Atlas Copco Infrastructureless Guidance System for High-Speed Autonomous Underground Tramming

J. Larsson *Atlas Copco Rock Drills AB, Sweden*

J. Appelgren *Atlas Copco Rock Drills AB, Sweden*

J. Marshall MDA Space Missions, Canada *(currently with Carleton University, Canada)*

T. Barfoot MDA Space Missions, Canada *(currently with University of Toronto, Canada)*

Abstract

During the last decade, mining companies and mobile equipment manufacturers have pursued improved efficiency, productivity, and safety in underground mining operations by automating some of the functions of underground vehicles. This paper describes the implementation and successful field testing of a new infrastructureless guidance system for autonomous tramming of centre-articulated underground mining vehicles (e.g., load-haul-dump and mine trucks). The project described in this paper is the result of a technical partnership between MDA, an experienced mining high-tech provider, and Atlas Copco, a world leader in the design of underground mining equipment.

1 Introduction

This paper describes the implementation of a new infrastructureless guidance system for autonomous tramming of underground load-haul-dump (LHD) mining vehicles. Also, included are the results of tests that have been performed on two different LHD models to verify the functionality of the system. The project described in this paper is the result of a technical partnership between MDA, an experienced mining robotics provider, and Atlas Copco, a world leader in the design of underground mining equipment. This work was first described in Marshall and Barfoot (2007).

Despite previous efforts by several parties to automate the tramming function of underground machines, widespread adoption of such technologies has yet to occur in the minerals industry. It has been reported by mine operators that this is, at least in part, due to poor reliability and lack of robustness in these technologies. Thus, the purpose of the described project was to design a fast, reliable, and robust "autotramming" technology that does not require the installation of fixed infrastructure throughout the mine.

There are several factors that make infrastructureless autotramming a challenging task. Firstly, the characteristic centre-articulated and hydraulically actuated steering mechanism makes such vehicles difficult to control at high speeds. In this paper, we describe a system architecture that effectively handles these substantive vehicle dynamics. Another significant challenge is the problem of precise and real-time underground localization. Underground mines are constantly changing, thus a system that requires fixed infrastructure to localize vehicles would necessitate the constant installation of this infrastructure as the mine advances. It is widely viewed that an infrastructureless system is preferable, but such a system must easily allow for possible changes/advancement in the environment. In this paper, we describe a robust localization method that fuses data from various sensors to determine the position and orientation of the underground vehicle with respect to a self-generated metric map of the underground mine. This approach contrasts existing systems that either require infrastructure or employ topological methods, which are considered by some in the robotics community to be non-robust when arbitrary tunnel geometries are possible.

During 2006 and 2007, extensive field trials were conducted using two different Atlas Copco LHDs at the Kvarntorp Mine, in Sweden. Performance results have been compared with previously recorded manual operator baseline times to establish an efficiency that effectively matches that of an experienced operator. The autotramming system was found to have remarkable reliability, based on a large number of repeated tramming operations. In summary, this paper describes what we believe to be the next-generation infrastructureless guidance system for underground mining vehicles.

1.1 Development Approach

The development approach used in this project can be divided in two main phases. In the first phase, a detailed model of the machine was designed in a simulation environment. This model was then extensively used for the development, initial testing and validation of the autotramming system. The intention was that the autotramming system's software should have settled to its final design upon completion of this phase and that only minor tuning of control parameters should be necessary when porting the system to a real vehicle. In the second phase, the autotramming system was integrated with the control system of the real vehicle, and the complete system was tested and verified in a real underground mine environment. This approach, which is a legacy of MDA's experience as a provider of equipment for space exploration, turned out to be very successful. The complete integration and acceptance tests were performed in only 5 weeks.

In this paper, we report on only the final design of the autotramming system and on the experiments and tests performed in the second phase of the project, i.e. the final tests and verification of the autotramming system on real LHD vehicles in an underground mine environment.

2 System Overview

At the user level, our autotramming system's operation consists of three steps: teaching, route profiling, and playback. During the teaching step, an operator drives the vehicle along a desired route, either on the vehicle, or by teleoperation. Simultaneously, sensor data is collected and stored in a binary log file (front and rear SICK laser rangefinders, a hinge angle encoder, and a drive shaft encoder for measuring displacement). During the route-profiling step, data logged during the teaching step is processed (offline) and converted into a format suitable for use by the estimation and control algorithms during playback. The output is referred to as a *route profile*, which contains information about the travelled path, a sequence of overlapping metric maps along the path, a record of any pause points (e.g., for dumping/loading material), as well as a vehicle speed profile to be tracked during playback. Finally, in the playback step, the system autonomously plays back a route profile generated by the teaching and route profiling steps. During playback, navigation and guidance algorithms use data from the specified route profile to estimate longitudinal, lateral, heading, and vehicle speed errors at discrete instants. The control system then stabilizes these errors so that the vehicle follows the profiled path at the desired speed.

Once profiled, a route can be played back many times. It is expected that re-profiling would only be necessary if significant changes to the environment were made; e.g., due to significant mine development. Since the vehicle is driven along a collision-free path during teaching, complex path planning is not required. However, the system does include short-range guidance algorithms designed to stop the vehicle should the profiled path be subsequently obstructed during playback.

The following subsections provide further details regarding some of the system's key features.

2.1 Route Profiling

A route profile consists of four components: a path profile, a pause profile, a sequence of locally-consistent metric maps, and a speed profile. Firstly, a sequence of locations that are equally spaced (typically 0.5 m) along the path are created, called path points. We then associate with each path point the configuration of the vehicle at that point during the teaching step by interpolating the pre-processed logged data. Thus, the sequence of path points and associated data comprise the path profile.

Locally-consistent metric maps of the mine environment along the path profile are generated using both odometry and rangefinder data. Each map is an occupancy grid (Elfes and Moravec, 1985). For localization in underground mines, this approach is much more flexible than a system that must classify tunnel topology in that it will work regardless of the shape of the walls, so long as the maps are of sufficient resolution. However, the use of a single monolithic map to represent the mine environment suffers from two key problems. Firstly, in some situations high memory usage is required. Secondly, map inconsistencies can result on longer traverses when a vehicle closes a loop or crosses its own path. To address these problems, we employ a sequence, or atlas, of metric maps attached along the path to form an overall route profile; which is to say, the system does not rely on one monolithic map and an absolute frame of reference. The underlying idea is to create a situation in which the vehicle's path exists in a high-dimensional space wherein it never intersects itself (Howard, 2004). Figure 1 shows an example monolithic and the corresponding atlas maps generated from real data acquired during one of our tests in a real underground mine. The monolithic map has a resolution of 0.3 m, while the atlas maps each have a resolution of 0.1 m.

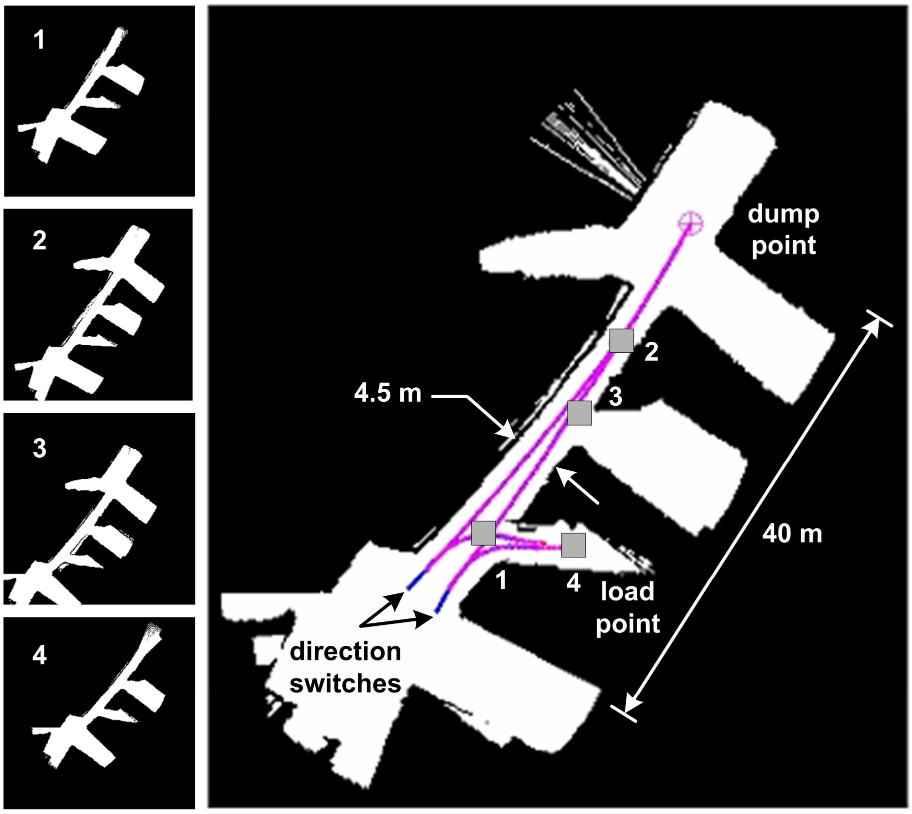

Figure 1 Example monolithic occupancy grid with atlas maps and vehicle path shown. The grey squares indicate the atlas map centres and the circle with cross hairs is a dump point.

2.1 Playback

Our objective was to create a system that permits a large articulated vehicle to robustly track the path specified by a route profile. During playback, this has been achieved through the design of a two-timescale control system. At the slower timescale, or outer loop, are localization and path-tracking algorithms that work to reject lateral and heading path errors. At the faster timescale, or inner loop, are rate estimators and two controllers that track reference steering rates and vehicle speeds. The underlying justification for this two-timescale design is founded on the assumption that we can specify sufficient bandwidth separation between the nested control loops. A schematic of the control architecture is shown in Figure 2.

At the outer loop, there exist two basic algorithms: localization and path tracking. The localization problem solved here is one of estimating the vehicle's pose as it travels through a (locally) known environment. Recently, a number of techniques have been developed in the mobile robotics community that globally localize a robot in a known environment. Many of these techniques use a particle filter representation of the vehicle's pose (Thrun et al., 2001). An initial design using a particle filter was shown to work in simulation, but required the use of too many particles (e.g., greater than 100) for convergence from a reasonable initial pose estimate. Variations requiring fewer particles and computational resources exist, but are complex to implement. Moreover, the task at hand does not actually require a solution to the global localization problem. Instead, we chose to implement a variation of the Unscented Kalman Filter (UKF) (Julier and Uhlmann, 1996; Wan and van der Merwe, 2000) to position with respect to the locally consistent submaps defined during route profiling. The inputs to the UKF algorithm are the laser rangefinder data as well as wheel odometry (i.e., hinge angle and wheel speed); the final outputs are heading and lateral errors of the vehicle with respect to the profiled path.

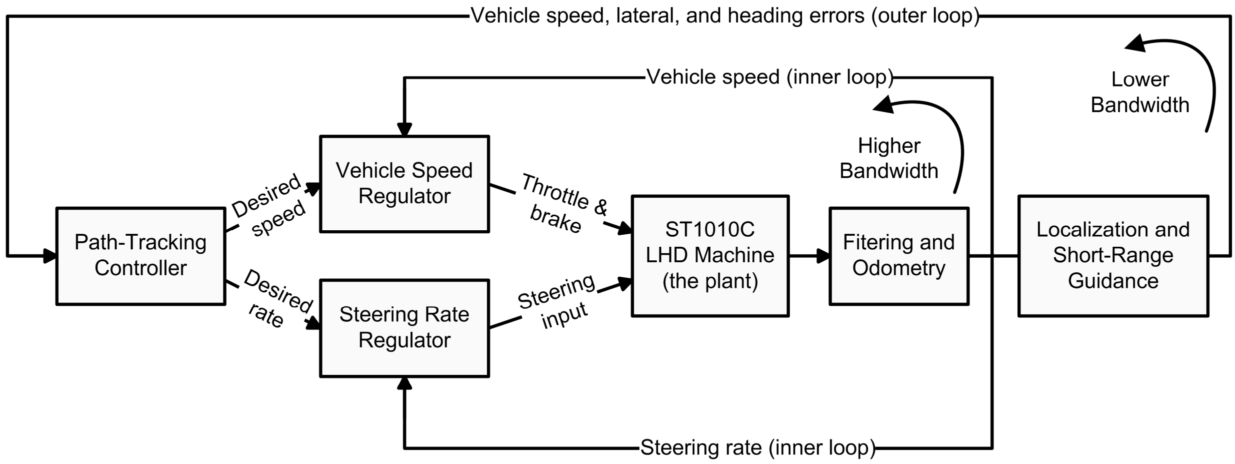

Figure 2 Two-timescale control system architecture.

Briefly, the unscented transformation works by parameterizing mean and covariance information in a way that allows for propagation through a nonlinearity. This is done by creating a discrete approximation that can be directly transformed, which has the same mean and covariance as the original Gaussian. This approximation takes the form of a set of $2n + 1$ so-called sigma points (where n is the dimension of the configuration space of the vehicle), which are symmetric and have the desired mean and covariance. In our case, $n = 3$ because we are interested in estimating the vehicle planar position and orientation.

A path-tracking controller is required to guide the vehicle along the path specified by the route profile. Path-tracking control for articulated vehicles has been extensively discussed in the engineering literature, yet there is some disagreement over the form such a controller should take. Some researchers argue that the position of both front and rear components of the vehicle should be tracked; others suggest that wheel slip should be explicitly accounted for (Ridley and Corke, 2001). We have found that, in practice, neither of these tasks is necessary under the two-timescale control architecture if inner-loop controllers are robust enough to handle these model uncertainties. Moreover, we have found that rejection of only the heading error and lateral error of the front component (when travelling forward) is necessary. We obtain these error signals from the UKF algorithm and then employ a nonlinear controller, based on feedback linearization, to drive them to zero, thereby forcing the vehicle to track the profiled path.

3 Integration and system verification tests

Our autotramming system design was initially integrated and tuned for verification on a 10-tonne capacity Atlas Copco ST1010C LHD. Initial tests of the autotramming system were carried out over a period of several weeks during March-April 2006 at the Kvarntorp Mine (an Atlas Copco testing facility created from a closed mining operation), in Sweden.

During 2007 the system was ported to, and tuned for, a newly released 14-tonne capacity ST14 LHD. This section also presents the result from experiments and tests performed on this platform, which focus on localization and path tracking during playback.

At the time of writing, a field test of the ST14 automation system, including autonomous tramming, is underway at a mine in northern Finland. Here, the machine is working in a backfilling operation. The initial results from these tests are very good, with flawless autotramming even though backfilling operations were not the intended task when the system was designed.

3.1 Atlas Copco Test Vehicles

The ST1010C is a one of a kind prototype machine equipped with Atlas Copco's standard CAN-based Rig Control System (RCS), Figure 3. The machine is based on an Atlas Copco Wagner ST1010 LHD which has an empty mass of 26.3 tonnes and takes a payload of 10 tonnes.

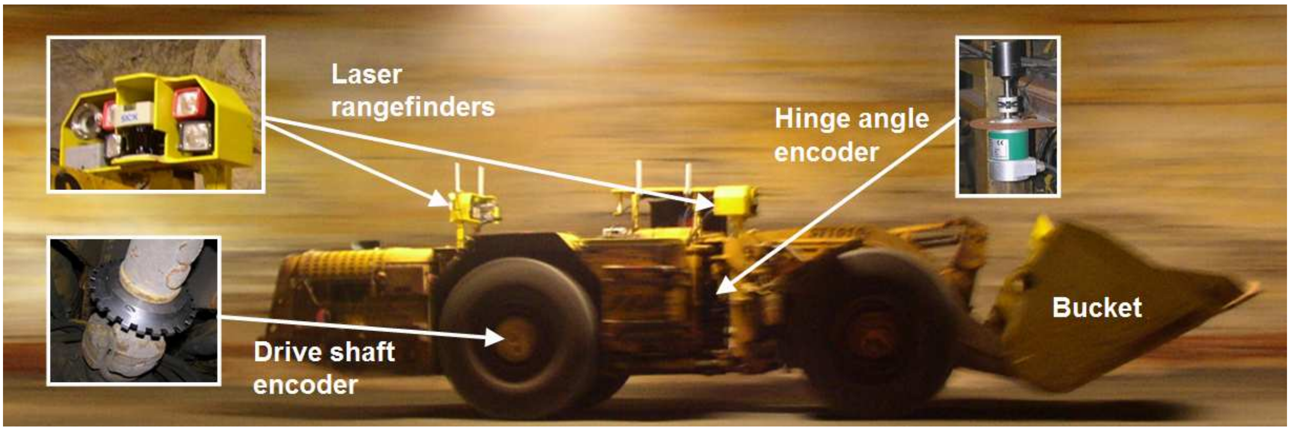

Figure 3 Atlas Copco ST1010C test vehicle, with sensor layout.

The standard ST1010 is equipped with direct hydraulics for control of steering, brakes and bucket movement, and analogue interfaces to the transmission and engine, e.g. accelerator pedal and gear selector switch. In the ST1010C, all vehicle functions are interfaced via a CAN-based computerized control system. Since this machine is of an old type and one of a kind, it is not optimal for research on future products. However, when the project started this was the only available Atlas Copco LHD that was already equipped with the RCS. To enable autonomous tramming, the machine was equipped with some additional sensors and an extra computational module. The main sensors added were: a drive shaft encoder to measure drive length; a hinge angle encoder; and two laser range finders, one for each direction of travel.

After the verification tests had been performed on the ST1010C, the newly developed Atlas Copco ST14 Scooptram was made available for automation, Figure 5. All ST14 LHDs come equipped with a CAN-based control system that also includes a hinge angle encoder as standard equipment. The sensor equipment used on the ST14 automation ready vehicle (ARV) is slightly different from the ST1010C in that the ST14 is also equipped with an Inertial Measurement Unit (IMU). Apart from this, the automation installations are the same with laser rangefinder sensors, a drive shaft encoder and an extra computational module.

The ST14 is not only of a newer design than the ST1010C, it is also a significantly larger machine. With an empty weight of 38 tonnes and the ability to carry a payload of 14 tonnes, the ST14 is nearly 50% heavier than the ST1010C. This is also reflected in its dimensions, as the ST14 is significantly higher, wider, and longer than the ST1010. Compared to the ST1010C, the ST14 also has a higher top speed in addition to an improved hydraulic system, where the later reflects in a more responsive steering.

Figure 4 Atlas Copco ST14 Scooptram test vehicle, the sensor layout is the same as for the ST1010C. The IMU is mounted in the housing of the front laser scanner.

It should be noted that both of the two test vehicles were used before they were engaged in the automation tests. The ST1010C had previously been used for more than 4000 hours in an extensive field test as the first LHD equipped with the Atlas Copco RCS system. Our ST14 was also somewhat worn, as it was the first prototype of the ST14 model and had been engaged in a field test for six months in a mine in northern Sweden. The fact that both vehicles were used is both a benefit and a drawback from an automation development point of view. It is a drawback since a worn vehicle can be more difficult to control. A worn vehicle is also more likely to suffer from breakdowns such as leaking hoses or engine and transmission failures. The latter did not directly affect our evaluation of the autotramming system, but failures delayed the overall testing and added uncertainty to the planning. On the other hand, worn vehicles can also be seen as a benefit when evaluating a new automation system. Eventually, any automation system employed on heavy machinery will have to deal with worn mechanics and hydraulics. In our case, since all tests have been performed on used vehicles the autotramming system has already proven to be tolerant to some amount of degradation of the vehicles as compared with brand new machines.

3.1 ST1010C Automation Testing

In this paper we report on 36 of the trial runs that were performed under strict observations as a part of the acceptance test of the delivery of the autotramming system from MDA to Atlas Copco. The main requirements of the acceptance tests were two: firstly, the autotramming system should not perform significantly slower than a manual operator; secondly, the reliability should be equal to or better than for a manual operator. The second requirement was quantified as a navigation precision requirement of ±0.5 m at the end of a route, and that the vehicle was not allowed to hit the walls of the tunnels during any of the tests.

For the purpose of evaluating the autotramming system, a set of test routes were designed to cover different aspects and situations that are common in the operation of LHDs. This includes tramming in wide and narrow drifts with and without intersections (T1 and T2), as well as curved drifts with both perpendicular and long sweeping curves (T3). Also paths including dumping at a dump point (T4 and T5), and direction switches (T5) were evaluated. Figure 5 display the monolithic maps for the different routes.

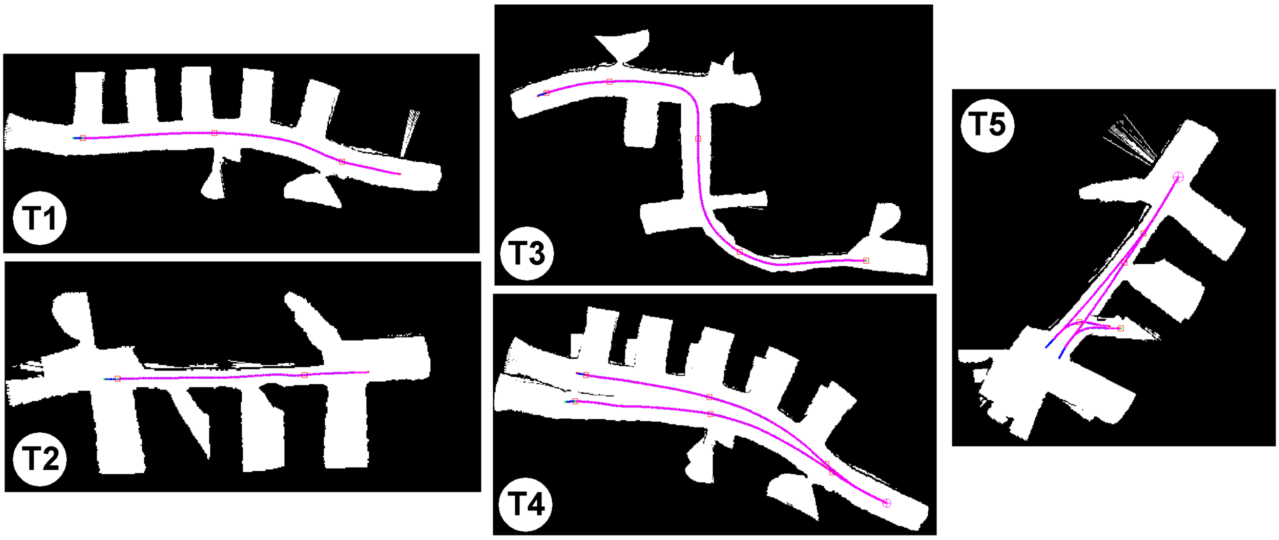

Figure 5 Monolithic maps displaying the routes used in the verification of the autotramming system. The drift widths in the wider parts of the routes are nominally 11 m, while they are 4.5 m wide in the narrow parts.

Manual baseline times were obtained by timing a human operator driving the same route as the autonomous tramming system. The navigation precision of the tramming system was measured by comparing the position of the vehicle after stopping at the endpoint of a route, to markings on the ground representing the location where the vehicle was stopped during teaching, see Section 2. Several of the test scenarios were conducted both with, denoted ‘L’, and without, denoted ‘E’ load in the bucket.

Table 1 displays the average times for each test scenario for both the human operator as well as for the autotramming system. From these results, the autotramming system was calculated to have an efficiency of approximately 0.97. During the tests of the ‘T4E’ scenario the vehicle consistently refused to switch to fourth gear due problems with the transmission, both in manual and autotramming mode. The timing of these runs is therefore not representative and the ‘T4E’ scenario was excluded from efficiency calculation, even though displayed in Table 1. Notable is also that the entire area in which the tests were performed is illuminated. This is of great help for the manual operator in achieving precision and speed, while it has no effect on the autotramming system.

To measure the navigation accuracy and repeatability of the autotramming system, the final position at the end of each run was directly compared to markings on the ground at the endpoint of each manually taught route; see Figure 6. The results of these measurements are presented in Table 2.

Table 1 Average times (in seconds)

Test Scenario	T1	T2	T3	T4E	T4L	T5E	T5L	Total
Human Operator	41.8	27.0	57.0	85.0	90.0	81.0	83.0	464.8
Autotramming	41.3	25.8	59.5	97.3	89.8	81.8	85.0	480.5

As can be seen in Table 2 the navigation errors at both dump points and end points of the routes are consistent within a given test scenario, and that the precision requirements are fulfilled with large margin in most cases. It should also be noted that no route specific tuning was performed to achieve this precision. By applying route specific tuning, the navigation errors could for sure be reduced even further.

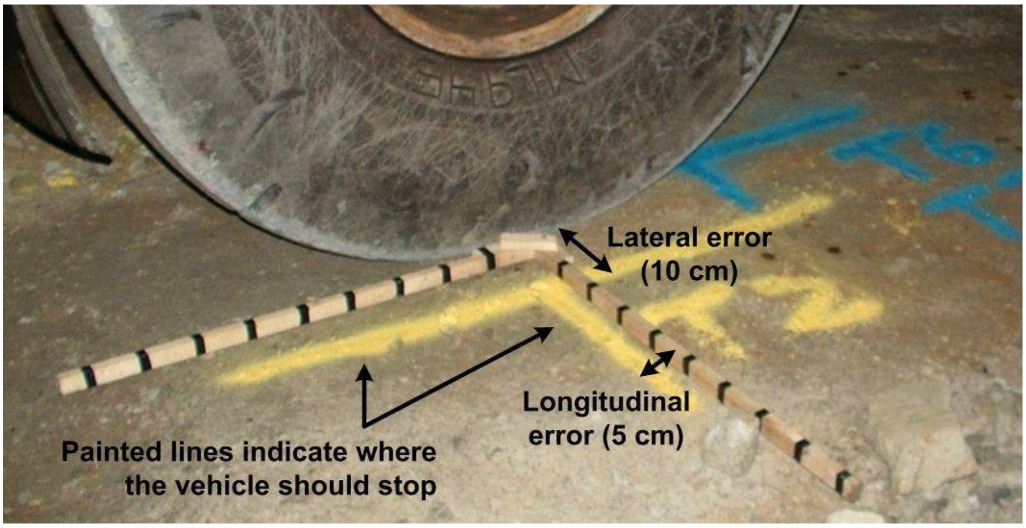

Figure 6 Example positioning measurement from a test run, where the orthogonal lines marked 'T2' show the manually taught position of the wheel's outer edge.

Table 2 Longitudinal and lateral positioning errors (cm) covering all measured stops for all relevant underground test runs.

Error / Test	T1	T2	T3	T4E	T4L	T5E	T5L	T7
Longitudinal	0	5	-10	-22	-20	-20	-25	0
	0	0	-10	-20	-30	-25	-25	5
	0	5	-10	-25	-15	-20	-25	5
	-5	0	-10	-30	-20	-25	-20	5
Lateral	45	10	0	-15	-10	20	-10	15
	40	10	0	-15	-15	25	-15	5
	40	10	0	-15	-10	20	-10	5
	45	15	0	-10	-10	20	-10	0

3.2 ST14 Automation Testing and Commercialisation

Following successful verification on the ST1010C, the autotramming system was ported to the ST14 vehicle and its control system during 2007. In this step, the autotramming system was commercialised along with a new tele-operation system. As a complement to the autotramming system, the automation system on the ST14 also features automatic dumping functionality. This however, is outside the scope of this paper and will therefore not be further described.

During integration and initial tuning of the autotramming system on the ST14, it was noted that the expected performance improvement of the inner loop, due to better steering response of the ST14 compared to the ST1010C, did not occur. All attempts to increase the bandwidth of the steering rate controller compared to the ST1010C resulted in oscillations in the steering, but this was not seen as a large problem since the ST14 outperformed the ST1010C in the most important test scenarios T4 and T5. However, as new test cases were run and the focus of the tests turned towards the overall path tracking performance instead of merely the easily measured precision at the endpoint of the routes, it appeared as though the autotramming system on the ST14 had a major problem. One of the new test cases consisted of a loop where the machine travelled forward in a nearly rectangular path. When playing back this route at high speed (> 4 m/s) it was found that the tracking algorithm not only cut the corners of the path, but also severely failed to track the path on the straight sections in-between the corners; see Figure 7. In large parts of the path the lateral path tracking error exceeded 1 m with a maximum of 1.39 m, something that is completely unacceptable.

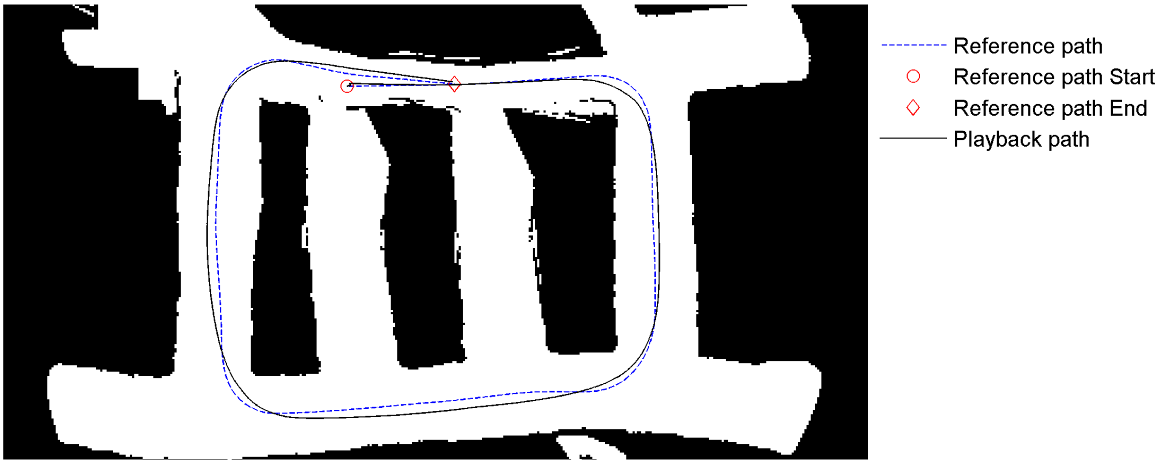

Figure 7 **Monolithic map with estimated path superimposed on the reference path. Nominal drift width is 11m. The path tracking error exceeds 1m in large parts of the path.**

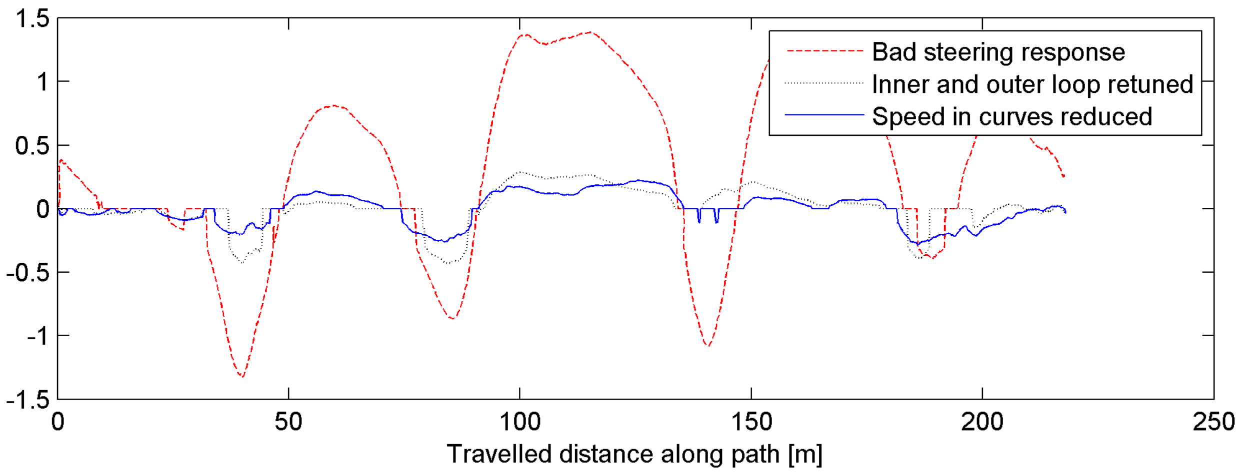

Figure 8 **Estimated lateral path tracking errors for three different runs of the route displayed in Figure 7**

Simultaneously, it was discovered that some small but highly significant changes to the hydraulic system had been made since the initial modelling of the ST14; see Section 1.1. After remodelling the hydraulic system according to the new specifications and performance of the steering, new simulations were made. From the simulations, a slightly changed steering controller design was derived and new controller parameters were evaluated. The results from the simulations were encouraging and, during the following re-tuning of the inner and outer loop on the machine, significant improvements in performance were noticed. When playing back the same route again the path tracking errors were greatly reduced to a maximum value of 0.44 m. Finally, the speed profile of the route was slightly adjusted to avoid saturation of the steering in the corners. This reduced the path tracking error even further to a maximum of 0.28 m, while still reaching a speed of close to 5 m/s on the straight parts of the route. The speed reduction in the corners resulted in a modest increase of the total time to playback the path from 69 s to 74 s. Figure 8 displays the path tracking errors for all three of these described runs.

4 Conclusions

By applying map-based localisation using an atlas of metric maps together with a suitable control architecture, our autotramming system was designed with the aim of matching the performance of a manual operator. The results presented in this paper are encouraging, and indicate that the system fulfils the performance and robustness requirements of an autonomous tramming system for underground mines.

The system has been tested on two different Atlas Copco LHD vehicle models in underground mine environments for more than 18 months (at the time of writing), and is currently being field tested in a real operating mine. Furthermore, the porting of the system from our original prototype ST1010C demonstrator to the new ST14 commercial vehicle was made easy by the modular architecture of the autotramming system's design. This is important because the system is likely to be ported to many of Atlas Copco's new and upcoming LHD and mine truck models.

Nevertheless, there remain both challenges and opportunities. One of the major challenges, from a commercialization point of view, is to develop tools and interfaces that allow trained LHD operators to not only use, but to also administrate the autotramming system; including to record, generate and validate new routes. From a technical point of view, the development of an auto loading system is a great opportunity, as this would allow for fully autonomous operation of the complete Load – Haul – Dump cycle.

5 Acknowledgements

Many people were involved in the autonomous guidance project described in this paper. At the Atlas Copco Rock Excavation Technology Center in Örebro, Sweden, we would like to acknowledge in particular Tobias Furuholm, Richard Hendeberg, Michael Krasser, Olav Kvist, Fredrik Nilsson, Lars Sandstrom, Casper Swart, and Oscar Tryggvesson. At MDA in Brampton, Canada, we would like to acknowledge in particular Joseph Bakambu, Leif Bloomquist, Rob Corcoran, Roy Jakola, Raja Mukherji, Andrew Ogilvie, Dave Parry and Robert Ward.

6 References

A. Elfes and H. Moravec (1985) 'High resolution maps from wide angle sonar', Proceedings of the IEEE International Conference on Robotics and Automation, St. Louis, MO, 116-121.

A. Howard (2004) 'Multi-robot mapping using manifold representations', Proceedings of the IEEE International Conference on Robotics and Automation, New Orleans, LA, 4198-4203.

S. Julier and J. Uhlmann (1996) 'A general method for approximating nonlinear transformations of probability distributions', Technical report, Department of Engineering Science, University of Oxford, Oxford, UK.

J. A. Marshall and T. D. Barfoot (2007) 'Design and field testing of an autonomous underground tramming system', Proceedings of the 6th International Conference on Field and Service Robotics, Chamonix, France, 393-402.

P. Ridley and P. Corke (2001) 'Autonomous control of an underground mining vehicle', Proceedings of the 2001 Australian Conference on Robotics and Automation, Sydney, Australia, 26-31.

S. Thrun, D. Fox, W. Burgard, and F. Dellaert (2001) 'Robust Monte Carlo localization for mobile robots', Artificial Intelligence, Volume 128, 99-141.

E. A. Wan and R. van der Merwe (2000) 'The unscented Kalman filter for nonlinear estimation', Proceedings of the IEEE AS-SPCC, Lake Louise, AB.