2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

A Probabilistic Method for Detecting Impending Vehicle Interactions

Stewart Worrall and Eduardo Nebot
Australian Center for Field Robotics
University of Sydney, Australia
Email: {s.worrall, nebot} @acfr.usyd.edu.au

Abstract— **In mining operations it is advantageous to be able to predict the future movements of nearby vehicles. For autonomous mining, this can be used for localised, short term path planning and risk assessment. For semi-autonomous or non-autonomous mining, this can be used for collision avoidance, situational awareness and risk assessment of maneuvers between a human operated vehicle, and another vehicle (operated by a human or otherwise). This paper introduces a probabilistic approach to predicting vehicle movements, in particular, the time until two vehicle paths intersect. Results are shown using real data collected from the operation of two separate fleets of vehicles.**

I. INTRODUCTION

Predicting the future movements of vehicles is important in many safety and planning applications, both autonomous and non-autonomous. This paper presents a new approach to predicting the movements of surrounding vehicles, and estimating a timeframe for the interactions between vehicles. A complete system is introduced that includes detection of the nearby vehicles, and the probabilistic algorithms to obtain an estimate of the time when the paths of the vehicles will intersect.

The first section of this paper examines current approaches to finding a measure of Time to Intersection (TTI) between vehicles. In this context, the intersection between vehicles refers to the point where the two vehicles meet, or pass each other on a road. To measure the TTI between vehicles, it is first necessary for the vehicles to be detected. A discussion into the different approaches to this problem is presented. Once the surrounding vehicles are detected, it is necessary to determine the path of the vehicle, and to calculate the TTI. Vector analysis is the most common form of predicting TTI, usually used for predicting collisions between vehicles. In this work, some examples of the shortcomings of this method are provided.

Section III describes the requirements of the system. One requirement is to provide a measure of TTI, including the confidence bounds of the predictions. The confidence is necessary for determining the estimated worst case scenario, which is the earliest possible time that the vehicle paths will intersect. Communication between the vehicles is another important system requirement. The communication is necessary to transfer motion information between vehicles such as position, velocity and heading.

The algorithm to predict the vehicle motion and TTI is defined in Section IV. The first part of this is the modelling of the system. To include map information into the algorithm, the velocity is defined as a function of distance. A velocity profile is generated for each section of road, and this is used to propagate the vehicle model.

Once the system is modelled, the time to intersection can be predicted. When it is calculated that two vehicles have intersecting paths, the position PDF for each vehicle is calculated. A time based convolution is performed on these two position PDFs, and the result is a time based probabilistic measure of the TTI. The outcome of this is a probabilistic measure of how confident the system is of the vehicles motion, including the estimated first possible chance of the vehicles intersecting. This can be used for planning vehicle maneuvers, or for measuring safety through calculating a risk of collision.

Results of the new system in operation are provided. Data was collected from two separate fleets of vehicles, each operating over several days. A performance analysis of the new system over many hundreds of vehicle interactions is provided. Statistics are provided for the performance of the two main measures of TTI, these being the first possible time, and the most likely time of an intersection between vehicle paths.

II. PREDICTING MOVEMENTS OF OTHER VEHICLES

To predict the movement of surrounding vehicles, it is first necessary that they are detected. Once the other vehicles are detected, it is then necessary to predict whether the paths of these vehicles will intersect with the vehicle in consideration. If it is determined that the paths will intersect, the time until intersection is calculated, which is useful for predicting potential collisions, or planning maneuvers in autonomous vehicles.

Detection of Other Vehicles

It is possible to detect surrounding objects, such as other vehicles, by fitting sensors such as cameras or radar to a vehicle ([1], [2], [3]). The sensor data can be processed to extract inferences about the surrounding vehicles motion properties. This usually includes range or position, and can sometimes provide heading and/or velocity to some extent. This can be considered passive detection, because it is not necessary to fit any equipment to the objects being detected. The advantage of this method is that all objects visible to the sensor can potentially be detected, without any special equipment on the surrounding objects. There are however limitations with this method, such as finding a unique identifier for individual objects, accurately predicting dynamic properties of the objects, handling obstructed/partially obstructed objects and dealing with false positives. Also, sensors such as cameras do not perform well in environments such as dusty roads, night operation and rain/foggy weather.

An alternative method is to use active detection ([4], [5], etc). This method involves fitting all vehicles with a form of communication, and broadcasting the identity, position, velocity and other information to surrounding vehicles. This method provides a clear indication of the properties of the surrounding agents with very low probability of false positives. An advantage of this system is that wireless communiation can provide significant range of detection ($>$ 500m). To some extent, these systems are not bounded by direct visibility since wireless communication has some ability to work out of a direct line of sight. The main problem with this method is that a system must be installed and working in each vehicle that is to be detected. This is a problem for operation in public roads, but is feasible in closed environments such as mines, quarries, ports, private enterprises, etc. This method was selected for the mining application outlined in this paper.

Predicting Time To Interaction

Depending on the sensor information available and the range of prediction required, there are several possible approaches to the problem of predicting the Time to Intersection between vehicles. The first approach reviewed here is by projecting the current vehicle state vector ([4], [5], etc). This involves considering the heading, position and velocity of each vehicle, and projecting the vectors to find the point of intersection, and the time until intersection. This is not sufficient in many cases, since the road is not taken into consideration. Figure 1 illustrates several cases where the vector analysis fails. In some cases (for example, top left in the Figure 1) the velocity vectors intersect even though the vehicles are on separate roads, and will not meet. Another problem (illustrated in the remaining images from Figure 1) is that the velocity vectors indicate a time to intersection that does not accurately represent the real situation. Without integrating map information into the equations, there is no way to consider many possible road scenarios.

Without the inclusion of map information in the prediction of future state vectors, the properties of roads that affect the velocity of the vehicle are ignored. This means that the models do not consider that vehicles drive at different speeds depending on the properties of the road, such as gradient, curvature and visibility. As an example, vehicles travelling into a sharp curve will slow down, so analysing the vehicle TTI will give an incorrect result using a constant velocity model, or other simplistic models.

The second potential approach considered in this paper regards learning vehicle trajectories, examples of which can be found in [6] and [7]. This body of research involves using statistical motion patterns that are learnt using many different techniques, a comprehensive list of these can be found in [7]. These techniques have been used to learn vehicle trajectories using fixed cameras to track the vehicles. These learned trajectories have then been used to forward predict the behaviour of the vehicle. In the context of the problem outlined in this paper, the focus of this research to date has been more towards monitoring traffic at intersections, detecting anomalies, and tracking vehicles around a site.

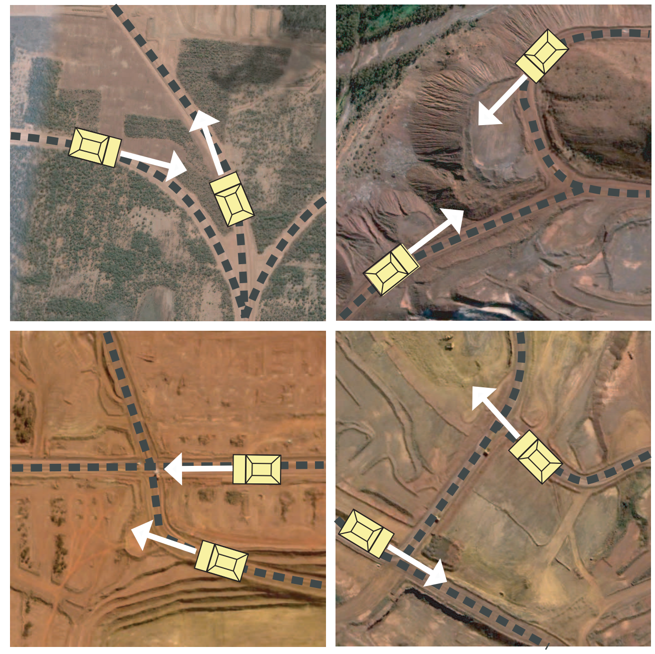

Fig. 1. Several scenarios are presented here to highlight that a simple vector analysis to predict the interactions between vehicles is not sufficient. It is necessary to have map information to correctly predict the outcome. The dotted lines mark where the roads are in the image, and the arrows represent the velocity vectors of each vehicle.

This paper introduces an alternate vehicle model where the trajectory of the vehicle is fixed to a known map. The map is in the form of a graph, resulting in single dimensional roads and intersections. A velocity model of the roads is learnt using position/velocity data collected from a fleet of vehicles, and this road model is integrated into the vehicle model. A prediction algorithm is used to provide an estimated Time to Intersection for vehicles over any distance as a PDF.

III. System Requirements

Measuring TTI

To determine whether a vehicle maneuver is safe, it is necessary to have a probabilistic measure of TTI. This measure can be used to plan maneuvers, or calculate the risk of collision for safety applications. An estimated TTI is not useful without the inclusion of confidence bounds. The most important confidence bound is the lower bound, which is the estimated first possible time to interaction between vehicles. This can be considered as a "worst case scenario", and is the most important measure for determining the safety of maneuvers.

Communicating between Vehicles

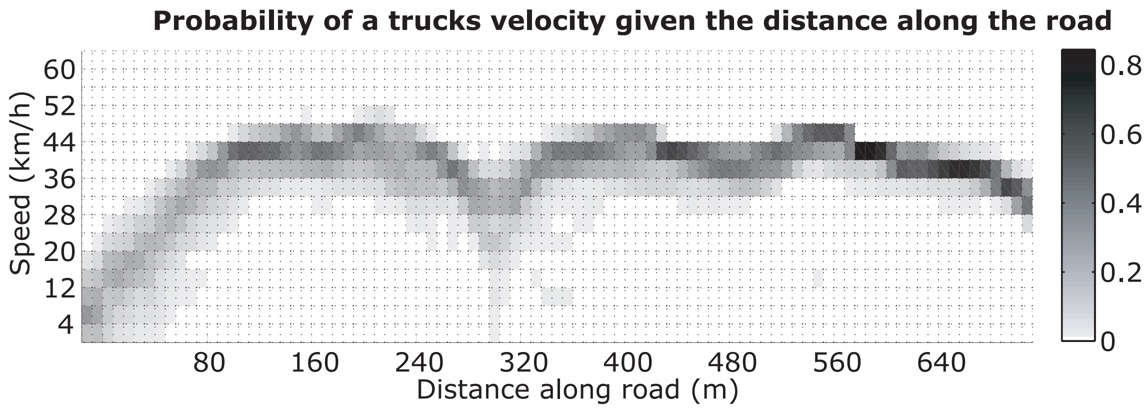

Fig. 2. An example of the velocity PDF generated, with the probability colour scale shown on the right. On the left of the graph, the vehicle starts from a very slow speed, then acceleration occurs. At around 300m, there is a section where the vehicles slow down, most likely a sharp corner.

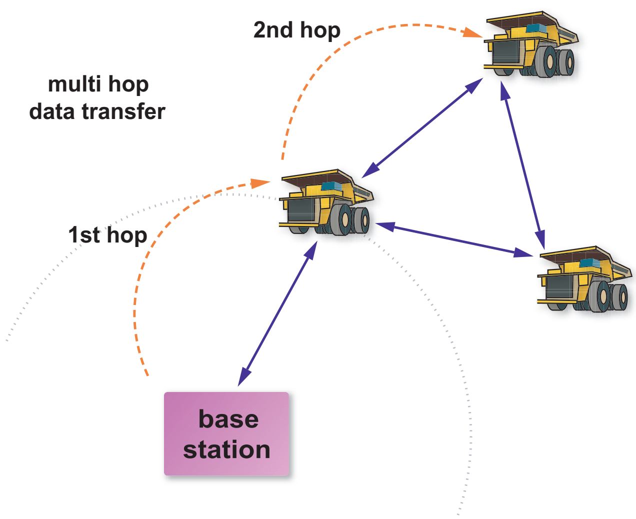

Fig. 3. An overview of the mesh network implemented in the vehicles, and in a base station.

A mesh network is used for communication between vehicles, as illustrated in Figure 3. This network allows the vehicles to transmit their position, velocity and other data. The mesh network was designed for long range, robust connections, which is essential for planning and estimation applications. A base station computer is also fitted to be a node on the mesh network. This allows the data to be collected periodically from the vehicles, and also allows for the transferring of programs, configuration files and updated digital maps to the vehicles.

A secondary, redundant wireless network working at a different frequency has been included in the system to increase the reliability of the transmission of information. The secondary system is a lower bandwidth radio implemented for location information only, the high bandwidth data (such as maps, etc) is transmitted only through the mesh network.

IV. Prediction Algorithm

Modelling the System

A process of incorporating map information into the vehicle model was introduced in [8]. The following subsection provides an overview to this process, put into the context of the application defined in this paper.

The digitised map is in the form of a directed graph (digraph). Each edge of the graph represents a road that connects the graph nodes, which are the intersections. Each edge of the graph is unidirectional, meaning that a normal two lane road is represented by two edges, one for each direction. An illustration of this can be seen in in [8].

The graph representation means that the state space can be considered as a single dimension, the distance along the road. The state space is assigned the value S, and a discrete representation can be created by dividing the state space into the set of divisions (s_i) as described in Equation 1. The motivation behind using a discretised space is to be able to represent non-linear, non-parametric velocity functions [8].

The high level vehicle plan (future sequence of graph edges) is considered to be unknown. This means that for a vehicle approaching an graph node (intersection), the prediction algorithm considers that the vehicle could travel in any valid direction, following any possible graph edge leading away from the node. For a safety application, it is necessary to assume that the human driver could take any direction at an intersection, even if the driver was instructed to go in a particular direction. This could lead to an indication of a potential interaction between vehicles even if the vehicles do not eventually meet. In the case of an autonomous system, it is possible to transmit high level plans between the vehicles over the network to eliminate this possibility.

$$S = s_1 \cup s_2 \cup \cdots \cup s_m \tag{1}$$

A velocity profile for each road division (s_i) is generated using historical GPS data [8]. This defines the historical probability (PDF) of vehicle velocity for a given section of road, meaning that velocity is represented as a function of distance. This is illustrated in Figure 2, where the combined velocity PDF for an entire section of road is shown. The road speed is usually limited by features such as road gradient, road curvature, etc, and these properties are built into the velocity model for the road. Considering velocity as a function of distance leads to Equation 2.

$$\begin{aligned} \text{velocity } V &= f(S) \\ P(V|S) &= P(V|s_i) \text{ for } s_i \in S \\ &= \text{PDF shown in Figure 2} \end{aligned} \tag{2}$$

The motion of the vehicle can then be described as shown in Equation 3;

$$\begin{aligned} \text{distance } S &= \int V\ dt \\ S_{t+1} &= S_t + V\ \Delta t \quad \text{(discrete model)} \end{aligned} \tag{3}$$

The propagation model for the prediction algorithm comes from the combination of Equation 2 and 3, shown below;

$$\begin{aligned} S_{t+1} &= S_t + V\ \Delta t \quad \text{(discrete model)} \\ P(S_{t+1}) &= P(S_t) * P(V|s_i) \text{ for } s_i \in S \end{aligned} \tag{4}$$

Equation 4 shows that the posterior is given by the convolution ($*$) of the prior distance PDF with the velocity PDF. This model is then implemented as a histogram filter, as described in Worrall [8].

Predicting the TTI

A model has been defined in the previous section for predicting the location of a vehicle by considering velocity as a function of distance. Consider now two vehicles both using this model, where the position of each vehicle is given by $P(S^1)$ and $P(S^2)$ respectively.

The probability of the vehicles being at the same location for a given time is defined in Equation 5. Here, $P(I)$ is the probability of the vehicles being in the same location over time, and $P(i_{t=1})$ represents the probability of co-location for a discrete time sample (k). Considering $P(I)$ as a discretised function, as the time divisions approach zero, the function approaches the true value. For accuracy, the optimal time division between samples is dependant on the speed of the vehicles. For faster vehicles it is necessary to use closer divisions to accurately represent $P(I)$. The reason for this will become apparent later in this section.

$$P(I) = P(i_{t=0}) \cup P(i_{t=1}) \cup \dots \tag{5}$$

For a specific time, the probability that the two vehicles will be in the same location is given by the integration of the two position PDFs multiplied together. This is shown in Equation 6

$$\begin{aligned} P(i_{t=k}) &= \int_S P(S^1|t_k) \bullet P(S^2|t_k) \\ &\approx \sum_S P(S^1|t_k) \bullet P(S^2|t_k) \end{aligned} \tag{6}$$

This process is a convolution where $P(I)$ is the convolution of the two vehicle position PDFs over time. Each $P(i_{t=k})$ is a discrete sample from this function, and so the more samples taken, the closer the result will be to the true convolution. The reason why more samples are required for faster vehicles is that the peak value will be higher over the time domain, and the spread less. With fewer samples, it is possible to miss the peak value. In any case, the convolution will not sum to one because the vehicle PDFs are modified for each time step in the convolution calculation.

The resulting convolution $P(I)$ is the probability that two vehicles are in the same location over time. The first significant value in this function represents the earliest time that the vehicles will meet. The highest value represents the most likely time that the vehicles will meet. These are the two important measures, as defined in Section III.

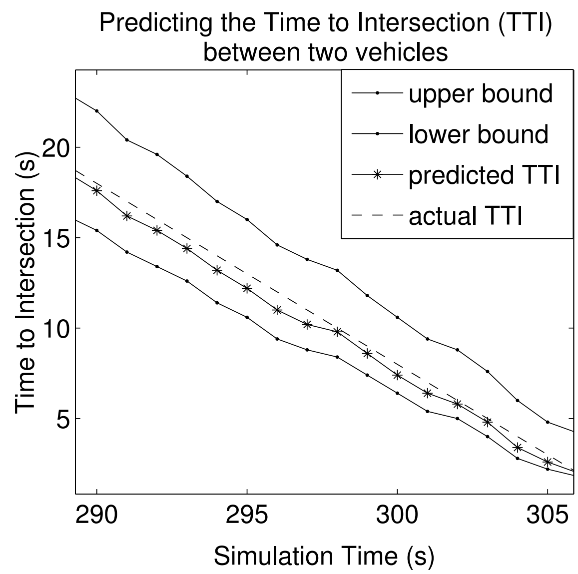

Fig. 4. This graph shows the result of two trucks moving closer together. For each time step, the calculation of TTI is plotted with the filter bounds. The actual measured time until the vehicles intersect is also shown.

When two vehicles establish communication and transmit their positions and other data, the algorithms can be initialised, and the times can be calculated. Each corresponding communication can be used to recalculate the times, and update the predicted TTI. This will reduce the variance on the probability $P(I)$. This can be seen in Figure 4, which demonstrates the output of the algorithms for two trucks moving closer together. The estimated TTI is updated each second, and as the vehicles get closer, the variance is reduced.

V. Results

The following results were taken from two datasets collected from vehicle fleets in two different mines. The data from several days operation of both vehicle fleets was used to give the following statistics. The data encompasses many hundreds of vehicle interactions, at varying speeds and at different locations in the mines. It is important to point out that even with such a variety of locations, headings, distances and speeds, the results are consistently close to the actual measured TTI from the data.

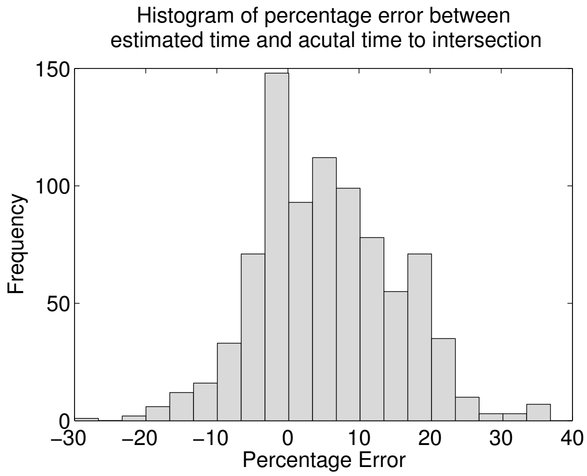

Fig. 5. A histogram showing the error between the estimated TTI and the actual measured TTI. This includes samples from many different situations, including different areas of the mine, different speeds, etc.

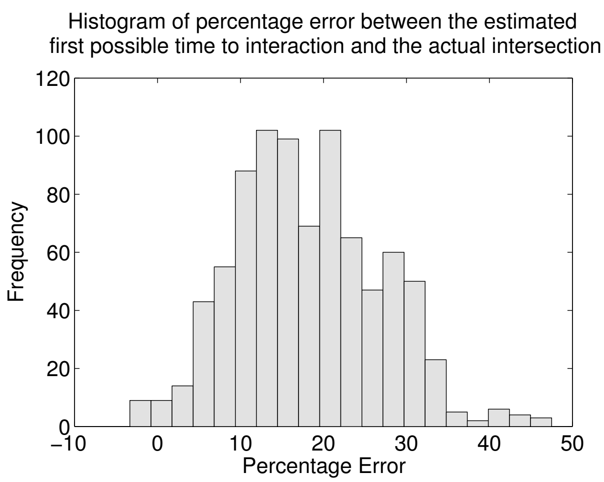

Fig. 6. A histogram showing the error between the estimated first possible TTI and the actual measured TTI. The graph shows that all interactions measured in the data occured on, or after the time of the estimated first possible TTI. Again, this graph includes many samples from vehicles at different speeds, distances and also in different areas of the mine.

The first result shows a comparison between the measured TTI and the actual TTI, illustrated in 5. This shows the spread of results from many vehicle interactions at different speeds, and in different parts of the mine. The measure is a percentage error, i.e. an error of +10 means that the actual time to vehicle intersection was 10% greater than the estimated time. The data was considered only for times where the vehicles passed after less than 20 seconds.

Figure 6 shows the percentage difference between the estimated "first possible time" to intersection (as shown in Figure 4) and the actual TTI. Again, an error of +10 means that the actual time to vehicle intersection was 10% greater than the estimated first possible time. The results show that even with the variety of situations presented in the datasets, the actual TTI is consistently later than the first possible TTI. This is important if this measure is to be used for planning vehicle maneuvers, or detecting unsafe driver behaviour.

VI. CONCLUSION

This paper presents a new approach to predicting future movements of vehicles, and a method of predicting the time until the vehicles intersect. Existing methods are inaccurate in many situations because they do not take into account the properties of a road, and the velocity that vehicles will travel at different points along a road. The approach presented in this paper involved building a velocity profile for each road using collected GPS data. A model for vehicle motion was given with velocity as a function of distance along a road.

Results of the new system were provided, showing the performance of the algorithms implemented. Data collected from two separate fleets of vehicles, in different locations was used in the algorithms, and a statistical analysis of the performance was provided. The new approach was shown to provide a useful measure of the estimated time until intersection, and more importantly, a measure of the estimated first possible time until intersection. These measures can be used for path planning, risk assessment, collision avoidance and many other applications, both autonomous and non-autonomous.

REFERENCES

[1] A. Eidehall and F. Gustafsson, "Combined road prediction and target tracking in collision avoidance," in *Intelligent Vehicles Symposium, 2004*, June 2004, pp. 619–624.

[2] D. Sinclair, B. Boufama, and R. Mohr, "Independent motion segmentation and collision prediction for road vehicles," in *Proceedings from Computer Vision and Pattern Recognition, 1994*, June 1994, pp. 958–961.

[3] N. Srinivasa, Y. Chen, and C. Daniell, "A fusion system for real-time forward collision warning in automobiles," in *Proceedings from Intelligent Transportation Systems, 2003*, vol. 1, October 2003, pp. 457–462.

[4] Y.-J. Chen, C.-C. Chen, S.-N. Wang, H.-E. Lin, and R. C. Hsu, "Gpsensecar -a collision avoidance support system using real-time gps data in a mobile vehicular network," *ICSNC*, pp. 71–76, 2006.

[5] Q. Huang, R. Miller, P. MacNeille, G.-C. Roman, and D. DiMeo, "Development of a peer-to-peer collision warning system," *Ford Technical Journal*, vol. 5, no. 2, March 2002.

[6] C. Stauffer, W. Eric, and L. Grimson, "Learning patterns of activity using real-time tracking," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, August 2000, pp. 747–757.

[7] W. Hu, X. Xiao, Z. Fu, D. Xie, T. Tan, and S. Maybank, "A system for learning statistical motion patterns," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 9, September 2006, pp. 1450–1464.

[8] S. Worrall and E. Nebot, "Using non-parametric filters and sparse observations to localise a fleet of mining vehicles," in *IEEE International Conference on Robotics and Automation, 2007*, April 2007, pp. 509–516.