Available online at www.sciencedirect.com

SCIENCE DIRECT®

Transportation Research Part C 13 (2005) 299–319

TRANSPORTATION RESEARCH PART C

www.elsevier.com/locate/trc

Case studies on the Braess Paradox: Simulating route recommendation and learning in abstract and microscopic models

Ana L.C. Bazzan [a,*,1], Franziska Klügl [b]

[a] Instituto de Informática, UFRGS, Caixa Postal 15064, 91.501-970 Porto Alegre, RS, Brazil

[b] Department of Artificial Intelligence, University of Würzburg, Am Hubland, 97074 Würzburg, Germany

Abstract

The Braess Paradox is a well-known phenomenon: adding a new road to a traffic network may not reduce the total travel time. In fact, some road users may be better off but they contribute to an increase in travel time for other users. This situation happens because drivers do not face the true social cost of an action. Some measures have been proposed to at least minimize the effects of the paradox. However, it is not realistic to assume that the drivers would have all the necessary knowledge in order to compute their rewards from a point-of-view which is not their own, i.e. it cannot be expected that drivers would consider the global performance of the system. Therefore this paper discusses the effects of giving route recommendation to drivers in order to divert them to a situation in which the effects of the paradox are reduced. Two contributions are presented: a generalized cost function for the abstract model, which is valid for any number of drivers, and the calibration and results for a microscopic simulation, where the cost functions are not necessary anymore. These are replaced in the microscopic simulation by the real commuting time perceived by each driver. In all cases we use a learning mechanism to allow drivers to adapt to the changes in the environment. Different rates of drivers receive route recommendation with different rates of acceptance. We show that it is useful to manipulate the route information given to the agents.

* Corresponding author. Tel.: +55 51 3316 6823; fax: +55 51 3316 7308.

E-mail addresses: bazzan@inf.ufrgs.br (A.L.C. Bazzan), kluegl@informatik.uni-wuerzburg.de (F. Klügl).

1 Author partially supported by CNPq.

0968-090X/$ - see front matter

doi:10.1016/j.trc.2005.07.003

Keywords: Braess Paradox; Traffic simulation; Route choice; Agent-based simulation

1. Introduction

The Braess Paradox was originally presented by Braess (1968). It consists of a phenomenon which contradicts the common sense: in a traffic network, when a new link connecting two points (e.g. origin and destination) is constructed, it is possible that there is no reduction regarding the time necessary to commute from the origin to the destination. Actually, frequently this time increases and so the costs for the commuters.

Several authors have investigate the original problem (Akamatsu, 2000; Arnott and Small, 1994; Pas and Principio, 1997; Penchina, 1997; Turner and Wolpert, 2000 among others) as it will be detailed in the next section. The present work describes a model already partially used by Turner and Wolpert (2000) who have investigated the use of a multi-agent system to control routing of packages in a computer network using the so-called Collective Intelligence (COIN) formalism. The authors conclude that the network is also sensitive to the paradox in some cases. It happens because agents, by trying to reduce their individual routing times in a greedy way, end up increasing the global time.

We depart from Tumer and Wolpert's work as to what concerns the scenario and the determination of private goals for each agent. We use the classical scenario proposed by Braess, i.e. road traffic, where a new link added to the network is very attractive to the users since the commuting time there is low. Also, the COIN formalism assumes that agents can be aligned with the global objective. This is only possible in the computer network scenario in which router nodes have an aggregate knowledge that drivers in the traffic network do not have.

Therefore we continue our previous studies (Bazzan and Klügl, 2003; Klügl and Bazzan, 2004) in which the motivation is to let drivers learn the best route to select, based on past selections. However, here we also add a route recommendation element—the recommendation is computed by a control center based on traffic information. Variable rates of drivers receive this information with different rates of acceptance.

Furthermore, we simulate a network where the paradox happens using a microscopic traffic simulator which permits the actual measure of the travel times, thus eliminating the need for the cost functions similar to the one proposed by Braess in his seminal work. We show how to configure and calibrate a network so that the paradox is observed, as well as change some parameters in the microscopic model to investigate what happens with the distribution of drivers among the routes. In all cases we let drivers learn the best route option and investigate the effects of learning and receiving recommendation (with and without the presence of non-informed drivers).

A discussion about the Braess Paradox and previous approaches to it are presented in the next section. Section 3 describes how to generalize the cost function in order for the paradox to happen for an arbitrary number of drivers. Section 4 describes the abstract approach, i.e. the one close to the original formulation, based on cost functions for the links. We discuss different scenarios varying the type of information recommendation, rate of informed drivers, acceptance of recommendation, etc. (Section 4.1), as well as the results of these simulations (Section 4.2). Section 5 shows the simulation of the paradox in a cellular-automata based microscopic simulator, a trial to bring the paradox closer to the reality. Section 6 outlines the conclusions and future work.

2. The Braess Paradox and some previous approaches and results

2.1. Basic Braess scenario

In the scenario proposed in Braess (1968), drivers or agents can select between two or three available paths to commute from origin O to destination D, as depicted in Fig. 1(a) and (b), respectively. Commuting times are computed by means of the functions depicted for each link. They are all function of the flow or number of vehicles e.g. ($T_{OQ} = 10 * f$), where f is the number of vehicles in that particular link.

In the configuration depicted in Fig. 1(a), it is clear that if we have 6 vehicles (original number in Braess paper), the user equilibrium occurs when routes OQD and OPD carry 3 vehicles each. No one would be better off changing route.

In order to understand the paradox which occurs in the configuration depicted in Fig. 1(b), let us consider an increasing number of vehicles starting with only one. For a single vehicle, route OQPD is much cheaper: it takes only $10 + 11 + 10 = 31$ time units to commute from O to D. Taking OPQ or OQD would take 61 time units. It is easy to see that for more than 2.58 drivers, the new route is not advantageous anymore, so drivers do not use OQPD at all.

Table 1 shows some distribution of the six drivers and, following the cost functions given in Fig. 1, the costs per link and total cost (over all links). The lowest *total cost* (176) is observed in three situations: when drivers avoid the link QP (hence OQPD), and when all six drivers use either OQD or OPD. Distributing drivers evenly among the three routes causes a higher cost (even if less than 2.58 drivers use OQPD).

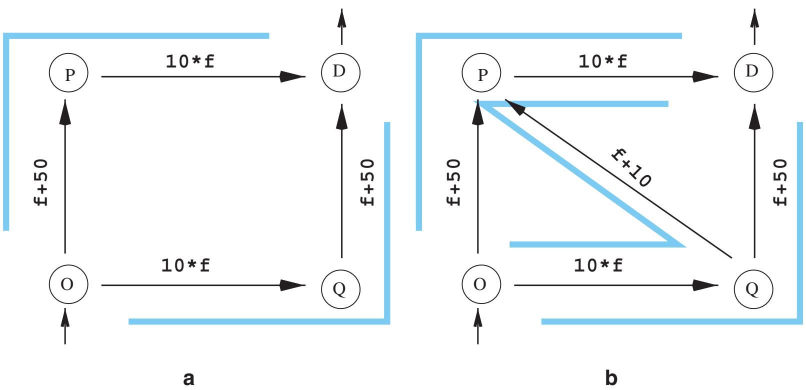

Fig. 1. Two configurations of the network in the classical scenario of the Braess Paradox: (a) routes are OPD and OQD (four-link network); (b) with additional route OQPD (five-link network).

Table 1
Cost for links and total cost for different distributions of drivers ($N = 6$)

No. of drivers			Cost of link					Cost
OP	QD	QP	OP	QD	QP	OQ	PD	Total
3	3	0	53	53	10	30	30	176
2	2	2	52	52	12	40	40	196
6	0	0	56	50	10	0	60	176
0	6	0	50	56	10	60	0	176
0	0	6	50	50	16	60	60	236

2.2. Related work

Following the seminal work by Braess, many authors have been proposing applications and modifications on the formulation of the problem in order to avoid the paradox. A paper by Smith (1978) shows how, in a particular simple case, total journey time varies with the travel time along an uncongested link. This result is of particular relevance for towns with a good bypass or a good outer ring road.

From the perspective of the economics of traffic, Arnott and Small (1994) analyse three paradoxes in which the usual measure for alleviating traffic congestion, i.e. expanding the road system, is ineffective. The resolution of these paradoxes—among them the Braess Paradox—employs the economic concept of externalities (when a person does not face the true social cost of an action) to identify and account for the difference between personal and social costs of using a particular road. For example, drivers do not pay for the time loss they impose on others, so they make socially-inefficient choices. This is a well-studied phenomenon more generally known as The Tragedy of the Commons (Hardin, 1968). Regarding the Braess Paradox, in the scenario analysed by Arnott and Small the travel time for each of the two original routes is 20 min, while after the addition of the new route the travel time for the equilibrium situation rises to 22.5 min for each route.

Returning to the traffic engineering and physics communities, in Penchina (1997) the "simplest anti-symmetric" two-path network is described which exhibits the Braess Paradox. A discussion of the good effects (non-paradoxical) of a bridge (especially a two-way bridge) is also included. Their Minimal Critical Network and graphical solution technique gives a clear understanding of the paradox for this network. They are also especially useful for several analysis of sensitivity related to changes in parameters, elastic demand, general non-linear (even non-continuous) cost functions, two-way bridges, tolls and other methods to control the paradox, and diverse populations of users. It is shown that the paradox occurs in a simpler network and with a larger Braess penalty than previously noticed.

Yang and Bell (1998) deal with network design via the Braess Paradox and show how this capacity paradox can be avoided by introducing the concept of network reserve capacity into network design problems. Pas and Principio (1997) examine properties of the paradox and show that whether the paradox does or does not occur depends on the conditions of the problem (link congestion function, parameters and the demand for travel). Akamatsu (2000) explores the properties of dynamic flow patterns on two symmetrical networks.

Although the literature from the transportation branch also proposes some ways of avoiding the paradox, ail of them concentrate on the parameters of the network, not on the driver itself, thus relegating the human component which plays an important role. One can argue that even if the "right" network is constructed (i.e. one which avoids the paradox from the point of view of the mathematics of the problem as it was proposed in the literature), drivers will always seek to maximize their individual payoffs, frequently leading to sub-optimal global distribution of traffic in the roads of the network. Besides, even if the "right" parameters are taken into account in the design of the network, the increasing demand for mobility in our society leads to a rapid obsolescence of the network.

2.3. Multi-agent approaches

Multi-agent systems approaches are interesting in the Braess scenario since one can focus on the issues related to the driver and its decision-making process.

2.3.1. Collective intelligence

Turner and Wolpert (2000) have investigated the use of a multi-agent system to control routing of packages in a computer network, as well as the sensitivity of the network to the Braess Paradox using the Collective Intelligence (COIN) formalism. Performance using COIN is compared to the case in which each agent in a set of agents estimates the "shortest path" to its destination. Since each agent decision (which is based on this estimation) ignores the effects of the decision of other agents on the overall traffic, shortest path based performances are badly sub-optimal. The authors conclude that the network is also sensible to the paradox in some cases. It happens because agents, by greedily trying to reduce their individual times, end up increasing the global time.

In the COIN approach, the new goals are tailored so that if they are collectively met the system maximizes throughput. The world utility, $G(\zeta)$, is an arbitrary function of the state of all agents across all times. The utility for an agent is given by the difference between the total cost accrued by all agents in the network and the cost accrued by agents when all agents sharing the same destination are "erased". The problem with the use of this approach in a traffic network is that we cannot expect the agents (drivers) to have any idea about the global goal or cost. This is possible in the computer network scenario because the agents are router nodes which have knowledge about the throughput (packages to be routed to different destinations).

2.3.2. Iterated route choice

Although the Braess Paradox is not about binary decision, the basic conditions from the iterated route choice (IRC) scenario (Klügl and Bazzan, 2004) are valid. This is a model for adaptive choice in which each agent has no information about other agents. Agents decide which route to select based on a local inference about the costs or rewards of each available action from the action set. This inference is based on the update of the probability according to which an agent selects each alternative action.

Let i be the index of an agent, and j the index of an action available to agent i. In a set of actions $A_i = (a_1, \ldots, a_j)$, a learning rule for time t is a rule which specifies the probabilities $P_{i,t} = (p_{i,\tau}(a_1), \ldots, p_{i,\tau}(a_j))$ as a function of the rewards obtained by selecting actions in the past $\tau < t$. In the adaptive scenario the agent i updates these probabilities with a certain periodicity according to the rewards it has obtained selecting route j up to that point. In the Braess scenario, this means that each agent remembers how many times s/he used each route and the total amount of time spent traveling on each of them. The reward is obviously inverse to the time spent commuting from O to D, for each route:

$$p_{i,\tau}(a_j) = \frac{\frac{1}{\sum_t T_{\mathrm{OD},j}}}{\sum_j \left(\frac{1}{\sum_t T_{\mathrm{OD},j}} \right)} \tag{1}$$

With this information in mind, the agent then calculates the average travel time spent in each route and selects the one with the shortest time. The average travel time in each route is actually computed by means of a discount factor δ, in order to allow us to play with the fact that it is not desirable that agents remember *all the past* with the same weight they remember the more recent outcomes. Thus, the last commuting time measured at a given route is multiplied by δ while the previous average time (i.e. a measure of the past) is multiplied by $(1 - \delta)$.

Bazzan and Klügl (2003) depart from Turner and Wolpert work as to what regards the scenario and the determination of private goals for each agent. The classical scenario proposed by Braess is used, i.e. road traffic, in which a new link added to the network is very attractive to the users since the commuting time there is low. However, since it overlaps (see Fig. 1) with the existing paths, the global commuting time increases.

The approach in Bazzan and Klügl (2003) is based on learning and adaptation in order to try to minimize both the global and local performance losses. Also, it is implemented using the multi-agent paradigm that requires little processing, making it scalable. It is also important to notice that agents do not need to explicitly communicate in order to coordinate their choices. Thus, the learning heuristics proposed—called A2B (Adapt to Braess)—improves the performances because it implicitly includes some factors of the global performance in the individual ones.

3. General cost function for abstract scenarios

The original work by Braess discusses an abstract scenario with six drivers. Later, Arnott and Small (1994) proposed cost functions for N drivers. These functions are given in Table 2. However, it is necessary to point out that, in those functions, the authors assume that the additional route (OQPD) can be traversed in a fixed number of time units, *regardless of traffic volume*. Since this is not very realistic, we depart from their assumption and show a set of equations and inequalities (Eqs. (2)–(6)), which arise from the paradox constraints, and allow us to derive the parameters $\alpha_1, \alpha_2, \ldots, \beta_4$ as a function of N.

Table 2
Functions to compute the time to traverse each route in the network (adapted from Arnott and Small (1994))

Links	Function
OQ (similar for PD)	$T_{\mathrm{OQ}} = \alpha_1 + \beta_1 \times f_{\mathrm{OQ}}$
QD (similar for OP)	$T_{\mathrm{QD}} = \alpha_2 + \beta_2 \times f_{\mathrm{QD}}$
QP	$T_{\mathrm{QP}} = \alpha_4$

Eq. (2) comes from the fact that, for the paradox to exist, when $N/2$ drivers use each OQD and OPD, this is cheaper than when $N/3$ use each one of the three routes (remember that OQ and PD are shared by $2N/3$ drivers). Eq. (2) is derived for OQD but symmetrical inequalities exist for OPD. In the example of the six drivers (Table 1, lines 1 and 2), $30 + 53 < 40 + 52$. Eq. (3) states more or less the same but this time from the point of view of drivers using OQPD (again, see Table 1: $30 + 53 < 2 * 40 + 12$).

Eq. (4) is the most important one because it represents the user equilibrium: when $N/3$ use each route, no driver is better off changing routes. For six drivers we have: $40 + 52 = 2 * 40 + 12$. Eq. (5) comes from the fact that link OP is more expensive than OQ in a two-route situation ($53 > 30$). For a three route situation (Eq. (6)), this is still valid and additionally we have that OQ is more expensive than PD ($52 > 40 > 12$).

$$(\beta_1 * (N/2) + \alpha_1) + (\beta_2 * (N/2) + \alpha_2) < (\beta_1 * (2N/3) + \alpha_1) + (\beta_2 * (N/3) + \alpha_2) \tag{2}$$

$$(\beta_1 * (N/2) + \alpha_1) + (\beta_2 * (N/2) + \alpha_2) < 2 * (\beta_1 * (2N/3) + \alpha_1) + (\beta_4 * N/3 + \alpha_4) \tag{3}$$

$$(\beta_1 * (2N/3) + \alpha_1) + (\beta_2 * (N/3) + \alpha_2) = 2 * (\beta_1 * (2N/3) + \alpha_1) + (\beta_4 * N/3 + \alpha_4) \tag{4}$$

$$\beta_2 * N/2 + \alpha_2 > \beta_1 * N/2 + \alpha_1 \tag{5}$$

$$\beta_2 * N/3 + \alpha_2 > \beta_1 * 2N/3 + \alpha_1 > \beta_4 * N/3 + \alpha_4 \tag{6}$$

Using these equations, it is possible to determine the parameters $\alpha_1, \alpha_2, \ldots, \beta_4$ as a function of N. Once these are determined, costs can be computed for any scenario, i.e. for any N. For instance, solving this set of equations for $N = 6$ return us the original scenario by Braess (Fig. 1 and Table 1). Solving for $N = 1500$—our scenarios here—the parameters are as in Table 3.

Table 4 is similar to Table 1 but here the distributions and costs are over $N = 1500$ drivers, and parameters are as given in Table 3. Table 5 shows more computations for the same situation: the cost of each route, the total cost (sum over all links), and cost over all drivers.

Table 3
Values for parameters when $N = 1500$

Parameter	α_1	α_2	α_4	β_1	β_2	β_4
Values	70	500	5	0.2	0.05	0.5

Table 4
Cost for links for different distributions of drivers

No. of drivers			Cost of links					Cost over
OP	QD	QP	OP	QD	QP	OQ	PD	All links
750	750	0	537.5	537.5	5	220	220	*1520*
749	749	2	537.45	537.45	6	220.2	220.2	1521.3
500	500	500	525	525	255	270	270	*1845*
400	500	600	520	525	305	290	270	1910
400	600	500	520	530	255	290	250	1845
500	600	400	525	530	205	270	250	1780
1500	0	0	575	500	5	70	370	1520
0	1500	0	500	575	5	370	70	1520
0	0	1500	500	500	755	370	370	2495

Table 5
Cost for links and total cost for different distributions of drivers ($N = 1500$)

No. of drivers			Cost of route			Cost over
OP	QD	QP	OPD	OQD	OQPD	All drivers ($\times 10^6$)
750	750	0	757.5	757.5	445	*1.14*
500	500	500	795	795	795	*1.19*
400	500	600	790	815	865	1.24
400	600	500	770	820	795	1.2
500	600	400	775	800	725	1.16
1500	0	0	945	570	445	1.42
0	1500	0	570	945	445	1.42
0	0	1500	870	870	1495	2.24

4. Abstract modeling

4.1. Approaches and scenarios

The basic scenario presented in Section 2.3.2 can be extended to include a traffic control center giving traffic information to the drivers. This information can be the actual traffic situation, route recommendation, or "manipulated information". The recommended route is the one which had the best last travel time. Manipulated recommendation is different: it is given to agents to try to force an equilibrium distribution between the alternatives, thus resembling both the approaches in Sections 2.3.1 and 2.3.2. To understand the manipulation process, let us take an example: if the current distribution is $\text{number}_{\text{OQPD}} = 500$ and $\text{number}_{\text{OPD}} = \text{number}_{\text{OQD}} = 500$, it is clear that the system could do better in terms of *total sum* of travel times (1845 vs. 1520). Also, the rewards of agents are sub-optimal (1.19 vs. 1.14). When the control system perceives such a situation, it tries to induce agents to come closer to the global optimum by given manipulated recommendation to them. In the particular case above, the control system can give a bad forecast for route *OQPD* to try to divert drivers to other routes. Apart from this information element, we also analyse the impact of non-informed drivers. Thus, in some scenarios, we have different rates of informed to non-informed drivers.

The implementation was done using SeSAm, a shell for developing agent-based simulation.[2] The situations we have simulated are:

(I) *No random drivers:* drivers receive recommendation but there is no *manipulation*; and recommendation is followed by different rates of drivers: 0% (Ia), 25% (Ib), 50% (Ic), 75% (Id), and 100% (Ie).

(II) *No random drivers:* drivers receive recommendation which is *manipulated*; and recommendation is followed by different rates of drivers: 25% (IIb), 50% (IIc), 75% (IId), and 100% (IIe).

(III) *Random drivers:* drivers receive recommendation but there is no manipulation; and recommendation is followed by different rates of drivers: 25% (IIIb), 50% (IIIc), and 100% (IIIe).

[2] Available for download at www.simsesam.de

(IV) *Random drivers:* drivers receive recommendation which is manipulated; and recommendation is followed by different rates of drivers: 25% (IVb), 50% (IVc), and 100% (IVe).

In all the above situations, drivers receive recommendation. However, different rates of them actually follow it (different rates are associated with different letters). Besides, we have different situations (marked with roman numerals) for: presence (situations (III) and (IV)) or absence ((I) and (II)) of random drivers, under manipulation ((II) and (IV)) or not ((I) and (III)). Random drivers just select a route randomly, while the others use the learn-and-adapt mechanism discussed in Section 2.3.2. When the traffic control center gives "false" forecast to prevent drivers from using a route which is expected to increase the global costs, an interesting issue is that agents may learn that the information cannot be trusted. This happens because, besides receiving false information, they also adapt their heuristics.

4.2. Description of the experiments and results

In this section we discuss the results regarding the situations presented above. For the experiments, we use $N = 1500$ agents, $\delta = 0.8$, and learning frequency of 0.2 (one in each five trips, on average). Cases with less agents and in particular with $N = 6$ are reported in Bazzan and Klügl (2003). As for the cost functions to compute the commuting time, they are shown in Table 2, except that we have $T_{\mathrm{QP}} = \alpha_4 + \beta_4 * f_{\mathrm{QP}}$. For $N = 1500$, the paradox happens with parameters shown in Table 3. With these parameters, the global costs can be computed for the particular situations of interest:

(1) When the $N = 1500$ drivers are distributed equally among the 3 routes, the route cost is 795 for OQPD, OQD, and OPD; this is also the average cost per driver, while the global cost is 1845 (over all links).
(2) When all drivers avoid OQPD and 750 use each remaining route, the cost is 445 for OQPD and 757.5 for the other two, thus adding up 1520.

In the simulations, the quantities measured were mainly distribution of drivers in each route, time in OQPD, time over all drivers, time over all routes, and the probability of selecting each route (averaged over N agents).

In principle, the distribution of drivers between the three routes tells roughly how good drivers perform because we have shown the analytical computation of times per link, total time, etc. for $N = 1500$ (Tables 4 and 5). However, due to the random decisions and the changing in route choices due to the learning mechanism, there are high variations in those distributions. For instance, Table 6 depicts the number of drivers (third column), time in route OQPD (fourth), time over all drivers (fifth), and time over all links (sixth), for different rates of drivers actually following the recommended route. When nobody follows it (first line), everybody acts as a random driver given that there is no other information available. Thus, the travel time (using OQPD) is around 800 as predicted (Table 5). The time over all drivers is around 1.2×10^6 and the time over all links is 1845, hence not the optimal one.

Increasing the rate of people actually following the recommended route, the most visible result in an increase in the variation of all measured quantities: during the whole simulation there are

Table 6
Situations with recommendation, no random driver, no manipulation

Case	(%) follow recomm.	No. of drivers OQPD	Time OQPD	Time over all drivers ($\times 10^6$)	Time over all links
Ia	0	≈500	≈800	1.2	1845
Ib	25	350–700	700–1000	1.15–1.35	1747–2007
Ic	50	200–1000	600–1150	1.2–1.55	1650–2170
Id	75	100–1250	500–1300	1.3–1.85	1585–2332
Ie	100	0	445	1.42	1520

between 350 and 700 drivers in OQPD when 25% of the drivers follow the recommendation. This of course causes variations in all other quantities. In an extreme case, when 75% follow the recommendation, there are between 100 and 1250 drivers using OQPD. This happens because drivers basically select the recommended route given by the traffic control center. The learning component plays a lesser role the more the driver actually relies on the recommendation. When everyone follow the recommendation, after a short time no one uses OQPD because the travel time is higher there *when everyone is there* as it appears in the last line of Table 5. One can argue that this kind of recommendation is not clever but from the point of view of the control center, the only information available is the best route in the near past. Any other type of information (e.g. diverting drivers) is treated here under manipulation, thus situations II and IV that are discussed below.

We also show the evolution of the average probability of selecting OQPD along time for situation I (Fig. 2, top). When nobody follows the recommendation, the average probability is around $\frac{1}{3}$ and therefore on average 500 drivers use OQPD. Increasing the rate of drivers following the recommendation causes the average probability to decrease.

Next, Table 7 depicts the same measures as before, this time with drivers receiving manipulated information (Situation II). We do not simulate the case when nobody follows recommendation because it would make no difference here. Comparing the corresponding lines in Tables 6 and 7, we notice a decrease in the variation in the number of drivers in OQPD (e.g. from 100–1250 to 100–150 when 75% follow the recommendation). This also reflects in the other quantities. This is explained by the fact that manipulation is a more clever form of recommendation, even if it is based on the interest of the global optimum, not in the interest of the user optimum. In any case the distributions of drivers are much more stable (and so times) and this is also better for the learning mechanism of the drivers. When only 75% follow recommendation (hence only 25% use the learn mechanism) we have a good performance from the point of view of the global system (with time over all links between 1585 and 1617). The evolutioo of the average probability of selecting OQPD is given in Fig. 2, bottom. We see that in all cases these probabilities are higher than in situation I. This is so because the manipulation actually is aimed at avoiding the use of OQPD, hence those who do not follow the recommendation and use OQPD are better off (hence the probability of using it increases for these).

Now, it is interesting to investigate what happens if not all drivers are informed, i.e. when random drivers perform the route selection together with drivers who learn and get recommendation. We have simulated this situation with different shares of random drivers: 25%, 50%, and 75%, with and without manipulation. Results for the former (Situation III) are depicted in Table 8. Henceforth we only show number of drivers in OQPD. This number generally decreases with the increasing rate of random drivers, and the variation in that number decreases with increasing

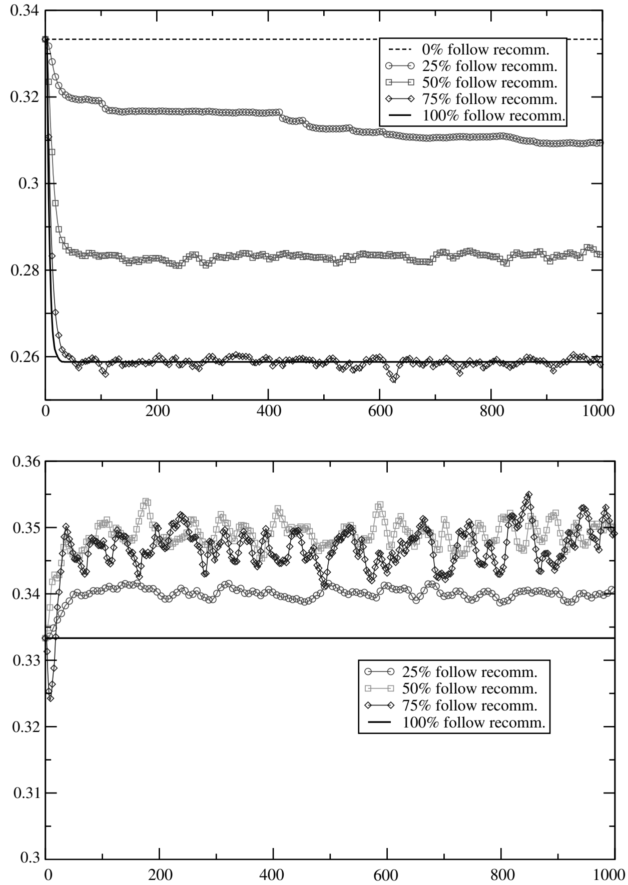

Fig. 2. Average probability of selecting route OQPD along time: No random drivers, without manipulation (top), with manipulation (bottom).

Table 7
Situations with recommendation, no random driver, with manipulation

Case	(%) follow recomm.	No. of drivers OQPD	Time OQPD	Time over all drivers ($\times 10^6$)	Time over all links
IIb	25	350–400	700–725	1.1–1.18	1747–1780
IIc	50	225–275	600–650	1.05–1.2	1666–1700
IId	75	100–150	500–550	0.95–1.28	1585–1617
IIe	100	0	445	0.85–1.42	1520

Table 8
Situations with recommendation, with random driver, no manipulation

Case	(%) follow recomm.	No. of drivers OQPD (25% random)	No. of drivers OQPD (50% random)	No. of drivers OQPD (75% random)
IIIb	25	500–650	370–670	400–590
IIIc	50	300–900	320–770	400–670
IIIe	100	100–1250	220–1020	350–760

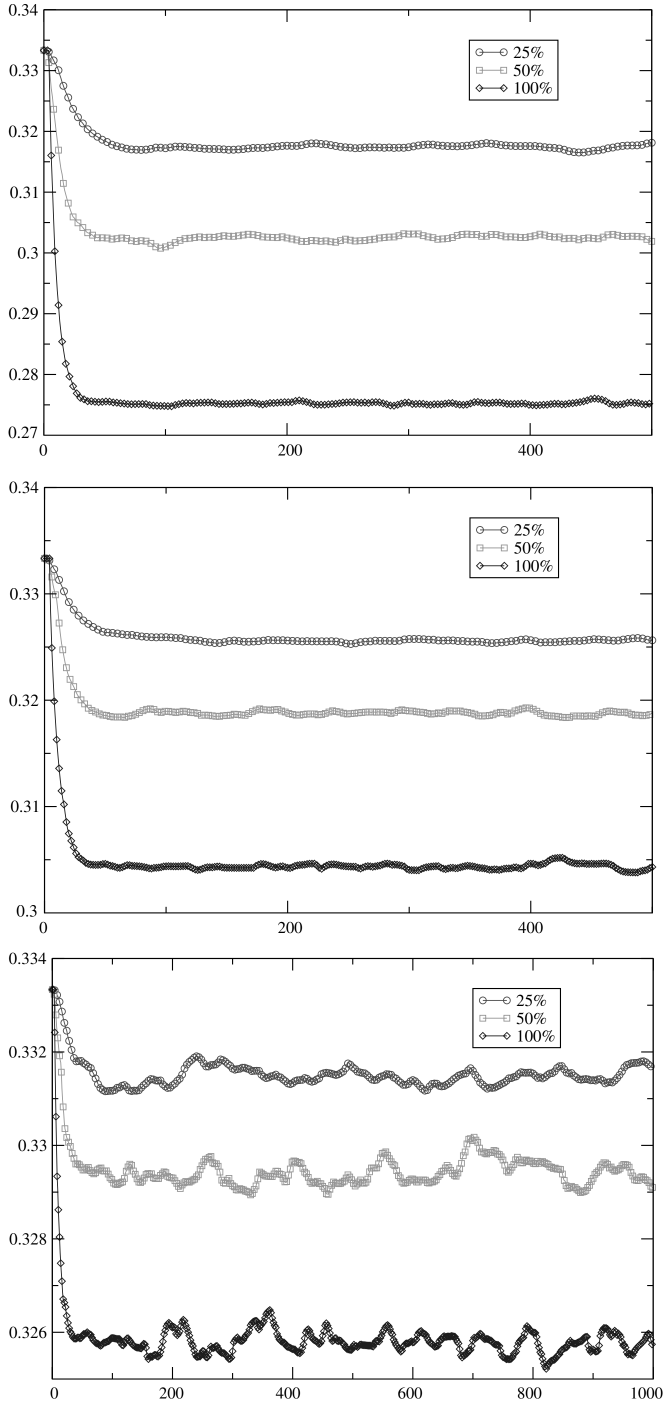

Fig. 3. Average probability of selecting route OQPD along time, no manipulation, with different rates of random drivers: 25% (top), 50%, and 75% (bottom).

rate of random drivers because these can select any of the three routes. In all cases the variation of the number of drivers in OQPD is high (although the latter is not as high as in situation I), and the number of drivers is higher than in situation II.

As for the probability of selecting OQPD, Fig. 3 depicts its evolution with time for the three rates of random drivers. The variation increases with increasing rate of random drivers, no matter the rate of people following recommendation.

Finally, Table 9 showns the number of drivers in OQPD for different rates of random drivers, this time under manipulation (Situation IV). As we noticed for situation II, the manipulation reduces the variation in those numbers. Here we can also observe that the number of drivers in OQPD increases slightly with the rate of random drivers (Fig. 4).

Summarizing, the main conclusions of the abstract modeling using recommendation and manipulation are:

(1) manipulation reduces the *variance* in the distribution of drivers (number of drivers in each route);
  - *without* random drivers, the reduction is from 100–1250 (when 75% follow recommendation) to 100–150 (again, 75% follow recommendation);
  - *with* random drivers, the reduction is less significant because the presence of random drivers already reduces the variance (although it increases the number of drivers using OQPD because random drivers select each route with $\frac{1}{3}$ of probability); the reduction in the variance is from 100–1250 (100% follow recommendation) to 100–150 (idem);

(2) manipulation also reduces the *number of drivers in the* OQPD; in fact the lowest number of drivers selecting OQPD is the one in which there are few random drivers (if any) and almost everybody follows the recommendation (which is manipulated); there is a reduction in all cases (25%, 50%, and 75% follow recommendation), with or without random drivers;

(3) the more random drivers, the more variance when there is no manipulation; under manipulation this is not significant;

(4) special cases:
  - no recommendation (i.e. 0% follow recommendation): in all cases the distribution is $\frac{1}{3}$ in each road due to the incentive given to individual drivers to use OQPD (link QP costs only 445);
  - when everyone gets and follows recommendation, since this is the last best route, drivers switch from one route to another; eventually everybody uses the same route, thus all drivers end up with the worst time possible (last lines in Table 5);

Table 9
Situations with recommendation, with random driver, and manipulation

Case	(%) follow recomm.	No. of drivers OQPD (25% random)	No. of drivers OQPD (50% random)	No. of drivers OQPD (75% random)
IVb	25	350–450	380–480	380–480
IVc	50	250–390	320–400	310–410
IVe	100	100–150	170–300	150–290

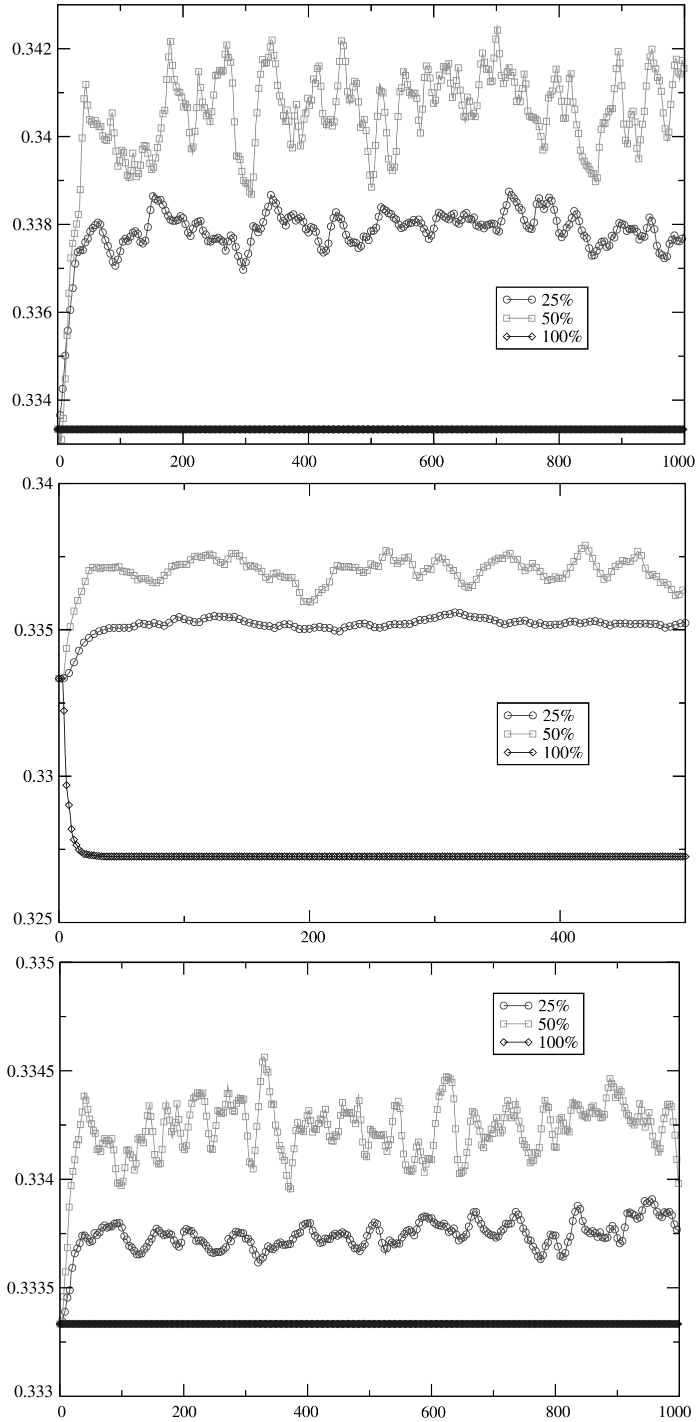

Fig. 4. Average probability of selecting route OQPD as a function of time, with manipulation, for different rates of random drivers: 25% (top), 50%, and 75% (bottom).

(5) the increasing use of recommendation:
  - *without* manipulation: increases the time over all drivers *but* increases the variance for all quantities so that nothing can be concluded;
  - in general: reduces the number of drivers on OQPD.

5. Microscopic simulation of the Braess Paradox

5.1. The cellular automaton model for microscopic simulation

Microscopic models of simulation allow the description of each road user as detailed as desired (given computational restrictions), thus permitting a model of drivers' behaviors. Multi-agent simulation is a promising technique for microscopic traffic models, since drivers' behavior can be described incorporating complex individual decision-making.

In the present paper, we use the microscopic traffic flow model proposed in Nagel and Schreckenberg (1992), based on a cellular automaton (CA). It represents a minimal model in the sense that it is capable of reproducing basic features of real traffic. For completeness, the definition of the model for single-lane traffic is briefly reviewed. The road is subdivided in cells with an arbitrary length. Each cell is either empty or occupied by only one vehicle with an integer velocity $v_i \in \{0, \ldots, v_{\max}\}$. The movement of the vehicles is described by the following rules (*parallel dynamics*):

R1 Acceleration: $v_i \leftarrow \min(v_i + l, v_{\max})$.
R2 Deceleration to avoid accidents: $v_i' \leftarrow \min(v_i, \text{gap})$.
R3 Randomization: with a certain probability p_b do $v_i'' \leftarrow \max(v_i' - 1, 0)$.
R4 Movement: $x_i \leftarrow x_i + v_i''$.

The variable *gap* denotes the number of empty cells in front of the vehicle at cell *i* and restricts the velocity. The first two rules (Rl, R2) describe a somehow optimal driving strategy. The driver accelerates if the vehicle has not reached the maximum velocity $v_{\max}$ and brakes to avoid accidents. This can be summed up as follows: *drive as fast as you can and stop if you have to!* Such a cellular automaton is deterministic and the stationary state depends only on the initial conditions. But drivers do not react in this optimal way: they vary their driving behavior without any obvious reasons, reflected by the *braking noise* p_b (R3). It mimics the complex interactions between the vehicles and is also responsible for spontaneous formation of jams.

5.2. Topology and calibration

The first step in the microscopic simulation of the Braess scenario is to find a network topology in which the paradox occurs. Contrary to the abstract scenario which uses cost functions, the topology is, of course, particular for a given number of drivers. If the network has too much capacity, then drivers are expected to drive always at free flow and the paradox may not happen; when the network has too little capacity, then the effects of the paradox may again not happen due to oversaturation. Therefore, it is necessary to design the network for a specific number of drivers. We keep using $N = 1500$ here. For this number, we designed the network to hold $N/2$ drivers in OQ, QD, OP, and PD (the original two route version of the network, before a third link is built).

Due to the CA modeling, vehicles must have a gap of at least v_{max} cells between them in order to travel at v_{max}. We allow different v_{max} at different links in order to set the paradox (Eqs. (5) and (6)). Also, the basic CA model can be extended to allow overtaking, thus we also allow different number of lanes in the links. For $N = 1500$, the design parameters are shown in Table 10 (calibration column).

In order to check whether the constraints on the cost functions (Eqs. (3) and (4)) hold for this network, for each link we plot the density (vehicles/metric unit) (which in our case is the same of

Table 10
Network and CA model parameters

Parameter	Calibration	Other scenarios
Length OQ (PD)	1500	1500
v_{max} OQ (PD)	5	5
#No. of lanes OQ (PD)	3	3
Length OP (QD)	1500	1500
v_{max} OP (QD)	3	3
#No. of lanes OP (QD)	2	2
Length QP	$1500 * \sqrt{2}$	$1500 * \sqrt{2}$
v_{max} QP	7	7
#No. of lanes QP	5	5
Breaking prob. p_b	0	0/0.1/0.25
Learning prob. p_l	–	0.2/1.0
Discount factor δ	–	0.8

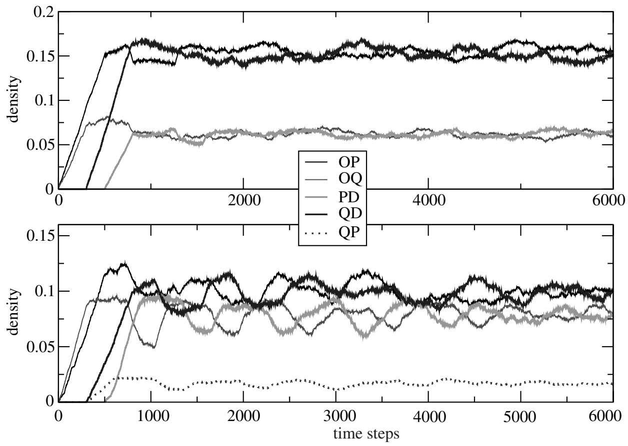

Fig. 5. Density of vehicles at each link in the calibration (without learning): upper and lower plots are for 2 and 3 routes, respectively.

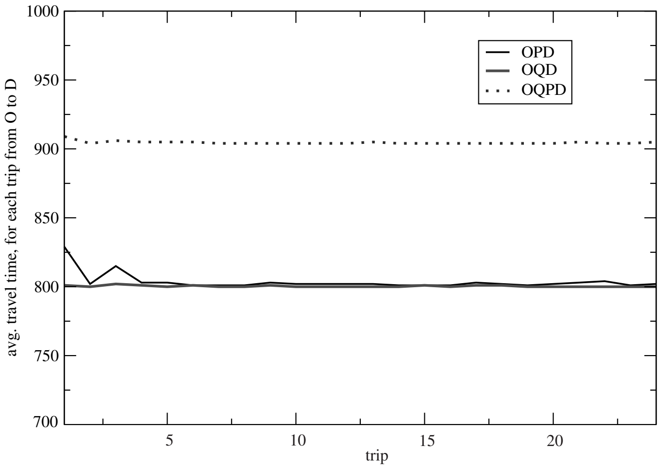

Fig. 6. Average time taken for drivers arriving at node D, without learning.

occupancy (vehicles/cells) because each vehicle occupies only one cell). This appears in Fig. 5. Densities are low here due to the need of a gap at least equal to $v_{\max}$ ahead of vehicles. The upper plot is for the situation with only two routes, while the bottom plot includes link QP. One sees that there are more vehicles on links OP and QD than in OQ and PD in both situations, and that the density in QP is the lowest. Thus the constraints of the paradox hold.

As for travel time, it is easy to see that the expected commuting time (at free flow) is 800 time steps (1500/3 + 1500/5) if only OQD and OPD are used. If we have three routes and 500 drivers in each, then the time is 903 (300 + 303 + 300) for those using OQPD (with the extra link). If we plot the time taken for each driver when reaching the destination node (D), no matter which route was taken, we see that using two routes takes drivers around 800 time steps most of the time; using three routes it takes longer. Fig. 6 confirms those numbers and that the paradox in fact occurs for this topology and set of parameters. It shows the average time over all drivers *using the same route*, when reaching the destination node (D), for the first 24 trips.

5.3. Parameters, scenarios, and results

Once the calibration has shown that this network topology is affected by the paradox, we can now simulate situations with drivers learning and adapting to the rewards received. In the microscopic simulation, the rewards are similar to those discussed in Section 2.3.2 with the only difference that now the travel time is generated directly in the simulation, not via abstract cost functions as in Section 4. The parameters for the further simulations are also given in Table 10. The network topology parameters are the same, while we vary the breaking probability (rule R3 of the CA)—$p_b = \{0.0, 0.1, 0.25\}$—, and the learning probability—$p_i = \{0.2, 1.0\}$.

Also, the microscopic simulation allows us to introduce floating cars (FC) that can provide us with punctual information about the links. FC information was also used by us in Wahle et al. (2000) with good results. We have defined three FCs, one for each route, and since they cannot change route, they do not learn. For each FC we measure position and speed. Due to limitation

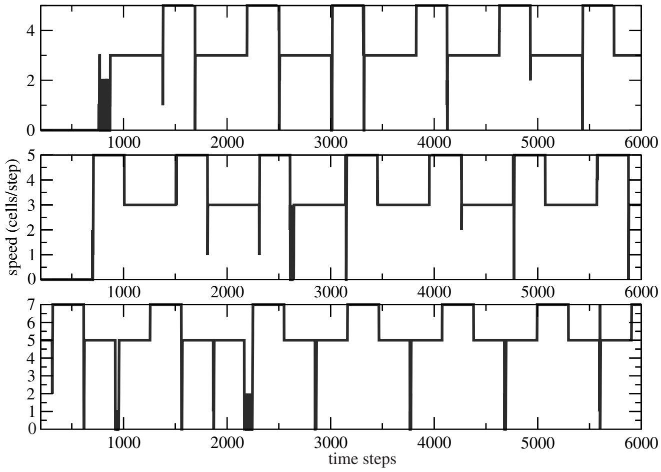

Fig. 7. Speed of FCs in the 3 routes (OPD, OQD, OQPD), with learning, high initial bias to select OQPD.

of space we only show here one plot of speed for the FCs (Fig. 7). This plot is for $p_l = 1.0$, $p_b = 0$, and the initial route probability for all non-FC drivers is $p_{OPD} = p_{OQD} = \frac{100}{1500}$ and $p_{OQPD} = \frac{1300}{1500}$ so to try an extreme condition in which initially most of the drivers would use OQPD thus jamming the link QP. The curves show the FCs in OPD (top), OQD (middle), and OQPD (bottom). The plot shows that the FCs travel at each link's v_{max} most of the time. However, sometimes congestions do happen in OQ and, especially, in QP, at the beginning of the simulation when the drivers did not start learning to use the routes OQD and OPD.

Regarding the time taken by the drivers to reach D, Fig. 8 shows average travel time in two situations (in both $p_l = 1.0, p_b, = 0$): the upper plot is for initial probabilities $p_{OPD} = p_{OQD} =$

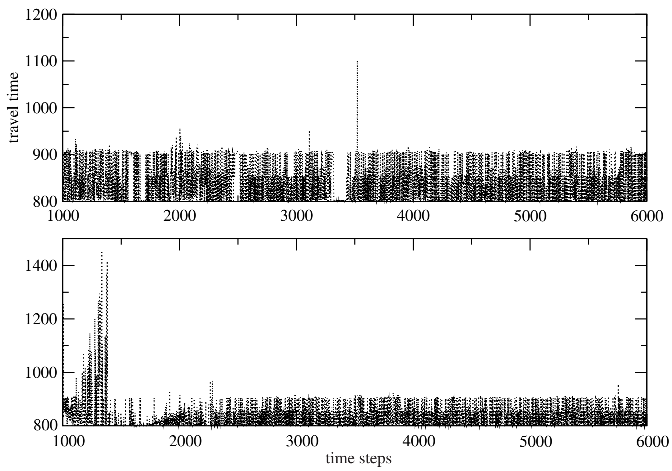

Fig. 8. Time taken for drivers arriving at node D, with learning; upper plot: no bias for the routes; lower plot: initial bias to select OQPD.

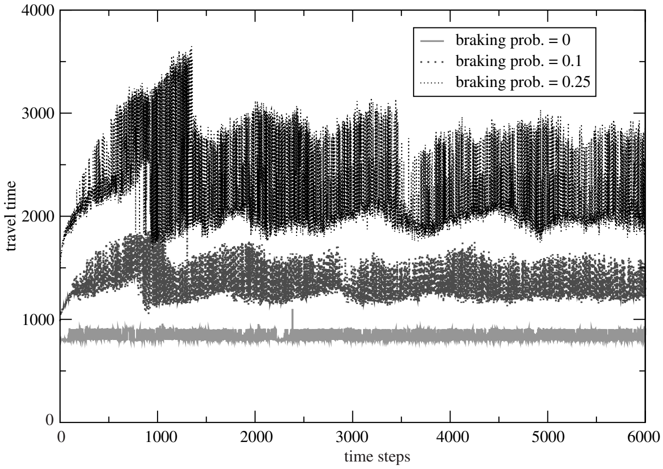

Fig. 9. Average time taken for drivers arriving at node D; learning frequency is $p_l = 1$, for different braking probabilities.

$p_{OQPD} = \frac{500}{1500}$ and the bottom plot for $p_{OPD} = p_{OQD} = \frac{100}{1500}$ and $p_{OQPD} = \frac{1300}{1500}$. In the former, since drivers already start with the user equilibrium probability, they only do fine adapting with time. However, in the latter case they start with a far-from-user-equilibrium probability: on average drivers select OQPD with probability $\frac{1300}{1500}$ causing jams and increasing the overall travel time. After step 1500, the travel time decreases sharply because many drivers completely avoid this route once the travel has taken too long (thus the selection probability becomes too low). Due to learning, after some time people start using it again so that the user equilibrium is reached by time step 2500, and the travel time remains around the expected 800–1000 time steps.

Regarding the distribution of drivers in the three routes, with learning, the results are as expected from the abstract scenario: 500 go to each route, on average. Fluctuations are smaller than in the abstract scenario (450–550 drivers in each route). The user equilibrium is reached no matter which learning and braking probability is used, but the higher the latter, the more it takes for that equilibrium to be reached. Also, increasing the braking probability p_b has an impact in the average travel times, as expected. Fig. 9 depicts a plot of travel times when p_b is changed. The higher this probability, the more noise is added due to the braking vehicles, increasing travel times.

6. Conclusion and future work

In the classical scenarios for the Braess Paradox reported in the literature, if drivers or agents of any kind act to maximize their own profits, the global performance of the system may decrease. This happens because the global goal opposes the individual goals in most cases. Thus the Braess Paradox is especially interesting in the study of issues such as learning and adaptation, and the effects of forecast and information manipulation in order to maximize both the global and local performances. Moreover, the scenario is challenging when moving from an abstract to a realistic simulation model, for the constraints of the paradox must apply to the network topology.

This paper has contributions in three directions: First it shows that the abstract model applies regardless the number of drivers, provided the system of equations and inequalities derived from the paradox constraints is solved. Second, it presents a series of simulations for the abstract model (using arbitrary cost functions) under several conditions: learning; presence or absence of uninformed drivers; presence of traffic guidance (route recommendation), with different rates of acceptance; presence or absense of information manipulation. In all cases, agents do not need to explicitly communicate in order to "coordinate" their choices. Third, microscopic simulations are performed using a cellular automaton model, thus permitting the study of individual drivers, including floating cars.

Regarding the abstract model, the simulations show that it is useful to manipulate the route information given to agents. By doing so, the control system is able to divert them to the more convenient alternative from the point of view of both the overall system and the individual agents (as this is the situation in which individual rewards are the highest). This holds for several rates of uninformed drivers. Thus, an important conclusion is that, in networks suffering the effects of the Braess Paradox, having agents provided with the most accurate information (information about the *actual* state of the system) is not necessarily good. This is so because when learning, users can at most reach the user equilibrium, which is sub-optimal.

Regarding the microscopic simulation, it successfully reproduced the previous results achieved with the abstract scenario with basic learning: drivers converge to selecting the three routes with equal probability (on average).

Future directions can be pursued both for the abstract as well as for the microscopic models. In the former, we made the assumption that the central component acts in the interest of the highest system performance. However, when we have other interests involved (for instance when there are several competing "central" controller components), some of them may act primarily to disturb the others. This remains to be investigated, as well as the simulations including random drivers and information in the microscopic version.

References

Akamatsu, T., 2000. Dynamic traffic equilibrium assignment paradox. Transportation Research B 6, 515–531.

Arnott, R., Small, K., 1994. The economics of traffic congestion. American Scientist 82, 446–455.

Bazzan, A.L.C., Klügl, F., 2003. Learning to behave socially and avoid the Braess Paradox in a commuting scenario. In: Parsons, S., Gmytrasiewicz, P. (Eds.), 5th Workshop on Decision Theoretic and Game Theoretic Agents, held together with AAMAS 2003, pp. 107–112.

Braess, D., 1968. Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 12, 258.

Hardin, G., 1968. The tragedy of the commons. Science 162, 1243–1248.

Klügl, F., Bazzan, A.L.C., 2004. Route decision behaviour in a commuting scenario. Journal of Artificial Societies and Social Simulation 7 (1). Available from: <http://jasss.soc.surrey.ac.uk/7/1/1.html>.

Nagel, K., Schreckenberg, M., 1992. A cellular automaton model for freeway traffic. Journal de Physique I 2, 2221.

Pas, E.I., Principio, S.L., 1997. Braess' Paradox: some new insights. Transportation Research B 31 (3), 265–276.

Penchina, C.M., 1997. Braess Paradox: maximum penalty in a minimal critical network. Transportation Research A 31 (5), 379–388.

Smith, M.J., 1978. In a road network, increasing delay locally can reduce delay globally. Transportation Research 12 (6), 419–422.

Turner, K., Wolpert, D., 2000. Collective intelligence and Braess' Paradox. In: Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI). AAAI Press, pp. 104–109.

Wahle, J., Bazzan, A.L.C., Klügl, F., Schreckenberg, M., 2000. Decision dynamics in a traffic scenario. Physica A 287 (3–4), 669–681.

Yang, H., Bell, M.G.H., 1998. A capacity paradox in network design and how to avoid it. Transportation Research A 32 (7), 539–545.