2008 IEEE/RSJ International Conference on Intelligent Robots and Systems
Acropolis Convention Center
Nice, France, Sept, 22-26, 2008

A Perception Mechanism for Supporting Autonomous Intersection Handling in Urban Driving

Young-Woo Seo
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh PA 15213, USA
young-woo.seo@ri.cmu.edu

Chris Urmson
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh PA 15213, USA
curmson@ri.cmu.edu

***Abstract*— Knowledge of the driving environment is essential for robotic vehicles to comply with traffic rules while autonomously traversing intersections. However, due to limited sensing coverage and continuous changes in driving conditions, rigidly-mounted sensors may not guarantee coverage of all regions of interest, all the time. Unobserved regions around intersections increase uncertainty in driving conditions. This paper describes a dynamic sensor planning method that searches for the optimal angles of two pointable sensors to maximally cover relevant unobserved regions of interest. The obtained angles are used to adjust the orientations of the pointable sensors and hence reduce uncertainty around intersections. Simulation results show that the sensor planning method increases the percentage of covered area. In addition to the sensor pointing problem, we provide an initial discussion of how to reason about occlusions in an urban environment. Occlusions caused by structures and other environmental features can increase uncertainty in driving environments. An occlusion handling method is used to detect these occlusions and enable our vehicle to model the presence of occluded regions. An awareness of occlusions enables safer driving decisions.**

I. INTRODUCTION

Driving through an intersection is a complex task. For safe driving, a human driver must perform various activities in parallel: interpret traffic rules based on the type of intersection; continuously observe oncoming traffic; judge when it is safe to merge; steer the vehicle along the lane; prepare for the unexpected, and so on.

These parallel and complex tasks make it challenging for robotic vehicles to handle intersections autonomously. In particular, when a robotic vehicle approaches an intersection, it needs to figure out the structure of the intersection and the corresponding traffic rules. The robotic vehicle must comply with traffic rules and needs to know when and where to look, in order to effectively deal with traffic. It must pay attention to this relevant information and steer itself through the intersection. Two types of information are necessary to autonomously accomplish intersection driving: road network information and moving obstacle information. Road network information is generally static, and may be provided a priori. It provides vehicles with geometric and topological information expressing how vehicles are expected to travel through an intersection. Moving obstacle information must be generated in real time, and will generally come from onboard sensors such as ladar, radar, vision, etc. It provides more dynamic information about moving obstacles (e.g., other vehicles) or unexpected events (e.g., stalled vehicles, accidents, etc.). The fields of view of rigidly mounted sensors are generally limited and driving conditions are continuously changing. Thus fixed sensors may not guarantee coverage of all regions of interest, all the time. Unobserved regions increase uncertainty in driving conditions. Limited knowledge about the environment can make it difficult for a robotic vehicle to sufficiently analyze the current state of the environment and consequently drive safely.

In this paper, we describe a perception mechanism that reduces uncertainty around intersections for autonomous urban driving. It consists of two components: a dynamic sensor planning algorithm and a method for modeling occlusions. Our dynamic sensor planning method is designed to cover relevant regions, which are not observed by fixed sensors, by dynamically pointing movable sensors. When regions of interest are occluded by urban structures, uncertainty in occluded regions is inevitable. An occlusion handling method is used to handle uncertainty caused by these occlusions.

Methods for adjusting the orientation of sensors to obtain better estimates of an environment have been intensively investigated in the robotics community. A common way of implementing this idea is to change sensor poses by utilizing degrees of freedom offered by robot mobility. For instance, Spletzer and Taylor present a team formation algorithm that changes the pose of an individual mobile robot equipped with an omnidirectional vision sensor [5]. Their primary interest is to find an optimal team formation in order to reduce the object tracking error, whereas our interest is to find the optimal orientation of a rotational sensor to reduce uncertainty in driving conditions. The "next best view" (NBV) problem [4] is similar to the problem of our dynamic sensor planning, in that both concern planning of sensor pose for a better estimation of a given environment. The primary difference is that given previous measurements of an object, a NBV solution attempts to determine the next best pose of a sensor for obtaining the complete surface geometry of the object. By contrast, our method does not utilize a history of sensor measurements since we are not interested in building high resolution models of fixed geometry, instead we are interested in observing regions of interest for moving vehicles. Another approach was proposed by Stamos and Allen in which an interactive sensor planning system suggests a set of admissible camera orientations satisfying given constraints [6]. This work is similar to ours in that both approaches seek the optimal sensor orientation for the maximal coverage of environment. However, their approach addresses sensor placement onto a particular configuration of a world whereas our approach concerns the orientation of sensors for continuously changing driving conditions. An alternative approach is the use of active vision techniques which share the same aspect of our dynamic sensor planning problem. In particular, in active vision techniques, algorithms are intended to find the optimal next viewpoint to reduce the error of reconstruction of objects' 3D shape [8] or scenes' 3D structure [2]. Specifically, Wenhardt and his colleagues proposed a method of finding the optimal next viewpoint to reduce reconstruction error. The optimality of a viewpoint is evaluated by its reduction of uncertainty given previously obtained images. In other words, the optimal viewpoint is one that minimizes the conditional entropy [8]. Information entropy has thus been utilized to address uncertainty in environment. For our problem, information entropy is used to model uncertainty in our ability to detect vehicles approaching in intersection. A *good* sensor configuration is one that minimizes the number of interesting regions (in bound road lanes) that are not observed by the vehicles' sensors.

Numerous urban structures (e.g., buildings, large traffic signs, tunnels, etc.) complicate autonomous intersection handling. Depending on the size of buildings and the structure of intersections, a vehicle's field of view may be partially or completely blocked for a period of time as the vehicle approaches the intersection. The presence of these static obstacles generates occlusions in a vehicle's view. For instance, when a vehicle stops at a "cross" type intersection, buildings may block all or some of its view of the crossroad. Thus the traffic condition of the crossroad is completely uncertain. In order to reduce uncertainty, a human driver would creep their vehicles toward the center of an intersection until they could obtain a sufficient view of the lane. An implementation of such a behavior is hard to achieve in practice as it requires considering hard to quantify factors such as the cost of not completing the mission (through inaction) and the risk of causing a traffic accident. Alternatively the use of a special range finder, which can see through manmade structures, might reduce uncertainty caused by occlusions. However we are unaware of an appropriate sensing technology in reality. Hence, we present an occlusion handling method that allows our robotic vehicle to be aware of the existence of occluded or unseen regions. Such an awareness of occlusions can help the robot make conservative and therefore safe judgements about entering intersections.

To the best of our knowledge, handling occlusion while traversing an intersection has not been investigated intensively. In the realm of developing autonomous ground vehicles or driving assistance technology, occlusions have been addressed in the application of vehicle tracking. Fayad and Cherfaoui propose a method of tracking partially occluded objects from laser scans [1]. They predefine five different classes of objects, which are partially observed due to occlusion or sensing limits, they use to estimate the dimension (i.e., width and length) of objects and track them using a Kalman filter. Zhang and his colleagues propose an appearance-based object tracking method that combines three different appearance models: *fixed*, *fast-change*, and *eigenspace* to cope with variations due to illumniation, pose changes, and occlusions [9]. In addition, they utilize robust statistics to tackle occluded objects. While their work addresses occlusions to improve the performance of object tracking, our work addresses occlusions for safe driving.

The remainder of this paper is organized as follows. Section II-A describes how we model instantaneous occlusions around the vehicle by using sensors with fixed poses. Section II-B presents a dynamic sensor planning method designed to cover regions of interest not observed by fixed sensors. Our occlusion handling method appears in section II-C. Conclusions and future work are presented in section III.

Fig. 1. The view of our robotic vehicle equipped with various sensors. Three different types of sensors are mainly discussed in this paper: A Velodyne HDL-64 at the top, two IBEO Alasca XT on both sides of the hood, and two ARS-300 installed on the panhead platforms on the driver and passenger side.

II. Perception Mechanism for Supporting Autonomous Intersection Handling

A. Modeling Instantaneous Occlusions

This section describes how we model instantaneous occlusions in the vicinity of the robotic vehicle based on data from fixed-orientation onboard sensors. The vehicle (Figure 1) was developed and tested for the DARPA Urban Challenge [7]. The Urban Challenge is a robotic car competition in which competitors need to finish missions that include parking and traversing intersections as part of driving a 60 mile-long course through an urban environment while abiding by California traffic rules. [1] We installed various sensors on our robotic vehicle for effective traffic handling and safe city driving. This paper considers three different types of these sensors: a Velodyne HDL-64 lidar, two IBEO Alasca XT lidars and five ARS-300 radar sensors (three with fixed poses and two on movable platforms). The sensor measurements

[1] Visit the following for more information, http://www.darpa.mil/grandchallenge/index.asp

(a) A satellite image of our Robot City (RC) test site. The polygon in white depicts a road network.

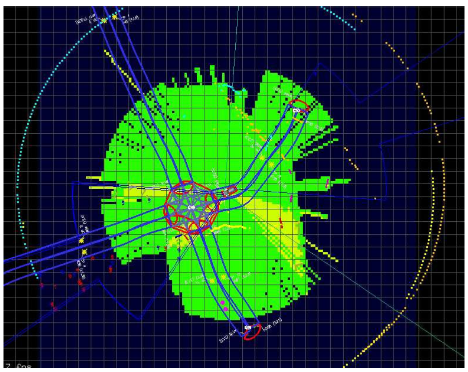

(b) A snapshot taken from our simulator shows the visibility map. A red rectangle at the center of the circular region is our robotic vehicle whereas the transparent (dark blue) square grid represents the range of the map. The radius of the circular region is 50 meters. The green regions are updated by information from Velodyne whereas yellow regions are updated by information from IBEOs. The size of the visibility map is 150 × 150 meters and the size of a cell is 1 × 1 meter. This visibility map information is updated at 10Hz.

Fig. 2. The visibility map.

from two different types of lidars (Velodyne and IBEO) are used to provide the vehicle with *visibility map*: a planar representation of the environment that tells the vehicle where it can see and where it cannot see at any given time. We present the visibility map as a rectangular grid in which we need to determine whether individual cells are visible at a given time. To this end, we first obtain the range measurements from three lidar sensors: a Velodyne and two IBEOs. These measurements are centered on the current pose of the vehicle. Sub-meter accurate pose estimates are provided by a combination of GPS and inertial data by an Applanix POS-LV200 at 100Hz. The IBEO sensors provide a set of raw range measurements while the Velodyne data is provided in the form of an *obstacle map* that represents which cells in the grid are occupied. By tracing a ray from the measurements to the vehicle's current location, we know whether the region, represented by a series of world coordinates on the ray, is visible from the current location. As the vehicle drives through the world, the visibility map is continuously updated, generating an instantaneous model of where the vehicle can see. The satellite image in the Figure 2(a) shows one of our test sites; Figure 2(b) shows visibility information generated at that location.

B. Dynamic Sensor Planning

The visibility map describes the area over which fixed sensors provide coverage. This coverage is sufficient for our robotic vehicle to accomplish most of its tasks. However, due to variability of intersection geometry, the rigidly-mounted sensors cannot guarantee coverage of all regions of interest all the time. Hence it is necessary to have another field of view that is flexibly adjustable with respect to the varying structure of intersections so as to maximally cover unobserved regions. To this end, we installed two pointable (or panhead) platforms for long-range sensors.

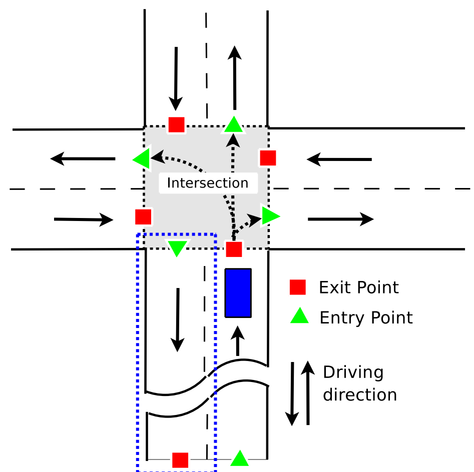

Fig. 3. An illustration of the internal representation of a road network. A square (gray) box represents an intersection that connects more than one road. A dashed-line (blue) rectangle depicts a road lane that is represented by a number of features such as the locations of entry/exit points and a driving direction. Three dashed arrow lines depict all the possible trajectories of our robotic vehicle represented in a (blue) filled rectangle.

These pointable sensors are intended to ensure maximum coverage of regions of interest. The regions of interest, called *observation zones*, are defined by polygonal boundaries of road lanes for oncoming traffic and all exit points of an intersection. As shown in Figure 3, an exit point is the location where a car exits its lane to enter another road. In Figure 3, our vehicle, depicted as a rectangle (blue), has three possible entry points to go to from the exit point, which it is currently approaching. From the road network information, our behavioral software system identifies regions of interest as a set of observation zones. Through this mechanism the behavioral reasoning layer informs the perception system, prior to reaching the intersection, where it should focus attention to enable the vehicle to safely traverse the intersection. Note that we assume that geometrical errors in identifying observation zones are negligible because the pose estimate provided to by the perception system is sub-meter accurate with respect to the road. Figure 4 illustrates the fixed sensing coverage around our robotic vehicle. As we described earlier, the omnidirectional sensing coverage is built by combining sensor measurements from two types of lidar sensors (Velodyne and IBEOs) and five ARS-300 radar sensors which are capable of detecting moving obstacles at long ranges. In Figure 4, there are seven observation zones defined around the intersection: four over four exit points and three of them over road lanes approaching the intersection.

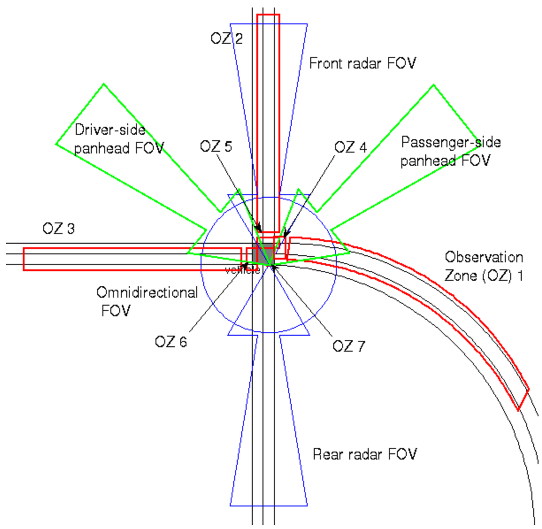

Fig. 4. Three different types of fixed sensors provide our robotic vehicle with an integrated field of views. There are five radar sensors: two (solid blue line) at the front bumper, one (solid blue line) at the rear bumper, and two (thick green line) installed on the panhead platform. The initial orientation of two radars on the movable platform is 45 degree from the vehicle's forward direction. The 50 meter circle represents the omnidirectional sensing coverage and provided by the Velodyne and the IBEOs. There are 7 different observation zones (red polygons, annotated OZ $x, (x = 1, ..., 7)$ defined over this intersection.

Due to variability in possible intersection shapes and sizes, a fixed sensor arrangement may not be able to cover all necessary observation zones all of the time. For example, in Figure 4, the observation zones numbered 4 to 7 are located inside the fixed-sensor coverage whereas most of the regions in observation zones 1 and 3 remain uncovered. Thus the goal of dynamically changing the orientation of the pointable sensors is to find orientations that maximally cover the observation zones. We approach this task in terms of reducing uncertainty. If a configuration of sensor orientations does not cover some observation zones, there is uncertainty. The fewer regions covered the more uncertain we are about approaching traffic. To address this problem, we need first to quantitatively measure the uncertainty of a driving condition. Although complex driving conditions constitute uncertainty, we consider that uncertainty is primarily caused by the current configuration of sensors. In other words, a good arrangement of sensor orientations would cover most of the regions of interest and hence the robot would be given a sufficient amount of information about the environment.

Let us denote $O_k = \{o_{k,1}, o_{k,2}, ..., o_{k,n}\}$ a set of observation zones in an intersection k and $A(O_k) = \{\alpha_{k,1}, \alpha_{k,2}, ..., \alpha_{k,n}\}$, the covered area of corresponding observation zones by fixed sensors. In particular, $\alpha_{k,i} \in [\frac{1}{n}, 1]$ is the area of the ith observation zone in the kth intersection which our fixed sensing coverage can observe partially or completely. According to this covered area value, we can measure the uncertainty in an observation zone by using $\log_2\left(\frac{1}{\alpha_{k,i}}\right) \in [0, \log_2(n)]$. In particular, if an observation zone is not observed at all (i.e., $\alpha_{k,i} = \frac{1}{n}$), the uncertainty in that observation zone is $\log_2(n)$. By contrast, it is not uncertain at all if the observation zone is completely covered by our fixed sensing coverage, $\log_2(\frac{1}{1}) = 0$.

Given a set of observation zones and coverage of those regions by fixed sensors, information entropy is utilized to measure the average amount of uncertainty around an intersection.

$$H(A(O_k)) = E\left[\log_2 \frac{1}{\alpha_{k,i}}\right] = -\sum_{i=1}^{n} \alpha_{k,i} \log_2 \alpha_{k,i} \quad (1)$$

where $0 \leq H(A(O_k)) \leq \log_2(n)$ [3]. $H(A(O_k))$ is zero when observation zones of an intersection are completely covered by sensing coverage whereas H has its maximum value if none of the observation zones are covered. From this formulation, we can restate our problem in a more concise way: find the angles of two pointable sensors that maximally cover unobserved regions of a set of observation zones. Figure 5 illustrates the fixed sensing coverage provided by combining omnidirectional coverage with the front and rear radars' fields of view. The domain of angles for our dynamic sensor planning method is defined as

$$\Theta = I_d \cup I_p = [\theta_0, -\pi] \cup [\theta_0, \pi]$$

where $I_d = [\theta_0 - \frac{\delta}{2}, -\pi + \frac{\delta}{2}]$ is the angle range available on the driver-side, $I_p = [\theta_0 + \frac{\delta}{2}, \pi - \frac{\delta}{2}]$ is the range of the available angles on the passenger-side, θ_0 is the orientation angle of the fixed radar sensor at the vehicle front, and δ is the width of radar cone. For example, in the Figure 5, suppose that θ_0 is $\frac{\pi}{3}$ and δ is $\frac{\pi}{3}$. I_d is defined as $I_d = [\frac{\pi}{3} - \frac{\pi}{6}, \pi + \frac{\pi}{6}]$ and I_p is defined as $I_p = [\frac{\pi}{3} + \frac{\pi}{6}, \pi - \frac{\pi}{6}]$, respectively. Once the domain of angles is defined, the "panhead planner" searches for a pair of angles, $\theta_d^* \in I_d$ and $\theta_p^* \in I_p$ that minimize the entropy of the unobserved regions by pointing the sensors to maximize the sensing coverage of observation zones.

In summary, when our robotic vehicle is approaching an intersection and its distance to an exit point is less than a threshold (e.g., 5 meters), the panhead planner begins to search for the pair of angles that minimizes the entropy in equation 1. First, the planner selects a candidate angle from the domain of angles. Second, it computes the covered areas of observation zones by the candidate angle and calculate the entropy. Finally, it returns the angle that minimizes the entropy.

To verify the usefulness of our sensor planning method, we test it with four different simulated driving scenarios at

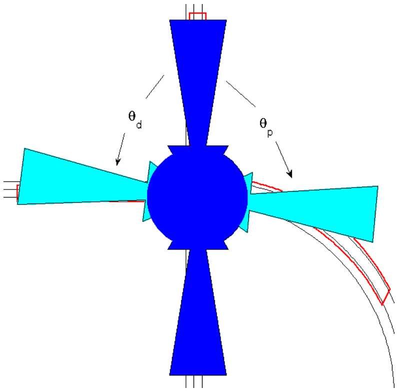

Fig. 5. The union of omnidirectional sensing coverage and the front and rear radar fields of view provides a fixed sensor coverage polygon depicted in the filled (blue) shape. The complement of the union of the observation zones and the fixed sensor coverage polygon is the domain of angles for our dynamic sensor planning method.

different sites. Through the simulations, we would like to see how much uncertainty is decreased and what percentage of the covered area is increased by the proposed sensor planning method. Thirty-one different intersections from four different driving scenarios were tested in terms of two measures: entropy for measuring intersection uncertainty and percentage of covered area for measuring the covered area of observation zones. For each driving scenario, we simulate the car driving autonomously on a test course. While the vehicle traverses each of the simulated intersections, we compared intersection uncertainty and percentage of covered area before and after applying the proposed sensor planning. Table I presents simulation results. Each row represents the result from a test site simulation. For example, there are 15 observation zones defined over 4 different intersections of the "BeaveRun (BR)" test site. The sum of individual intersections' initial entropies is 7.49 which was reduced to 2.76 by fixed sensing coverage and further to 0.28 by sensor planning. For the percentage of covered area, sensor planning was able to cover 2493.79 (square meters, m^2) out of 2700.03 (m^2) whereas the fixed sensors only covered 1395.81 (m^2). Since all the observation zones in the "DPG (DP)" test site are completely covered by fixed sensing coverage, there is no need for planning sensors' orientations dynamically. By contrast, in the "CastleAFB (CA)" simulation, fixed sensing coverage observed only 34.31% of the union of observation zones, thus sensor pointing was particularly important here. Our method increased the covered region of observation zones from 2649.90 (m^2) to 5313.96 (m^2) out of 7724.52 (m^2). Not surprisingly fixed-sensor coverage was useful to reduce uncertainty in intersections. In fact, 45.48% (6939.90 / 15259.00 m^2) of all the observation zones were covered by fixed sensors. However, when fixed sensing coverage was combined with our sensor planning method, the coverage increased greatly. Figure 6 summarizes these results. Sensor planning was able to decrease uncertainty to 6% of its original value while increasing sensing coverage from 45% to 81%.

	Uncertainty			Coverage		
	before	after	initial	before	after	total
BR (4/15)	2.76	0.28	7.49	1395.81	2493.79	2700.03
RC (14/39)	1.64	0.13	19.75	2530.74	4314.14	4471.01
DP (6/16)	0.00	0.00	0.00	363.45	363.45	363.45
CA (7/39)	5.37	2.96	16.12	2649.90	5313.96	7724.52
(31/109)	9.77	3.37	43.36	6939.90	12485.00	15259.00

TABLE I

A COMPARISON OF UNCERTAINTY AND PERCENTAGE OF COVERED AREA BEFORE AND AFTER DYNAMIC SENSOR PLANNING.

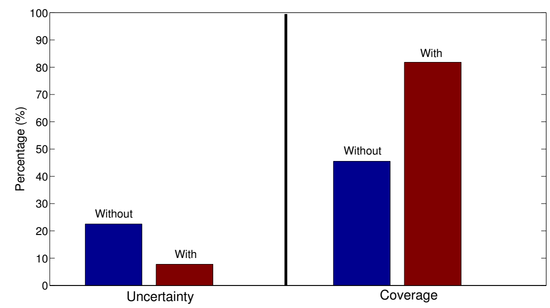

Fig. 6. A summary of simulation results shows a comparison of the relative performance with and without applying our dynamic sensor planning method.

C. Simple Method of Handling Permanent Occlusions

Urban structures around intersections can generate occlusions that lead to unsafe driving conditions. These occlusions increase uncertainty in intersections because occluded regions can not be observed and thus unidentified vehicles may be in these regions. Since it is difficult to determine the correct value (or range of values) of uncertainty for detecting occlusions, we develop a simple, but practical method for detecting occlusions. We analyze a time history of the visibility map. At each timestep, we build a visibility map at the current pose of our robotic vehicle from range sensor measurements. As our vehicle drives, the list of updated cells changes. If a cell is observed more than once, its value is accumulated. The value of the ith cell in the temporal visibility map at the given time t can be occupied, empty or unobserved, $cell(t, i) \in \{\text{occupied}, \text{empty}, \text{unobserved}\}$. This accumulated value of an individual cell indicates how long the cell has been in a particular state. For example, if we count how many unit times an individual cell has not been observed within the range of the visibility map, we would know how long that cell has been occluded.

Consider Figure 7 as a canonical example of occlusion caused by static obstacles. There is a building between our vehicle and an oncoming car. It is a yield intersection where our vehicle slowly approaches a yield line. Since the relative location of the building occludes the lanes behind it, we

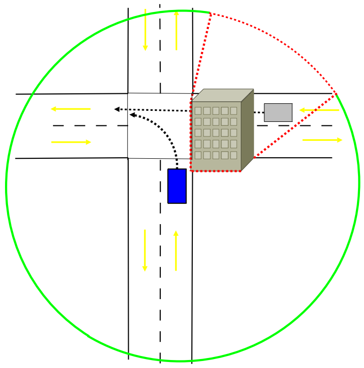

Fig. 7. A canonical example of an occlusion at an intersection.

would not be sure of the existence of another vehicle until having a clear line of sight on that lane.

A function, $\psi(T, i)$, is used to determine whether the ith cell is occluded at a given time, T.

$$\psi(T, i) = \begin{cases} 1, & \text{if } \sum_{t=t_0}^{T} I(cell(t, i) = \text{unobserved}) \geq \tau \\ 0, & \text{otherwise} \end{cases}$$

where t_0 is the time that a cell begins to update, $I(x = c)$ is an indicator function which is 1 if the condition is satisfied, and 0 otherwise. This temporal information is designed to be updated at 10Hz. We set the value of τ to 10 empirically, meaning that an individual cell has not been observed more than 1 second and it is regarded as occluded.

We again utilize observation zones for detecting occlusions. As mentioned previously, observation zones are defined around an intersection and include all exit locations and a polygonal region around the incoming traffic lane. For detecting an occlusion, we identify occluded cells using the function $\psi(T, i)$. Once we detect an occlusion, we place a phantom vehicle in the corresponding region. This ensures that our vehicle applies safe-driving features, such as distance keeping, to the phantom vehicle instead of immediately proceeding through an intersection.

III. Conclusions and Future Work

To autonomously navigate an intersection successfully, it is essential for robotic vehicles to obtain a set of relevant information about driving conditions. Otherwise a robotic vehicle would not be able to safely interact with other vehicles and hence would be unable to comply with traffic rules. However it is difficult to maintain relevant information due to frequent changes in driving conditions and structures in the driving environment. Insufficient information increases uncertainty in driving conditions. To handle uncertainty, this paper presents two different methods. To decrease uncertainty, we propose a sensor planning method to maximize the benefit of pointable sensors. Information entropy is utilized to model uncertainty in intersections. Our robotic vehicle looks for a pair of optimal angles, which maximize the coverage of its sensors to minimize uncertainty, and uses the obtained angles to change the orientations of pointable long-range sensors. To tackle uncertainty caused by occlusions, we develop an occlusion-handling method. Instead of considering the geometry of occluding obstacles, we utilize temporal visibility information to determine the presence of occlusions. Such an awareness of occlusions enhances the safe driving capability of our robotic vehicle.

As future work, we aim to utilize the temporal visibility information to improve the performance of moving obstacle tracking. Specifically, tracked vehicles might disappear temporarily from our vehicle's view due to occlusions instantaneously caused by other vehicles. Instead of reinitiating the state of the tracked object whenever it appears in our view, it would be useful if a tracker could know whether a seemingly new track is the reappearance of a previously tracked vehicle. This temporal information might be useful in combination with a local visibility map in order to estimate global traffic patterns by accumulating these pieces of information over time. For instance, an analysis of this saved local information could tell us how crowded a driving environment is and could help predict what the condition of upcoming roads might be. As an extension of our occlusion handling method, we would like to investigate statistical machine learning methods that are capable of capturing spatial and temporal relations among sensor measurements. By understanding those relations, we might be able to predict what is behind occluded regions based on the history of sensor measurements.

IV. ACKNOWLEDGMENTS

This work would not have been possible without the dedicated efforts of the Tartan Racing team and the generous support of our sponsors including General Motors, Caterpillar, and Continental. This work was further supported by DARPA under contract HR0011-06-C-0142.

References

[1] Fadi Fayad and Veronique Cherfaoui, Tracking objects using a laser scanner in driving situation based on modeling target shape, In *Proceedings of IEEE Intelligent Vehicle Symposium* (IV-07), 2007

[2] Eric Marchand and Francois Chaumette, Active vision for complete scene reconstruction and exploration, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(1): 65-72, 1999.

[3] David J.C. Mackay, *Information Theory, Inference, and Learning Algorithms*, Cambridge Univ. Press, 2003.

[4] Richard Pito, A solution to the next best view problem for automated surface acquisition, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(10): 1016-1030, 1999.

[5] John R. Spletzer and Camillo J. Taylor, Dynamic sensor planning and control for optimally tracking targets, *The International Journal of Robotics Research*, 22(1): 7-20, 2003.

[6] Ioannis Stamos and Peter K. Allen, Interactive sensor planning, In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pp. 489-494, 1998.

[7] Chris Urmson *et al.*, Autonomous driving in urban environments: Boss and the Urban Challenge, *Journal of Field Robotics*, to appear.

[8] S. Wenhardt, B. Deutsch, J. Hornegger, H. Niemann, and J. Denzler, An information theoretic approach for next best view planning in 3-D reconstruction, In *Proceedings of 18th International Conference on Pattern Recognition* (ICPR-06), Vol. 1, pp. 103-106, 2006.

[9] Bo Zhang, Weifeng Tian, and Zhihua Jin, Efficient hybrid appearance model for object tracking with occlusion handling, *Optical Engineering*, Vol. 46, Issue 8, 2007.