

# Collaborative Transportation Systems

Michał Piórkowski

LCA, School of Computer and

Communication Sciences

EPFL, 1015 Switzerland

Email: [michal.piorkowski@epfl.ch](mailto:michal.piorkowski@epfl.ch)

**Abstract**—We propose a new class of applications for Intelligent Transportation Systems (ITSs), called *collaborative transportation applications* that aim at solving transportation problems such as congestion and parking. Specifically, we define two applications: *SmartPark* and *SmartRide* that leverage short-range wireless communication. We quantify the potential benefits these collaborative transportation applications can offer to an individual and to the public. To this extent, we conduct both the realistic simulations and the analysis of the performance of a taxi cab fleet from San Francisco. Our analysis shows that both collaborative transportation applications can provide with significant savings in travel times, fuel consumptions, etc. Finally, we discuss the functional requirements of collaborative transportation applications and we present the challenges that these applications are facing.

## I. INTRODUCTION

The economic loss caused by traffic jams in most countries is measured in billions of US dollars yearly. For example, in the year 2007 congestion caused urban Americans to travel 4.2 billion hours more and to purchase an extra 2.9 billion gallons of fuel for a congestion cost of 78 billion US dollars [12]. In many cities 30% of the traffic is caused by drivers looking for parking places [13]. In the year 2008 Transportation Alternatives<sup>1</sup> conducted a twelve-month study on a 15-block area of Columbus Avenue in New York, US (commercial corridor on Manhattans Upper West Side). The data reveals that only while looking for parking, in the survey area, motorists drive 590000 km, wasting 130000 US dollars in gas and generating 325 tons of CO<sub>2</sub> yearly.

Youn et al. in [14] show that drivers seeking the shortest route to a given destination eventually reach a Nash equilibrium, in which no single driver can do any better by changing his strategy unilaterally. It is less efficient than the equilibrium reached when motorists coordinate their movements to benefit the entire group. Relying on the results of realistic simulations, Youn et al. found that the inefficiency caused by selfish drivers can be high - selfish drivers waste 30% more time than they would under socially optimal conditions.

This observation clearly suggests that in order to utilize the public infrastructure optimally and to reduce the travel time, mobile end users need to collaborate. This brings us to the definition of the collaborative transportation applications that are based on the interaction and collaboration between pedestrians, motorists and traffic operators, in order to accomplish

the objective of improving the efficiency of traveling in urban environments.

In this work we first discuss the idea behind collaborative transportation applications. Next, we present two such applications called *SmartRide* and *SmartPark*. Afterwards, we focus on the quantitative benefits of the collaborative transportation applications. Our results are based on the realistic simulations as well as on the analysis of the mobility trace of a taxi fleet from the city of San Francisco in the USA. Finally, we present the research and implementation challenges and discuss the functional requirements of collaborative transportation applications.

## II. FROM COLLECTIVE TO COLLABORATIVE TRANSPORTATION

The typical approach to solving transportation problems related to congestion and parking is to reduce the number of vehicles on the roads. Usually this is achieved by relying on collective transportation (e.g. public transportation), which addresses the general transportation needs. Such generalization may result in certain travel inconveniences, because of multiple transfers and detours or long wait times at off-peak hours. Clearly the collective transportation does not meet the door-to-door transportation demand of an individual.

The road traffic efficiency and safety systems, envisioned by the Intelligent Transportation Systems (ITSs) community, may provide a better solution to such transportation problems. ITSs vary in technologies applied, from basic management systems, such as car navigation, to advanced applications that integrate real-time data from a number of other sources, such as parking systems and weather information. In the last decade, it was the information technology that was the driving force of ITSs. The architecture of the currently available ITSs very often separates sensing from computation, as well as decision making from actuation. For instance, many ITSs collect road traffic information from immobile roadside sensors and process the data off-line at a central server. These road traffic statistics are used afterwards to decide, for example on the optimal traffic-light schedules. In such systems the time between sensing and actuating is long, i.e., on the order of days or weeks. The new wireless and computing technologies (e.g. DSCR and/or ZigBee) provide designers with new ways to create distributed traffic safety and efficiency applications. In fact such technologies allow for mixing sensing, decision making, actuation, computation, and networking. Thanks to this, such

<sup>1</sup><http://www.transalt.org>

novel distributed ITSs can drastically reduce the feedback time by orders of magnitude. As a result the traffic related problems can be reduced, which translates into huge savings.

In the following section we present two instances of the distributed ITSs, which fall into the class of collaborative transportation applications. They leverage short-range wireless technology to solve the transportation problems related to parking and congestion. We describe the way they operate and we present the potential benefits they offer.

### III. THE SMARTPARK APPLICATION

The main goal of SmartPark is to collect information about parking space availability and to coordinate drivers in order to guide them to free parking spots. There already exist centralized solutions to the parking problem that offer similar functionalities as SmartPark. The closest solution is the system offered by the Streetline Networks<sup>2</sup> operating in San Francisco. In the fall of 2008 the company, together with the municipality, started a transportation experiment. They equipped 6'000 on-street parking spots with wireless sensors that are connected to a central server, which is monitoring the parking occupancy. Drivers can be alerted to empty parking places either by displays on street signs, or by looking at maps on screens of mobile phones<sup>3</sup>.

SmartPark aims to take the parking guidance a few steps further. Notably, the SmartPark application should operate on a completely autonomous basis, without any need for a centralized element. To this end, every parking spot is equipped with a sensor that tracks the occupancy and communicates over a short-range wireless interface with other nearby sensors and drivers. Each vehicle is equipped with a wireless communication device that provides a driver with information about parking space availability and guides them by turn-by-turn instructions.

We envision the typical scenario for SmartPark as follows:

- 1) While driving towards his destination, the driver announces to the sensors in the destination area that he wants to find an empty parking spot in the neighborhood of the destination.
- 2) The parking spot sensors nearby the destination have information about parking spot availability in the neighborhood and can react to the incoming requests.
- 3) If there is an empty parking spot nearby the destination, SmartPark provides the driver with the directions to this spot.

#### A. Potential Benefits

Here we investigate to what extent a system such as SmartPark can reduce the urban traffic and the travel time of an individual. For this purpose we conduct a realistic simulation study of a centralized SmartPark system, where an oracle is aware of the current parking availability and provides drivers, willing to park, with turn-by-turn instructions. We compare

this idealized SmartPark system to the reference scenario, where no parking guidance system is available. In other words, there are two different simulation scenarios.

- 1) *coordinated scheme* - In this scenario drivers are aided by an oracle. Whenever a motorist wants to park, the oracle looks up the empty parking spot nearest to the chosen destination and guides the motorist to that spot.
- 2) *non-coordinated scheme* - In this scenario, drivers look for parking spots by themselves.

We consider two metrics. The first metric indicates the overall performance of the SmartPark system and measures its usefulness. Specifically, *time to park* measures the elapsed time between driver's decision to park and actually getting to an available parking spot. The second metric indicates the indirect influence of the SmartPark system on smoothing the urban traffic. For this purpose, we monitor the instantaneous *total traffic density* defined as the number of all active (not parked) vehicles at a given time.

#### Methodology

The goal is to evaluate both schemes in a realistic setting. This requires not only a realistic road network layout and mobility model, but also a realistic parking search strategy. The latter defines the behavior of a driver in search of a parking space. Scientists from the transportation research community identify different types of parking search strategies [2]. In our study we consider the following two.

For *coordinated scheme* the driver knows the position of an available parking space in advance and navigates directly to it. We implement this in our simulation as follows. If a driver wants to park, he sends a request to the server, which responds with a reservation for the parking spot nearest to the destination. The navigation system guides the driver to the parking spot, and when the vehicle approaches the parking spot, it uses the reservation to check into the parking spot.

For *non-coordinated scheme* the driver does not know the position of an available parking spot and has to look for it. He starts circling around the destination to find an empty spot (equivalent to Strategy VI from [2]). We implement this in our simulation as follows. If a driver wants to park, he moves to his destination and broadcasts a *request for parking* within a short range distance. We take this distance to be 30 meters, which in real life is an average distance to spot visually a free parking place from behind a windshield, given the obstacles and road layout. If the driver finds an available parking spot, he parks in it immediately. If he does not find one, he will begin a new parking search by checking all nearby road segments (beginning with the nearest one) and periodically rebroadcasting his *request for parking*.

What is common for the coordinated and non-coordinated schemes is that each driver obeys the driving rules and, if she wants to park and finds a free parking spot, she then occupies it immediately.

<sup>2</sup><http://www.streetlinenetworks.com>

<sup>3</sup>Technology Review Find a Parking Space Online:  
<http://www.technologyreview.com/web/21123>



Fig. 1: Road network topology (Boston, Massachusetts) used for evaluating the benefits of the SmartPark application. Picture from OpenStreetMap [8].

TABLE I: Simulation parameters for evaluating the benefits of the SmartPark application

Number of parking spots	50
Parking-to-vehicle ratios	0.5, 0.34 (typical for Zürich), 0.25
Region size	750m x 940m = 0.705km <sup>2</sup>
Parking decision	Within first 20 min, each driver decides to look for a parking spot

### Simulation

In order to simulate both schemes, we use the JiST/SWANS<sup>4</sup> with STRAW<sup>5</sup> simulation environment. STRAW was designed to enhance SWANS with simple but realistic road network based mobility model. Specifically, it provides two mobility models: (i) *Random Street Mobility* - upon arriving at an intersection, the node chooses randomly the road segment to take next, and (ii) *Random Origin-Destination Mobility* - a destination road is chosen randomly, and then the shortest path between the current location and the destination is followed.

For the purpose of the simulation, the initial placements for all drivers are generated once for both scenarios. The same applies to the simulated wireless parking spot sensors. We had to modify the SWANS environment such that each simulated network node is able to take autonomous decisions about its own mobility.

The road network topology used for the simulation is depicted in Figure 1. In Table I we give the main simulation parameters and their corresponding values. In our simulations we use the Random Origin-Destination Mobility model to implement both parking strategies presented earlier. Also, in our simulation scenario there are no transit vehicles, i.e., vehicles that just traverse the simulated area. In this respect, we consider the worst-case transportation scenario, when all drivers in a certain region are looking for a free parking space.

### Results

<sup>4</sup>Java in Simulation Time/Scalable Wireless Ad hoc Network Simulator: <http://jst.ece.cornell.edu>

<sup>5</sup>Vehicular mobility model for network simulations: <http://www.aqualab.cs.northwestern.edu/projects/STRAW/index.php>

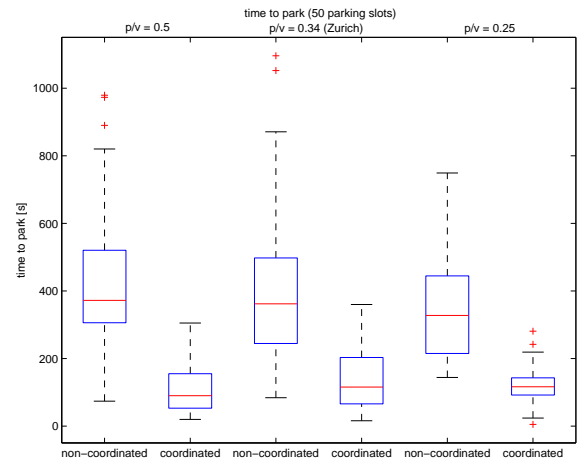


Fig. 2: The box-plot of the population-based average time to park for non-coordinated and coordinated schemes

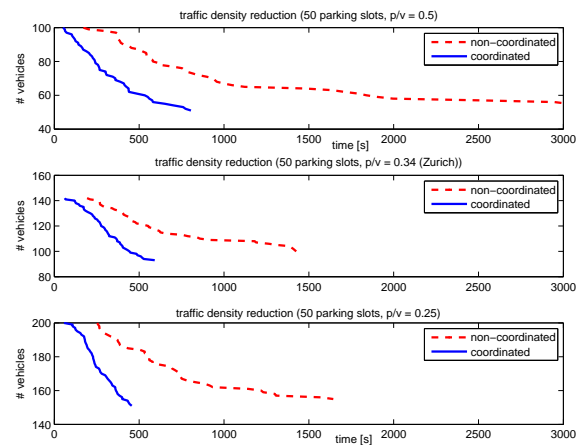


Fig. 3: Traffic density reduction - the total number of active vehicles present on the road network after a certain time. Comparison between non-coordinated and coordinated schemes.

In Figure 2 we depict the quantitative comparison of the population-based average time to park for both the non-coordinated and coordinated schemes. The simulations show that the population-based average time to park is reduced by a factor of up to 3. As a consequence, thanks to the coordinated scheme, cars are removed from the urban traffic at least two times faster than in the case of the non-coordinated scheme. We visualize this urban traffic smoothing effect in Figure 3.

The results of this realistic simulation clearly support the intuition that a collaborative transportation system would be beneficial to the urban environment as a whole.

### IV. THE SMARTRIDE APPLICATION

The main goal of the SmartRide application is to match drivers, who offer their empty seats for a ride, with pedestrians that want to get to their destinations. There exist centralized solutions that facilitate hitchhiking and ridesharing already, e.g. Carlos<sup>6</sup>. Also, with the appearance of location-based services for mobile phones many web-based ridesharing applications are available for the end users, for instance the

<sup>6</sup><http://www.carlos.ch>

Carticipate<sup>7</sup> or OpenRide<sup>8</sup>. The main drawback is that users of such systems have to plan in advance their trips. This works well when the users commute frequently on the same route. Also, such centralized, web-based systems are suited well for long-distance trips. In contrast, the SmartRide application aims mostly at urban, opportunistic trips. As mentioned before, it is complementary to the public transportation and, to some extent, it is similar to the cab service combined with the ridesharing feature. Similarly to the SmartPark application, SmartRide aims to take the ridesharing idea a few steps further, i.e., the SmartRide application should operate in an autonomous fashion, without any need for a centralized element. To this extent, each system participant is equipped with a mobile device that communicates with other devices in vicinity via short-range wireless interface. When the pedestrian, who is looking for a ride, finds an appropriate vehicle, both the pedestrian and the driver are offered with the location of a meeting point. We envision the typical scenario for the SmartRide application as follows:

- 1) The pedestrian announces a ride request to the drivers traveling towards the same destination as the pedestrian.
- 2) The drivers with at least one available seat can accept the incoming request.
- 3) If there is an empty seat nearby the pedestrian, SmartRide provides both the driver and the pedestrian with the directions to the meeting point.

#### A. Potential Benefits

For any transportation system, the proper allocation of resources is an essential condition for a robust and effective operation. The performance of such a system as a whole can be investigated by comparing the capacity of the system to its actual load.

We want to understand what is the use pattern of existing transportation systems similar to SmartRide. This will help us to quantify the benefit the SmartRide application could offer. For this, we investigate the relationship between capacity and load of a transportation system that meets the door-to-door transportation demand but does not offer ridesharing. Notably, we study the occupancy pattern for a taxi fleet of 536 cabs operating in San Francisco over a period of 26 days [10]. This dataset was obtained through the periodic mining of the website of the cabspotting project that provides an open API for accessing the taxi information in real time<sup>9</sup>. Apart from the regular spatiotemporal information about a taxi, each mobility trace contains basic information also about taxi occupancy - every time the taxi is occupied an appropriate notification is also sent to the central server. Thus it is possible not only to track cab movements but also the duration, the origin and the destination of each trip.

The individual capacity  $C_i$  offered by taxi  $i$  is expressed as the total time  $i$  operates, i.e., the time it actually moves around

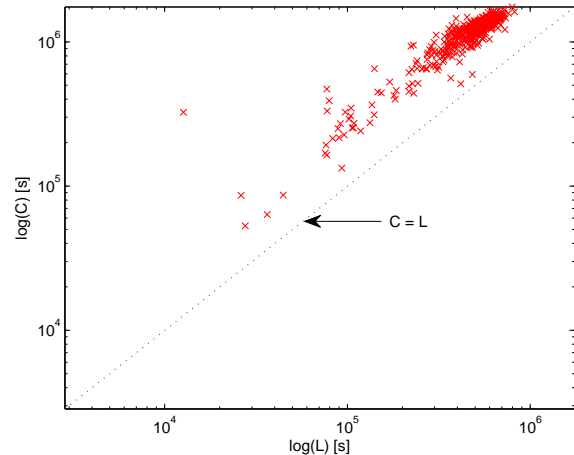


Fig. 4: Capacity (total time cab was *operating*) versus load (total time cab was *occupied*) for the San Francisco cabs in loglog scale. Each marker corresponds to one cab. The diagonal represents the line of the maximum efficiency.

in the city. The individual load  $L_i$  of taxi  $i$  is the total time the taxi is occupied. By comparing the two, we are able to quantify the efficiency of such transportation system. Figure 4 shows the relation between the load and capacity of a taxi fleet from San Francisco in a loglog scale.

We find that there is a certain gap between the actual system use and the line of maximum efficiency ( $C = L$ ). It indicates that the allocation of physical resources (seats) is far from effective. In addition, we define the efficiency coefficient for a transportation system as the ratio between the total load and the total capacity:  $\alpha = \frac{\sum_i L_i}{\sum_i C_i}$ . We find that, for the San Francisco cab system,  $\alpha$  is 0.418, which suggests that for almost 60% of the time, cabs are moving without passengers.

Even if we assume that the same amount of time is required for cabs to get to the passengers waiting for a ride, it is still striking that for almost 20% of the time cabs circle around the city with empty seats, contributing to the congestion. This analysis shows that there is a certain room for improvement that an application such as SmartRide could offer.

#### B. Costs

In contrast to SmartPark, the SmartRide application induces certain costs on drivers and passengers that decide to share their ride. This is because in order to pick up passengers for ridesharing additional distances need to be traveled by the drivers. It means that certain trajectory modifications will occur.

As shown in the previous subsection the potential for fuel savings and pollution reduction from increased ridesharing can be significant. However this benefit is offset by the need for extra travel such that ridesharing passengers can get on-board. Jacobson and King in [6], using driving data for car and small truck fleets, have studied the tradeoff between saving fuel and spending time to pickup additional passengers. They find that, on average, ridesharing is unlikely to be attractive to many travelers even at high fuel prices. But they observe

<sup>7</sup><http://www.carticipate.com>

<sup>8</sup><http://open-ride.com>

<sup>9</sup><http://cabspotting.org/api>



however that even small increase in parking fees and road toll costs makes ridesharing much more attractive to travelers than driving alone. This suggests that ridesharing can be very popular in urban environments, where congestion charges are imposed on drivers, e.g. in cities like London or Stockholm.

## V. FUNCTIONAL REQUIREMENTS AND CHALLENGES

The attractiveness of the proposed applications stems from the fact that nodes forming an ad hoc network can function without any fixed infrastructure, e.g. a cellular network. This however presents new challenges, related to both, to the technology used and to the application design, which do not exist in traditional, centralized systems.

We believe that the first real-life implementation of collaborative transportation applications will combine vehicular, pocket-switched and sensor networks. This is because none of the available short-range wireless technologies exclusively allows us to deploy a fully distributed system that enables communication between drivers, pedestrians and parking sensors. Specifically, the DSRC wireless technology envisioned for vehicular networks was designed to support high data rate communications for long distances (up to 1000 meters). The most popular short-range wireless technology for handsets, i.e., Bluetooth supports moderate connection speeds (up to 3Mbps) and has the communication range typically around 30 meters. The ZigBee standard has the typical communication range of up to 100 meters and provides connection speeds between 10kbps and 250kbps. In contrast to Bluetooth, ZigBee was designed to target large-scale deployments such as citywide monitoring and control, e.g. Göteborg Energi<sup>10</sup>. We believe that all three can be successfully combined to take advantage of their complementary strengths.

Concerning the application design, both SmartPark and SmartRide must ensure that any free seat/parking spot accepts only one request at any point in time. Also, both applications need to provide certain mechanism for the *producer* (a parking spot sensor or a driver with an empty seat) and the *consumer* (a pedestrian looking for a ride or a driver looking for a parking spot) to authenticate each other when they meet. For the SmartRide application there is an additional need for a specific negotiation protocol between passengers onboard and potential passengers that allows to put a price on ridesharing depending on the fuel costs, extra travel time, etc. (cf. subsection IV-B).

We observe that for such collaborative transportation applications there is a need for a robust communication service that ensures that the relevant information is delivered to end users. The traditional MANETs pay more attention to relatively long-distance point-to-point routing. In fact, if two nodes want to communicate via MANET with the IP-like architecture, then before the communication takes place the following happens. First, valid IP addresses must be assigned to all mobile nodes in the network. Second, the sender node has to discover and locate the destination node and its corresponding IP addresses

(e.g. by flooding the network). Third, a route has to be setup from the source node to the destination node for example by an active routing protocol such as AODV [9] or DSR [7]. This approach appears to be problematic, in the context of distributed collaborative transportation applications, as the identity and location is combined in an IP address and thus forces mobile nodes to change their identity when moving between subnets. For the collaborative transportation applications, in order to avoid such problems, the addressing and naming might be a hybrid, multi-level scheme, where the context information plays an important role. For instance, in the MAC layer, the geo-location of nodes can be used for directional broadcast, whereas in the networking layer, the road network topology can be used to establish a forwarding path, similar to the source routing scheme in IP [11]. In fact the geo-location information is a necessary component of applications like SmartRide and SmartPark, therefore networking services such as geocasting or geomulticasting are a natural choice for collaborative transportation applications.

In a distributed setting, when only short-range communication is in use, the end-to-end communication is often not possible due to the high mobility of the system elements. It implies that the networking services supporting distributed collaborative transportation applications will have to cope with the high dynamics of network topology. In urban environments the average vehicle speed is around 40 km/h (11.1 m/s). Human speed is around 5 km/h (1.4 m/s). Considering a realistic transmission range of 100 meters for a short-range communication device in an urban environment, we find that the communication interval varies from approximately 11 seconds (for two vehicles traveling in opposite directions on a straight road) to 140 seconds (for a human passing by an immobile sensor on a straight road). The lack of end-to-end connectivity, relatively short duration of contacts between nodes and the high dynamics of network topology gives rise to a number of specific research problems, while designing effective communication schemes for distributed collaborative transportation applications. For instance, we suggest a change in the network service model where the mobility can no longer be abstracted away. Instead such networking services should exploit the mobility of system's elements. Therefore one of the challenges is to design robust *mobility-centric* communication protocols for such applications.

## VI. RELATED WORK

The Mobile Century group demonstrates in [1] that GPS-enabled mobile phones can be used as sensors for vehicle traffic monitoring purposes. Their experiment shows the possibility of reconstructing traffic conditions using a penetration rate of equipped vehicles less than 5%. The proposed system is a centralized system that comprises four entities: the probes (i.e. GPS-equipped mobile phones traveling onboard vehicles), a cellular network operator, an ID proxy server, and a traffic monitoring and reconstruction system. It would be possible to implement both SmartRide and SmartPark relying on the architecture proposed by the authors of [1], e.g. drivers would

<sup>10</sup>[www.zigbee.org/imwp/download.asp?ContentID=16120](http://www.zigbee.org/imwp/download.asp?ContentID=16120)

notify the system through mobile phones about events such as starting a journey towards a specific destination or vacating a place in the car or a parking spot. Note that our approach is orthogonal to theirs since we postulate to use fully distributed system.

The idea for solving the parking problem by leveraging the wireless sensor network technology has been discussed in the literature frequently in the past decade. In [3] Bi et al. present a parking management system based on wireless sensor network that consists of three types of nodes. The *monitoring node* is installed near each parking space to detect its occupancy. *Guiding nodes* control the LEDs display on the turnoffs of the road to help drivers finding an empty parking space. *Sink node* takes charge of collecting data from the whole network and transmitting the data to a management station, which is set in the monitoring room of the park. According to the topological structure of the parking lot and the occupied status of each parking space, the management station finds the shortest path to an empty spot and transmits it to proper guiding nodes. Caliskan et al. propose in [4] a topology independent, scalable information dissemination protocol for spatio-temporal traffic information such as parking place availability using a vehicular network solely.

Concerning the SmartRide-like architectures, the closest is the one proposed by Zhou et al. in [15]. They present an application called EZCab, which allows pedestrians to book nearby cabs using their mobile phones or PDAs equipped with short-range wireless interfaces. This application upon a query issued by a pedestrian discovers and books free cabs using the VANET network. It does not provide any ride-sharing features.

In [5] Gidofalvi and Pedersen address the problem of cab-sharing. Specifically, they propose a centralized cab-sharing system that provides door-to-door transportation service, while minimizing the transportation costs and increasing the efficiency of cab space utilization. They address the problem of grouping nearby cab requests into cab-shares. They report that in exchange for a short (up to 15 minutes) wait time, their centralized cabsharing service can group together requests in a way that efficiently utilizes resources and provides significant savings to the user.

## VII. CONCLUSIONS AND FUTURE WORK

This paper introduces a new class of applications for distributed ITSs. We quantify the potential benefits these applications can offer. For this purpose we perform realistic simulation of the first proposed application. Then we analyze the performance of a real-life transportation system aiming to achieve the same goal as our second application, i.e., offering door-to-door transportation service for pedestrians. Our results clearly show that the collaborative transportation applications are beneficial to all the participants of the urban life. We also discuss the functional requirements and present some of the research and technical challenges the collaborative transportation applications will have to solve.

In the future work we intend to study the collaborative transportation applications in more detail. Specifically, we

want to focus on the information dissemination protocols that would support such applications, operating in fully distributed environment. For this purpose we need to understand both the mobility characteristics of different urban traffic participants, as well as the city travel patterns, e.g. the spatiotemporal distribution of city-trips.

## ACKNOWLEDGMENT

The authors are indebted to Matthias Grossglauser for help in improving the manuscript. The work presented in this paper has been supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5006-67322.

## REFERENCES

- [1] S. Amin, S. Andrews, S. Apte, J. Arnold, J. Ban, M. Benko, A. M. Bayen, B. Chiou, C. Claudel, C. Claudel, T. Dodson, O. Elhamshary, C. Flens-Batina, M. Gruteser, J.-C. Herrera, R. Herring, B. Hoh, Q. Jacobson, M. Kumar, T. Iwuchukwu, J. Lew, X. Litrico, L. Luddington, J. Margulici, A. Mortazavi, X. Pan, T. Rabbani, T. Racine, E. Sherlock-Thomas, D. Sutter, A. Tinka, K. Tracton, O.-P. Tossavainen, T. West, A. Wiedmer, D. B. Work, and Q. Wu. Mobile Century - Using GPS Mobile Phones as Traffic Sensors: A Field Experiment. In *The 15th World Congress on Intelligent Transportation Systems*, New York, November 2008.
- [2] K. Axhausen. Parking Demand and its Characteristics, 1994. <http://rouages.met.wallonie.be/metpub/src/trafic06/p27.html>.
- [3] Y.-Z. Bi, L.-M. Sun, H.-S. Zhu, T.-X. Yan, and Z.-J. Luo. A Parking Management System Based on Wireless Sensor Network. *ACTA AUTOMATICA SINICA*, 32(6):968–977, November 2006.
- [4] M. Caliskan, D. Graupner, and M. Mauve. Decentralized Discovery of Free Parking Places. In *The 3rd International Workshop on Vehicular Ad hoc NETWORKS, VANET*, pages 30–39, New York, NY, USA, 2006. ACM.
- [5] G. Gidofalvi, T. B. Pedersen, T. Risch, and E. Zeidler. Highly Scalable Trip Grouping for Large-Scale Collective Transportation Systems. In *EDBT '08: Proceedings of the 11th International Conference on Extending Database Technology*, pages 678–689, New York, NY, USA, 2008. ACM.
- [6] S. H. Jacobson and D. M. King. Fuel saving and ridesharing in the US: Motivations, limitations and opportunities. *Transportation Research Part D: Transport and Environment*, 14(1):14–21, January 2009. <http://dx.doi.org/10.1016/j.trd.2008.10.001>.
- [7] D. Johnson, Y. Hu, and D. Maltz. (RFC4728) The Dynamic Source Routing Protocol (DSR) for Mobile Ad Hoc Networks for IPv4, February 2007. <http://www.ietf.org/rfc/rfc4728.txt>.
- [8] OpenStreetMap - a repository of free geospatial data. <http://www.openstreetmap.org>.
- [9] C. Perkins, E. Belding-Royer, and S. Das. (RFC3561) Ad hoc On-Demand Distance Vector (AODV) Routing, July 2003. <http://www.ietf.org/rfc/rfc3561.txt>.
- [10] M. Piorkowski, N. Sarafjanovic-Djukic, and M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). <http://crawdad.cs.dartmouth.edu/epfl/mobility>.
- [11] (RFC791) Internet Protocol Specification, September 1981. <http://www.ietf.org/rfc/rfc791.txt>.
- [12] D. Schrank and T. Lomax. Urban mobility report. Texas Transportation Institute, The Texas A and M University System, 2007. [http://tti.tamu.edu/documents/ums/mobility\\_report\\_2007\\_wappx.pdf](http://tti.tamu.edu/documents/ums/mobility_report_2007_wappx.pdf).
- [13] D. C. Shoup. Cruising for parking. *Transport Policy*, 13(6):479–486, November 2006.
- [14] H. Youn, M. T. Gastner, and H. Jeong. Price of Anarchy in Transportation Networks: Efficiency and Optimality Control. *Physical Review Letters*, 101(12), September 2008.
- [15] P. Zhou, T. Nadeem, P. Kang, C. Borcea, and L. Iftode. EZCab: A Cab Booking Application Using Short-Range Wireless Communication. In *The 3rd IEEE International Conference on Pervasive Computing and Communications, PerCom'05*, pages 27–38, 2005.