A Multiagent Solution to Overcome Selfish Routing In Transportation Networks

Mohammad Rashedul Hasan
University of Nebraska-Lincoln
Lincoln, NE 68588, USA
Email: hasan@unl.edu

Ana L. C. Bazzan
PPGC / UFRGS
CP 15064, 91501-970, Brazil
Email: bazzan@inf.ufrgs.br

Eliyahu Friedman and Anita Raja
The Cooper Union
NY 10003
Email: {friedm3,araja}@cooper.edu

***Abstract*—It is well-known that selfish routing, where individual agents make uncoordinated greedy routing decisions, does not produce a socially desirable outcome in transport and communication networks. In this paper, we address this general problem of the loss of social welfare that occurs due to uncoordinated behavior in networks and model it as a multiagent coordination problem. Specifically we study strategies to overcome selfish routing in traffic networks with multiple routes where a subset of vehicles are part of a social network that exchanges traffic related data. We investigate classic traffic flow paradoxes that are ubiquitous in various types of networks leading to severe congestion. We present a novel distributed traffic coordination algorithm that alleviates congestion by harnessing the real-time information available through the driver's online social network. We also propose a utility computation mechanism for route choice that generates near-optimal flows. Our extensive simulation results show that social network based multiagent traffic route coordination contributes to mitigate the effects of these paradoxes and significantly reduces congestion.**

I. Introduction

Selfish routing, where individual agents make uncoordinated routing decisions only in the interest of their own performance, is known to be inefficient. It leads to congestion problems in networked systems [7], [13], [14].

We have several motivations for the work reported in this paper. First, since the publication of the aforementioned (and other) works on selfish routing, the panorama has changed. There is an increasing prevalence in the use of social networks where communication of information among individuals using the network is real-time and seamless. This trend is giving traction to mobile apps for transport networks, such as Waze and Roadify, designed to enable commuters to share realtime traffic information. This work is aimed at opening a thread of inquiry to determine the underpinnings of Waze-like algorithms, which are not accessible to the public, with the goal of providing improvements where possible. It has also led to recent interest [9], [10], [5] in studying the potential benefits of the social communication aspects of these mobile apps and how they can be effectively applied towards minimizing traffic congestion. Our work is motivated by such initiatives and we seek to take advantage of these new technologies to use them to better distribute vehicles in the network with the simulation goals of reducing congestion.

Second, the increase in average travel time due to selfish routing has been extensively studied with respect to two classic traffic flow paradoxes: Pigou [11] and Braess [2]. Using a simple two route network, Pigou's paradox demonstrates the principle that selfish behavior need not produce a socially optimal behavior. Braess' paradox reinforces the sub-optimality of selfish routing by showing that network improvements (adding a zero or low cost route to reduce congestion) can result in increased congestion when drivers choose routes selfishly. In real transportation networks, it has been observed [12] that this problem could be resolved by restructuring the network (shutting down a road). Another strategy is to introduce tolls on faster routes that not only incur extra costs for drivers but also change the cost structure of the routes. Our goal is to solve these paradoxes without restructuring the networks or by enforcing tolls, both of which could be very expensive. Instead we use the driver's social network data to facilitate coordinated routing.

Third, traffic congestion is closely related to two fundamental concepts in traffic assignment, namely the user equilibrium (UE) and the system (or social) optimum (SO), formulated by Wardrop [16]. The former states that "under equilibrium conditions, traffic arranges itself in congested networks such that all used routes between an origin-destination (OD) pair have equal and minimum costs, while all those routes that were not used have greater or equal costs." This is Wardrop's first principle, also known as Wardrop's user, or Nash equilibrium. Wardrop's second principle corresponds to the SO and states that under social equilibrium conditions, traffic should be arranged in congested networks in such a way that average (or total) travel cost is minimized. Further, the concept of price of anarchy [7] measures the effect of decentralized decision-making in a system in which its components act selfishly, i.e., it is defined as the ratio between the worst equilibrium (worst UE) and the optimal solution (SO). Therefore, in order to achieve a socially beneficial outcome, it is important to reduce the price of anarchy by establishing optimal flows. In the event of congestion, all drivers may strictly prefer a coordinated outcome to the flow at Nash equilibrium to reduce the price of anarchy. Our approach seeks to achieve the SO in which the Nash flow is strictly Pareto-dominated by the optimal flow.

In order to create optimal flows in transportation networks, we model this problem as a multiagent system (MAS) distributed coordination problem. Drivers traveling through the roads act as agents that make route choice decisions in a distributed fashion. The goal is to create a mechanism for coordinated decision making. We envision a scenario in which a certain fraction of drivers in transportation networks are connected via an online social network traffic app. When a driver encounters congestion in the route she is currently traveling or just finished traveling, she will report the level of congestion in the form of a post in the traffic app. Creating a post can be constrained by the GPS readings so that a driver can only post congestion levels about a route in which she is traveling. Based on the collection of congestion reports, the traffic app continuously computes the utility of each route in real-time. The utility of a route is a measure of traffic moving smoothly on that route and is computed using the collective reported congestion information. Unlike conventional electronic map apps, there is bi-directional information flow in our app, i.e., drivers connected via the app have access to the route utility values which they can choose to use to make route choices.

We propose to augment the traffic app with a traffic route preference function (TRPF) that computes the route utility values. The TRPF is based on real-time congestion data collected from the drivers. Drivers are able to choose a moving time window to compute the route utility based on data from the past *n* hours so that the app is able to improve its prediction over the long run. Our goal is to improve the travel time for individual drivers while also reducing the average travel time of all the drivers. TRPF not only allows users to share relevant information with each other to improve their travel experience, but also offers valuable data that could ultimately help transportation agencies upgrade the way they structure initiatives aimed at improving and streamlining commutes.

To gain a deeper understanding about how this traffic app could improve the system behavior, we perform an extensive simulation study. We choose two simple traffic networks [11], [2] for investigation that are well studied in the literature for resolving the congestion problem. As discussed below, one of the challenges of providing global information about the routes to drivers is that it might lead to oscillation. Using simulations, we determine the optimum parameter settings based on these two simple network scenarios that reduce oscillation and ensure a stable traffic flow across routes. Our primary goal for this paper is to identify the fundamental aspects of TRPF to be able to implement it for larger network scenarios in the future. Simulation results show that the social network based route coordination approach successfully resolves these paradoxes while reducing both congestion and the average travel time.

To summarize, in this paper we emphasize the benefits of harnessing the online social network platform to reduce congestion using a couple of well-known networks as case studies . The main contributions of this paper are:

- For a given number of vehicles traveling from a source to destination, our route-preference approach achieves the socially optimal (or near-optimal) distribution of drivers.

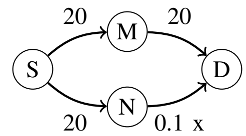

(a) 2 route network

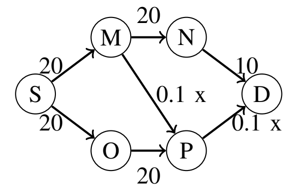

(b) 3 route network

Fig. 1. Modified Pigou's Networks

- We present a MAS coordination model that leverages drivers' social network to facilitate establishing the optimal traffic distribution.
- We show that our model is able to resolve the classic traffic flow paradoxes such as *Pigou* and *Braess* that result from selfish routing.

II. Classic Traffic Scenarios

We describe two networks used to illustrate our approach. The simple network in Fig. 1(a) is a variant of Pigou's example [11] on the effects of selfish routing. Both SMD and SND routes lead from a source S to a destination D. Consider that a flow of 200 want to use the network. The individual costs (latency/travel time) associated with each edge are either constant ($c = 20$) or variable (a function of the flow x using an edge). The cost of SMD and SND are 40 and $20 + 0.1x$ respectively. Route SND is cheaper from each individual driver's perspective. Thus, under the assumption that each driver aims to maximize her utility, we can expect that the entire flow uses SND. Moreover, no driver has incentive to shift to SMD since it also incurs a cost of 40 there. Hence, under UE only SND is used. This is certainly not the best distribution from the overall point of view: if the traffic flow could be coordinated (instead of being a product of selfish routing), the best distribution from the social point of view (SO), would be assigning 100 to each route. The average cost in this case would be 35, which is smaller than that at UE.

Similar conclusions can be reached in the case of three routes (Fig. 1(b)): because SMPD is the cheapest route from the individual point of view, the entire flow of 200 will use SMPD with an average cost of 60. However, the social optimum occurs when 75 use SMPD and 125 use SMND, causing a reduced average cost of 44.375.

Another example vastly used in the literature on selfish routing is due to Braess [2]. An instance of it is depicted here in Fig. 2. Assume flow is 4000 and $c = 45$. The paradox occurs due to the following: Initially, there were just two routes to go from S to D, one via v and one via w (Fig. 2(a)). Since costs are identical, the whole flow splits equally producing an individual cost of $20 + 45$. The traffic authority then decides to construct a high capacity link from v to w (Fig. 2(b)). This new link has cost or latency of 0. Obviously, now, all flow will deviate to route SvwD. This causes high cost in both Sv and wD edges. The individual cost is now $40 + 0 + 40$, which is higher than the one experienced prior to when the additional edge was included.

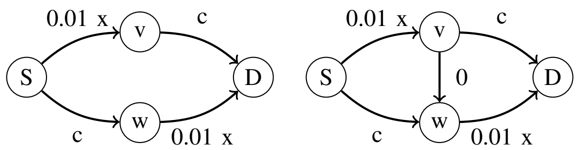

(a) initial network (b) extended network

Fig. 2. Braess' Paradox

Worse, if a single driver decides to take for instance SwD, its travel time increases to 85. Thus it has no incentive to do so. Like the previous examples, we could do better if we could route the flow to the SO in the following way: 1745 take SvD, 1745 take SwD, and 510 take the route via vw. The cost in the latter group is just 45, while it is 67.5 for the others. The average cost is obviously lower than the one obtained by selfish routing (64.68 instead of 80).

III. Related Work

The problem of routing vehicles (or packages in a communication network) has received great attention in the literature. For example, the aforementioned literature on selfish routing and the price of anarchy [7], [13] aims at bounding the worst possible severity of phenomena related to selfish routing. Roughgarden [12] suggests three ways to cope with selfishness: increasing the capacity of the network, influencing traffic with edge taxes (tolls), and routing a small amount of traffic centrally. The first involves social, environmental, and economic costs that lie in the area of political decision-making. Regarding the second direction, Buriol et al. [3] proposed an approach that finds a traffic flow that is both UE and the SO. However, they solve this optimization problem by finding a set of edges in which a number of toll booths should be placed, i.e., they solve a different problem, namely imposing additional costs to make the user equilibrium and the system optimum coincide. As for the third direction, a solution via Stackelberg routing is discussed in [12]. This solution was formulated for communication networks; it is unclear how a portion of the *vehicular* traffic can be *centrally* routed. In the present paper, we argue that traffic routing can be best implemented using route recommendation in conjunction with collaborative congestion data collection.

In the area of guidance and recommendation, there are classical as well as agent-based works that investigate the effects of providing information to drivers. In the interest of space, we just mention that oscillations are seen when everyone receives the same information.

A simple form of social network is reported in Bazzan et al. [1] for a commuting scenario (restricted to a simple binary route choice) in which information can be provided by office mates, by acquaintances, by malicious agents, or by route guidance. It was found that information is beneficial.

Vasserman et al. [15] investigate whether Waze can, in general, implement the socially optimal solution. It differs from ours in the sense that it does not explicitly consider the drivers as source of information; it assumes that all drivers use Waze; and its focus is on incentive compatible policies to compute a mediated equilibria. In the VANET community, there has been some early attempts [4], [8] to consider social relationships between vehicles or drivers. Most of them focus on reproducing the mobility models. However these models aim mainly at detecting communities of vehicles and they do not map the vehicles on a real topological space.

In the same line, other works advocate for providing knowledge to vehicles and for the use of data collected from social networks, in order to help to improve traffic prediction, as in [5], [10]. These however do not discuss the use of the methods in routing scenarios.

IV. Our Approach

In the New Cities Foundation (Connected Commuting) study [9] on the urban commuting problems in San Jose, the central question posed is whether a new level of networking between commuters could enhance the overall commuting experience. One of their key observations is that commuter comments collected by smartphone applications provide valuable high-quality real-time data about commuter sentiment in relation to their commutes. In our approach, users report the congestion level of the route they traveled. Each driver then calculates the Traffic Route Preference Function (TRPF), which fuses the congestion reports into a representation of the utility of each route. The driver can then choose the route that maximizes the TRPF. The goal of this method is for each vehicle to use the congestion reports to choose a route such that average travel time for all vehicles is minimized.

A. Problem Model

We model a commuting scenario. Each edge has an associated cost function, which represents the travel time, or latency, a driver would experience traveling along that edge. In our scenarios, the cost function is linear with respect to the number of drivers on that edge. Given N drivers and k routes from which they can choose, we want drivers to distribute themselves in such a way as to achieve the optimal flow. The optimal number of drivers along a route, O_i $i = 1, 2, ..., k$ can be found by solving a convex optimization problem using CPLEX (quadratic programming).

Next, we introduce the three main parameters of our model: G, p, and T. To model drivers' decisions, we assume that without any external stimulus, drivers will choose the route that has previously given them the lowest travel time. This causes the driver distribution to converge to the UE. However, with probability $1 - G$, a driver will stick with the same route as last time, which is a way for us to model a person's desire to exploit their previously gained experience rather than exploring new options. Parameter G allows us to control the number of agents who change their route each round to determine the effect of this number on the performance of TRPF; especially how it influences oscillation.

We assume that a fraction p of the drivers have access to a Waze-like app creating a dynamic social network. This is the fraction of drivers who will be using the TRPF. When drivers run into congestion, if they are part of the social network, they will report it using one of two possible congestion levels: Level 1: moderate congestion (traffic moving slowly) or Level 2: heavy congestion (standstill traffic). We define a weight metric called the congestion multiplier associated with each congestion level. For Level 1, a congestion multiplier $w = 3$ is used, and for Level 2, $w = 5$. Each vehicle in the social network receives reports from the other vehicles in the social network and calculates the TRPF. Drivers may use a number of rounds, represented by the parameter T, to determine their individual utility value for each route.

B. Traffic Route Preference Function (TRPF)

The TRPF is a measure of whether more drivers are on a route than the optimal distribution of drivers would suggest. It fuses the traffic reports from the drivers in the social network to result in a function that achieves its maximum when fewer drivers are on a route than should be in the SO. The TRPF for route i is a function of the traffic reports for the last T rounds. If t is the current round, $n_i^{(t)}$ is the number of drivers reporting congestion for route i at time t, and $w_i^{(t)}$ is the congestion multiplier for route i at time t, then the TRPF is given by:

$$TRPF_i(t) = 1 - \frac{\sum_{\tau=t-T}^{t} n_i^{(\tau)} w_i^{(\tau)}}{\sum_{\tau=t-T}^{t} n_i^{(\tau)}}$$

This express the total congestion report for route i over the last T rounds, normalized by the number of drivers reporting congestion. The purpose of the subtraction is so that the function is maximized when fewer drivers use route i.

If fewer than O_i drivers (the optimal number of drivers on route i) choose route i, then the TRPF for route i will be 1. The TRPF will decrease as the difference between the number of drivers on a route and the SO of that route increases.

Our approach proceeds in the following steps:

1) Initialization: Each driver will use the TRPF with probability p.
2) Each driver chooses a new route. With probability $1 - G$, a driver will choose the same route as last time. Of the drivers who can change their route, those who use the TRPF will choose the route with the maximum TRPF value and those who do not use the TRPF will choose the route that has historically resulted in the minimum travel time.
3) Drivers travel along the route and receive the travel cost.
4) Drivers who use the TRPF report the congestion levels for their routes.
5) Drivers who use the TRPF receive the reports from other drivers and calculate the new TRPF based on the last T rounds.
6) Repeat, starting from Step 2.

V. Simulation and Results Analysis

We conduct simulations with the following goals: (i) investigate the performance of TRPF for reducing congestion and average cost (travel time) (ii) determine the conditions under which use of TRPF provides optimum performance, i.e., minimum-possible average cost with no/small oscillations in the route choices. We use the previously introduced Modified Pigou's Network with three routes (Fig. 1(b)) and Braess' Network (Fig. 2) with a constant flow of 200 and 4000 drivers respectively to study the effectiveness of our approach. In our experiments, the total number of vehicles on the combined routes for the two networks sometimes exceed the total number of vehicles since these are averages over a number of rounds.

We vary the three parameters G (fraction of drivers that change routes in each round), p (fraction of drivers that use TRPF to choose routes), and T (number of previous rounds used to calculate TRPF) to observe the effect of TRPF on traffic distribution (average number of drivers in each route) and average cost. Each simulation reports the results of 200 rounds. These results are generated by taking the average of last 100 rounds once the traffic distribution becomes stable.

Table I reports the results of simulations for four scenarios in Pigou's and Braess' network. Rows 1-5 of Table I represent scenario 1, in which 20% drivers ($G = 0.2$) may change their routes every journey. Also, drivers having the TRPF app use data from the past one hour ($T = 1$) to compute TRPF values. Among the 20% drivers (who may change their routes), those who do not have the app will choose the historical lowest-cost (selfish) route. If they are already using the historical lowest-cost route, then they do not change routes. However, if some (or the majority) of these drivers have access to the TRPF app, then they switch to the route with highest TRPF value. Parameter p controls the fraction of drivers that use the app to change routes. When the value of p is large, a majority of the drivers use TRPF for choosing a new route. With G and T set at fixed levels (second and third columns), we vary the value of p to observe how it influences the traffic distribution and the average cost.

When $p = 0$, the majority of the drivers choose the historical lowest-cost route, i.e., SMPD in Pigou's network and SvwD in Braess' network. As anticipated in Section II, we observe that a majority of the drivers choose these routes, creating heavy congestion with an increased average cost. As the value of p increases, the traffic distribution in the three routes approaches the SO with a decrease in average cost (as described in section II, at SO the average cost for Pigou's and Braess' network would be 44.38 and 64.68 respectively).

However, Fig. 3 shows that with $p \geq 0.5$ traffic oscillation gradually starts increasing, peaking at $p = 0.9$. The reason is that the majority of the drivers that change their routes use the TRPF app as compared to small values of p (≤ 0.5), in which only a small fraction changes route and the majority stick to their lowest-cost route. That is why higher values of p (≥ 0.5) lead to oscillation. We remark however that the average cost does not decrease at higher values of p. This indicates that oscillation, while not preferable, does not degrade system performance (i.e., does not increase average cost). The standard deviation of the traffic distribution increases with the increase of oscillation for larger p values. We observe a relatively stable traffic distribution for the values

TABLE I

MODIFIED PIGOU'S NETWORK WITH THREE ROUTES & BRAESS' NETWORK: EFFECT OF THE VARIATION OF G, T & P.

				Modified Pigou's Network with 3 Routes								Braess' Network							
	G	T	p	# SMND drivers		# SMPD drivers		# SOPD drivers		Cost		# SvD drivers		# SwD drivers		# SvwD drivers		Cost	
				Avg	Std	Avg	Std	Avg	Std	Avg	Std	Avg	Std	Avg	Std	Avg	Std	Avg	Std
1	0.2	1	0	49.25	2.06	147.47	2.71	3.28	1.70	49.96	0.32	118.02	22.79	116.99	16.51	3764.98	34.54	78.01	0.28
2	0.2	1	0.3	51.56	3.62	137.13	2.45	11.31	2.65	49.29	0.43	627.35	23.43	639.08	19.07	2733.57	22.55	70.93	0.13
3	0.2	1	0.5	85.66	4.34	101.62	2.54	12.71	3.64	45.83	0.28	1041.97	24.37	1055.90	24.68	1902.12	34.01	67.15	0.12
4	0.2	1	0.7	100.25	6.95	85.48	5.07	14.28	5.30	45.13	0.32	1399.68	31.71	1411.03	39.51	1189.28	48.87	65.29	0.09
5	0.2	1	0.9	107.86	7.13	78.61	6.77	13.52	4.42	44.92	0.20	1681.57	139.08	1679.59	144.94	638.83	112.68	64.81	0.14
6	0.2	3	0	49.95	1.77	147.28	1.90	2.77	1.62	49.87	0.24	111.89	13.67	113.7	20.38	3774.4	30.57	78.09	0.25
7	0.2	3	0.3	48.87	3.56	141.93	2.64	9.19	2.82	49.75	0.44	648.46	16.36	649.16	21.00	2702.38	26.37	70.75	0.15
8	0.2	3	0.5	94.51	5.43	95.72	4.37	9.76	3.82	45.32	0.30	990.29	26.74	989.92	25.32	2019.79	33.77	67.58	0.13
9	0.2	3	0.7	110.27	8.26	78.95	7.41	10.72	5.34	44.82	0.25	1419.7	34.72	1416.16	31.73	1164.14	47.76	65.23	0.09
10	0.2	3	0.9	112.87	14.54	73.27	12.79	13.86	8.19	44.98	0.41	1642.97	300.3	1633.96	305.08	723.07	221.18	65.20	0.42
11	0.2	5	0	48.71	2.00	148.13	2.06	3.16	1.47	50.04	0.28	114.01	14.36	112.79	22.80	3773.2	34.15	78.08	0.28
12	0.2	5	0.3	57.17	4.21	134.58	3.17	8.25	3.46	48.67	0.47	615.88	18.11	622.21	20.65	2761.91	21.92	71.09	0.12
13	0.2	5	0.5	72.16	4.19	120.17	3.01	7.67	3.51	47.00	0.37	1015.54	30.73	1013.88	28.43	1970.57	37.82	67.40	0.14
14	0.2	5	0.7	115.77	9.42	75.39	8.86	8.84	4.71	44.72	0.20	1420.65	33.96	1416.98	32.43	1162.36	45.08	65.23	0.08
15	0.2	5	0.9	113.94	20.61	72.09	18.02	13.96	9.53	45.16	0.65	1653.28	498.13	1670.09	494.33	676.63	215.26	65.95	1.18
16	0.5	1	0	49.09	2.29	147.65	2.60	3.26	1.85	49.98	0.33	77.42	11.37	77.83	7.80	3844.75	15.41	78.67	0.13
17	0.5	1	0.3	44.32	6.96	143.13	5.26	12.54	5.90	50.30	0.90	622.95	35.36	629.35	34.22	2747.70	60.17	71.01	0.35
18	0.5	1	0.5	80.15	12.56	103.24	8.89	16.61	9.47	46.31	0.96	999.93	57.22	1006.78	53.95	1993.29	102.33	67.48	0.41
19	0.5	1	0.7	90.72	14.61	90.59	12.09	18.63	12.18	45.73	0.80	1361.34	69.81	1358.16	73.37	1280.50	137.09	65.48	0.31
20	0.5	1	0.9	101.15	22.91	76.92	18.80	21.93	15.63	45.65	0.87	1580.13	440.75	1595.21	417.39	824.66	321.33	65.73	1.01

of p between $0.3 \sim 0.5$ in both Pigou's (Fig. 3(a) & 3(c)) and Braess' (Fig. 3(b) & 3(d)) networks.

In order to observe the effect of parameter T, scenarios 2 and 3 are created by increasing the value of T to 3 and 5 (rows 6-15 of Table I). In other words, the drivers who have the TRPF app use social network data (congestion reports) from the past three and five hours to compute their TRPF values. In rows 6-10 of Table I (scenario 2), when T is set to 3, no further decrease in average cost can be observed.

Similarly, rows 11-15 of Table I (scenario 3) show that with T set to 5, there is no performance improvement when compared with the first scenario reported in rows 1-5 ($T = 1$). This indicates that the past one hour's congestion report is enough to identify the less-congested routes.

Rows 16-20 of Table I show the results of simulations for scenario 4 in which 50% drivers ($G = 0.5$) change their routes every time they begin their journey. We set T to a fixed value of 1 so that drivers having the TRPF app use data from the past one hour to compute their TRPF values. With these two parameters set at fixed values, we vary the value of the parameter p. For Pigou's network, Fig. 3(g) shows that the amplitude of the oscillation is highest at a relatively low value of p (≥ 0.5). In other words, if 50% of the drivers change their routes in the beginning of their journey, then the use of TRPF app results in decreased average cost with increased oscillation. This is reflected by the increased standard deviation values for the average number of vehicles across all the routes for larger values of p ($0.5 \sim 0.9$). We observe a similar situation in Braess' network (Fig. 3(h)).

For the above-mentioned four scenarios, we performed simulations multiple times to observe whether the average number of vehicles (for the three routes) and average cost differ. However, the standard deviations were small. For space constraint, we do not report results from statistical analyses.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we solved the problem of selfish routing in transportation networks by leveraging the driver's social network data. We presented the Traffic Route Preference Function (TRPF) that computes utility values of the alternative routes for a given source-destination pair. TRPF uses posts on congestion levels (moderate or heavy congestion) from the driver's social network for this computation. We emphasized two standard traffic networks, Pigou's and Braess' network, that are used to investigate the phenomenon of selfish routing. Using extensive simulations we have demonstrated that TRPF is able to establish coordination among the drivers for choosing alternative routes. Our main findings are as follows:

- TRPF achieves system optimum (SO) by reducing congestion and decreasing the average travel time.
- TRPF is shown to perform especially well in networks where the Nash flow is strictly pareto-dominated by the system optimum flow.
- Braess' paradox (as well as Pigou's) can be resolved without restructuring or enforcing tolls.
- TRPF approaches SO even when less than 20% of the drivers use TRPF to choose routes.
- Oscillation becomes prominent when 50% or more drivers change routes at the beginning of their journeys. However, oscillation does not exacerbate the performance. This indicates that there is a network-specific range for parameter p that optimizes congestion control.

As future work, we intend to extend our investigation to multi-commodity networks that have more than one origin-destination pair. Also we plan to use learning algorithms so that TRPF can learn how to minimize oscillations.

REFERENCES

[1] A. L. C. Bazzan, M. Fehler, and F. Klügl. Learning to coordinate in a network of social drivers: The role of information. In K. Tuyls, P. J. Hoen, K. Verbeeck, and S. Sen, editors, *Proceedings of the International Workshop on Learning and Adaptation in MAS (LAMAS 2005)*, number 3898 in Lecture Notes in Artificial Intelligence, pages 115–128, 2006.

[2] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258, 1968.

[3] L. S. Buriol, M. J. Hirsh, P. M. Pardalos, T. Querido, M. G. Resende, and M. Ritt. A biased random-key genetic algorithm for road congestion minimization. *Optimization Letters*, 4:619–633, 2010.

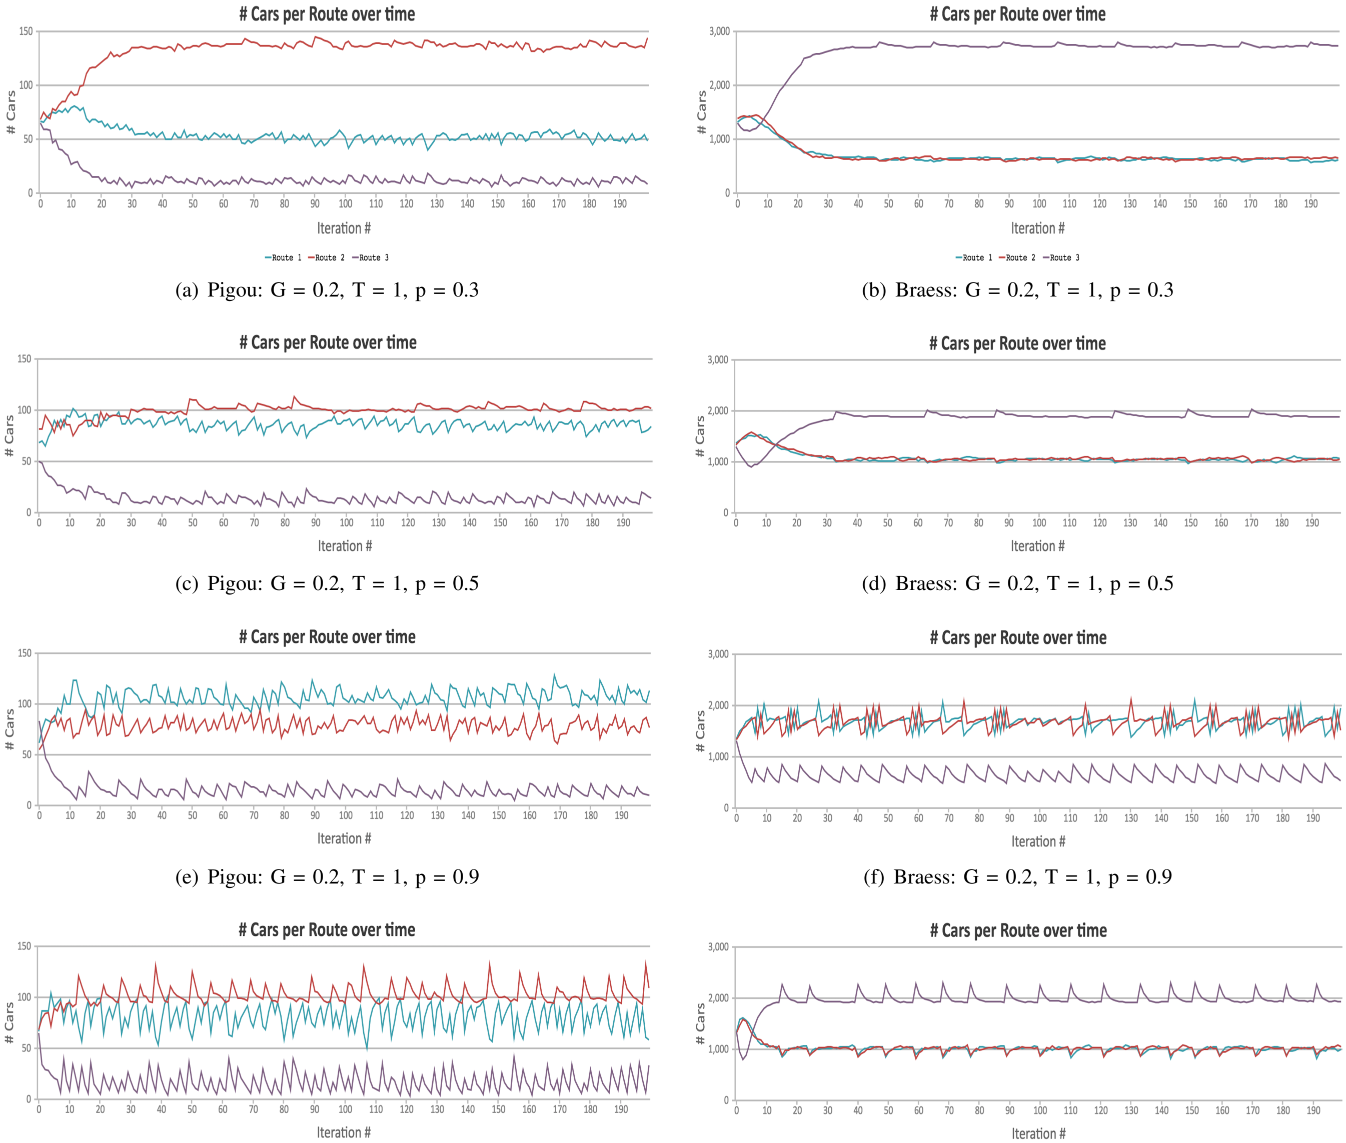

(a) Pigou: G = 0.2, T = 1, p = 0.3

(b) Braess: G = 0.2, T = 1, p = 0.3

(c) Pigou: G = 0.2, T = 1, p = 0.5

(d) Braess: G = 0.2, T = 1, p = 0.5

(e) Pigou: G = 0.2, T = 1, p = 0.9

(f) Braess: G = 0.2, T = 1, p = 0.9

(g) Pigou: G = 0.5, T = 1, p = 0.5

(h) Braess: G = 0.5, T = 1, p = 0.5

Fig. 3. (best seen in color) Traffic Distribution in Modified Pigou's (Route 1 = SMND, Route 2 = SMPD, Route 3 = SOPD) and Braess' (Route 1 = SvD, Route 2 = SvwD, Route 3 = SwD) Network.

[4] A. Gainaru, C. Dobre, and V. Cristea. A realistic mobility model based on social networks for the simulation of vanets. In *VTC Spring*. IEEE, 2009.

[5] J. He, W. Shen, P. Divakaruni, L. Wynter, and R. Lawrence. Improving traffic prediction with tweet semantics. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, pages 1387–1393. AAAI Press, 2013.

[6] Y. A. Korilis, A. A. Lazar, and A. Orda. Achieving network optima using stackelberg routing strategies. *IEEE/ACM Trans. Netw.*, 5(1):161–173, 1997.

[7] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science (STACS)*, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.

[8] M. Musolesi and C. Mascolo. A community based mobility model for ad hoc network research. In *Proceedings of the 2nd International Workshop on Multi-hop Ad Hoc Networks: From Theory to Reality*, pages 31–38, New York, NY, USA, 2006. ACM.

[9] New Cities Foundation. Connected commuting: Research and analysis on the new cities foundation task force in San Jose. Technical report, 2012.

[10] D. Pathania and K. Karlapalem. Social network driven traffic decongestion using near time forecasting. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1761–1762, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

[11] A. Pigou. *The Economics of Welfare*. Palgrave Classics in Economics. Palgrave Macmillan, 1920.

[12] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.

[13] T. Roughgarden and É. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[14] K. Tumer and D. Wolpert. Collective intelligence and Braess' paradox. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 104–109. AAAI Press, 2000.

[15] S. Vasserman, M. Feldman, and A. Hassidim. Implementing the wisdom of waze. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, pages 660–666, 2015.

[16] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institution of Civil Engineers*, volume 1, pages 325–362, 1952.