

On the Benefits of Cheating by Self-Interested Agents in Vehicular Networks*

Raz Lin and Sarit Kraus
Computer Science Department
Bar-Ilan University
Ramat-Gan, Israel
{linraz,sarit}@cs.biu.ac.il

Yuval Shavitt
School of Electrical Engineering
Tel-Aviv University, Israel
shavitt@eng.tau.ac.il

ABSTRACT

As more and more cars are equipped with GPS and Wi-Fi transmitters, it becomes easier to design systems that will allow cars to interact autonomously with each other, e.g., regarding traffic on the roads. Indeed, car manufacturers are already equipping their cars with such devices. Though, 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 achieved by self-interested agents if no counter-measures are implemented.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Intelligent agents*

General Terms

Experimentation

Keywords

agent-based deployed applications, artificial social systems

1. INTRODUCTION

As technology advances, more and more cars are being equipped with devices, which 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

*This work was supported in part under ISF grant number 8008.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawaii, USA.
Copyright 2007 IFAAMAS.

between cars, e.g., for locating road congestions [1] and optimal routes [15] or improving traffic safety [2].

Vehicle-To-Vehicle (V2V) communication is already onboard by some car manufactures, enabling the collaboration between different cars on the road. For example, GM's proprietary algorithm [6], called the "threat assessment algorithm", constantly calculates, in real time, other vehicles' positions and speeds, and enables messaging other cars when a collision is imminent; Also, Honda has begun testing its system in which vehicles talk with each other and with the highway system itself [7].

In this paper, we investigate the attraction of being a selfish agent in vehicular networks. That is, we investigate the benefits achieved by car owners, who tamper with on-board devices and incorporate their own self-interested agents in them, which act for their benefit. We build on the notion of *Gossip Networks*, introduced by Shavitt and Shay [15], in which the agents can obtain road congestion information by gossiping with peer agents using ad-hoc communication.

We recognize two typical behaviors that the self-interested agents could embark upon, in the context of vehicular networks. In the first behavior, described in Section 4, the objective of the self-interested agents is to maximize their own utility, expressed by their average journey duration on the road. This situation can be modeled in real life by car owners, whose aim is to reach their destination as fast as possible, and would like to have their way free of other cars. To this end they will 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, and causing them to choose a different route.

The second type of behavior, described in Section 5, 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 vandalism or terrorists, who aim to cause as much mayhem in the network as possible.

We note that the introduction of self-interested agents to the network, would most probably motivate other agents to try and detect these agents in order to minimize their effect. This is similar, though in a different context, to the problem introduced by Lamport *et al.* [8] as the *Byzantine Generals Problem*. However, the introduction of mechanisms to deal with self-interested agents is costly and time consuming. In this paper we focus mainly on the attractiveness of selfish behavior by these agents, while we also provide some insights

into the possibility of detecting self-interested agents and minimizing their effect.

To demonstrate the benefits achieved by self-interested agents, we have used a simulation environment, which models the transportation network in a central part of a large real city. The simulation environment is further described in Section 3. Our simulations provide insights to the benefits of self-interested agents cheating. Our findings can motivate future research in this field in order to minimize the effect of selfish-agents.

The rest of this paper is organized as follows. In Section 2 we review related work in the field of self-interested agents and V2V communications. We continue and formally describe our environment and simulation settings in Section 3. Sections 4 and 5 describe the different behaviors of the self-interested agents and our findings. Finally, we conclude the paper with open questions and future research directions.

2. RELATED WORK

In their seminal paper, Lamport *et al.* [8] 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 cannot send a message to all the other agents. Dolev *et al.* [5] has built on this problem and has analyzed the number of faulty agents that can be tolerated in order to eventually reach the right conclusion about true data. Similar work is presented by Minsky *et al.* [11], who discuss techniques for constructing gossip protocols that are resilient to up to t malicious host failures. As opposed to the above works, our work focuses on vehicular networks, in which the agents are constantly roaming the network and exchanging data. Also, the domain of transportation networks introduces dynamic data, as the load of the roads is subject to change. In addition, the system in transportation networks has a feedback mechanism, since the load in the roads depends on the reports and the movement of the agents themselves.

Malkhi *et al.* [10] present a gossip algorithm for propagating information in a network of processors, in the presence of malicious parties. Their algorithm prevents the spreading 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 is static and also that the agents are static (they discuss a network of processors). This is not true for transportation networks. For example, in our model, agents might gossip about heavy traffic load of a specific road, which is currently jammed, yet this information might be false several minutes later, leaving the agents to speculate whether the spreading agents are indeed malicious or not. In addition, as the agents are constantly moving, each agent cannot choose with whom he interacts and exchanges data.

In the context of analyzing the data and deciding whether the data is true or not, researchers have focused on distributed reputation systems or decision mechanisms to decide whether or not to share data.

Yu and Singh [18] build a social network of agents' reputations. 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 obtained from reliable referral chains. After a bad experi-

ence 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 update their ratings accordingly. This approach might be implemented in our domain to allow gossip 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. This is mainly due to the dynamic nature of the list of neighbors in transportation networks. Thus, not only does it require maintaining the neighbors' list, since the neighbors change frequently, but it is also harder to build a good reputation system.

Leckie *et al.* [9] focus on the issue of when to share information between the agents in the network. Their domain involves monitoring 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 sufficient likely he 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 might also vary (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 cheating.

Some work has been done in the context of gossip networks or transportation networks regarding the spreading of data and its dissemination.

Datta *et al.* [4] focus on information dissemination in mobile ad-hoc networks (MANET). 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 which node is interested in the information is made by the data item itself, using heuristics. However, their work concentrates on the movement of the data itself, 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 the agent itself roams the road and is responsible (and has the capabilities) for spreading the data to other agents in the network.

Das *et al.* [3] propose a cooperative strategy for content delivery in vehicular networks. In their domain, peers download a file from a mesh and exchange pieces 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 among themselves.

Shibata *et al.* [16] propose a method for cars to cooperatively and autonomously collect traffic jam statistics to estimate arrival time to destinations for each car. The communication is based on IEEE 802.11, without using a fixed infrastructure on the ground. While we use the same domain, we focus on a different problem. Shibata *et al.* [16] mainly focus on efficiently broadcasting the data between agents (e.g., avoid duplicates and communication overhead), as we focus on the case where agents are not cooperative in

nature, and on how selfish agents affect other agents and the network load.

Wang *et al.* [17] also assert, in the context of wireless networks, that individual agents are likely to do what is most beneficial for their owners, and will act selfishly. 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 (a mechanism of *incentive compatibility*). However, the domain of wireless networks is quite different from the domain of transportation networks. In the wireless network, the wireless terminal is required to contribute its local resources to transmit data. Thus, Wang *et al.* [17] 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. Unlike this, 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 regarding the credibility of the data they exchange.

In the next section, we will describe our transportation network model and gossiping between the agents. We will also describe the different agents in our system.

3. MODEL AND SIMULATIONS

We first describe the formal transportation network model, and then we describe the simulations designs.

3.1 Formal Model

Following Shavitt and Shay [15] and Parshani [13], 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 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 points (vertices). A *journey* is defined as the (not necessarily simple) path 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 this path. Every agent has to travel between its origin and destination points and aims to minimize its journey length.

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 inter-vehicle 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: (a) between gossip agents passing one another, or (b) between gossip agents located at the same junction. We assume that each agent stores the most recent information it has received or gathered around the edges in the network.

A subset of the gossip agents are those agents who 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 4 and 5.

3.2 Simulation Design

Building on [13], the network in our simulations replicates a central part of a large city, and consists of 50 junctions and 150 roads, which are approximately the number of main streets in the city. Each simulation consists of 6 iterations. The basic time unit of the iteration is a step, which equivalents to about 30 seconds. Each iteration simulates six hours of movements. The average number of cars passing through the network during the iteration is about 70,000 and the average number of cars in the network at a specific time unit is about 3,500 cars. In each iteration the same agents are used with the same origin and destination points, whereas the data collected in earlier iterations is preserved in the future iterations (referred to as the history of the agent). This allows us to simulate somewhat a daily routine in the transportation network (e.g., a working week).

Each of the experiments that we describe below is run with 5 different traffic scenarios. Each such traffic scenario differs from one another by the initial load of the roads and the designated routes of the agents (cars) in the network. For each such scenario 5 simulations are run, creating a total of 25 simulations for each experiment.

It has been shown by Parshani *et al.* [13, 14] that the information propagation in the network is very efficient when the percentage of gossiping agents is 10% or more. Yet, due to congestion caused by too many cars rushing to what is reported as the less congested part of the network 20-30% of gossiping agents leads to the most efficient routing results in their experiments. Thus, in our simulation, we focus only on simulations in which the percentage of gossip agents is 20%.

The simulations were done with different percentages of self-interested agents. To gain statistical significance we ran each simulation with changes in the set of the gossip agents, and the set of the self-interested agents.

In order to gain a similar ordinal scale, the results were normalized. The normalized values were calculated by comparing each agent's result to his results when the same scenario was run with no self-interested agents. This was done for all of the iterations. Using the normalized values enabled us to see how worse (or better) each agent would perform compared to the basic setting. For example, if an average 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 be $60/50 = 1.2$.

More details regarding the simulations are described in Sections 4 and 5.

4. SPREADING LIES, MAXIMIZING UTILITY

In the first set of experiments we investigated the benefits achieved by the self-interested agents, whose aim was to minimize their own journey length. The self-interested agents adopted a cheating approach, in which 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.

4.1 Modeling the Self-Interested Agents' Behavior

While the gossip agents gather data and send it to other agents, the self-interested agents' behavior is modeled as follows:

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

Basically, the self-interested agent acts the same as the gossip agent. It collects data regarding the weight of the roads (either by traversing the road or by getting 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 in 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 in its route are heavily congested. This is in order to free the path for itself, by making other agents recalculate their paths, this time without including 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 in its route and sends them to the other agents.

While an agent can also divert cars from its route by falsely reporting congested roads in parallel to its 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 congestion on 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 its own lies and other lies spreading in the network, all self-interested agents will ignore data received about roads with heavy traffic (note that data about roads that are not heavily traffic will not be ignored)¹.

In the next subsection we describe the simulation results, involving the self-interested agents.

4.2 Simulation Results

To test the benefits of cheating by the self-interested agents we ran several experiments. In the first set of experiments, we created a scenario, in which a small group of self-interested agents spread lies on the same route, and tested its effect on the journey length of all the agents in the network.

¹In other simulations we have run, in which there had been several real congestions in the network, we indeed saw that even when the roads are jammed, the self-interested agents were less affected if they ignored all reported heavy traffic, since by such they also discarded all lies roaming the network

Table 1: Normalized journey length values, self-interested agents with the same route

Iteration Number	Self-Interested Agents	Gossip - SR	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	2.93	1.35	1.16
5	1.13	2.00	1.40	1.17
6	1.08	2.02	1.43	1.18

Thus, several cars, which had the same origin and destination points, were designated as self-interested agents. In this simulation, we selected only 6 agents to be part of the group of the self-interested agents, as we wanted to investigate the effect achieved by only a small number of agents.

In each simulation in this experiment, 6 different agents were randomly chosen to be part of the group of self-interested agents, 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 that road at each time step. About 8,000 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 as opposed to the average journey length of other regular gossip agents traveling along the same route. Table 1 summarizes the normalized results for the self-interested agents, the gossip agents (those having the same origin and destination points as the self-interested agents, denoted *Gossip - SR*, and all other gossip agents, denoted *Gossip - Others*) and the regular agents, as a function of the iteration number.

We can see from the results that the first time the self-interested agents traveled the route while spreading the false data about the roads did not help them (using the paired t-test we show that those agents had significantly lower journey lengths in the scenario in which they did not spread any lies, with $p < 0.01$). This is mainly due to the fact that the lies do not bypass 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 he is about to travel, in the first iteration.

Only when the self-interested agents had repeated their journey in the next iteration (iteration 2) did it help them significantly ($p = 0.04$). The reason for this is that other gossip agents received this data and used it to recalculate their shortest path, thus avoiding entrance to the roads, for which the self-interested agents had spread false information about congestion. It is also interesting to note the large value attained by the self-interested agents in the first iteration. This is mainly due to several self-interested agents, who entered the jammed road. This situation occurred since the self-interested agents ignored all heavy traffic data, and thus ignored the fact that the road was jammed. As they started spreading lies about this road, more cars shifted from this route, thus making the road free for the future iterations.

However, we also recall that the self-interested agents ignore all information about the heavy traffic roads. Thus,

Table 2: Normalized journey length values, spreading lies for a beneficiary agent

Iteration Number	Beneficiary Agent	Gossip - SR	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	1.14
5	1.05	1.32	1.05	1.12
6	0.92	1.40	1.06	1.11

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 increased above 1.00. Using the paired t-test to compare these values with the values achieved by these agents when no lies are used, we see that there is no significant difference between the two scenarios.

As opposed to the gossip agents, we can see how little effect the self-interested agents have on the regular agents. As compared to the gossip agents on the same route that have traveled as much as 193% more, when self-interested agents are introduced, the average journey length for the regular agents has only increased by about 15%. This result is even lower than the effect on other gossip agents in the entire network.

Since we noticed that cheating by the self-interested agents does not benefit them in the first iteration, we devised another set of experiments. In the second set of experiments, the self-interested agents have the objective to help another agent, who is supposed to enter the network some time after the self-interested agent entered. We refer to the latter agent as the beneficiary agent. Just like a self-interested agent, the beneficiary agent also ignores all data regarding 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 2 summarizes the normalized values for the different agents. As in the first set of experiments, 5 simulations were run for each scenario, with a total of 25 simulations. In each of these simulation 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 8,000 and 32,000 agents were designated as regular gossip agents and regular agents, respectively.

We can see that as the number of iterations advances, the lower the normalized value for the beneficiary agent. In this scenario, just like the previous one, in the first iterations, not only does the beneficiary agent not avoid the jammed roads, since he ignores all heavy traffic, he also does not benefit from the lies spread by the self-interested agent. This is due to the fact that the lies are not yet incorporated by other gossip agents. Thus, if we compare the average journey length in the first iteration when lies are spread and when there are no lies, the average is significantly lower when there are no lies ($p < 0.03$). On the other hand, if we compare the average journey length in all of the iterations, there is no significant difference between the two settings. Still, 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, virtually are 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 just one self-interested car, we can see that both the gossip agents that follow the same route as the lies spread by the self-interested agents, and other regular agents, increase their journey length by more than 14%.

In our third set of experiments we examined a setting in which there was an increasing number of agents, and the agents did 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 the cars were repeating their journeys (that is, more than one iteration was made). 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 that they will all have the same origin and destination points.

As in the previous sets of experiments we ran 5 different simulations per scenario. In each simulation 11 runs were made, each run with different numbers of self-interested agents: 0 (no self-interested agents), 1, 2, 4, 8, and 16. Each agent adopted the behavior modeled in Section 4.1. Figure 1 shows the normalized value 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. Also, we have seen in the previous set of experiments and we have provided explanations as to why the self-interested agents do not gain much from their behavior in the first iteration.

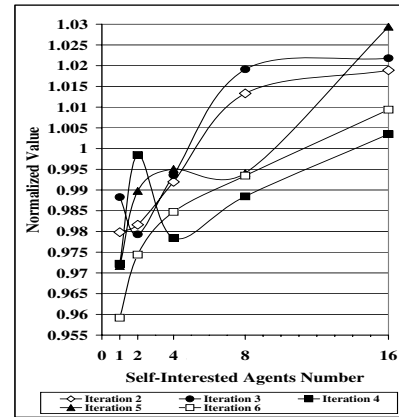


Figure 1: Self-interested agents normalized values as a function of the number of self-interested agents.

Using these simulations we examined what the threshold could be for the number of randomly selected self-interested agents in order to allow themselves to benefit from their selfish behavior. We can see that up to 8 self-interested agents, the average normalized value is below 1. That is, they benefit from their malicious behavior. In the case of one self-interested agent there is a significant difference between the average journey length of when the agent spread lies and when no lies are spread ($p < 0.001$), while when

there are 2, 4, 8 and 16 self-interested agents there is no significance 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 the next subsection we analyze the scenario as a game and show that when in equilibrium the exchange of gossiping between the agents becomes inefficient.

4.3 When Gossiping is Inefficient

We continued and modeled our scenario as a game, in order to find the equilibrium. There are two possible types for the agents: (a) regular gossip agents, and (b) 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 to them by recalculating their routes [14]. 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 4.1. We also 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 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, yet it may currently be jammed due to heavy traffic on the road).

Let T_{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 takes to traverse an edge. We will investigate the game, in which the self-interested and the 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 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 the self-interested agents spread as T_s , such that $T_{avg} \leq T_s \leq T_{max}$. Motivated by the results of the simulations we have described above, we saw 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 above a certain threshold, T_g . Obviously, $T_{avg} \leq T_g \leq T_{max}$. In order to prevent the gossip agents from detecting the lies and just discarding those values, the self-interested agents send lies in a given range, $[T_s, T_{max}]$, 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 agents, which is defined by the values of T_s and T_g . If the self-interested agents spread traversal times higher than or equal to the regular gossip cars' threshold, they will not benefit from those lies. Thus, the utility value of the self-interested agents in this case is 0. On the other hand, if the self-interested agents spread traversal time which is lower

than the threshold, they will gain a positive utility value. From the regular gossip agents point-of-view, if they accept messages from the self-interested agents, then they incorporate the lies in their calculation, thus they will lose utility points. On the other hand, if they discard the false values the self-interested agents send, that is, 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:

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

$$u^g = \begin{cases} T_g - T_{avg} & \text{if } T_s \geq T_g \\ T_s - T_g & \text{if } T_s < T_g \end{cases} \quad (2)$$

We are interested in finding the *Nash equilibrium*. We recall from [12], that the Nash equilibrium is a strategy profile, where no player has anything to gain by deviating from his strategy, given that the other agent follows his 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 $i \in \{\text{regular gossip, self-interested}\}$ 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, that is, if for all i , $u^i(T^*) \geq u^i(T_i, T_{-i}^*)$. That is, (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 have the following theorem.

THEOREM 4.1. (T_{avg}, T_{avg}) is the only Nash equilibrium.

Proof. First we will show that (T_{avg}, T_{avg}) is a Nash equilibrium. Assume, by contradiction, that the gossip agents choose another value $T_{g'} > T_{avg}$. Thus, $u^g(T_{avg}, T_{g'}) = T_{avg} - T_{g'} < 0$. On the other hand, $u^g(T_{avg}, T_{avg}) = 0$. 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 $T_{s'} > T_{avg}$. Thus, $u^s(T_{s'}, T_{avg}) = 0$, while $u^s(T_{avg}, T_{avg}) = 0$.

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

We have three cases. In the first $T_{avg} < T_g < T_s \leq T_{max}$. Thus, $u^s(T_s, T_g) = 0$ and $u^g(T_s, T_g) = T_g - T_{avg}$. 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: $u^g(T_s, T_g + 1) = T_g + 1 - T_{avg}$.

In the second case we have $T_{avg} < T_s < T_g \leq T_{max}$. 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 we have $T_{avg} < T_s = T_g \leq T_{max}$. 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: $u^s(T_g - 1, T_g) = T_g - 1 - T_{avg} + 1 = T_g - T_{avg} > 0$. ■

Table 3: Normalized journey length values for the first iteration

Self-Interested Agents Number	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	1.1
64	1.14	1.4	1.13
80	1.15	1.5	1.14
100	1.17	1.63	1.16

Table 4: Normalized journey length values for all iterations

Self-Interested Agents Number	Self-Interested Agents	Gossip Agents	Regular Agents
1	0.98	1.02	1.06
2	1.0	1.04	1.07
4	1.0	1.08	1.07
8	1.01	1.33	1.11
16	1.02	1.89	1.17
32	1.06	2.46	1.25
50	1.13	2.24	1.29
64	1.21	2.2	1.32
80	1.21	2.13	1.27
100	1.26	2.11	1.27

The above theorem proves that the equilibrium point is reached only when the self-interested agents send the time to traverse certain edges equals 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 any anomalies in the network.

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

5. 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 5 different simulations for each scenario. In each simulation different agents were randomly chosen (using a uniform distribution) to act as gossip agents, among them self-interested agents were chosen. Each self-interested agent behaved in the same manner as described in Section 4.1.

Every simulation consisted of 11 runs with each run comprising different numbers of self-interested agents: 0 (no self-interested agents), 1, 2, 4, 8, 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.

Tables 3 and 4 summarize the normalized journey length for the self-interested agents, the regular gossip agents and the regular (non-gossip) agents. Table 3 summarizes the data for the first iteration and Table 4 summarizes the data for the average of all iterations. Figure 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 4.2, we can see that randomly selected self-interested agents who 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 146% higher than in the setting with no self-interested agents at all, as opposed to an increase of only 25% for the regular agents). In contrast, these results also indicate that the self-interested agents do not succeed in causing a significant load in the network by their malicious behavior.




Figure 2: Gossip and regular agents normalized values, as a function of the iteration.

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 benefits, the question arises as to why would the behavior of the self-interested agents 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 regu-

Table 5: Normalized journey length values for all iterations. Network with congestions.

Self-Interested Agents Number	Self-Interested Agents	Gossip Agents	Regular Agents
1	1.04	1.02	1.22
2	1.06	1.04	1.22
4	1.04	1.06	1.23
8	1.07	1.15	1.26
16	1.09	1.55	1.39
32	1.12	2.25	1.56
50	1.24	2.25	1.60
64	1.28	2.47	1.63
80	1.50	2.41	1.64
100	1.69	2.61	1.75

lar 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 saw that although the false data was spread, it 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 this data was "too old" to be incorporated in the other agents calculations.

We also examined the impact of sending lies in order to cause chaos when there are already congestions 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 as described in Section 4.1, and the self-interested agents spread lies about their own route. The simulation results, detailed in Table 5, 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 the other agents in the network. For example, whereas the average journey length of the regular agents increased only by about 15% in the original scenario, in which the network was not congested, in this scenario the average journey length of the agents had increased by about 60%.

6. CONCLUSIONS

In this paper we investigated the benefits achieved by self-interested agents in vehicular networks. Using simulations we investigated two behaviors that might be taken by self-interested agents: (a) trying to minimize their journey length, and (b) trying to cause chaos in the network. Our simulations indicate that in both behaviors the self-interested agents have only limited success achieving their goal, even if no counter-measures are taken. This is in contrast to the greater impact inflicted by self-interested agents in other domains (e.g., E-Commerce). Some reasons for this are the special characteristics of vehicular networks and their dynamic nature. While the self-interested agents spread lies, they cannot choose which agents with whom they will interact. Also, by the time their lies reach other agents, they might become irrelevant, as more recent data has reached the same agents.

Motivated by the simulation results, future research in this field will focus on modeling different behaviors of the self-interested agents, which might cause more damage to the network. Another research direction would be to find

ways of minimizing the effect of selfish-agents by using distributed reputation or other measures.

7. REFERENCES

- [1] A. Bejan and R. Lawrence. Peer-to-peer cooperative driving. In *Proceedings of ISCSIS*, pages 259–264, Orlando, USA, October 2002.
- [2] I. Chisalita and N. Shahmehri. A novel architecture for supporting vehicular communication. In *Proceedings of VTC*, pages 1002–1006, Canada, September 2002.
- [3] S. Das, A. Nandan, and G. Pau. Spawn: A swarming protocol for vehicular ad-hoc wireless networks. In *Proceedings of VANET*, pages 93–94, 2004.
- [4] A. Datta, S. Quarteroni, and K. Aberer. Autonomous gossiping: A self-organizing epidemic algorithm for selective information dissemination in mobile ad-hoc networks. In *Proceedings of IC-SNW*, pages 126–143, Maison des Polytechniciens, Paris, France, June 2004.
- [5] D. Dolev, R. Reischuk, and H. R. Strong. Early stopping in byzantine agreement. *JACM*, 37(4):720–741, 1990.
- [6] GM. Threat assessment algorithm. <http://www.nhtsa.dot.gov/people/injury/research/pub/acas/acas-fieldtest/>, 2000.
- [7] Honda. <http://world.honda.com/news/2005/c050902.html>.
- [8] Lamport, Shostak, and Pease. The byzantine generals problem. In *Advances in Ultra-Dependable Distributed Systems*, N. Suri, C. J. Walter, and M. M. Hugué (Eds.). IEEE Computer Society Press, 1982.
- [9] C. Leckie and R. Kotagiri. Policies for sharing distributed probabilistic beliefs. In *Proceedings of ACSC*, pages 285–290, Adelaide, Australia, 2003.
- [10] D. Malkhi, E. Pavlov, and Y. Sella. Gossip with malicious parties. Technical report: 2003-9, School of Computer Science and Engineering - The Hebrew University of Jerusalem, Israel, March 2003.
- [11] Y. M. Minsky and F. B. Schneider. Tolerating malicious gossip. *Distributed Computing*, 16(1):49–68, February 2003.
- [12] M. J. Osborne and A. Rubinstein. *A Course In Game Theory*. MIT Press, Cambridge MA, 1994.
- [13] R. Parshani. Routing in gossip networks. Master's thesis, Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel, October 2004.
- [14] R. Parshani, S. Kraus, and Y. Shavitt. A study of gossiping in transportation networks. *Submitted for publication*, 2006.
- [15] Y. Shavitt and A. Shay. Optimal routing in gossip networks. *IEEE Transactions on Vehicular Technology*, 54(4):1473–1487, July 2005.
- [16] N. Shibata, T. Terauchi, T. Kitani, K. Yasumoto, M. Ito, and T. Higashino. A method for sharing traffic jam information using inter-vehicle communication. In *Proceedings of V2VCOM*, USA, 2006.
- [17] W. Wang, X.-Y. Li, and Y. Wang. Truthful multicast routing in selfish wireless networks. In *Proceedings of MobiCom*, pages 245–259, USA, 2004.
- [18] B. Yu and M. P. Singh. A social mechanism of reputation management in electronic communities. In *Proceedings of CIA*, 2000.