

Utilizing the Infrastructure to Assist Autonomous Vehicles in a Mobility on Demand Context

B. Rebsamen*, T. Bandyopadhyay*, T. Wongpiromsarn*, S. Kim*,
Z. J. Chong†, B. Qin†, M. H. Ang Jr.†,
E. Frazzoli*‡, D. Rus*‡

* Singapore-MIT Alliance for
Research and Technology,
Singapore

† National University of
Singapore,
Dpt. of Mech. Eng.

‡ Massachusetts Institute
of Technology,
Cambridge, MA., USA.

Abstract—In this paper we describe an autonomous vehicle that aims at providing shared transportation services in a mobility on demand context. As the service is limited to a known urban environment, prior knowledge of the environment can be exploited, as well as existing infrastructure sensors such as security cameras. We argue that utilizing infrastructure sensors yields greater safety of operation and allows reduction in the number of sensors required on-board, hereby reducing the cost of the vehicle. We describe the role that infrastructure sensors can play and show the resulting improved performances of the system, supported by simulation and field experiment results.

I. INTRODUCTION

As cities become bigger and more crowded, a new paradigm for urban transportation is required. Mobility-on-demand (MoD) systems operating on shared resources are becoming a viable alternative [6]. While successful in increasing resource utilization, these systems usually require a mean to re-balance the fleet, which remains an open problem for large vehicles such as cars.

Vehicle autonomy is considered as an efficient solution to automatically re-balance vehicles among stations, and enable a one-way vehicle sharing option in MoD systems. The ability of vehicles to drive autonomously in urban environments has been demonstrated, for example, in the 2007 DARPA Urban Challenge (DUC). However, the vehicles that competed in the DUC depend on a variety of sensors, some of which are prohibitively expensive while others are highly specialized, causing the deployment to be economically infeasible.

In this work, we aim at developing an autonomous vehicle that uses minimal sensing and off-the-shelf components to attain the same level of operational ability while making the system economically feasible. We rely on the assumption that the vehicle is always operating in the same environment, which is a-priori known. This allows the vehicle to learn important features of the environment and take advantage of the infrastructure. Infrastructure cameras at an intersection or pedestrian crossing can also help the autonomous vehicle detect oncoming vehicles or pedestrians, allowing the autonomous vehicle to make the decision (whether to go or stop) more efficiently.

The main contribution of this paper is the evaluation of exploiting infrastructure sensors in autonomous driving. In

particular, we show, using both simulations and field testing, that under some conditions, employing infrastructure sensors can significantly improve the performance of the system whereas under some other conditions, the improvement is not clear. We characterize both sets of conditions.

The remainder of the paper is organized as follows. Section II describes the important components of our autonomous vehicle. Section III discusses the use of existing infrastructure sensors to improve safety and efficiency of the autonomous vehicles. Results from both a simulation and a physical experiment are provided in Section IV and Section V. Finally, Section VI concludes the paper and discusses future work.

II. SYSTEM DESCRIPTION

Currently we have a single vehicle providing MoD service between four stations, covering a total distance of about 500m. It is based on an electric golf car, shown on Figure 1, equipped with a minimalistic sensor configuration. We use three 2-D LIDARs, a simple webcam to achieve required perception, two wheel encoders and a dual axis gyroscope to measure the vehicle's speed and orientation.

The vehicle localizes itself using the adaptive Monte Carlo localization (AMCL) approach [8], which uses a particle filter to track the pose of a robot against a known map. While this is traditionally done by matching features from the surrounding buildings, we augmented our map with curb information on those sections of the road with no or too few surrounding buildings. The tilted LIDAR is used to detect the curbs. This yields a precise localization (accuracy in the order of the centimeter). The interested reader is encouraged to consult [7] for more information.

Since the vehicle navigates on a known road network, the navigation problem can be simplified by following predefined routes. Navigation is thus reduced to path following, which is implemented by a pure pursuit controller [5], and speed control. This was described in details in [3].

III. EXPLOITING INFRASTRUCTURE SENSORS

The design philosophy of the DUC vehicles was to provide sufficient sensing capability to the vehicles such that they are able to navigate in a completely unknown environment with



Fig. 1. Our vehicle testbed is built upon an electric golf car, augmented with a drive-by-wire system, and a small amount of sensors and computers.

some degree of reliability. This leads to vehicles being almost completely covered with sensors (see for an example the MIT’s entry in DUC shown on Figure 2) and economically infeasible vehicle costs.

The domain of our work is where the autonomous vehicle operates in a known environment. In this context, the vehicle should use as much a-priori information as possible. For instance, it is standard practice for a robot to use a detailed map of the environment to localize itself and navigate. In this work we propose that the vehicle could also utilize some infrastructure sensors, such as traffic cameras.

In many cases, information from infrastructure sensing is not only beneficial, but also becomes a necessity. Take the example of Figure 3(a) where the view of the vehicle is almost completely occluded by the bus. The knowledge of the vehicles and pedestrians in the occluded region is absolutely important for safe navigation. Being too conservative may cause the traffic flow to freeze up causing unacceptable travel time. In such situations, an infrastructure sensor provides a more complete view of the environment for the vehicle to make decisions, which would be very difficult to achieve from on-board sensors alone.

Another example is show in Figure 3(b), where the pedestrian crossing is almost entirely occluded by the pillars and barriers. At this particular location, as in many other urban locations, a security camera is monitoring the area (in red on the picture), and so it could also be used as an extra sensor for our autonomous vehicle. The most difficult part would only be to obtain the proper authorization to access the information from the camera.

We argue that the main advantages of an infrastructure sensor are that it is shared among all vehicles and that it can be fine tuned to deal with a particular location (e.g. provide an optimal viewpoint). As a consequence, utilizing

infrastructure sensors would reduce the sensing burden on the individual vehicles, since on-board sensors can thus be limited to the strict minimum required for basic perception. Moreover, it would increase the safety and reliability of the system at difficult and dangerous locations, such as pedestrian crossings and intersections. Infrastructure sensors are also easily available in many growing urban communities; such as in Singapore where electronic road pricing is implemented.



Fig. 2. The MIT’s vehicle at the Darpa Urban Challenge.

IV. PEDESTRIAN CROSSING EXPERIMENT

To show the effect of the additional information from the infrastructure, we performed an experiment in which our autonomous vehicle was heading toward a pedestrian crossing. Due to the limitation of on-board sensors and occlusions caused by environmental features and other vehicles on the



(a) Example of occlusion by other vehicles: the bus is blocking the view of the right side of the intersection.



(b) Example of occlusion by infrastructure features: from the vehicle point of view, it is very difficult to detect pedestrians coming out of the building.

Fig. 3. In such situations, an infrastructure camera would provide the autonomous vehicle with useful information on the state of the intersection or pedestrian crossing, which would allow the vehicle to make a safer decision, without being too conservative and stalling the traffic.

road, the pedestrians could not be reliably detected from the vehicle while approaching the crossing (see Figure 3(b)). Thus, the vehicle had to slow down at the pedestrian crossing each time, irrespective of whether there were pedestrians or not.

An infrastructure sensor, in the form of a security camera, was installed above that pedestrian crossing (in red on Figure 3(b)). Pedestrian were detected using the default HOG-based people detection algorithm implemented by OpenCV [1], [4]. This infrastructure sensor then sent a binary information to the golf car whether there were pedestrians about to cross the road or if the region was pedestrian free. Any pedestrian detection would trigger the autonomous vehicle to slow down in anticipation for the pedestrian to cross the road.

The details of that experiment were published in [2]. We observed that when some pedestrians were detected, the vehicle would slow down. Besides, when no pedestrian was detected the vehicle would go through the crossing without slowing down. Thus it appeared that using an infrastructure sensor

improved the traffic flow.

V. PEDESTRIAN CROSSING SIMULATION

In the previous section, we saw that utilizing an infrastructure sensor allowed for a higher motion efficiency: when it had access to the infrastructure sensor, the vehicle only stopped when a pedestrian was detected, whereas it had to stop each time when it had to rely solely on its on-board sensors. Although the difference on the total traveling time was only of a few seconds, it would cause unnecessary inconvenience and discomfort to the passenger.

In this section we will explore what would have happened if we considered a flow of vehicles rather than a single unit. To explore this question, we designed an experiment similar to the experiment described previously but with a flow of vehicles heading towards the pedestrian crossing. Since we only have one vehicle, and in order to explore the impact of some parameters, we conducted this investigation in simulation.

A. Description of the Simulation

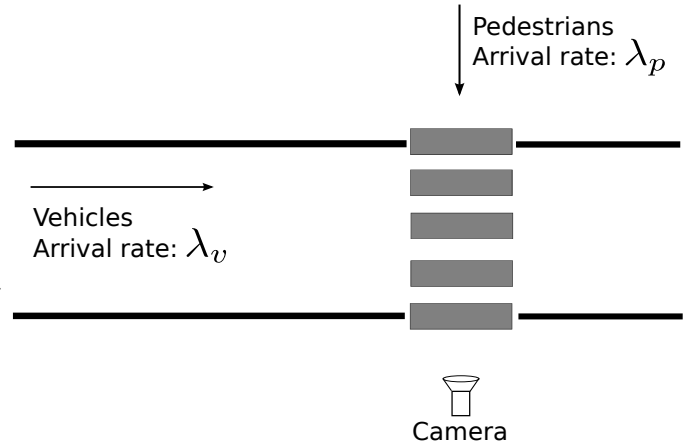


Fig. 4. Schematic of the pedestrian crossing simulation.

The simulation setup is made of a straight road heading towards a pedestrian crossing, as illustrated on Figure 4.

Mobiles (pedestrians and vehicles) are modeled as points moving in one dimension following the second order kinematic equation 1, where a can be either 0, a_{max} or $-a_{max}$.

$$\begin{cases} \ddot{x} = a \\ 0 < \dot{x} < v_{max} \end{cases} \quad (1)$$

Pedestrians appear a few meters before the crossing, with a rate of arrival that follows a Poisson distribution with parameter λ_p . Pedestrians travel at a constant velocity of $1 \text{ m}\cdot\text{s}^{-1}$ ($a = 0, v_0 = v_{max} = 1$).

Vehicles appear 100m before the pedestrian crossing with a rate of arrival that follows a Poisson distribution with parameter λ_v . Vehicles modulate their velocity so as to avoid collision with the vehicle in front and the pedestrians, while trying to travel as fast as possible. Their acceleration profile

TABLE I
SIMULATION RESULTS

V_{max}	λ_p	λ_v													
		0.01		0.10		0.20		0.30		0.40		0.60		0.80	
6	0.00	19	17	19	17	20	17	20	17	22	17	-	17	-	17
	0.01	19	17	19	17	20	17	20	17	21	17	-	17	-	18
	0.04	19	18	20	18	20	18	22	18	27	18	-	18	-	18
	0.07	20	18	20	18	21	19	23	18	37	19	-	20	-	21
	0.10	20	18	21	19	22	19	26	20	-	21	-	21	-	-
	0.20	21	20	24	21	31	23	63	22	-	27	-	-	-	-
	0.40	26	25	42	32	-	-	-	-	-	-	-	-	-	-
10	0.00	13	10	13	10	14	10	14	10	15	10	-	10	-	10
	0.01	13	10	13	10	14	10	14	10	16	10	-	10	-	10
	0.04	14	11	14	11	15	11	16	11	29	11	-	12	-	12
	0.07	14	11	14	11	16	11	18	11	25	12	-	13	-	14
	0.10	14	12	14	11	20	12	-	13	-	13	-	16	-	19
	0.20	15	13	18	14	-	14	-	15	-	19	-	-	-	-
	0.40	19	17	-	23	-	-	-	-	-	-	-	-	-	-
14	0.00	11	7	11	7	12	7	12	7	12	7	-	7	-	7
	0.01	11	7	11	7	12	7	12	7	15	7	-	7	-	7
	0.04	11	8	12	8	13	8	14	8	-	8	-	8	-	8
	0.07	12	8	12	9	13	9	-	9	-	9	-	10	-	12
	0.10	12	9	12	9	14	9	-	9	-	10	-	10	-	12
	0.20	14	10	-	11	-	10	-	13	-	15	-	-	-	-
	0.40	18	15	-	12	-	-	-	-	-	-	-	-	-	-

is set to $a_{max} = 3m.s^{-2}$, and their maximum velocity v_{max} is a parameter of the simulation.

As with the experiment described above, we assume that because of occlusion from some environmental features (e.g. walls, pillars, etc.), the vehicles on the road cannot see the pedestrians coming to the pedestrian crossing. However an infrastructure sensor is available and can communicate to the vehicles the position of the pedestrians. We simulated two types of vehicles.

The first type of vehicles do not have access to the infrastructure sensor. Irrespective of whether a pedestrian is approaching or not, those vehicles have to slow down before the pedestrian crossing until they come to a stop, and resume their motion only if no pedestrian is crossing. These vehicles will hereafter be referred to as base vehicles.

The other type of vehicles communicate with the infrastructure sensor. Based on this information, the vehicles infer how long it will take for an approaching pedestrian to reach the beginning of the pedestrian crossing and stop only if they have to. These vehicles will hereafter be referred to as infra vehicles.

Optionally, the simulation can be run in real time, allowing simultaneous graphical visualization. Figure 5 shows a snapshot of the simulation for $\lambda_p = 0.01$, $\lambda_v = 0.6$ and $v_{max} = 6m.s^{-1}$, showing that the base vehicles (in green) form a traffic jam, whereas the infra vehicle (in blue) move smoothly through the pedestrian crossing.

B. Results

We measured the vehicles' transit time, defined by the time taken to cross the pedestrian crossing, starting 100m before. Table I shows, for various combinations of λ_p , λ_v and v_{max} , the average transit time of base vehicles (first column) and infra vehicles (second column). In cases where the vehicles form a traffic jam, the average transit time is meaningless




Fig. 5. Snapshot of the simulation. The thick gray line at the bottom right corner represents the pedestrian crossing. The red ball represents a pedestrian, which has just crossed the road. Vehicles are coming from the top left corner of the image. The green balls represent the vehicles that do not have access to the infrastructure sensor (base vehicles), the blue balls those that do (infra vehicles).

since individual transit times keep increasing as the simulation progresses (unstable condition). These cases are marked by a dash in place of the average transit time.

From Table I it can be observed that transit times increase with λ_p and λ_v , until a point where the traffic jam condition occurs. Besides, we can see that transit times are always higher for base vehicles than for infra vehicles, and that the traffic jam condition occurs at lower values of λ_p and λ_v for the base vehicles than for the infra vehicles. It can also be observed that

the traffic jam condition occurs at lower values of λ_p and λ_v for higher values of v_{max} .

These observations confirm our idea that having access to the infrastructure sensor improves the traffic flow, with a substantial gain for heavy traffic conditions (high values of λ_v) with few pedestrians (low values of λ_p).

VI. CONCLUSION

In this paper we argued that for autonomous vehicles operating in a known environment, the complexity of the system can be brought down by relying on a-priori information and infrastructure sensors. Firstly, disposing of an accurate map and a network of predefined routes simplifies the localization and navigation tasks. Secondly, integrating information coming from infrastructure sensors allows to negotiate more efficiently some of the most difficult situations, such as intersections and pedestrian crossings.

This results in safer and cheaper vehicles, hence easier to deploy in a Mobility-on-Demand scenario. Moreover, as the number of vehicles increases, the additional cost incurred by the modification of the infrastructure would represent a smaller fraction of the total cost.

Currently we are working toward more sophisticated ways to integrate the information coming from different sources, both on-board and from the infrastructure, taking into account various sources of uncertainties such as transmission errors, or damaged sensors.

REFERENCES

- [1] G. Bradski, "The OpenCV Library," *Dr. Dobbs's Journal of Software Tools*, 2000.
- [2] Z. J. Chong, B. Qin, T. Bandyopadhyay, T. Wongpiromsarn, E. S. Rankin, M. H. Ang Jr., E. Frazzoli, D. Rus, D. Hsu, and K. H. Low, "Autonomous Personal Vehicle for the First- and Last-Mile Transportation Services," in *IEEE International Conference on Robotics, Automation and Mechatronics (RAM)*, 2011.
- [3] Z. J. Chong, B. Qin, T. Bandyopadhyay, T. Wongpiromsarn, B. Rebsamen, P. Dai, E. S. Rankin, and M. H. Ang Jr., "Autonomy for mobility on demand," in *IEEE International Conference on Intelligent Autonomous Systems (to be published)*, 2012.
- [4] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 1, 2005, pp. 886–893.
- [5] Y. Kuwata, J. Teo, G. Fiore, S. Karaman, E. Frazzoli, and J. P. How. Real-Time Motion Planning With Applications to Autonomous Urban Driving. <http://dspace.mit.edu/openaccess-disseminate/1721.1/52527>.
- [6] W. J. Mitchell, C. E. Borroni-Bird, and L. D. Burns., *Reinventing the Automobile: Personal Urban Mobility for the 21st Century*. Cambridge, MA: The MIT Press, 2010.
- [7] B. Qin, Z. J. Chong, T. Bandyopadhyay, M. H. Ang Jr., E. Frazzoli, and D. Rus, "Curb-Intersection Feature Based Monte Carlo Localization on Urban Roads," in *IEEE International Conference on Robotics and Automation*, 2012.
- [8] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT Press, 2005.