recent one is chosen). If all the data in the history are characterized as having an unknown value, then an average of the road's load is calculated. Then, the data item in the history that is closest to the average load is chosen as the believed data about the road.

2) When Maintaining the History Is Inefficient: In Section IV-C, we proved that there is an equilibrium in which gossiping is inefficient when no countermeasures are implemented against the malicious agents. We will demonstrate now that gossiping is inefficient when maintaining a history as well. To do so, we model our scenario as a game to find the equilibrium. Two possible types of agents participate in the game—regular gossip agents and malicious agents. Each of these agents is a representative of its group, and thus, all agents in the same group have similar behavior. The gossip agents can choose the size of the history that they maintain, whereas the malicious agents can choose the size of the coalition that they form to try to manipulate the entire history so that it will consist of only false data. If the coalition size is larger than the history that is maintained by the gossip agents, then the malicious agents can gain control over the history. In this case, the malicious agents gain utility, whereas the gossip agents lose. However, the larger the coalition's size, the larger the overhead and the coordination that are required by the malicious agent. Thus, the larger the coalition is, the lower the utility value they gain. Similar considerations apply to the gossip agents. If the history size is larger than the coalition's size, then the gossip agents can use the history to minimize the effects of malicious agents, and they gain a higher utility value. On the other hand, the larger the history size is, the more computation required by the agents; thus, they gain lower utility values. Given these considerations, we can generate the payoff matrix, shown at the bottom of the page, in which the rows represent the coalition size, and the columns represent the history size.

From the payoff matrix, we can observe that as the coalition size (history size) increases, the utility value of the malicious agents (gossip agents) decreases. In addition, whenever the coalition size (history size) is larger than the history size (coalition size), the utility value of the malicious agents (gossip agents) is positive, whereas the utility value of the gossip agents (malicious agents) is negative. Furthermore, the highest utility of the malicious agents (gossip agents) is gained when the coalition size (history size) is minimal yet larger than the history size (coalition size), i.e., a coalition size (history size) of two and a history size (coalition size) of one.

It is easy to see that a *Nash equilibrium* exists in which the history size and the coalition's size are of size n. Following our results in Section IV-C, in this situation, gossiping is inefficient.

B. Trusted Agents

In the second mechanism that we implemented, we assume that a subset of the gossip agents that roam the network can be characterized as trusted agents. This can be modeled, for example, by ambulances or police cars, which are known to be trustworthy and have no incentive to spread misinformation. Data that are received from the trusted agents are always presumed to have a *true value* and, thus, receive a higher priority when updating the data about the road. The updating of the history (whether there is no history, i.e., the history size is one, or the history size is larger than one) and the generation of the belief about the roads are similar to the algorithms described above. Note that we assume that the network infrastructure

$$
\left(
\begin{array}{c|cccccc}
 & 1 & 2 & \dots & n-1 & n \\
\hline
1 & (n-1,n-1) & (-1,n) & \dots & (-1,3) & (-1,2) \\
2 & (n,-1) & (n-2,n-2) & \dots & (-2,3) & (-2,2) \\
3 & (n-1,-1) & (n-1,-2) & \dots & (-3,3) & (-3,2) \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
n-1 & (3,-1) & (3,-2) & \dots & (1,1) & (-(n-1),2) \\
n & (2,-1) & (2,-2) & \dots & (2,-(n-1)) & (0,0)
\end{array}
\right)
$$

TABLE VI
NORMALIZED JOURNEY LENGTH VALUES WHEN SIX SELF-INTERESTED AGENTS, WITH THE SAME ORIGIN AND DESTINATION, SPREAD LIES ABOUT THEIR ROUTE; ONE ROAD, ON THE ROUTE OF THESE AGENTS, WAS PARTIALLY BLOCKED. GOSSIP AGENTS ARE IMPLEMENTED WITH THE HISTORY MECHANISM ONLY (HISTORY SIZE OF THREE)

Iteration Number	Self-Interested Agents	Gossip - Same	Gossip - Others	Regular Agents
1	1.02	1.10	1.03	1.06
2	1.02	1.04	1.01	1.10
3	1.09	1.07	1.04	1.12
4	1.07	1.02	1.01	1.09
5	0.99	1.02	1.01	1.07
6	1.05	1.03	1.02	1.06

TABLE VII
NORMALIZED JOURNEY LENGTH VALUES WHEN SIX SELF-INTERESTED AGENTS, WITH THE SAME ORIGIN AND DESTINATION, SPREAD LIES ABOUT THEIR ROUTE; ONE ROAD, ON THE ROUTE OF THESE AGENTS, WAS PARTIALLY BLOCKED. GOSSIP AGENTS ARE IMPLEMENTED WITH BOTH A HISTORY MECHANISM (HISTORY SIZE OF THREE) AND 1% OF TRUSTED AGENTS

Iteration Number	Self-Interested Agents	Gossip - Same	Gossip - Others	Regular Agents
1	1.10	1.07	1.03	1.05
2	1.03	1.04	1.01	1.10
3	1.04	1.04	1.04	1.12
4	0.93	0.97	1.00	1.10
5	1.01	1.01	1.01	1.08
6	1.01	1.02	1.01	1.07

supports this mechanism. That is, it provides a way to detect the messages of trusted agents and prevents other agents from disguising themselves as trusted agents (for example, using private and public key encryptions). Section VI-C describes our simulation results using both mechanisms.

C. Simulation Results

We ran two sets of experiments. In each set, we implemented our mechanisms to decrease the effect that is caused by the malicious agents. In both experiments, the history size was set at three. In one set, no trusted agents were present, whereas in the other, 1% of the gossip agents (approximately 80 agents) were trusted agents. We believe that there would not be a higher proportion of trusted agents in real settings.

In the first set of experiments, we created a scenario in which a small group of self-interested agents spreads lies about the same route and tested their effect on the journey length of all the agents in the network while implementing our mechanisms to overcome their effect. Thus, several cars, which had the same origin and destination points, were designated as self-interested agents. We selected only six self-interested agents in an attempt to investigate the effect that is achieved by only a small number of agents.

In each simulation in this experiment, six different self-interested agents were chosen randomly. In addition, one road on the route of these agents was randomly selected to be partially blocked, allowing only one car to go through at each time step. About 8000 agents were randomly selected as regular gossip agents, and the other 32 000 agents were designated as regular agents. When implementing the trusted agent mechanism, a random number of 80 agents of the 8000 gossip agents were randomly selected to act as trusted agents.

We analyzed the average journey length of the self-interested agents as opposed to the average journey length of other regular gossip agents traveling along the same route. Tables VI and VII summarize the normalized results for the self-interested agents, the gossip agents, and the regular agents as a function of the iteration number. The two tables list the results when the history size was three without trusted agents and with 1% trusted agents, respectively. These results can be compared with Table I, in which neither of the two mechanisms to overcome the malicious agents was implemented.

The results clearly illustrate the benefit of implementing the history mechanism. For example, in the last iteration, when neither of the two mechanisms was implemented, the gossip agents with the same original route as the malicious agents doubled their journey length (a normalized value of 2.02). However, when the history mechanism was implemented, the effect on the gossip agents significantly decreased to a normalized value of just 1.03 in the last iteration. These results reveal that maintaining a history helps to minimize the effects of the malicious agents. This can be attributed to two main reasons. The first is that true data are given a priority. Thus, even if several malicious agents spread data on the same road, the false data cannot override the true data that exist about the road. The second reason is the fact that an agent can only attribute one instance to the history of a given road. Thus, a malicious agent cannot aggregate the data and fill the history of a given road with its own misinformation.

Adding the trusted agents' mechanism together with the history mechanism does not help the gossip agents to further decrease their journey length, which has already significantly decreased due to the use of the history mechanism. To clarify this, we also ran experiments (which are not presented in this paper) in which the history was set to one, and no trusted agents existed. In these experiments, the results also revealed that our new history update mechanism enables a significant decrease in the effects caused by the malicious agent, and thus, the benefit of the trusted agents in the system is minimized.

In the second set of experiments, we tested the effect of our mechanisms when the malicious agents aim to cause disorder in the network. This can be achieved, for example, by malicious agents causing an increase in the average journey length of all agents, even at the cost of increasing their own journey length. We ran two sets of simulations: In the first set, 32 malicious agents were present, and in the second set, 100 malicious agents were present. The malicious agents spread lies about the same 13 main roads in the network. Table VIII is a snapshot of Table IV, which summarizes the average results of all size iterations when no mechanism is used, whereas Tables IX and X summarize the average results of all six iterations with a history

TABLE VIII
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS. THIRTY-TWO AND 100 SELF-INTERESTED AGENTS WITH THE OBJECTIVE OF MINIMIZING THEIR AVERAGE JOURNEY LENGTH. NO OVERCOMING MECHANISM WAS IMPLEMENTED

Malicious Agents Number	Malicious Agents	Gossip Agents	Regular Agents
32	1.06	2.13	1.25
100	1.26	2.10	1.27

TABLE IX
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS. THIRTY-TWO AND 100 SELF-INTERESTED AGENTS WITH THE OBJECTIVE OF MINIMIZING THEIR AVERAGE JOURNEY LENGTH. GOSSIP AGENTS ARE IMPLEMENTED WITH THE HISTORY MECHANISM ONLY (HISTORY OF SIZES ONE AND THREE)

History	Malicious Agents Number	Malicious Agents	Gossip Agents	Regular Agents
H = 1	32	1.01	1.03	1.05
H = 3		1.00	1.00	1.06
H = 1	100	1.01	1.04	1.05
H = 3		1.00	1.00	1.05

TABLE X
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS. THIRTY-TWO AND 100 SELF-INTERESTED AGENTS WITH THE OBJECTIVE OF MINIMIZING THEIR AVERAGE JOURNEY LENGTH. GOSSIP AGENTS ARE IMPLEMENTED WITH BOTH THE HISTORY MECHANISM (HISTORY OF SIZES ONE AND THREE) AND 1% OF TRUSTED AGENTS

History	Malicious Agents Number	Malicious Agents	Gossip Agents	Regular Agents
H = 1	32	1.00	1.04	1.05
H = 3		1.00	1.00	1.05
H = 1	100	1.02	1.04	1.05
H = 3		1.00	1.01	1.06

of size 1 $(H = 1)$ and a history of size 3 $(H = 3)$ when only the history mechanism is implemented and when both the history mechanism and the trusted agents' mechanisms are implemented, respectively.

Again, in this experiment as well, we can see the significant decrease in the journey length for the gossip agents due to the incorporation of the history mechanism. We can also see that the addition of the trusted agents' mechanism when the history mechanism is already implemented has no significant effect on the results.

VII. COALITIONS OF MALICIOUS AGENTS

In the previous section, we demonstrated how the history mechanism allows the gossip agents to minimize the effect of the malicious agents. The question arises as to what will happen if the malicious agents are aware of the protection method that is implemented by the gossip agents. Can the malicious agents manipulate this mechanism to their own benefit?

In Section VI-A2, we have shown that gossiping is inefficient under some assumptions of maintaining a history and a coalition formation by the malicious agents. In this section, we examine whether the coalition formation by the malicious agent can also assist the malicious agents in increasing their effect on the gossip agents in the network, while the gossip agents maintain a history mechanism. The main goal is to check whether the malicious agents can form coalitions that will enable them to take control of the different roads upon which they spread false data and, thus, make the gossip agents believe that the actual road load is the false one.

TABLE XI
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS, WITH A HISTORY OF SIZE ONE AND A COALITION OF SIZE TWO

Number of Malicious Agents	Gossip Agents	Regular Agents
32	1.03	1.05
100	1.04	1.05

TABLE XII
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS, WITH A HISTORY OF SIZE THREE AND A COALITION OF SIZE TWO

Malicious Agents Number	Gossip Agents	Regular Agents
32	1.00	1.06
100	1.02	1.06

TABLE XIII
NORMALIZED JOURNEY LENGTH VALUES FOR ALL ITERATIONS, WITH A HISTORY OF SIZE THREE AND A COALITION OF SIZE FOUR

Malicious Agents Number	Gossip Agents	Regular Agents
32	1.02	1.06
100	1.01	1.06

To test this, we ran two sets of experiments. In each experiment, the gossip agents used the history mechanism as a mechanism to decrease the effect that is caused by the malicious agents. In addition, two runs were made in each experiment. The first consisted of 32 malicious agents being present in the network, and the second consisted of 100 malicious agents. The malicious agents were randomly selected and followed the same strategy—spreading lies about the same 13 main roads in the network. We defined a coalition of K cars to be a set of K agents that have the same route (same source and destination nodes) and enter the network at approximately the same time. For example, if the coalition size is set to four, and the network consists of 100 malicious agents, then they form 25 different coalition groups.

In the first set of experiments, the malicious agents were grouped into coalitions of size two, and we conducted two simulations. In the first, the history size of the gossip agents was set at one, whereas in the second simulation, it was set at three. This allowed us to examine the effect of a coalition of size two, both when the history size is smaller than the coalition size and when it is larger than the coalition size. In the second set of experiments, the malicious agents were grouped into coalitions of size four, and we had a single simulation in which the history size was set at three. Tables XI and XII summarize the average results of all six iterations of the first experiment, whereas Table XIII summarizes the results of the second experiment. Note also that in all of the results, the standard deviation was lower than 0.002. Since the goal of the malicious agents is to cause chaos in the network and not minimize their own journey length, we omit the results concerning the malicious agents themselves. The results of the previous experiments in which no coalitions were formed are presented in Table IX.

When we observe the normalized journey length of the regular gossip agents and the regular agents (a maximal increase of 2% and 6%, respectively), we can deduce that the coalition formation did not help the malicious agents in achieving disorder in the network. One reason for this could be the way the coalition was formed and the way the history is updated. The coalition is formed by grouping malicious agents that are traversing the same route at about the same time. However, the malicious agents themselves do not spread false data about the roads they traverse but, rather, about 13 main junctions in the network. We hypothesized that by going the same route, the coalition will be able to take control of the history of other gossip agents on that route. Yet, it seems that the way in which the history is updated proffers no advantage to the coalition groups. Although the malicious agents in the coalition *can* gain monetary control over the history, if the gossip agents receive new data regarding the same roads, they will override the false data. The chances of agents, on the route of the 13 main junctions in the network, receiving other data about these roads are quite high, as it takes time until the malicious data are propagated to them, and, in addition, when they are propagated, only one instance of the data is communicated, and the history list can recover quickly. Simulating coalitions that spread false data regarding their own route is similar to the results presented in Section IV-B in which six self-interested agents spread lies regarding their own route. Table I summarizes the results, which, indeed, reveal how the self-interested agents can benefit from the lies, while causing harm to other gossip agents in the network, i.e., mainly the gossip agents on the same route as the self-interested agents. Based on the latter experiments, it seems that implementing the history mechanism will significantly decrease the harm that is inflicted by the self-interested agents in that scenario.

VIII. Discussion and Conclusion

In this paper, we have investigated the benefits that are achieved by self-interested agents in vehicular networks and whether mechanisms can help gossip agents in overcoming malicious agents in transportation networks. Using simulations, we have investigated two behaviors that might be taken by self-interested agents: 1) trying to minimize their journey length and 2) trying to cause chaos in the network. Our simulations indicate that in reference to both behaviors, the self-interested agents have only limited success in achieving their goal, even if no countermeasures are taken. This is in contrast to the greater impact that is inflicted by self-interested agents in other domains (e.g., e-commerce). Several reasons for this are the special characteristics of vehicular networks and their dynamic nature. Although the self-interested agents spread lies, they cannot choose the agents with which they will interact. Also, by the time their lies reach other agents, they might become irrelevant, as more recent data have reached the same agents.

The importance of implementing mechanisms to overcome malicious agents cannot be overrated, as we have seen the effect of malicious agents on other agents in the network when no countermeasures are implemented. However, it is also important that these mechanisms not be costly, nor time consuming, due to the dynamic nature of the transportation network and in light of the fact that the interaction is range and bandwidth limited. Furthermore, the fact that the agents cannot choose the agents with which to interact might affect the efficacy of these mechanisms. Our simulations indicate that for both behaviors that are implemented by the malicious agents in the experiments, our mechanisms enabled gossip agents to significantly overcome the effects of malicious agents. In addition, we show that even a short history mechanism can suffice to overcome the effects of malicious agents. We also demonstrate that malicious agents cannot take advantage of the history mechanism by simply grouping into coalitions.

Motivated by the simulation results, future research in this field will focus on modeling different behaviors of self-interested agents, which might cause more damage to networks. Another direction would be to focus on the benefits of distributed reputation mechanisms in this model, as well as using this type of mechanism to penalize malicious agents.

References

[1] A. Bejan and R. Lawrence, “Peer-to-peer cooperative driving,” in *Proc. ISCIS*, Oct. 2002, pp. 259–264.
[2] Y. Shavitt and A. Shay, “Optimal routing in gossip networks,” *IEEE Trans. Veh. Technol.*, vol. 54, no. 4, pp. 1473–1487, Jul. 2005.
[3] I. Chisalita and N. Shahmehri, “A novel architecture for supporting vehicular communication,” in *Proc. VTC*, Sep. 2002, pp. 1002–1006.
[4] GM, *V2V*, 2005. [Online]. Available: http://www.gm.com/company/gmability/safety/newsissues/releases/sixthsense102405.html
[5] Honda, 2005. [Online]. Available: http://world.honda.com/news/2005/c050902.html
[6] U.S. Dept. Transp., *VII Initiative*, 2006. [Online]. Available: http://www.its.dot.gov/vii/docs/vii factsheet.pdf
[7] L. Lamport, R. Shostak, and M. Pease, “The Byzantine Generals Problem,” in *Advances in Ultra-Dependable Distributed Systems*, N. Suri, J. Walter, and M. Hugue, Eds. Los Alamitos, CA: IEEE Comput. Soc. Press, 1982.
[8] G. H. Bham and R. F. Benekohal, “A high fidelity traffic simulation model based on cellular automata and car-following concepts,” *Transp. Res., Part C, Emerg. Technol.*, vol. 12, no. 1, pp. 1–32, Feb. 2004.
[9] D. Helbing, A. Hennecke, V. Shvetsov, and M. Treiber, “Micro- and macro-simulation of freeway traffic,” *Math. Comput. Model.*, vol. 35, no. 5, pp. 517–547, Mar. 2002.
[10] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, “SUMO (simulation of urban mobility): An open-source traffic simulation,” in *Proc. 4th MESM*, Sharjah, U.A.E., Sep. 2002.
[11] S. Kraus, R. Parshani, and Y. Shavitt, “A study of gossiping in transportation networks,” *IEEE Trans. Veh. Technol.*, vol. 57, no. 4, pp. 2602–2607, Jul. 2008. DOI:10.1109/TVT.2007.912339.
[12] D. Dolev, R. Reischuk, and H. R. Strong, “Early stopping in Byzantine agreement,” *J. ACM*, vol. 37, no. 4, pp. 720–741, Oct. 1990.
[13] Y. M. Minsky and F. B. Schneider, “Tolerating malicious gossip,” *Distrib. Comput.*, vol. 16, no. 1, pp. 49–68, Feb. 2003.
[14] D. Malkhi, E. Pavlov, and Y. Sella, “Gossip with malicious parties,” School Comput. Sci. Eng., Hebrew Univ. Jerusalem, Jerusalem, Israel, Tech. Rep. 2003-9, Mar. 2003.
[15] B. Yu and P. Singh, “A social mechanism of reputation management in electronic communities,” in *Cooperative Information Agents IV—The Future of Information Agents in Cyberspace*. Berlin, Germany: Springer-Verlag, 2000.
[16] C. Leckie and R. Kotagiri, “Policies for sharing distributed probabilistic beliefs,” in *Proc. ACSC*, 2003, pp. 285–290.
[17] A. Datta, S. Quarteroni, and K. Aberer, “Autonomous gossiping: A self-organizing epidemic algorithm for selective information dissemination in mobile ad-hoc networks,” in *Proc. IC-SNW*, Jun. 2004, pp. 126–143.
[18] S. Das, A. Nandan, and G. Pau, “Spawn: A swarming protocol for vehicular ad-hoc wireless networks,” in *Proc. VANET*, 2004, pp. 93–94.

[19] N. Shibata, T. Terauchi, T. Kitani, K. Yasumoto, M. Ito, and T. Higashino, "A method for sharing traffic jam information using intervehicle communication," in *Proc. V2VCOM*, 2006, pp. 1–7.
[20] W. Wang, X.-Y. Li, and Y. Wang, "Truthful multicast routing in selfish wireless networks," in *Proc. MobiCom*, 2004, pp. 245–259.
[21] R. Parshani, "Routing in gossip networks," M.S. thesis, Bar-Ilan Univ., Ramat Gan, Israel, 2004.
[22] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–291, Mar. 1979.
[23] D. Kahneman, P. Slovic, and A. Tversky, *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, U.K.: Cambridge Univ. Press, 1982.
[24] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.

Sarit Kraus received the Ph.D. degree in computer science from The Hebrew University of Jerusalem, Jerusalem, Israel, in 1989.

She is currently a Professor of computer science with the Bar-Ilan University, Ramat-Gan, Israel, and an Adjunct Professor with the Institute for Advanced Computer Studies, University of Maryland, College Park. She has extensively worked in the following areas: the development of intelligent systems, negotiation and cooperation among agents, large-scale systems, networks, personalization, training, optimization of complex systems, machine learning, information agents, user interfaces, and decision support tools. She has published over 230 papers in leading journals and major conferences. She is the author of the book *Strategic Negotiation in Multiagent Environments* (MIT Press, 2001) and is a coauthor of the book *Heterogeneous Active Agents* (MIT Press, 2000). She is an Associate Editor of the *Artificial Intelligence Journal* and the *Annals of Mathematics and Artificial Intelligence Journal*. She is on the editorial board of the *Journal of Autonomous Agents and Multi-Agent Systems*, the *Journal of Applied Logic*, and the *Journal of Web Semantics*.

Dr. Kraus was elected as a Fellow with the Association for the Advancement of Artificial Intelligence in 2002 and the European Coordinating Committee for Artificial Intelligence in 2008. She was the recipient of the 1995 International Joint Conferences on Artificial Intelligence Computers and Thought Award (the premier award for a young AI scientist), the 2001 IBM Faculty Partnership Award, and the 2007 Association for Computing Machinery Special Interest Group on Artificial Intelligence Agents Research award. Her paper with Prof. B. Grosz received the International Foundation for Autonomous Agents and Multiagent Systems influential paper award (joint winner). She has joint projects with researchers from Harvard University, Cambridge, MA; University of Maryland, College Park; Stanford Research Institute, Menlo Park, CA; Tel-Aviv University, Ramat-Aviv, Israel; The Hebrew University of Jerusalem; University of Southern California, Los Angeles; and University of Liverpool, Liverpool, U.K.

Raz Lin (M'08) received the B.Sc. degree (*summa cum laude*) in mathematics and computer science, the M.Sc. degree (*magna cum laude*) in computer science, and the Ph.D. degree from the Bar-Ilan University, Ramat-Gan, Israel, in 2001, 2002, and 2008, respectively.

He is currently a Postdoctoral Fellow with the Department of Computer Science, Bar-Ilan University, where he investigates issues of automated negotiations, personalization, training, and learning, and where he also received a four-year President's scholarship for outstanding students.

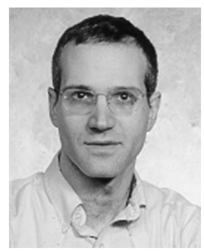

Yuval Shavitt (S'88–M'97–SM'00) received the B.Sc. degree (*cum laude*) in computer engineering, the M.Sc. degree in electrical engineering, and the D.Sc. degree from the Technion–Israel Institute of Technology, Haifa, Israel, in 1986, 1992, and 1996, respectively.

After graduation, he spent a year as a Postdoctoral Fellow with the Department of Computer Science, Johns Hopkins University, Baltimore, MD. Between 1997 and 2001, he was a member of the Technical Staff with Bell Labs, Lucent Technologies, Holmdel, NJ. Since October 2000, he has been a member of the faculty of the School of Electrical Engineering, Tel-Aviv University, Tel-Aviv, Israel. His current research focuses on Internet measurement, mapping, characterization, QoS in networks, and routing.