

# Automation of an Underground Mining Vehicle using Reactive Navigation and Opportunistic Localization.

Elliot S. Duff, Jonathan M. Roberts and Peter I. Corke  
CSIRO Manufacturing & Infrastructure Technology  
PO Box 883, Kenmore, Qld 4069, Australia  
<http://www.cat.csiro.au/cmit/automation>

## Abstract

*This paper describes the implementation of an autonomous navigation system onto a 30 tonne Load-Haul-Dump truck. The control architecture is based on a robust reactive wall-following behaviour. To make it purposeful we provide driving hints derived from an approximate nodal-map. For most of the time, the vehicle is driven with weak localization (odometry). This need only be improved at intersections where decisions must be made — a technique we refer to as opportunistic localization. The truck has achieved full-speed autonomous operation at an artificial test mine, and subsequently, at a operational underground mine.*

## 1 Introduction

Mining is an important global industry that has so far exhibited a slow uptake of robotics and automation technology — but this is beginning to change. The main driving forces for this industry are a need for increased productivity, societal demands for safer workplaces, an aging work-force, and declining value (in real terms) of its products. To date, productivity increase has been achieved through mechanisation, moving from human and animal power early last century to present-day electric and diesel powered machines. Machines have become progressively larger and more powerful but practical limits are reducing the rate of growth. Automation has been identified as the most likely means to attain the next quantum jump in productivity and safety since sensing, control and computing technologies are advancing rapidly.

One of the key issues for underground automation is navigation. A review of appropriate navigation techniques is presented in [1]. Both absolute and reactive navigation techniques have been, and are being developed to guide autonomous vehicles in mine tunnels. Absolute navigation techniques that rely on detailed *a priori* maps are probably not viable in the mine tunnel application. By their very nature, mines change every day and to rely on old

maps is hazardous. Simultaneous localization and mapping (SLAM) techniques provide a nice solution in that maps are generated and updated continuously. In this respect SLAM is bringing active perception concepts into the absolute navigation framework. However, SLAM systems are in their infancy and have not been tested in the mining environment. A very real problem with this technique is with long straight tunnels which provide no natural topographical features to robustly detect and track.

In contrast, reactive navigation is ideally suited to the underground mining environment. In the last decade there have been a number of attempts to use reactive navigation on LHDs, notably there have been systems that required “rail-guides” to assist with navigation [2]. A navigation system that relies upon wall-following has a significant advantage over “rail-guided” systems because it has the potential to operate with no guidance infrastructure (apart from the walls). This leads to lower installation and maintenance costs which can be significant over large areas of a mine. The wall-following technique used in our system is described in [3].

For the vehicle to perform useful tasks, the low-level reactive wall-following behaviour must be augmented with a purposeful strategic behaviour [4]. Typically this would be achieved by knowing the location of the vehicle in the mine. In this paper, we will demonstrate that it is not always necessary to know the exact location of the vehicle. The classical question “Where am I?” can be answered simply by “In a tunnel! Heading for the next intersection.” This makes sense from everyday human experience where we drive to an unfamiliar location, not by following a trajectory of spatial coordinates, but by a road-following behaviour combined with recognition of intersections at which appropriate decisions are made. In this scenario, the vehicle need only determine its location whenever the opportunity arises — a concept we refer to as *opportunistic localization*.

## 2 Implementation

In 1996 we conducted a number of field trials at a mine in Queensland, Australia [5]. Data were collected from a number of sensors mounted on an LHD with the aim of determining which sensors performed best underground. It was found that the 2D laser scanner was an ideal primary navigation sensor in the underground environment. In July 1998 we began an industry funded project to develop an autonomous navigation system for an LHD which was field tested in July 1999.

### 2.1 Automation architecture

We were fortunate in being able to “piggy-back” the automation system onto an existing tele-operation system which provides an excellent substrate for automation since:

- there is only one point of control;
- the electrical interface necessary to control the vehicle already exists;
- the safety systems built into the tele-operation systems can be used; and
- the vehicle can be switched back to the underlying tele-operation mode.

This last point is particularly important because it means the vehicle can still be used for operations that have not been automated yet (i.e digging), and it enables the recovery of the vehicle if the automation system fails in a hazardous environment. Our automation system comprises three functional software layers:

**Strategic** layer determines when and where the vehicle will perform an action based upon a list of hints — defined by the selected route. To do this, the strategic layer must be able to estimate the approximate position of the vehicle. This knowledge is used to influence the behaviour of the vehicle through hints (i.e. driving and turning strategies) which are passed down to the tactical layer. Since this layer has some expectation that it is approaching a node, e.g. an intersection, it can use this knowledge to assist in its recognition.

**Tactical** layer obeys strategic driving hints and actually “drives” the vehicle whilst avoiding the tunnel walls. The desired vehicle path is estimated with active contours that follow the walls. The tactical layer has no knowledge of the location of the vehicle with respect to a global coordinate frame, it simply senses and reacts to the walls. It sends steering and speed set-points to the operational layer.

**Operational** layer contains the control loops that convert steering and speed set-points into low-level machine input signals that control functions such as throttle, gear box, brake and articulation joint hydraulic rams.

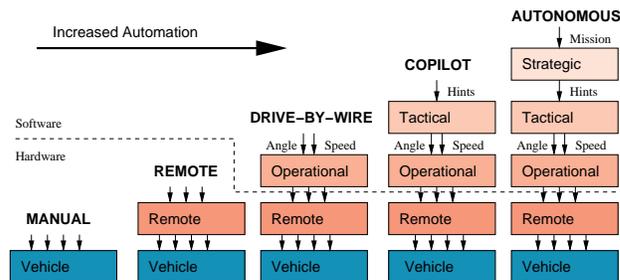


Figure 1: Operational modes of the system.

This control architecture allows various modes of operation (Figure 1). The original mode of operation is the **manual** mode, where the vehicle is controlled by an operator who is sitting in the driver’s seat. The second mode is the **remote** mode, where the vehicle is controlled with a joystick via a tele-operation system. In the first of the computer controlled modes, the **drive-by-wire** mode, the operational layer accepts speed and steering set-points from the operator and hides the actual dynamics of the machine. In the **co-pilot** mode, the tactical layer controls both speed and steering of the vehicle. The operator acts as a co-pilot and provides hints to the tactical layer which will influence behaviour at decision points. And finally, in the **autonomous** mode, the strategic layer interprets a mission and generates the appropriate hints to the tactical layer, which in turn generates the appropriate speed and steering demands to the operational layer. The vehicle is given a mission by the operator, who subsequently, has no influence over the vehicle’s driving behaviour.

### 2.2 Development environment

Our development environment includes a 30 tonne LHD and an artificial test mine constructed from shade-cloth (Figure 2). The roadway is 300m in length and contains curves, sharp corners, a loop and a large “room” (simulating an underground workshop). The shade cloth walls are opaque to the lasers and transparent to radio frequencies. Thus, it was possible to develop the system using a standard low-cost wireless high-bandwidth local area network (LAN).

### 2.3 Missions, routes and nodes

To perform useful tasks, the vehicle must perform a sequence of actions. A sequence of actions is defined by a

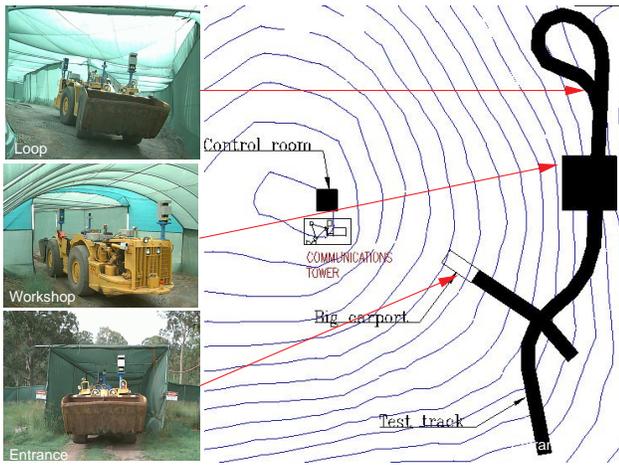


Figure 2: Layout of the test mine.

mission. Missions are loaded and monitored using a remote graphical user interface. There are three actions that the vehicle must perform; driving (along a route), dumping and digging.

A nodal-map of the mine can be constructed from a traditional map of the mine, or it can be constructed by driving the vehicle along the path and observing the local environment. Nodes are typically identified as points that have obvious natural features. A nodal-map of our test track is shown in Figure 3. Here the nodes are represented by circles and the lines that link nodes are called segments. A route is defined by a list of nodes that the vehicle must pass through. For example, the route from the start of the track to the stop-log via the loop, would be defined by the list of nodes: N1 N2 N3 N4 N5 N5 N4 N3 N2 N6.

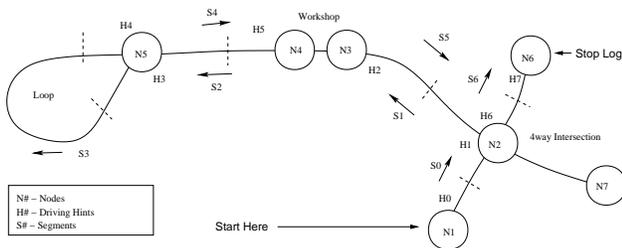


Figure 3: Nodal map of test mine.

To assist the tactical layer, the route contains a number of hints. These hints are only appropriate along specific sections of the track. After driving the LHD in co-pilot mode a few times it is possible to estimate the optimal driving strategies. This is analogous to creating pace notes for rally driving.

## 2.4 Odometry

We use odometry to provide an estimate of distance travelled along each segment (path between one node and the next). This is used to:

1. give some indication of distance remaining to the next node, and to flag an error if the expected node has not been found within a certain tolerance,
2. schedule driving hints such as maximum speed, keep left, keep right etc.

Odometry is a sensor that is poorly regarded within the outdoor mobile robotics literature. While the problems of slippage and wheel radius estimation are real, the difficulties are perhaps exaggerated. For our application, a 30 tonne vehicle, operating on dirt roads we have found odometry errors to be less than 1%. Our odometry is based on drive shaft rotation and calibrated to distance travelled. Clearly this calibration will change with time but a simple learning rule, based on odometer reset errors, could be used. This is in contrast to approaches such as [6] where 'wheel radius' is estimated online, but also includes lumped model error.

## 3 Experimental results

### 3.1 Test mine results

To test the performance of the autonomous LHD, a mission was constructed that simulated a typical task of an LHD in a real mine. The mission consisted of a route from a digging position to a dumping position, followed by a dumping action, which in turn was followed by the reverse route from dump to digging positions. The forward (dig-to-dump) route is shown by the arrows in Figure 3. Figure 4 shows the dig-to-dump route in a series of "snapshots" constructed from the laser data.

During a mission the odometry of the LHD was automatically corrected at each node. The size of the correction ranged from a few centimetres (when the LHD was travelling slowly) to a few metres (when the conditions were muddy with increased wheel slip, or when the vehicle was travelling quickly and became airborne).

To gauge the performance of the automated LHD it is useful to compare it to a human driver performing the same mission. A comparison between the speed of the vehicle under manual and automatic control is shown in Figure 5, where the range (x-axis) is defined as the distance from the starting position. In the upper plot the vehicle travels in the forward direction i.e. the speed is positive and the vehicle moved from 0 to 300m, while in the lower plot the vehicle travels in the reverse direction i.e. the speed is negative

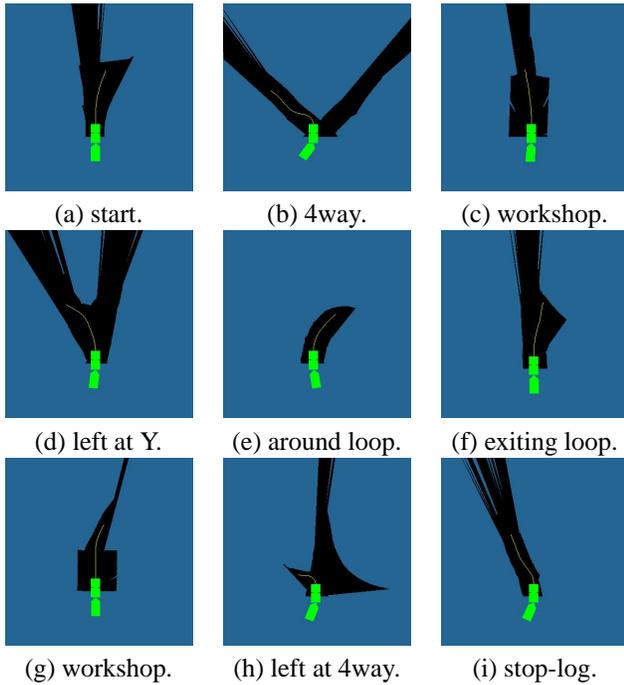


Figure 4: Laser profile of forward route.

and the vehicle moves from 300m back to 0m. From this comparison a number of points can be made:

- In the first 25m, as the LHD turns the first sharp left hand corner, its speed was fixed at 5km/h, while the human driver took a more aggressive speed of 10km/h.
- Once the corner was successfully negotiated, the LHD was driven at maximum speed. The actual speed was determined by the ground conditions. There was very little difference in the speed between manual and automated control.

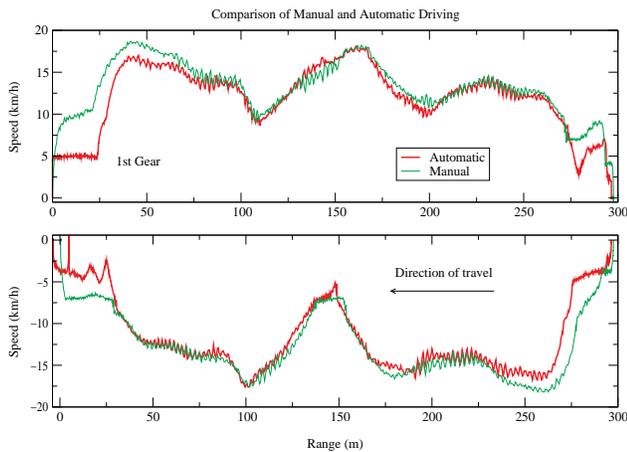


Figure 5: Speed of manual and automatic driving.

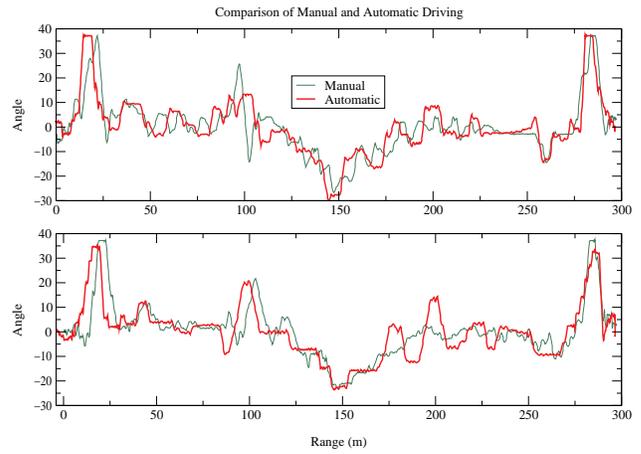


Figure 6: Steering at high speed.

- At the final corner, the LHD must be prematurely slowed for the front laser to see the corner and be given enough time to act.

A comparison between the steering of the vehicle is shown in Figure 6, from which the following observations can be made:

- At 20m and 280m the LHD articulated to full left lock as it negotiates the 4WAY intersection in both manual and automatic runs.
- At 150m the LHD articulated 20° to the right around the back of the loop in both manual and automatic runs.
- Along the straight sections there was considerable oscillation in both manual and automatic runs. This was due to the rough nature of the track.
- In the forward direction (upper plot) the first left hand turn was taken earlier by the human driver, while the last turn was taken at the same range.
- On the return route (lower plot) the situation was reversed. The first right hand turn was made at the same range, while the last right hand turn was made late.

The delay in turning is related to the interpretation of free space by the tactical (wall-following) layer and by the fact that the laser cannot see around the corner. This delay limits the speed at which the automation system can be used around such sharp corners.

Figure 7 shows the clearance between the LHD and the tunnel wall for the cases of manual and automation operation, from which the following observations can be made:

- The noise in the range data was due to either holes in the shade-cloth, or reflected light from dust.

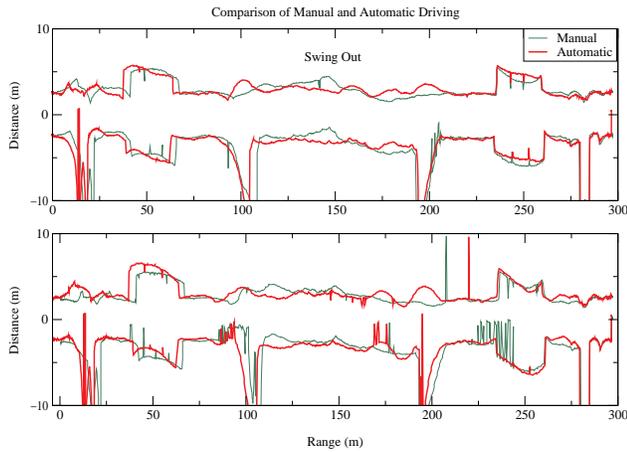


Figure 7: Distance to walls at high speed.

- The difference in the horizontal position of features (i.e. entrance to workshop) was due to slip and the orientation of the LHD.
- At 150m, in the forward direction, the LHD under manual control tended to hug the left hand wall of the loop, while the automated system remained in the centre of the tunnel.

To summarise, the LHD was reliably driven along a 300m route, which included two  $90^\circ$  corners and a sweeping loop with a radius of curvature of less than 8m, at speeds up to 18km/h. The vehicle has operated for over an hour at a time without any human intervention.

Under most conditions the LHD was driven autonomously at the same speed as a human operator. Subjectively, the LHD under autonomous control takes a better line, and reacts faster than a human operator. The only weakness occurs at sharp intersections where the autonomous LHD must travel slowly to “see” around the corner. An experienced human driver can drive more aggressively around a blind corner because they can remember the profile of the tunnel from a previous run. This is a small penalty to pay for having no map, however some form of local map of each intersection could be added and learnt online.

### 3.2 Underground results

The second phase of the project was to establish whether the vehicle would perform in the same manner under operational conditions. For this series of tests, the LHD was transported to an operational underground mine. The mission control computer was connected to the existing mine communications infrastructure (leaky feeder — a 1200 baud RF communications system). The system was up and running in less than two days. A section of tunnel, 150m

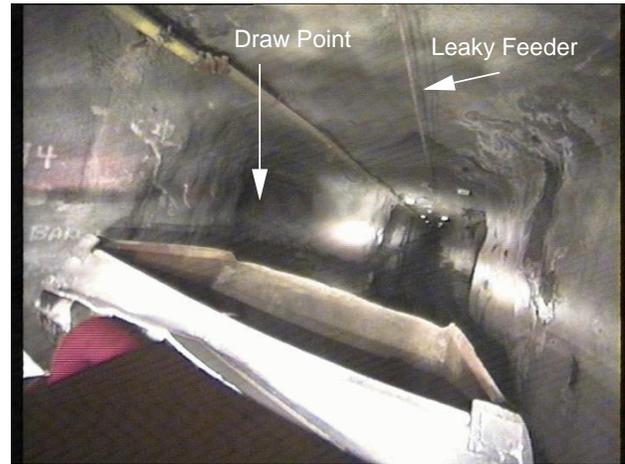


Figure 8: LHD at draw point.

in length was isolated for the test. For half a day the vehicle was driven up and down the test tunnel to acquire data to generate the relevant driving hints. In the main test the vehicle drove autonomously for an hour, periodically switching back to the tele-remote system to dig and dump ore. At one stage a Palm Pilot was used to enter driving hints, demonstrating the simplicity and independence of the control communications. During the demonstration a video record was taken from the LHD. In Figure 8, the LHD is shown stopping in front of a draw point (where ore is loaded). The route was over 100 m in length. At no stage during these underground trials did the LHD touch the walls.

## 4 Discussion

To reiterate, both absolute and reactive navigation techniques have been, and are being developed to guide autonomous vehicles in mine tunnels. Traditional approaches to absolute navigation (non-SLAM approaches) are blind to their environment — the control of the vehicle is inferred from the position of the vehicle, rather than what the sensors tell it about the position of the walls. Such techniques are therefore highly dependent on the accuracy of both the localization and the map — any error may cause the vehicle to collide with a wall.

Reactive navigation is far more robust since the vehicle is controlled by the actual free space perceived in front of the vehicle. The vehicle is able to move at high speeds without any knowledge of its global position, however such knowledge is essential if the vehicle is to be purposeful and choose the appropriate path at an intersection. Opportunistic localization implies the vehicle knows only the segment of the route on which it is travelling and the identity of the next node. Furthermore, knowledge of the vehicle’s

position allows the vehicle to operate at speeds higher (or lower) than the free-space would recommend (e.g. on long curves, or bumpy terrain).

The split between control and localization, highlights one of the main differences between absolute and reactive navigation. Using absolute navigation, vehicle control can only be achieved after the vehicle's absolute location has been established. The two routines are tied together; they are synchronous. The vehicle cannot be controlled if there is a problem with localization. In many situations, the task of localization can be very difficult. It can be computationally expensive and requires redundant information to make it robust. To improve the reliability of finding landmarks, infrastructure is added to the environment (e.g. reflecting beacons).

On the other hand, with reactive navigation, vehicle control and localization are "decoupled", the vehicle control can run independently, and at a much higher bandwidth than localization. This is critical for high-speed autonomous vehicles. In practice, vehicle control is a low-level process performed at a high bandwidth, while localization is a high-level process performed at a much lower bandwidth. Since localization is not critical to vehicle control, then its reliability and robustness is less important. It may even be possible to have no localization infrastructure, and in this application we use the intersections themselves as landmarks.

## 5 Summary

The results so far are very promising. Our autonomous LHD is close to the performance of a human operator around our test mine. The system has also been successfully demonstrated underground at a real mine. Again, the vehicle was able to operate at full-speed through a typical production cycle without localization infrastructure or physical changes to the mine tunnels. In the trials that we have conducted, we have shown that:

- reactive navigation can control a high-speed mining vehicle,
- opportunistic localization is sufficient to navigate underground,
- under most conditions our system can equal the performance of a human operator,
- it can operate with no localization infrastructure, and
- 2D scanning lasers are currently the sensor of choice.

## Acknowledgments

The authors would like to thank the rest of the project team: Graeme Winstanley, Stuart Wolfe, Andrew Castle- den, Leslie Overs and Reece McCasker who designed and built the hardware for our LHD. Thanks also to Pavan Sikka and Robin Kirkham who wrote key software modules, and to Jock Cunningham who managed the project. The experimental part of this work was conducted and partially funded through the CRC for Mining Technology and Equipment (CMTE). Funding from industry was provided through the Australian Mineral Industries Research Association (AMIRA) under project P517 by MIM Holdings Ltd, Normandy Mining Ltd, WMC Resources, North Limited, Pasminco Ltd, Western Metals, Automotive Industrial Mining Supplies and Caterpillar Elphinstone Pty Ltd. Research results from other CSIRO projects were also applied.

## References

- [1] J. Roberts, E. Duff, P. Corke, P. Sikka, G. Winstanley, and J. Cunningham, "Autonomous control of underground mining vehicles using reactive navigation," in *Proceedings of IEEE Int. Conf. on Robotics and Automation*, (San Francisco, USA), pp. 3790–3795, 2000.
- [2] G. Brophey and D. Euler, "The Opti-Trak System, a system for automating today's LHDs and trucks," *CIM Bulletin*, vol. 87, no. 984, pp. 52–57, 1994.
- [3] E. Duff and J. Roberts, "Wall-Following with Constrained Active Contours," tech. rep. Yet to be submitted.
- [4] R. C. Arkin, "Towards the unification of navigational planning and reactive control," Working Notes of the AAAI spring Symposium on Robot Navigation, Stanford University, March 1989.
- [5] S. Scheding, E. Nebot, M. Stevens, H. Durrant-Whyte, J. Roberts, P. Corke, J. Cunningham, and B. Cook, "Experiments in autonomous underground guidance," in *Proc. IEEE Conf. Robotics and Automation*, (Albuquerque, NM), pp. 1898–1903, 1997.
- [6] S. Scheding, G. Dissanayake, E. Nebot, and H. Durrant-Whyte, "Slip modelling and aided inertial navigation of an LHD," in *IEEE Conference on Robotics and Automation*, pp. 1904–1909, 1997.