MULTI AGENT LEARNING OF THE SURFACE QUALITY FOR FLEET OPERATION IN DAMAGE-PRONE ROADS

Dr.Jose E. Guivant
School of Mechanical and Manufacturing Engineering,
The University of New South Wales, Australia
j.guivant@unsw.edu.au

Abstract

This paper presents a cost effective approach for the cooperative learning of the quality of dirt roads in intensive hauling environments. The paper also introduces an efficient method to synthesize a 3D description of the roads based on commercial 2D laser and low cost dead reckoning resources. The result of the estimation process is a global belief about the road conditions, which is maintained by a fleet of mobile agents (e.g. mining trucks). Practical results obtained in realistic environments demonstrate the benefits of the system operation, allowing the vehicles to infer, in advance, the presence of risks on the road surface.

Keywords: 3D sensing, Cooperative Learning, Road Surface Modelling, Haul Trucks.

1 Introduction

The operation of fleets of trucks in dirt roads, typically in open mines, presents many issues related to safety and productivity. These machines usually operate almost continuously under hard conditions. One particular issue is the condition of the roads. These roads usually deteriorate due to heavy transit and adverse weather conditions. The hauling of ore and other bulk material are also the source for the presences of debris on the road. The machines although heavy and robust are affected by those road conditions. One usual problem is the damage in the tires when these hit rocks and other debris on the roads. A heavy truck (e.g. 300 tons) moving at 60Km/h and hitting a solid rock of 40 cm size would likely damage one of the tires. It is easy to think about some sensor that would scan the road ahead the truck and inform the driver (human or autonomous machine) to take adequate control actions (e.g. reduce speed or perform a slight turn to avoid the object). That warning should be given in advance, giving the driver (or even an autonomous system) enough time to perform a smooth avoiding action. Unfortunately, the necessary horizon of time should be in the order of seconds or its spatial equivalent: several tens of meters. This early warning process implies that the machine must detect objects that are located 100 meters ahead the vehicle. Even in good conditions of visibility and line-of-sight, a laser scanner (that is a commercially available, cost effective and accurate enough sensor) would not be able to scan the road with the necessary detail. By the time the machine realized about the obstacle it would be too late to perform any necessary and admissible avoiding maneuver. This situation means that the truck would need more powerful sensing capabilities. Sensors offering those better capabilities do exist, but those are more expensive and not reliable to operate in the dirt road conditions. Conditions such as vibration, dust, humidity, and temperature are too hard for these complex sensors.

Outdoor laser scanners such as the well-known Sick LMS211 ([15][15]) and LMS151, have successfully been being used in mining operation, e.g. in guidance systems ([2] [4]). Years of operation with acceptable maintenance issues have demonstrated that that sensing technology is adequate for mining environments. However, the problem of the limitation in the scanning range does still persist. The scanning range, ideally about 80 meters for that sensor, is when exposed to real conditions (dust clouds, opaque surfaces and reflecting from non-perpendicular surfaces) usually not better than 20 meters. There is no adequate control action, feasible of being generated based on just 20 meters of visibility, for a heavy vehicle moving at high speed. A truck moving at 20 m/s (72 Km/h) would have only one second to react. A human driver would need a big fraction of that first second just to get aware about the risk. After that warning is given, no feasible maneuver would be possible due to the non-holonomic dynamics of the machine itself. If the machine was in trajectory to hit the object it would likely hit it. This limitation would eventually be solved by installing good sensing capabilities on all the trucks. Expensive and sensible equipment in a high number of trucks is an equation that would discourage any technical or financial decision maker. The elevated initial cost and the permanent maintenance costs are the drawbacks of this option.

Alternatively, it would be extremely convenient if the sensing equipment, needed per truck, were cheap and robust (low maintenance) and also if there was no need for a large scale deployment, i.e. no strict needs to install the equipment in all the trucks. This option is feasible through a process involving collective learning of the road conditions. The collective learning is presented in this paper in addition to a method to synthesize the 3D representation of the road surface through relatively low cost and industrial grade sensing equipment.

2 Description of the Approach

The collective learning and sharing of belief about the condition of roads is feasible to be implemented in an inexpensive way. For this process to happen all the trucks would need, at least, a short range / low bandwidth communication system such as the one applied in [3] and [5]. A subset of the vehicles, denominated active vehicles, needs to be equipped with exteroceptive sensing capabilities such as laser scanners, (e.g. one scanner per active vehicle).

In this implementation each active vehicle has a laser scanner, low cost dead-reckoning resource (usually just speed measurements and low cost 3D gyroscopes) and a low cost GPS unit (operating in intermittent and low quality mode, i.e. the denominated autonomous mode). In fact the GPS measurements could be replaced by any low accuracy global localization module, able to provide position estimates of accuracy similar to a GPS properly operating in autonomous mode. The vehicle global localization could be done by GPS, GPS-SLAM fusion [1], constrained dead-reckoning [6] or any process able to provide absolute position estimates having an error in the order of meters (called in this context low quality *global position estimates*).

Based on the laser scans and on the local dead-reckoning estimates, it is possible to generate a high quality 3D local representation of the road and its surroundings. From that surface is it also possible to infer and detect terrain features such as protuberances on the road (usually rocks and debris) and aggressive terrain depressions (pot-holes, ditches). The size and other properties of these features are estimated as well, depending on the resolution of the sensing device. All those objects of interest (OOI) and their estimated characteristics are included in an *on-line database*, on board the vehicle's processing node. These objects are correlated to their absolute positions on the road (provided by the estimates of the vehicle absolute localization). More exactly, for the context of this approach, localization means the localization in the longitudinal coordinate of the road, quantized in sections, due to the fact that the roads are modeled through a set of one dimensional (1D) manifolds. The sections can have lengths of meters or about tens of meters. No higher localization accuracy is required in order to give an early warning to drivers about the presence of risks in certain parts of the road.

As some vehicles get in proximity to others they are able to share information (as the usual situation shown in **Figure 1**). Consequently, the vehicles can update their local beliefs (the on-board databases) with the latest information about the quality of the road sections, provided by other agents. This capability is independent of the sensing capabilities of the trucks, i.e. passive agents can also update their on-board beliefs based on the sharing of information with other agents. In this way, an efficient, reliable and inexpensive (in terms bandwidth usage and sensor deployment) process maintains the most accurate belief about the road condition, in each vehicle's local database.

This process (sharing of belief) is performed by any vehicle (passive or active). The active vehicles have the additional capability of being able to sense the road, in addition to give and receive fresh information to/from other vehicles.

Clearly, the maximum speed for learning the condition of the roads is by making all the vehicles to be active agents and increasing the range of the communication system. However, the approach does even work with just a small number of active units and operating with a short range communication resource, although this situation would decrease the convergence speed of the estimation process.

2.1 3D Representation of the Road Surface

Several fusion processes are involved in the synthesis and interpretation of the 3D imagery: The approaches for modelling the road and estimating the pose of the vehicle are achieved through sensor data fusion of on-board sensing resources: laser scanner, 3D inertial measurement unit (IMU) and speed measurements. Based on the estimates provided by the fusion processes, a number of client processes perform higher level tasks such as the feature extraction to detect obstacles.

Details about the perception processes are not included in this document; those are based on previous work of the authors. The 3D synthesis was treated in [8] and [9], and the 3D terrain modeling based on PWML (Piece Wise Multi Linear) was developed for different applications in [7] and [9].

In this particular application, where the main interest of the perception processes is to detect features of the road, the scanner is maintained at a fixed elevation angle, respect to the vehicle chassis, in order to scan the road surface linearly and generate a high quality 3D image just locally (there is no interest in generating large 3D maps). In other applications, developed by the authors such the works presented in [8],[13] and [9], the objective also considers the possibility of performing 3D scan matching, for generating global 3D maps. In those cases the scanner is periodically rotated in its elevation angle ([8],[13]) or in azimuth ([9]), in order generate sequences of good quality 3D images which have common areas. The overlapped 3D images allow a scan matching process to generate large 3D maps, a resource that is not required in the project presented in this paper, although it could be exploited for the global localization process.

It must be remarked that the 3D synthesis process, which is based on the mentioned sensor data fusion processes, allows the generation of high quality 3D imagery, in a local sense; consequently allowing an effective and reliable feature extraction for the detection of risk on the roads, nearby the active agent.

Figure 1. A usual context: Vehicles operating in a dusty environment that limits the visibility for human operators and sensing resources. The vehicles are able to establish short range communication, usually under Line of Sight (LoS) conditions, when some agents are temporarily in proximity (as the situation captured by the picture). They are briefly able to communicate in order to share their beliefs about the context of operation. A laser scanner of the LMS211 family can be seen at bottom right of the image, where the sensor was installed for testing the sensor, on a maintenance vehicle operating in a mine. This photograph was taken from inside the cabin of the vehicle.

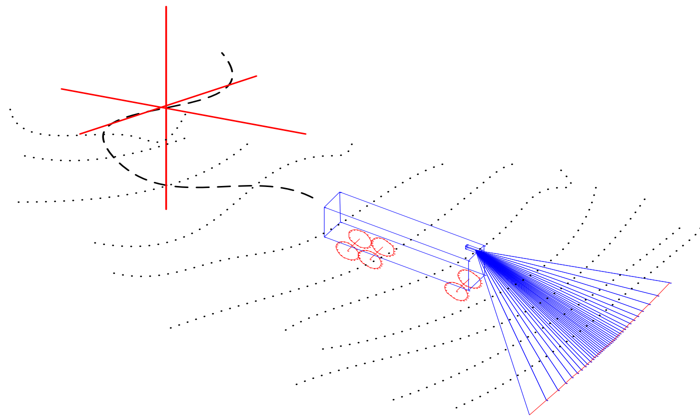

Figure 2. Sketch showing the scanning process: a sequence of laser scans and the associated 3D projection. In the figure the separation between consecutive scans, in the vehicle's longitudinal direction, is exaggerated. In a real case, e.g. for a truck travelling at 60km/h, the separation between scans would be, in average, approximately 22 cm (for a LMS211 operating at a scan rate of 75 Hz). Oscillations of pitch angle, mainly due to suspension effects, can temporarily increase and decrease the inter scan separation.

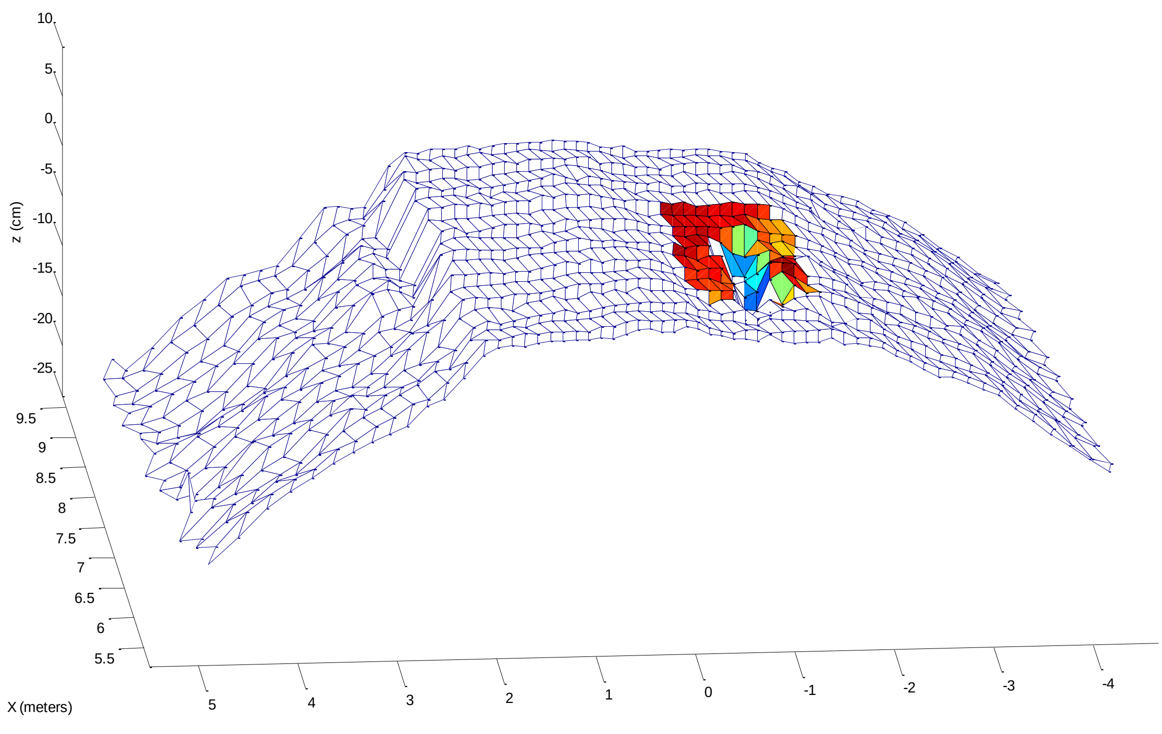

Figure 3. Example of a detected object of interest (OOI), in this case a pot-hole detected on the surface of a dirt road (Note that the scale of the z axis is in cm and the scale of X and Y is in meters). A small patch containing the OOI is extracted and published in the system database. A higher level client module would process the patch to classify and extract relevant features, in order to compress the information for sharing purposes. Those 'compressed' properties of the OOI are the pot-hole diameter, depth and low accuracy geographical

position (global localization). More details about these perception processes are available on-line in [11].

2.2 Learning of the Road Condition

The objective of the learning process is the synthesis of a belief for representing the condition of the road. It must be noted that the condition of the road is time varying, due to the fact that features can appear and disappear (e.g. removal of debris) through the time.

The learning process is based on observations that can be *positive* or *negative*. Positive observations are those that report the existence of an object on the surface of the road. Negative observations are those that report that an area of the road is clear of features (no objects of interest or OOI). Both types of observations contribute to the update of the belief about the road condition. Positive observations may produce the creation of a new item (in the population of OOIs, described in the current belief) or usually they increase the probability of existence of some previously defined OOI.

A negative observation would reduce the probability of existence of certain known OOI in a related area of the road. Typically, those cases happen when dangerous objects (e.g. debris) are removed from the road.

It must be noted that each observation involves a *Data Association* process, i.e. the new perceived OOI is usually associated to a currently known OOI.

For instance, if an OOI is detected by the agent perception processes, then the first step is associating it with a section of the road (global location). The second step verifies if one of the currently known OOIs (listed in the database that implements the belief) could be associated to the recently perceived OOI. If one of the currently known OOIs does match the new one, then its current assigned probability of existence is increased. If none of the current database's OOIs could be associated to the recently perceived OOI, then a new OOI is added to the database.

There are diverse OOI's properties that are used for performing the Data Association. One of the properties is the geographical position of the OOI, which is discretized in terms of *road sections*. The position is expressed by at least 2 integer parameters that identify the road and the section of the road where the point is located. Any global localization process, providing accuracies in the order of meters (e.g. 10 meters), would be adequate for allowing a reliable Data Association process. This is why the system only requires low accuracy estimates for the global localization.

A second relevant property is the *type* of OOI. The type is also an integer parameter, enough for labelling OOI as pot-hole, obstacle, ditch, etc.

A set of secondary properties, mostly dependent on the type of OOI, are also perceived and associated to each database's OOI. Those properties are used for diverse purposes, such as reporting and also Data Association as well.

2.3 Learning Policies

Two different policies for the learning process are proposed. One of them works in a centralized fashion by giving a particular node (or set of nodes) the task of collecting the observation of all the agents and performing the fusion in order to estimate the map of OOI for all the roads. That particular node is a permanent node and is termed *base station* (BS). The BS node may have other duties as well, although those are not of interest in this paper. Each agent that briefly communicates with the BS is able to receive the updated version of the belief.

This class of policy is effective in contexts of operation where all the agents usually visit the BS, typically in open mines operations.

A BS is not necessary implemented by a sole node; it could be implemented by set of nodes, geographically separated but permanently connected via network resources. Multiple BS nodes allow more flexible distribution of agents and itineraries, provided that each of the agents performs periodic visits to at least one of the BS nodes.

A second policy works in a decentralized way and allows all the nodes to contribute with on-board processing in the estimation stage, i.e. not just actuating as perception nodes. They collect observations like in the centralized policy but they additionally process these observations and produce an update of the private beliefs. Each time that two agents enter in brief radio contact, they perform a fusion of both beliefs. This policy is more expensive in terms of usage of communication resources, however it allows a more reactive update of the beliefs and also allows the agents to exploit the benefit of the cooperative learning without needing to visit any BS node.

3 Experimental Results

Comprehensive experiments, in simulation and real platforms were performed. For the experiments performed on a real fleet, the number of agents was low (3 agents) due to limitation in the available resources (mainly vehicles retrofitted with laser scanners).

Figure 4 shows the paths (i.e. coarse waypoints, planned on the roads) that were part of the real test. All the objects that were perceived to be risks for the vehicles were indicated on the road by yellow squares.

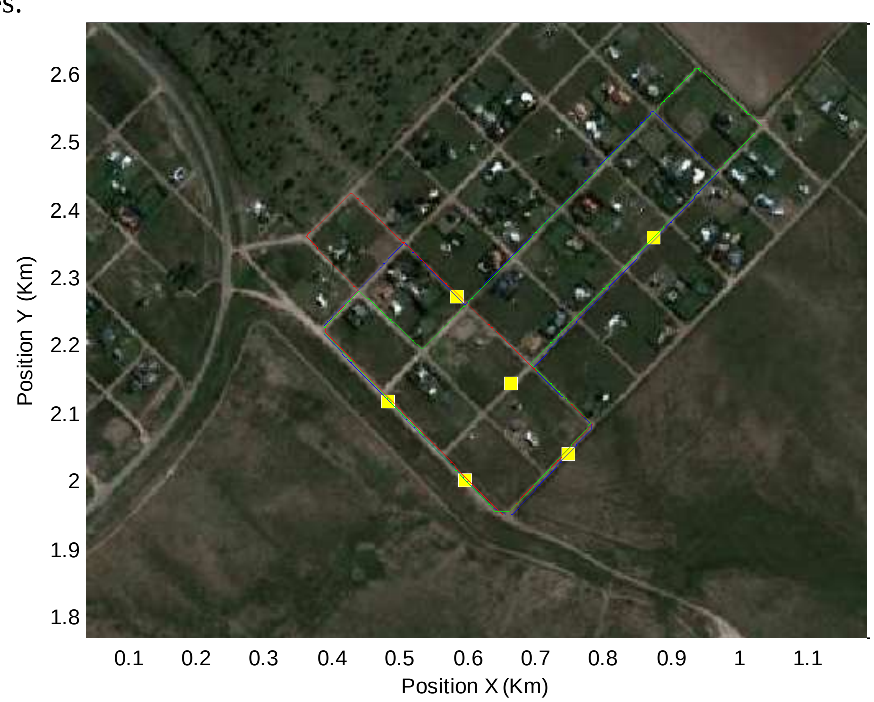

Figure 4. Satellite image showing the place of the experiments. The place included dirt roads in rural and residential areas. The superimposed polylines indicate the waypoints assigned to

each agent. The agents travelled some common sections of the roads but had different nominal speeds and sequences of waypoints. OOIs are indicated by yellow squares.

Some additional situations were tested such as processing negative information, i.e. the case where certain OOI had disappeared (an obstacle physically removed from road). The removal policy is conservative in the sense than more than one iteration is required in order to completely remove a risk from the belief (as false negatives are considered more critical than false positives). For the case of processing positive information, only one iteration is needed to accept the existence of a new OOI.

Some static agent was also included in some experiments. This type of agent behaves as a vehicle that is parked and still receives, updates its belief and shares information to temporarily nearby mobile platforms.

The implementation of the communication resource was based on standard narrow band radios (XBeePro [10]), operating data rates of about 128 Kbps.

Three mobile active nodes were implemented through standard cars, which were legally able to operate in the rural roads of the area where the experiment took place. Pictures of one the cars can be seen in [11], where an IMU, a laser scanner (a blue LMS200 ([15][15]), on the car's roof), a GPS and an XBee radio where installed, in addition to an on-board computer (a laptop). Figure 4 shows a satellite image of the test area, where the roads involved in the test are indicated. No all the roads were travelled by all the vehicles. Small yellow squares indicate the approximated location of the detected OOIs.

Figure 5 shows the events where the vehicles had an encounter, i.e. each time when a couple of them are in geographical proximity and establish a brief communication process. The event are shown against the time, that is labeled in minutes (is should be noted that the covered area is in order of kilometers, and the vehicles were moving at speeds of about 20Km/h).

Figure 6 shows the events where the active vehicles detect an OOI. Those are *positive observations*, i.e. real OOI being inferred from the perception processes.

Figure 7 shows the beliefs about the existence of the detected OOIs. The beliefs are estimated through probabilities implemented through counters, although the values, shown in the figure, are simplified to existence (value=1) or no existence (value=0).

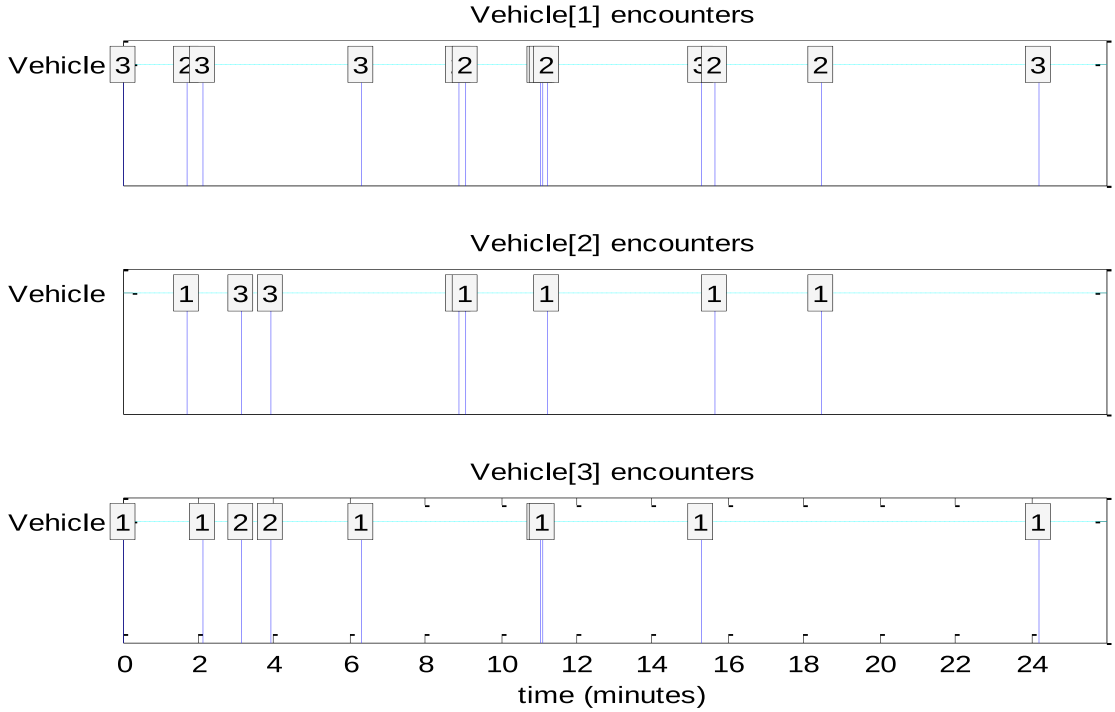

Figure 5 Encounters (brief data sharing action between nearby agent) experienced by each of the agents during part of the test. The top subfigure shows the encounters from the perspective of agent 1. Agent 1 establishes sharing sessions with agents number 2 and 3, as it is indicated in the figure. The other two subfigures present the encounters from the perspectives of agents 2 and 3 respectively.

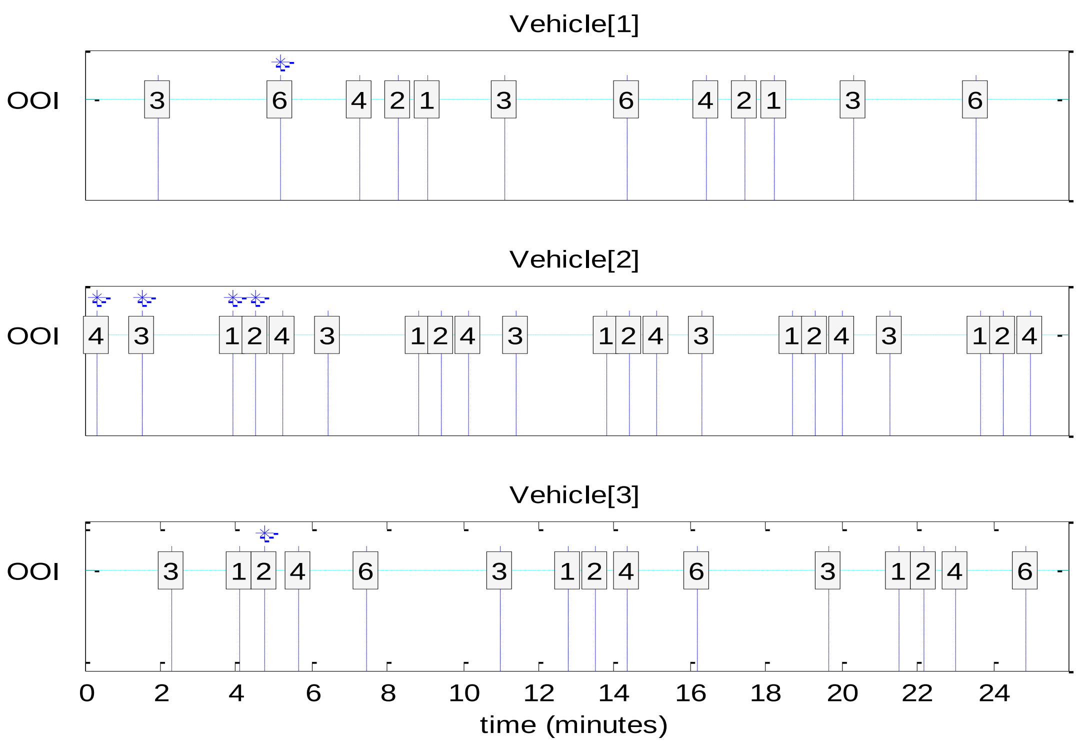

Figure 6 Risks individually observed by each of the agents. The symbol at the top of each vertical line is used to graphically indicate the ID of the risk. Naturally, the same agent may observe the same risk a number of times, as it revisits the same sections of the road when repeating its usual trip (a typical behavior of hauling trucks in mining contexts).

4 Conclusions

The performed tests validated the expected performance, which had been previously obtained in synthetic simulations. It showed that when an agent discovered a new threat the rest of the agents were usually able to know its existence well in advance to the time when they would physically encounter it. The main implication of the results is that this low cost system reliably operates in diverse weather conditions because is based on medium range perception capabilities but applied to short range perception, what gives the system the capability to still operate in low visibility conditions (e.g. dust, rain). It is also important to note that high accuracy perception is applied where it is needed and cost effective (local 3D imagery), while easily achieved low accurate estimates are generated for client processes because those just require that quality (e.g. global localization that is used for Data Association and Event Reporting).
In addition to the usual low accurate GPS (low cost autonomous or differential modes), the localization was achieved via Virtual GPS [6], intended for GPS denied contexts, what means this system could also operate in difficult contexts such as underground mines.

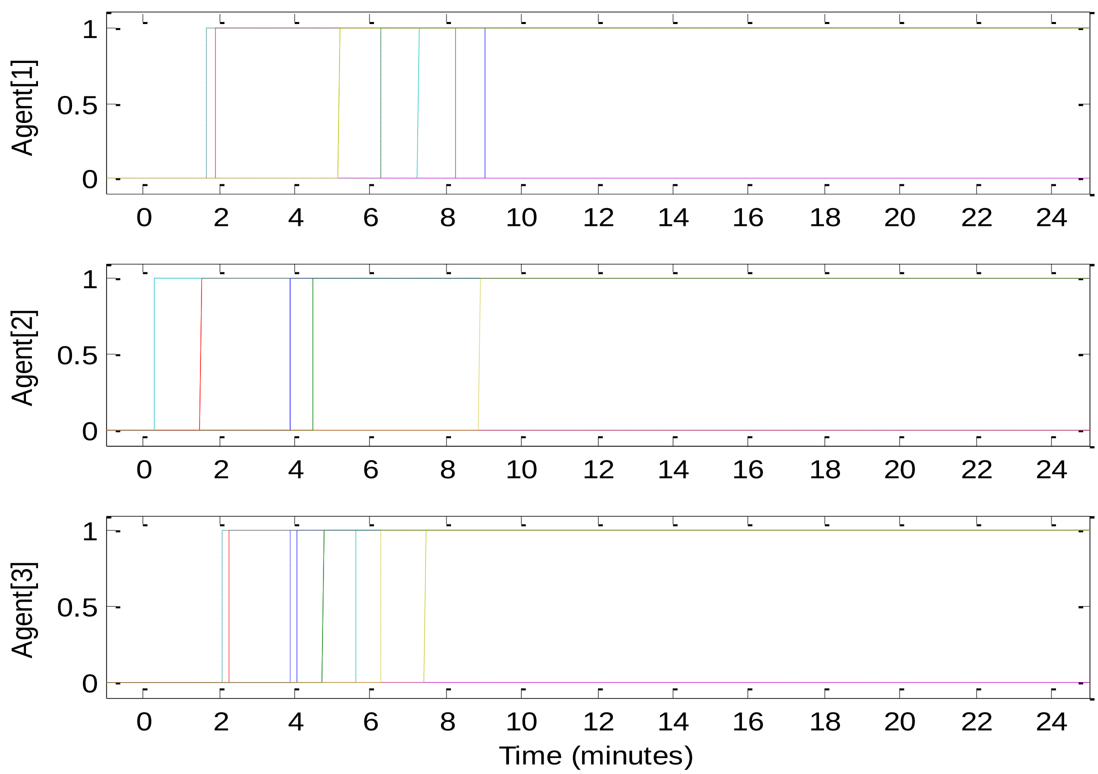

Figure 7. Both classes of individual beliefs: purely individual (based on on-board observations, indicated in solid lines) and the belief based on collective learning (shown by

broken lines). Value =1 means the agent knows about the existence of the risk. Each of the 6 risks is indicated by a different color. Values = 0 indicate that the agent does not know about the risk, at that time. Transitions, from 0 to 1, represent the moment when the risk is learned by the associated belief. For a particular color, it can be seen that in many cases, the broken lines show transitions in advance to the associated continuous lines (of the same color), what means that the agent, due to the collective belief, estimated the presence of those risks in advance to their physical detection.

Additional Resources

Additional resources, for the readers that may be interested in this work, are available on-line in [11]. These resources include videos and data used in this paper and future related work.

Acknowledgements

The author would like to thank the colleagues that helped in the field tests performed for the experimental validation of this work, in Bahia Blanca, Argentina, in 2012.
To Mr. Hernan Orta, for his technical advice and support.
To Ms. Alicia Robledo (Possum D&D), for the design of a number of components for the installation of sensors on the vehicles and for providing diverse software modules.

References and Previous Work

[1] Guivant J., Masson F., Nebot E. M. (2002) "Simultaneous Localization and Map Building Using Features and Absolute Information". Journal of Robotics and Autonomous Systems, Volume 40, Issues 2-3, 31 August 2002, Pages 79-90

[2] Nebot E., Guivant J, Worrall S. (2003) "Haul Truck Alignment Monitoring and Operator Warning System".Journal of Field Robotics, vol 23, no 2, 2006. pp 141-161.

[3] G Kloos, J Guivant, S Worrall, A Maclean, E Nebot, (2004) “Wireless Network for Mining Applications” Proceedings of the 2004 Australasian Conference on Robotics & Automation, Canberra, Australia, 2004

[4] E Nebot, J Guivant, (2005)- “Heavy Vehicle Guidance System”. WO Patent WO/2005/024,536,

[5] E Nebot, J Guivant (2005), “Virtual Network System”, WO Patent. WO2005032171 A1

[6] J Guivant, R Katz (2007) “Global urban localization based on road maps” Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ, 2007.

[7] Robledo, A. Cossell, S. and Guivant, J. (2011), “Outdoor ride: Data fusion of a 3d Kinect camera installed in a bicycle”. Proceedings of the 2011 Australasian Conference on Robotics & Automation, Melbourne, Australia, 2011.

[8] J. Guivant , (2008) “Real time synthesis of 3D images based on low cost laser scanner on a moving vehicle”, V Jornadas Argentinas de Robotica, Bahia Blanca, Argentina, 2008.

[9] Guivant, J., Cossell, S., Whitty, M. and Katupitiya, J., (2012) “Internet-based operation of autonomous robots: The role of data replication, compression, bandwidth allocation and visualization”. Journal of Field Robotics, Volume 29, number 5, pages 793-818, 2012,

[10] http://www.digi.com/gp/xbee

[11] http://possumrobot.com/Documents/HVTT13a.htm

[12] http://www.microstrain.com/inertial/3DM-GX3-25-OEM

[13] http://possumrobot.com/Documents/ManlyIn3D.htm

[14] https://www.mysick.com/

[15] http://sicktoolbox.sourceforge.net/docs/sick-lms-technical-description.pdf