An Approach to Predicting the Location of Moving Objects During On-Road Navigation

Craig Schlenoff, Raj Madhavan, Stephen Balakirsky
Intelligent Systems Division
National Institute of Standards and Technology
100 Bureau Drive, Stop 8230
Gaithersburg, MD 20899
Phone: 301-975-3456 F: 301-990-9688
craig.schlenoff@nist.gov, raj.madhavan@nist.gov, stephen.balakirsky@nist.gov

Abstract

For an autonomous vehicle to navigate in real-time within a dynamic environment, it must be able to respond to moving objects. In particular, it must be able to predict, with appropriate levels of confidence, where those objects are expected to be at times in the future. It must then capture this information internally in its world model in a format amenable for planners that intend to use it.

In this paper, we provide an overview of a framework to address the challenges involved in predicting and representing the future location of moving objects. This framework uses a multi-representational approach to model information about moving objects, thus allowing for planners that require different forms of knowledge representation. We then describe a probabilistic, logic-based algorithm to predict the future location of vehicles in an on-road environment. Included in this discussion are the factors that affect the probabilities associated with various actions that the moving object may take.

1 Introduction and Related Research

For an autonomous vehicle to navigate in real-time within a dynamic environment, it must be able to respond to moving objects. In particular, it must be able to predict, with appropriate levels of confidence, where those objects are expected to be at times in the future. It must therefore capture this information internally in its world model in a format amenable for planners that intend to use it.

In this paper, we introduce a framework to address the challenges involved in predicting and representing the future location of moving objects. This framework uses a multi-representational approach to model information about moving objects, thus allowing for planners that require different forms of knowledge representation. In addition, this framework accounts for different factors that would influence the future location of moving objects, including the environment it is in, *a priori* maps, mobility characteristics, the object's intention and indicators, environmental conditions, etc.

This framework is applicable to both on-road and off-road driving. However, for the remainder of this paper, we will be focusing on the more interesting problem of on-road navigation where road networks provide a constrained environment in which to navigate, and as such, introduce a number of additional factors that could influence the probability of other moving objects behaving in certain ways. These factors and their corresponding influences are the focus of this paper.

We are not aware of any efforts in the literature that have addressed the development of a framework for combined moving objects representation and prediction. However, there have been efforts focusing on individual components of this framework that could be leveraged, specifically in moving object representation. Very limited work exists in the representation of moving objects. Firby [4] uses NaTs (navigation templates) as a symbolic representation of static and dynamic sensed obstacles to drive a robot's motors to respond quickly to moving objects. Gueting [5] extends database structures to allow for the representation of dynamic attributes (i.e., ones that change over time) and also extends the database's query language to allow for simplified querying of the values of dynamic attributes. Singhal [10] introduces the concept of dynamic occupancy grids which allow each cell to have a state vector which contains information such as a probabilistic estimate of the entity's identity, location, and characteristics (such as velocity, acceleration) along with global probability distribution functions.

In the literature, it is common to find methodologies that predict where an object will be in the next one or two sensor images. This form of predicting the future location of moving objects for a relatively small number of time steps into the future (short-term prediction) is useful in determining where the object will be in the next sensor image so as to perform object tracking. This has been a well-researched area in which approaches including Kalman filters [3] and Bayesian-based methods [11] have shown good results. In this paper, we concentrate on long-term prediction of sensor images, i.e., predicting the position of objects 10's or 100's of sensor images in the future. For autonomous navigation, planners plan over time horizons anywhere from milliseconds to tens of seconds for path planning and obstacle avoidance. As such, prediction algorithms need to be able to generate both short and long term predictions to accommodate the needs of planners.

In this paper, we introduce an approach to predicting the future location of moving objects with the following characteristics:

- The predictions are received by the planner and are used for path planning and obstacle avoidance
- Predictions are made at longer time horizons, on the order of 10's or 100's of sensor images into the future
- Constraints on the environment are explicitly taken into account, such that only legal and possible actions are considered during prediction
- A logic-based approach is used to associate probabilities with various actions the moving object may take.

In Section 2, we provide an overview of the moving object framework. In Section 3, we describe the prediction algorithms and apply them to on-road driving. In Section 4, we discuss the influencing factors and constraints on motion that affect the probabilities associated with actions used in the prediction algorithms. In Section 5, we discuss the implications of applying this approach to an existing planner. In Section 6, we conclude the paper and discuss future work.

2 Moving Object Framework

The moving object framework provides a mechanism to apply appropriate prediction algorithms and representational approaches in order to fully capture the information needed to navigate in the presence of moving objects. This framework is shown in Figure 1.

We are assuming that the processed sensor data will be provided as input to the framework (as shown in the left-most box). Specifically, the information that we are assuming will be provided at pre-defined time intervals includes:

- The perceived dimensions of objects in the environment
- The location of objects in the environment
- The object's velocity and direction
- The color of the object

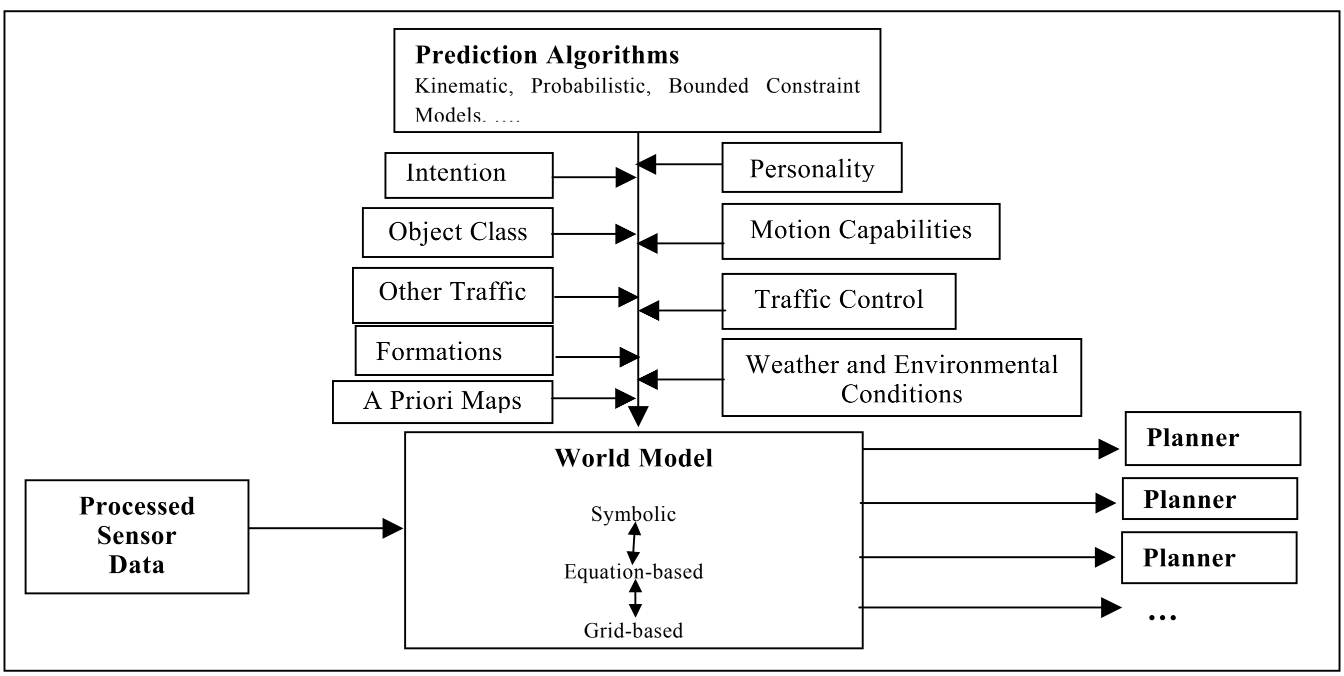

Figure 1: Moving Object Framework

We are also assuming knowledge about the environment in which the vehicle is navigating. This could take the form of *a priori* maps containing road networks and terrain characteristics, or could be dynamically generated based upon sensory input and processing.

In the framework, we attempt to classify the moving object based upon the processed sensor data. In the case of on-road driving, simply classifying the objects as vehicles (cars, motorcycles, trucks, buses, emergency vehicles), pedestrians, animals, or debris is enough for the purpose of motion prediction. We introduce a fairly simplistic object classification algorithm in [8] to provide the level of classification necessary to allow for informed moving object prediction.

Based upon the environment we are in, we employ different types of prediction algorithms. For off-road navigation, we have employed a bank of Kalman filters to predict the future location of the moving objects in the environment [6]. For on-road navigation, we are developing logic-based prediction algorithms that are intended to function in constrained environments. This logic-based approach is discussed in Section 3 of this paper.

Information about the moving object and its possible future locations are stored in the vehicle's world model in a multi-representational format. Information about the instances of object classes encountered in the environment (e.g., vehicle, animal, pedestrian, debris) is stored in a symbolic knowledge base with links to *a priori* detailed information about the corresponding object class. Based on the object class and the environment in which it is in, prediction algorithms, as discussed in the previous paragraph, are linked to the symbolic representations of the object. The results from these prediction algorithms are instantiated and represented in a time-based grid representation and provided to lower-level planners. A detailed discussion of these representation formalisms and their interactions can be found in [9].

One of the major advantages the proposed moving object framework provides is the ability to represent information about the moving object in many different, inter-related representations. It is expected that this moving object framework will provide information to planners that require fundamentally different kinds of underlying representations. For example, a planner that is planning for short time horizons (on the order of a few seconds) may require a grid-based representation describing occupancy probabilities of locations in space while a planner that plans at a longer time horizon may require a symbolic representation that describes characteristics of objects and equations governing their motion as opposed to locations in space. By using an interconnected multi-representational approach, we are able to provide information to the planner that is at a level of abstraction appropriate to its planning requirements. More information on planning in the presence of moving objects can be found in Section 5.

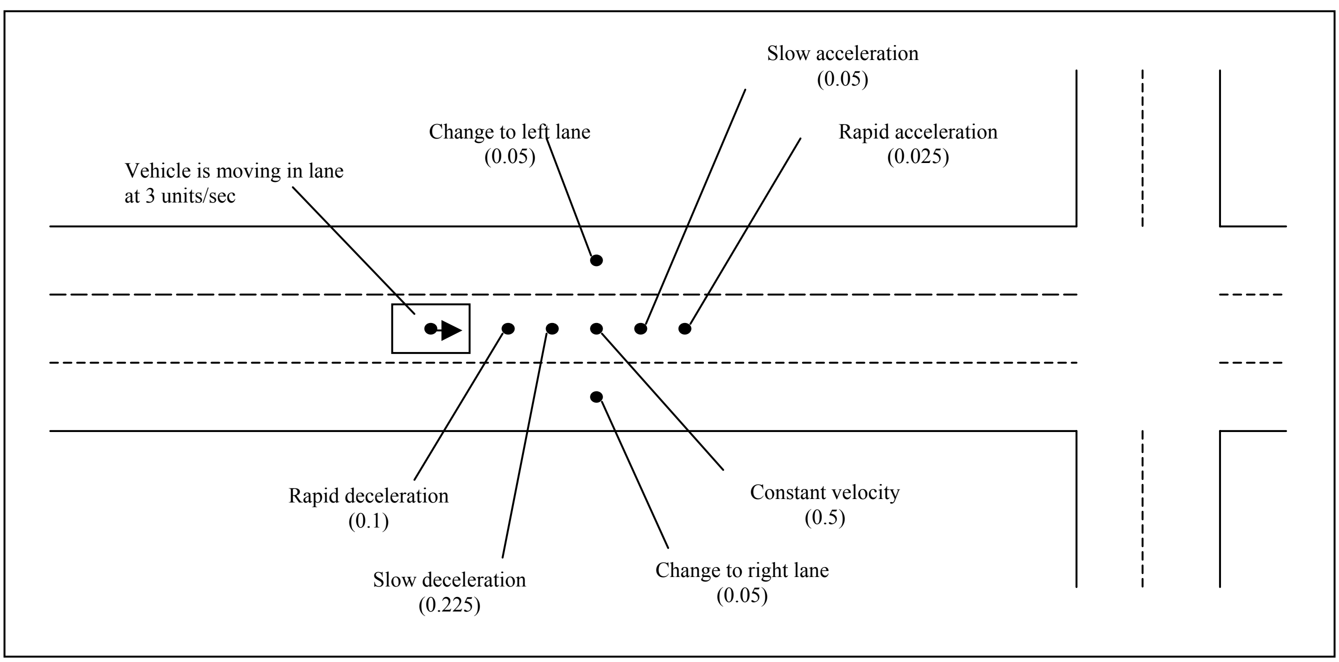

Figure 2: Driving Scenario

3 Logic-Based Motion Predictions in Constrained Environments

We are developing logic-based prediction algorithms for use in constrained environments. The purpose of these algorithms is to predict the probability that an object will occupy a given location in space at a given time by taking into account: 1) the constraints that are placed on the object's motion and 2) the influencing factors that would cause it to take a given action over another at specific times. These constraints and influencing factors are discussed in Section 4 of this paper.

In the case of on-road driving, vehicles must stay on the road and as such, the road network provides the constraints dictating the bounds in which a vehicle may travel. A database structure has been developed to capture detailed information about the road network, which includes information about the curvature of lanes, road interconnectivity, signage and traffic control, lane marking, etc. Equations representing the path of the roads can be inferred from the information in the database, and these equations serve as the basis for representing the possible paths the vehicle may take along the road network. Details about the database will be the topic of a future paper.

3.1. Discretizing Actions

The rule-based prediction approach requires that you discretize the possible actions that a moving object may take. In the case of a vehicle driving on-road, we limit the actions of the vehicle to be:

- Remain at a constant velocity in the current lane
- Slowly accelerate in the current lane
- Rapidly accelerate in the current lane
- Slowly decelerate in the current lane
- Rapidly decelerate in the current lane
- Change to a lane on the left
- Change to a lane on the right
- Turn to a lane on the left (at an intersection)
- Turn to a lane on the right (at an intersection)
- Make a U-Turn (at an intersection)

Figure 2 shows an example of a vehicle on a three-lane, one-way road. Each possible discretized action that the vehicle can take in this scenario is shown, along with the probability that the vehicle will take this action (represented by a value between zero and one in parenthesis). The point on the road that is referenced by each action shows the position the vehicle will be at if that action is performed. So, if we assume that the vehicle is at (0,0) to start and moving along its lane at 3 m/s, then the vehicle will be at (1,0) if a rapid deceleration action is performed, at (2,0) if a slow deceleration action is performed, etc.

Figure 3 shows how we can project these actions into the future to predict the position of the vehicle at longer time horizons. At time = t, the vehicle is at location (0,0). To get to time t+1, the vehicle may perform any of the discretized actions. The results of any of these actions will result in the vehicle occupying a location in the environment at one time step in the future (t+1). This location is shown as (x,y) coordinates next to each possible action. The probability that the vehicle will take any one of these actions over another is determined by the influencing factors described in Section 4.

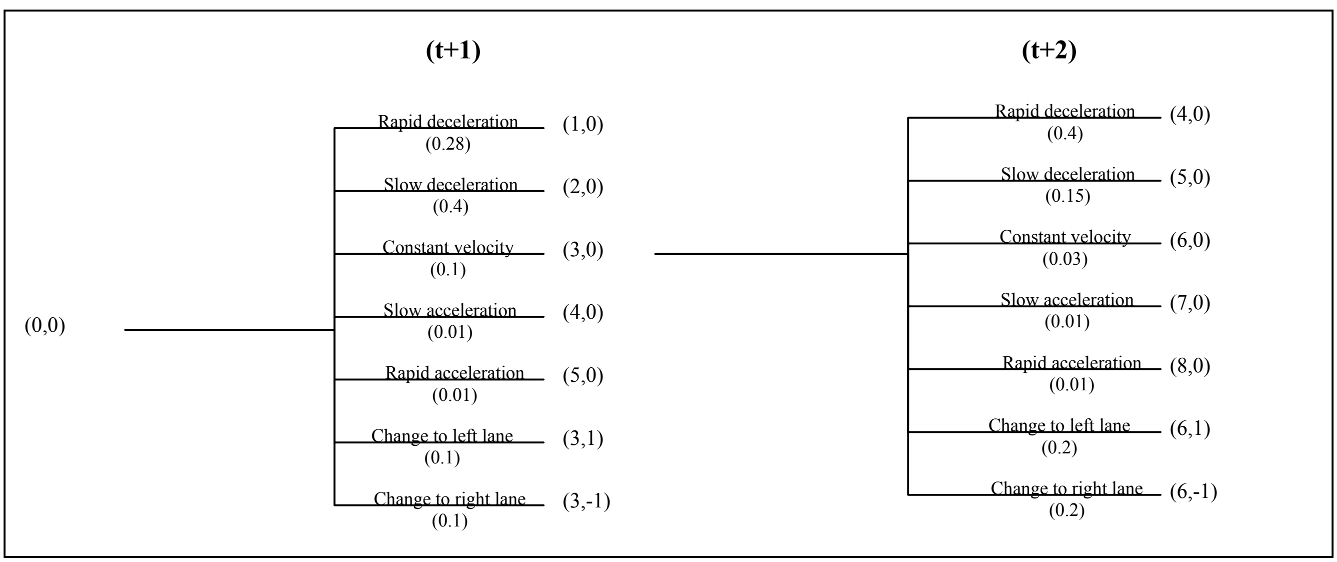

Figure 3: Probabilistic Prediction Over Multiple Time Steps

To get to the next time step (t+2), each of the possible actions that vehicle may take is again determined, and the probabilities are associated based upon the action it took at the previous time step. We continue this into the future as long as our planning horizon requires. Then to determine the probability of the vehicle occupying a given point in space at a given time, we determine if any of the paths result in the vehicle being in that location, and multiply together all of the probabilities along that tree branch to determine the overall probability.

In this case, we are assuming that the vehicle is driving on a straight, horizontal road and as such, the vehicle's location is simply moving in the x-direction. In reality, the location that the vehicle occupies will be derived from the information stored in an *a priori* road network database being developed at NIST.

3.2. Reducing Computation Time

One issue that arises with this approach is the possibility of a large amount of information that needs to be captured. If we have ten possible actions and we are projecting out 20 time steps into the future, we have 10^{20} values that need to be computed. This is an unrealistic expectation for any system that is expected to run in real-limiting the actions a vehicle may take at a given location, and then by limiting the time horizon of the prediction.

First, it is often the case that only a subset of the actions that the vehicle may take is possible at a given location on the road. For example, if the vehicle is not at an intersection, turning right, turning left, and making a U-turn is not possible. If the vehicle is driving in the right lane of a road, it is not possible for it to change to the right lane. If the vehicle is stopped, it may not perform either of the two deceleration actions. These limitations greatly limit the number of actions that need to be represented at any given time. In Figure 2, since the vehicle is not at an intersection, the turn right, turn left, and make a U-turn activities are not listed. It is expected that there will be between 5 and 7 possible actions, on average, at a typical position on the map.

Second, we will initially be implementing these algorithms in the 4D/RCS architecture [1]. 4D/RCS is a hierarchical architecture and limits the planning time horizon at each level of the architecture. Plans at each level typically have 5 to 10 steps between the anticipated starting state and a planned goal state at the planning horizon [1].

Third, we may wish to prune the tree from the onset to eliminate actions that have a very low probability of occurring. For example, if we assume that the probability that a vehicle will rapidly accelerate at time = t+1 is less than a certain percentage (say, 3 percent) then we may decide to ignore that action in the tree. By doing this, we would also ignore all branches of this tree that would follow from this action taking place, thus greatly reducing the size of the tree.

Considering these three factors, the number of values that need to be computed is greatly reduced (from 10^{20} to as little as 5^5) and as such, we believe that this approach should lend itself to real-time environments.

4 Constraints on Motion and Influencing Factors

This section discusses the factors that affect the probabilities associated with the possible actions that a vehicle may take while driving on-road. There are two classes of factors that we must consider. The first are factors that limit the possibilities of where the vehicle is *able* to reach. In other words, by considering these factors, we can eliminate certain portions on the maps that are not reachable by the vehicle. We call these 'constraints on motion'. The second are factors that influence which of the possible actions the vehicle *is likely to perform* out of those that are available to it. We call these influencing factors. These two categories of information are discussed below.

4.1 Constraints on Motion

As mentioned above, the constraints on motion limit the possibilities of the locations that the vehicle is able to reach. Below we discuss two constraints on motion:

- **A Priori Road Network Information:** Assuming that the vehicle is driving on-road and will remain on-road, the road network limits the possible locations that the vehicle can possibly attain.

- **Vehicle's Motion Capabilities:** Motion capabilities of a vehicle limit the possibilities of where it can possibly be in the future. For example, the vehicle's acceleration capabilities restricts the range of locations that are accessible by the vehicle in a given timeframe. Similarly, knowing a vehicle's minimum turning diameter as a function of its current velocity provides a limitation on how quickly it can change lanes and its ability to perform turns at an intersection. One of the ways this information may be used is to limit the possibility of a vehicle turning at an intersection as a function of its velocity approaching the intersection. That is, if a vehicle is approaching an intersection at a relatively high rate of speed, one may eliminate the possibility that the vehicle is turning at the intersection.

4.2 Influencing Factors

Influencing factors affect the probability that a vehicle will perform one action over another. Seven influencing factors are discussed below:

- **Weather and Environmental Conditions:** Weather and environmental conditions include rain, sleet, snow, fog, darkness, etc. and their effects on visibility and slickness of the road surfaces. As the weather and environmental conditions worsen, the probability often increases that the vehicle's velocity will decrease. Also in these conditions, vehicles often prefer to remain in their lane as opposed to switching lanes or performing passing maneuvers.

- **Vehicle's Intention and Indicators:** One of the strongest factors that play a role in human's ability to predict the future location of another vehicle is the vehicle's perceived intentions. Intention could be known *a priori*, such as knowing a vehicle is driving to the bank, and this knowledge could be used to determine the most probable path it will take to achieve that goal. More commonly, intentions could be derived from perception, such an indication that a vehicle is making a left turn based upon the vehicle moving into a turn lane or having its blinker on. As more information becomes available from the vehicle, this information can be used to either strengthen or weaken the perceived intentions, which in turn would increase/decrease the probabilities associated with the possible actions the vehicle may take in the future.

- **Class of the Vehicle:** Object classification provides information about the class of object that is being perceived. If we limit our scope to vehicles on the road, the class of vehicle could indicate the course the vehicle is expected to travel, or how it is expected to behave in certain situations. For example, if the vehicle was identified as being a city bus, we would most likely expect it to stop at a bus stop signs, and traverse primarily in the right most lane. Similarly, if the vehicle was identified as an emergency vehicle, we would expect it to travel at high rates of speed and not to necessarily stop at stop signs and traffic lights. If the vehicle was a motorcycle, we may not eliminate the possibility of it navigating in between vehicles stopped in a traffic jam.

- **Vehicle's Personality:** When humans drive on-road, they implicitly assign a personality to other vehicles. For example, if one sees another vehicle swerving in and out of traffic, and making unsafe lane maneuvers, one may assign a very aggressive personality to the vehicle. Conversely, if a vehicle is observed driving at or below the speed limit, keeping an extraordinarily far following distance, and rarely changing lanes, a low level of aggressiveness would be assigned. Based on these personality measures, one would expect different actions from that vehicle in the future. For example, an aggressive vehicle would be more prone to make lane changes, and as such, the probability assigned to the action of changing lanes would be greater for that type of vehicle.

- **Traffic Control Indicators / Rules of the Road:** The 'rules of the road' play a large role in predicting how a vehicle is expected to behave under certain situations. For example, if a vehicle is approaching a stop sign, one would expect that the vehicle would gradually decrease its speed until it reaches the stop sign, comes to a complete stop, and then proceeds when the intersection is safe to traverse. However, based on the perceived personality of the vehicle, we may expect that the vehicle only slows down but does not come to a complete stop, or traverses the intersection before most would consider it safe.

  Efforts at NIST have focused on encoding the rules of the road using finite state machines leveraging a driver's manual published by the Department of Transportation [7]. The document contains a comprehensive inventory of the behaviors involved in operating an automobile, along with the rated criticalities of these behaviors. The task descriptions are organized in terms of the situations giving rise to the behaviors; behaviors involved in controlling movement of the car without regard to specific situations; behaviors that must be performed continually or periodically while driving, rather then in response to a specific situation; and off-road behaviors that are performed before driving, to maintain the car in sound operating condition, and in compliance with legal regulations. The document organizes the task descriptions into the following categories:

  - basic control (situation-independent driving behaviors to control the movement of the vehicle),
  - general driving (continuously-performed driving behaviors in response to any specific situation),
  - situational behaviors (behaviors that are required in response to specific situations),

- pre-driving behaviors (behaviors taken prior to driving to assure safe and efficient operation),
- maintenance (behaviors directed toward the vehicle to assure safe and efficient operation), and
- legal responsibilities (legally imposed behaviors required to assure that drivers are responsible for the consequences of their actions).

- **Other Traffic:** In the same way that our vehicle is predicting the future locations of other vehicles in its vicinity, other vehicles are doing the same with vehicles in their vicinity. Hence, our vehicle needs to not only be cognizant of vehicles that run a risk of interfering with our path, but also of vehicles that could interfere with those vehicles' paths. This is analogous to a driver looking two cars ahead to try to predict what the car in front is going to do.

  Each vehicle on the road has a range of influence associated with it. The size of this range is a function of the vehicle's velocity and the presence of intersections, among other factors. Any vehicle within a defined range could be impacted by actions in which that vehicle performs. This, in turn, could cause a ripple effect. As shown in figure 4, Vehicle B is in Vehicle A's range of influence (denoted by the right-most oval). Similarly, Vehicle C is in Vehicle B's range of influence (denoted by the left-most oval). Even though Vehicle C is not in Vehicle A's range, Vehicle C still needs to be aware of Vehicle A's motions since these motions will affect vehicle B which in turn will affect Vehicle C.

  Information about the position and motion of other traffic will affect the probability of other vehicles taking certain actions. Constraints, such as maintaining safe following distance, play a strong factor in how a vehicle reacts to certain situations.

- **Formations:** Formations aren't as important for on-road driving as they are for off-road driving, but they are still worth mentioning here. If it can be determined that a vehicle is driving as a part of a larger formation, the rules that govern the formation play a large role in dictating where that vehicle will be in the future. In the case of a battlefield environment, the military has devised a number of formations that vehicles in a group follow (e.g., bounding overwatch, V-formation, etc.). Similarly, in Robocup competition, teams often implement different strategies that rely on different formations. Knowing the other team's strategy can help to predict the players' moves. Even in on-road environments, vehicles sometimes move in formations, such as funeral processions. Identifying a presidential procession can provide additional information about the future moves of the vehicles in the procession. For example, the vehicles in the

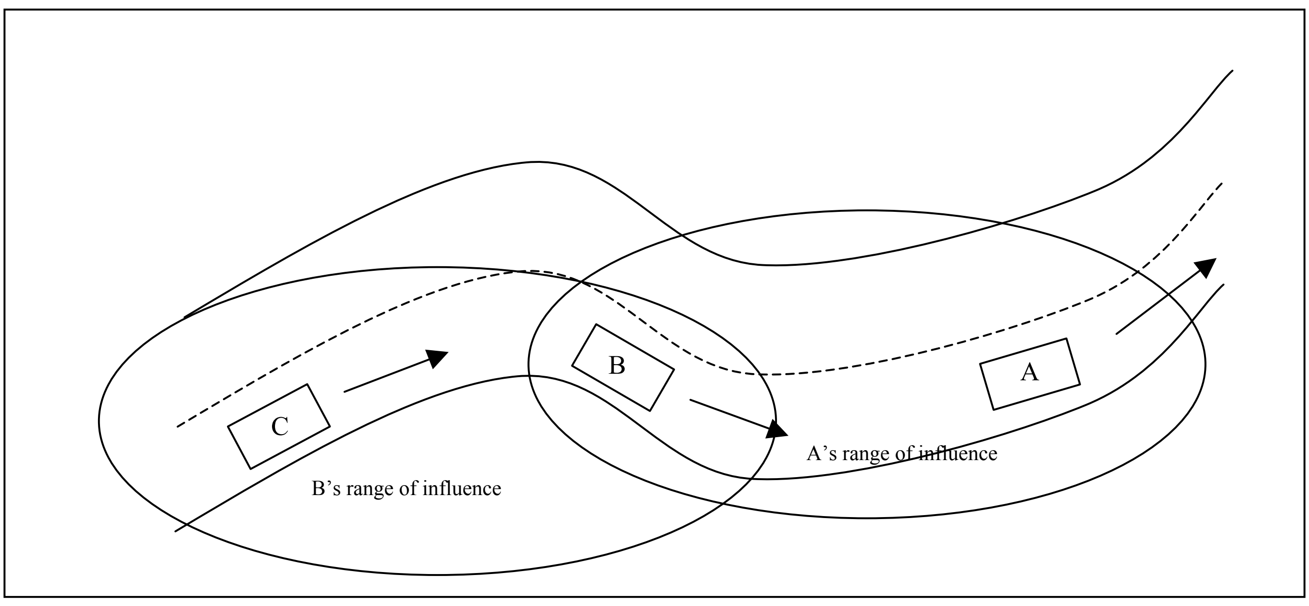

Figure 4: Vehicle's Range of Influence

procession will most likely change lanes when the vehicle in front of them changes lanes. Also, the vehicles will likely run red lights to keep up with the vehicle in front of them.

The factors mentioned above provide much of the input necessary to determine and refine the probabilities that predict the future location of moving objects in the environment. This information is then fed to the planners in the form of space/time probability distribution (in the planner's formalism of choice) to develop appropriate plans in the presence of moving objects.

5 Planning in a Dynamic Environment

As described in [2], the NIST planner utilizes incrementally created planning graphs to formulate potential vehicle trajectories. As part of the graph expansion/evaluation phase, a cost/benefit number must be assigned to each potential path segment. The dynamic obstacle layer of the planner's world model system determines a portion of this cost/benefit number.

If the trajectory of the moving object is known explicitly, the moving object prediction subsystem would produce a curve through space and time that represents the path of the moving object. The dynamic obstacle layer would then match this curve against the plan segment being evaluated to determine if an intersection exists. This collision information is passed onto a value judgment module that examines the predicted nature of the obstacle (e.g. is it a soda can or a tank) and the intent of the commander (e.g. allowed to run over soda cans, but not tanks) in order to formulate the dynamic obstacle portion of the overall cost/benefit number for the plan segment.

In the real world, moving obstacles seldom broadcast their exact trajectory ahead of time and predictive algorithms are necessary to compute a potential trajectory. This potential trajectory is made available to the dynamic obstacle layer in the form of equations that represent a volume in space/time for the expected location of the object. The volume is often a very tight circle at the current time (where the location of the vehicle is known with small uncertainty), and the size of the volume per unit time will gradually increase as one moves forward in time. This increase in volume represents the uncertainty in the location prediction. For a ground-based object, this bounding area may be viewed as a three-dimensional volume with axes of northing, easting, and time. The dynamic obstacle layer must now examine if a potential path segment lies inside the volume that represents any value over a pre-defined probability threshold. This information is sent to the value judgment module for use in the formulation of the final cost/benefit number. Through the use of this system, minimal collision or collision free paths may be planned.

6 Preliminary Results

We have implemented this approach in a simulated environment. In this environment, we place a vehicle on a roadway with various constraints on the environment, such as imposing speed limits and placing obstacles in the roadway. We then impose a cost on the vehicle based upon 1) the actions the vehicle takes (e.g., changing lanes, quickly accelerating, etc.), 2) not adhering to the driving rules of the road (e.g., not obeying the speed limit), and 3) coming within an unsafe stopping distance from other stationary or moving objects. We are initially basing the costs we associate with these actions on the criticality indexes that are documented in the DOT manual (ref.).

Based on initial experiments, the following was learned:

- We are able to predict out to 10 time steps in less than one second, without any substantial pruning of the graph. Once we implement some pruning techniques (most likely based on overall probability), we expect the time to be cut in half.
- The possible future location of a vehicle is based upon a reachability graph which is derived from the ten discrete actions which are described in section ??. The reachability graph accounts for the constraints that are imposed by the roadway.
- In our experiments, there are usually only two or three possible actions that a vehicle can legally take at any given point in the roadway at a given velocity profile. Assuming three velocity profiles, we are evaluating the cost and probability of 6-9 possible actions.
- Because the cost of the performing an action drive the probabilities of which action a vehicle may take, determining the costs becomes very important. Initially, we are basing the costs on the criticality indexes from the DOT manual mentioned above. We will them be refining these costs based on learning techniques that are still in the early stages of investigation.
- Although only applied to other vehicles in this paper, we firmly believe that a similar approach would be beneficial to determining the future location of other types of moving objects in the environment, such as pedestrians. Future efforts will explore this.
- The validity of the predictions will be determined based upon a set of experiments in which we model an existing road driving situation in the simulation package and then compare the results we get from our prediction algorithm with what actually occurs in the real environment.

7 Conclusion

In this paper, we have presented an overview of a framework for representing and planning the future location of moving objects. In our research, we quickly found that there was a clear void in the literature in areas focusing on long-range prediction of moving objects in a constrained environment. As such, we have developed an approach to predict the future location of objects in a constrained environment.

The approach explores applying probabilistic, logic-based algorithms to predict the future location of vehicles in an on-road environment. To apply this approach, the possible motions of the object are discretized and each action is assigned a probability based upon a series of 'constraints on motion' and influencing factors that are described at a high-level in this paper. Although these factors may not be exhaustive, we believe that they provide a good representative sample of the types of factors that would need to be applied.

The concepts in this approach are very new and a number of issues still need to be explored. We need to work out the details on how the influencing factors will contribute to the probabilities associated with the discretized actions. This paper discussed the contributions at a coarse level.

We also need to ensure that this approach can be performed in a real-time environment. Although we have proposed pruning mechanisms to curtail the unbounded growth of the probability trees, we still need to ensure that the pruned tree can still be developed and processed in the time constraints imposed on the planners in different types of on-road environments.

Before we use this approach, we need to decide which moving objects to apply it to. In the case of on-road driving, there are many moving objects in the environment, but not all of them need to have a detailed level of prediction associated with them. Only the vehicles that have the highest probability to affecting our path would be of concern. For other vehicles, we may employ a less accurate and a less computationally expensive approach.

Although this approach was only applied to on-road driving in this paper, it would be equally applicable to any other type of moving object provided that the actions of the object can be discretized. Future work will apply this technique to pedestrians and military vehicles.

References

1. Albus, J. and et.al., "4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems," NISTIR 6910, National Institute of Standards and Technology, Gaithersburg, MD, 2002.

2. Balakirsky, S. and Herzog, O., "Planning with Incrementally Created Graphs," NIST, 6895, Gaithersburg, MD, 2002.

3. Bar-Shalom, Y. and Fortmann, T. E., *Tracking and Data Association*, Academic Press 1988.

4. Firby, J., "Architecture, Representation, and Integration: An Example from Robot Navigation," *Proceedings of the 1994 AAAI Fall Symposium Series Workshop on the Control of the Physical World by Intelligent Agents*, New Orleans, LA, 1994.

5. Gueting, R. H., "A Foundation for Representing and Querying Moving Objects," *ACM Transactions on Database Systems (TODS)*, Vol. 25, No. 1, 2000, pp. 1-42.

6. Madhavan, R. and Schlenoff, C., "Moving Object Prediction and Tracking for Off-road Autonomous Navigation," *Proceedings of the SPIE Aerosense 2003 Conference*, Orlando, FL, 2003.

7. McKnight, J. and Adams, B., *Driver Education Task Analysis. Volume 1. Task Descriptions*, Human Resource Research Organization, Department of Transportation, National Highway Safety Bureau 1970.

8. Schlenoff, C., "Linking Sensed Images to an Ontology of Obstacles to Aid in Autonomous Driving," *Proceedings of the 18th National Conference on Artificial Intelligence: Workshop on Ontologies for the Semantic Web*, 2002.

9. Schlenoff, C., Madhavan, R., and Balakirsky, S., "Representing Dynamic Environments for Autonomouos Ground Vehicle Navigation," *Submitted to the IEEE/RSJ IROS 2003 Conference*, Las Vegas, NV, 2003.

10. Singhal, A., *Issues in Autonomous Mobile Robot Navigation*, Computer Science Dept, U. of Rochester 1997.

11. Stone, L., Barlow, C. A., and Corwin, T. L., *Bayesian Multiple Target Tracking*, Artech House 1999.