2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

An Evaluation of Local Autonomy Applied to Teleoperated Vehicles in Underground Mines

Johan Larsson, Mathias Broxvall and Alessandro Saffiotti

Abstract— **Autonomous vehicles are being increasingly introduced in the mining industry. While these may offer high safety and high productivity, fully autonomous solutions are not always applicable or economically viable. Teleoperation is an attractive option, since it increases safety and comfort of the drivers. Unfortunately, the difficulty to operate the vehicle remotely often results in reduced productivity. In this paper, we show that techniques from the field of mobile robotics can be used to mitigate this problem. We extend a commercial teleoperation system for use in underground mines with a local autonomy functionality, with the main purpose to evaluate if the achieved productivity improvement motivates development of general algorithms and a fully commercial implementation. We then describe a user study performed in an underground mine with a 38 tonne articulated wheel loader, which proves that local autonomy gives a significant improvement in productivity of the teleoperation system, while retaining or even reducing the maintenance costs.**

I. Introduction

In underground mines, LHD (Load-Haul-Dump) vehicles are used to transport ore and waste rock from one location to another. LHDs are typically manually controlled by an on board operator, a job that can be characterized as triple-D: Dirty, Dull and Dangerous. Due to this, it is desirable to automate the operation of the LHDs, and several commercial solutions for teleoperation and fully autonomous tramming of LHD vehicles are commercially available [1], [2], [3].

In many applications fully autonomous tramming is the best solution. However, there are situations where a human-in-the-loop solution is preferred. One of the most important reasons is the setup time for the autonomous tramming systems. Before the autonomous system can be run in production, the routes between the load and dump points need to be defined. This is often done by teaching a path which is then processed off-line and evaluated before it can be run in autonomous mode. Therefore the currently available systems for autonomous operation of LHDs are not cost-effective in dynamic environments, where the time to set up the system for autonomous operation is a significant part of the total operation time. These cases include for instance rapid progress room and pillar mining and stope backfilling where the operation time can be in the order of a few days. In these cases the total cost is less with a remotely operated, human-in-the-loop system, than with an autonomous one, even though the productivity during operation is lower.

Techniques for effective teleoperation have been extensively studied in the field of mobile robotics [4], [5], [6], [7], [8]. It is therefore natural to ask if these techniques can be used to improve the performance of teleoperated LHDs. For example, one would expect that the use of local autonomy [7], [8], which effectively decouples the remote operator from the control loop, has the potential to improve performance while reducing the cognitive load of the operator. Results from CSIRO [9] also indicate that a vehicle in local autonomy mode would be able to drive faster than the same vehicle in normal teleoperation. While these expectations are reasonable, there are no systematic studies that validate them.

The goal of this paper is to provide such a systematic study. More precisely, we set up an experiment in which ad-hoc implementations of two augmented teleoperation techniques are compared with an existing commercial teleoperation system. The comparison is based on a set of runs performed in a test mine by experienced operators. Our results show that both techniques significantly improve the productivity of the teleoperated LHD while reducing the wear of the machine compared to normal teleoperation. These results are especially important from a commercial point of view since they motivate the investment in time and money needed to develop a full scale implementation.

The rest of this paper is organized as follows. In the next section we discuss current solutions for LHD teleoperation, and suggest techniques from the domain of mobile robotics that hold promise to improve these solutions. In section III we describe our experimental methodology for evaluating the impact of these techniques. Sections IV and V describe the system built to perform our experiment and the results of the experiments, respectively. Finally, section VI concludes.

II. Teleoperation of Mining Vehicles

Commercial systems for teleoperation of LHDs in underground mines are available from several vendors. The work described here is based on a recent system from Atlas Copco, but the considerations below apply to most existing systems.

A. A Commercial System

Our system consists of an ST14 LHD vehicle prepared with additional sensors and communication capabilities (Figure 1) and an operator station (OPS, Figure 2). The additional sensors consist of an odometer measuring the rotation of the transmission drive shaft, an articulation angle sensor, and

This work was founded partly by the Swedish KK Foundation, and partly by Atlas Copco Rock Drills, Sweden.

Johan Larsson is with AASS, Örebro University, Örebro, Sweden and Atlas Copco Rock Drills, Örebro, Sweden johan.larsson@oru.se

Mathias Broxvall and Alessandro Saffiotti are with AASS, Örebro University, Örebro, Sweden {mbl,asaffio}@aass.oru.se

978-1-4244-5040-4/10/$26.00 ©2010 IEEE

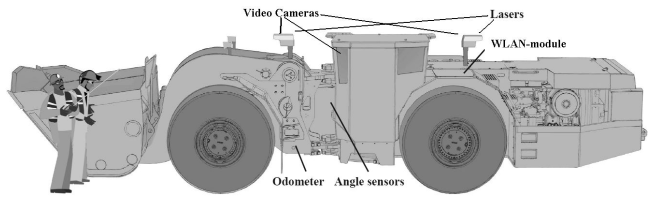

Fig. 1. ST14 LHD vehicle prepared for teleoperation. Two operators are shown on the left to indicate the scale.

Fig. 2. Operator's station, console and display.

three video cameras – two facing forward and one backwards. As an option to improve operator's awareness, the teleoperation system can be augmented with laser scanners providing information about the tunnel walls in the vicinity of the machine measured in a horizontal plane.

To enable communication between the LHD vehicle and the operator station, the teleoperation package includes a proprietary WLAN module enabling fast and reliable roaming between standard access points. This module is also responsible for streaming the video from the cameras.

The operator station consists of a display and a console. The display show the video stream from the selected camera, and has areas where system state information, warnings and errors are displayed. The left half of the screen is reserved for the optional laser range scanners. Here data from the front and rear laser scanners and from the articulation angle sensor are fused and displayed, creating a local track-up representation of the perceived environment. The operator console is equipped with two 2-axis proportional joysticks for controlling vehicle motion (left joystick) and the bucket (right joystick). It also has several push buttons for event based commands, e.g., to start the engine.

B. Problems with the Current Solutions

Teleoperation of a LHD has its drawbacks. Because of the difficulty to control the vehicle in teleoperation, the maximum allowed gear of the automatic transmission is set to second gear. This limits the speed to 10 km/h,[1] which leads to low tramming speed and hence to reduced productivity. Even at this speed, collisions are frequent in teleoperation, which lead to increased wear and tear of the vehicles and hence to increased maintenance costs.

[1]Typical tramming speeds for an LHD with an on-board operator are of the order of 15–20 km/h.

The difficulty experienced by an operator to control the vehicle mainly originates from the latencies in the control and feedback signals. Other negative factors include the limited awareness of the environment and machine state provided by the teleoperation interface, and the steering principle used in the interface.

The last point may require a clarification. On most vehicle types, steering is controlled through a steering wheel whose position corresponds to a specific turn radius. By contrast, and because of practical and historical reasons, on a typical underground mine LHD a 1-axis joystick is used as steering input device. The location of the joystick corresponds to the rate at which the articulation angle, and thus the turn radius, is changed. This means that the steering transfer function is not linear, but it includes an integrator. This makes it difficult for an operator to find the articulation angle at which the machine moves along a desired line, especially in teleoperation where the delay in control and feedback signals often result in wobbly steering even at rather low speed.

Due to the above drawbacks, there is a strong interest in developing functionalities that enable higher tramming speed of teleoperated LHDs, and preferably also reduce the wear and tear on the machine.

C. The Heritage of Mobile Robotics

In the literature of mobile robotics there has been extensive work on techniques for effective teleoperation, which might provide ways to improve teleoperation of LHD vehicles. Most of the proposed techniques fall within four categories: network control [4], [10]; graphical interface design [11], [5]; haptic feedback [12], [13]; and local autonomy [8], [14].

Network control approaches typically try to overcome network latency by presenting to the operator a predicted state of the robot at the remote location, based on a model of the robot and the control inputs from the operator. Unfortunately, prediction requires precise dynamic models of the robot, something that is difficult to achieve with a large, hydraulically controlled, articulated vehicle traveling in rough environments such as a mine. These approaches are therefore hard to apply in our domain.

Approaches that focus on graphical interface design and haptic feedback are also of limited interest in our domain, since they do not directly address the problems caused by network latency. Moreover they often require modifications to the existing hardware, which is undesirable for commercial reasons.

In this work we concentrate on the concept of local autonomy, which holds the greatest promise to improve performance in our domain while being commercially acceptable. This concept spans over a wide field of implementations, from direct control with collision avoidance [13] to heading control with obstacle avoidance [14], to supervisory control where the operator issues high level commands that the robot autonomously executes while the operator monitors the progress [8], [15]. The inclusion of local autonomy in teleoperation has the potential to increase the tramming speed and thus the productivity by reducing the problem of latency in the control streams, since the remote operator is decoupled from the control loop. It may also compensate for limited telepresence, since low level decisions can be handled on-board the robot, and lead to less collisions and hence to reduced maintenance costs. Finally, local autonomy may reduce the cognitive load of the operators since they only have to make low bandwidth high level decisions.

III. Evaluation Methodology

The above discussion suggests that the addition of local autonomy functionalities can improve teleoperation of LHD. More precisely, we form the hypotheses that these functionalities: (1) improve productivity by enabling higher tramming speed; (2) decrease maintenance costs by reducing the wear and tear of the machine; and (3) reduce the cognitive load of the operator and thus increase the productivity over a full eight hour shift.

We now detail our methodology to empirically validate these hypotheses. We first discuss the selected functionalities, then we go through the design of the experiment to evaluate them, and finally we present our performance metrics.

A. Functionalities

We decided to evaluate the effects of the following functionalities.

1) Augmented steering: As discussed in Section II-B, one of the sources of difficulty in teleoperating a LHD is how the steering works. To cope with this, we introduce and test a new functionality in this work, which we call *Augmented Steering*: an intermediate automation level between the direct control of ordinary teleoperation and the autonomous functionalities described above.

With augmented steering the operator controls the articulation angle, and thus the turn radius, instead of the rate of change of the articulation angle. This could potentially enable higher tramming speed in teleoperation as the loop for controlling the articulation angle is closed locally on the machine instead of via the remote teleoperation station. A benefit with augmented steering compared to the autonomous functionalities is that no exteroceptive sensors are needed. It can therefore be implemented on all teleoperated LHDs, without the need for the optional and relatively expensive laser range scanners.

2) Local autonomy: When it comes to local autonomy in mines, the set of possible functionalities is rather obvious. Given the structure of the majority of the mines, there are basically three situations that the machine can encounter during tramming: to follow a tunnel, to drive straight through an intersection, or to turn in an intersection.

Most tunnel following behaviors operate in one of two modes: relative to one of the walls, or relative to a virtual "lane" in the middle of the tunnel. In the case of underground mines, human LHD operators typically adopt an intermediate solution. Since many tunnels have ditches on one side, operators tend to drive near the wall opposite to the ditch in order to keep safely away from the ditch – whose edge is often difficult to see. However, operators do not strictly follow the surface of the wall, which is often interrupted by bays for infrastructures, intersections with other tunnels, or simply irregularities produced when building the tunnel. Rather, they follow the average direction of the wall.

Correspondingly, a local autonomy functionality for tunnel following should implement a suitable combination of wall following and virtual lane following. This functionality should maintain the vehicle at a predefined minimum distance to one wall in order to avoid threats that cannot be perceived by the on-board sensors, like a ditch, while maintaining a safe distance to the other wall. The resulting path should be as straight as possible within these bounds to allow high speed tramming.

Tunnel following should control both velocity and steering. However, it is desirable to give the operator the possibility to limit the maximum speed at any moment.

Driving straight at an intersection is directly captured by the previous functionality. When it comes to autonomous turning, the functional requirements are similar to the ones for tunnel following, with the additional one that the system should be able to detect intersecting tunnels and guide the machine into the new tunnel while maintaining the above safe distance constraints. Like for tunnel following, the operator should have the possibility to limit the maximum speed.

B. Experiment Design

The empirical way to evaluate the gain achieved by the above functionalities is to implement them and run a comparative case study. It should be noted, however, that our evaluation is aimed to assess the potential benefit of these functionalities *before* we decide to allocate the effort needed for their full scale implementation. Therefore, the evaluation must be based on a simplified, easy to implement version of these functionalities.

Our strategy has been to first decide a specific, fixed test environment, and then create an *ad-hoc* implementation of the target functionalities tailored to work on this environment. This allowed us to make a number of assumptions that greatly simplified the implementation. To make the results meaningful, the assumptions and the implementation should be such that: (1) the operator's experience of the system in this specific test is the same as if he had been using a general, full fledged implementation; and (2) the functionalities can in principle be implemented with the assumptions relaxed. This approach is somehow inspired by the "Wizard of Oz" methodology, extensively used for evaluations in the field of human-computer interaction [16].

Once we have decided a test environment and implemented the target functionalities, we run a user study in which several subjects run both the original teleoperated system, and the system augmented with the target functionalities, in the three standard driving situations for our domain: driving in a (nearly) straight tunnel, driving straight through an intersection, and turning at an intersection. We had initially intended to evaluate all possible combinations of functionalities and situations: e.g., manual operation inside a tunnel and autonomous turning at an intersection, versus autonomous tunnel following and manual turning. However, it soon became clear that mixing different modes would be confusing for the operator, so we restricted our experiment to use one type of functionality for each run.

C. Performance Metrics

Two independent variables were introduced: *Mode* (level of autonomy) and *Driving Situation*. The first has three levels: *Manual* (**Man**), *Augmented Steering* (**AS**), and *Local Autonomy* (**LA**). The second is either *Tunnel* (following the current tunnel) or *Turn* (at an intersection).

By measuring the time it takes for each subject to complete a test route, the augmented steering and local autonomy modes can be compared to the manual baseline to evaluate if higher tramming speed is achieved. The impact of the evaluated mode on the two driving situations is differentiated by separating the test path into *Tunnel* and *Turn* sections, and measuring the time to complete each section individually. For the evaluation, the dependent variables t_{Tunnel}, t_{Turn} and *wearEvents* were introduced. The first two variables correspond to the total time of each run spent driving along the straight parts and turning in the intersections. *wearEvents* is the number of wear-related events during the runs.

Data analysis is performed using the well known statistical tool ANOVA [17]. Two null hypotheses are formed for the first question mentioned at the beginning of this section: if the new functionality enables higher tramming speed.

(1) $H_{0,\mathrm{tunnel}}$: All t_{Tunnel} time samples are drawn from the same population.

(2) $H_{0,\mathrm{turn}}$: All t_{Turn} time samples are drawn from the same population.

A third null hypothesis is formed for the second question: if the new functionality reduces the wear.

(3) $H_{0,\mathrm{wear}}$: All *wearEvents* samples are drawn from the same population.

In addition to the above quantitative evaluation, we perform a *qualitative* evaluation by asking the subjects to fill up questionnaires after driving each mode (Figure 3). These aim at evaluating the subjective experience of the subjects, thus providing an indirect answer to the third question above: if the new functionality reduces cognitive load.

IV. Experimental Setup

We now describe the setup that we created for our evaluation study. We first describe the physical environment, and then the software system that implements the target functionalities in this specific environment. Changes to the OPS were kept as small as possible, complying with the requirement of no changes at all in the hardware.

A. Test Environment

The test area, shown in Figure 4, was a section of an abandoned mine that is used by Atlas Copco for development of LHD automation solutions. This test facility includes a complete teleoperation system as described in Section II.

Questionnaire for one test mode

With respect to the how the machine behaved during the test mode, please indicate how much you agre or disagree with the following opinions. Use the following scale:

Disagree completely	Disagree	Undecided	Agree	Agree completely
0	**1**	**2**	**3**	**4**

01	I feel that I safely can drive the machine at the speed that I achieved in this test while tramming in a tunnel	0	1	2	3	4
02	I feel that I safely can drive the machine at the speed that I achieved in this test while turning in an intersection	0	1	2	3	4
03	I have to be very focused to drive the machine at the speed that I achieved in this test while tramming in a tunnel	0	1	2	3	4
04	I have to be very focused to drive the machine at the speed that I achieved in this test while turning in an intersection	0	1	2	3	4
05	I believe that I would be able to drive the machine at the speed I achieved in this test for a full shift	0	1	2	3	4
06	I prefer the tunnel following mode used in this test case as compared to the standard tele-operation driving mode	0	1	2	3	4
07	I prefer the intersection handling mode used in this test case as compared to the standard tele-operation driving mode	0	1	2	3	4
08	I believe that I would be able to drive the machine at the speed that I achieved in this test for one hour straight	0	1	2	3	4
09	I believe that I would be able to drive the machine at the speed that I achieved in this test for two hours straight	0	1	2	3	4
10	I believe that I would be able to drive the machine at the speed that I achieved in this test for four hours straight	0	1	2	3	4
11	On a scale from 1 (I feel very relaxed) to 10 (I am completely focused) assess the level of attention needed to be able to drive the machine at the speed you achieved in the tunnel parts of the course					
12	On a scale from 1 to 10 assess the level of attention needed to be able to drive the machine at the speed you achieved in the intersection parts of the course					
13	I don't feel that I safely can drive the machine at the speed that I achieved in this test while tramming in a tunnel	0	1	2	3	4
14	I feel that I safely can drive the machine at the speed that I achieved in this test while turning in an intersection	0	1	2	3	4
15	I feel relaxed when driving the machine at the speed that I achieved in this test while tramming in a tunnel	0	1	2	3	4
16	I feel relaxed when driving the machine at the speed that I achieved in this test while turning in an intersection	0	1	2	3	4
17	I don't think that I would be able to drive the machine at the speed I achieved in this test for a full shift	0	1	2	3	4
18	I think that standard tele-operation is better suited for driving the machine in a tunnel than the mode used in this test.	0	1	2	3	4
19	I think that standard tele-operation is better suited for turning in an intersection than the mode used in this test.	0	1	2	3	4
20	I don't think that I would be able to drive the machine at the speed that I achieved in this test for one hour straight	0	1	2	3	4
21	I don't think that I would be able to drive the machine at the speed that I achieved in this test for two hours straight	0	1	2	3	4
22	I don't think that I would be able to drive the machine at the speed that I achieved in this test for four hours straight	0	1	2	3	4

Fig. 3. The questionnaire used to capture the subjects opinions after driving each mode.

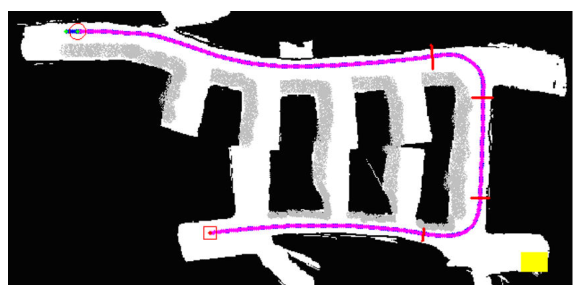

Fig. 4. Map of the test area including test path. Start point is marked with a circle and end point with a square, section limits are marked with red lines perpendicular to the path. The simulated ditches are marked in gray. The yellow rectangle is the operator's station cabin.

The test route was designed to get the longest possible *Tunnel* sections in the available test area and was approximately 2×225 m long, covering three *Tunnel* and two *Turn* sections that were driven in both directions. To reduce the test time the scenario was limited to the tramming part of the ordinary work cycle of a LHD, i.e., no loading or dumping was performed.

In several sections we simulated the presence of a wide ditch on one side, that requires the machine to travel close to the opposite wall. This also compensated for the fact that the drift width in our mine is larger than in most other mines. The simulated ditches were marked by a white line and by cones paced at about 10 m, see Figure 5. The cones and lines provided the operator with a visual perception of the location

Fig. 5. A tunnel with a painted white line indicating a simulated ditch.

Fig. 6. OPS display with local autonomy functionality active. The yellow shape with a black rectangle on the left represents the machine and the bucket. The dark gray areas are the free areas sensed by the laser scanners. The calculated path is shown by the thin line above the machine, and the detected side tunnel by a black line with x at its ends.

of the ditch at least as good as for a real ditch, since a real ditch edge would appear just as a shaded part of the roadbed.

B. Manual Mode

The original manual mode for teleoperating the vehicle has not been modified, but the speed limitation of 10 km/h has been removed by enabling the third and fourth gear to be used. The system thus allows the operator to reach the maximum speeds of 16 km/h and 25 km/h, respectively. This has been done to allow a fair comparison with the local autonomy modes, which do not have the 10 km/h restriction.

C. Augmented Steering

The first new mode is augmented steering. This mode changes the behavior of the left joystick in the OPS to control the desired articulation angle instead of the rate of change of the articulation angle. This was a straightforward implementation with an on-board PI-controller for controlling the hinge angle receiving reference values from the OPS steering joystick and the on-board articulation angle sensor. The steering of the machine was approximated to a FOLIPD system (First order lag plus integral plus time delay), and the model parameters were identified from recorded data. From the model parameters the controller parameters were calculated according to standard PI-tuning methods [18].

D. Local Autonomy

The second new mode enables local autonomy: the operator gives only high level commands for controlling maximum speed and the desired offset from tunnel walls, and for choosing which intersections to turn in.

1) Operator's Station: The changes to the OPS are limited to the laser view. Here the calculated path and detected side tunnels where the vehicle can turn are displayed superimposed to the laser data, as shown in Figure 6. In this mode, the right joystick is used to select offset, i.e., to keep to the center, left side or right side of the tunnel. The left joystick, which normally controls the steering, throttle and brake, is used to indicate to the intersection handling when to turn, and to provide the system with maximum allowed throttle and minimum allowed brake references. When the vehicle approaches an intersection, the operator can instruct it to turn by moving the joystick to the left or right.

2) Vehicle Controller: According to the experimental methodology discussed above, the implementation of this mode is not meant to be general, but it has the sole purpose to allow an evaluation that will motivate the development of a more general system. Therefore, we have relied on several characteristics that are true in our specific test environment and that allow a simpler implementation.

First, we reuse an existing path tracking system used in Atlas Copco's autonomous tramming system [19] to control the LHD. All that our system has to do, then, is to dynamically generate a suitable path for the vehicles to follow. Second, since in the test environment all the tunnels are relatively straight, we only consider slightly curved tunnels in our implementation of tunnel following. Finally, we assume that pre-recorded local paths are available for each turn at each intersection.

For the path generation part, we consider three cases: (1) following a tunnel or traversing an intersection without turning; (2) approaching an intersection with the intention of turning; and (3) turning at an intersection.

For case (1) we extract the wall positions from laser data, and dynamically generate a path from these while taking into account variable offsets. Lasers are used because vision is not suitable in underground mines [20] and because they are available as an option to the Atlas Copco's teleoperation system. Case (3) is handled by retrieving the relevant pre-recorded path for the specific intersection, since this was the easiest solution for our test. For case (2) we dynamically generate a path as for case (1), but with the added constraint that the paths ends at a point that is the start of the pre-recorded path for the intersection. This provides a smooth transition from tunnel following to the intersection handling.

Figure 7 shows the pseudo-code of the function implementing the top level control loop of this system. This function, running at 25 Hz, relies on a pre-existing laser-based global localization system, which is part of the

```
1  Inputs: offset ∈ {Center,L,R}, turn ∈ {Forward,L,R}
2  Outputs: path, refSpeed
3  Data: paths {P_1, ..., P_n}
4  laserdata ←range scans from the two lasers
5  pos ←updateLocalization(laserdata)
6  if turn = Forward  then
7     wayPoint ←projectWall(laserdata, offset)
8     path ←generatePath(laserdata, wayPoint, offset)
9  else
10    find P_i mininimizing : |pos − first(P_i)|
11    P = transform(pos,P_i)
12    if first(P)_x > 0  then
13       path ←generatePath(laserdata, first(P), offset)
14    else
15       path ←{p ∈ P|p_x > 0}
16    end if
17 end if
18 refSpeed ←safePathSpeed(path)
```

Fig. 7. Top level loop for generation of path and reference speed for low-level controller.

used path tracking system, and a set of pre-recorded paths for implementing turns in different intersections. It uses a simple path generation function `generatePath` based on traversable areas extracted from laser data for tunnel following (line 8) and preparing to turn in intersections (line 13). It expects a set of laser range readings covering 180 degrees in the driving direction of the vehicle, and a `waypoint` constituting the end point of the generated path. For the case of pure tunnel following, the waypoint is calculated by `projectWall` (line 7) which computes a point in front of the vehicle with a given offset orthogonal to the walls indicated by the laser data. For the case of turning at an intersection, the waypoint is the first point in the path implementing the turn (line 13). GeneratePath also uses an `offset` coming from the OPS representing the minimum allowed distance from the left or right wall. The result of this function is a smooth and continuous path with a safe distance from the detected walls.

For the case of approaching and turning in intersections, the algorithm finds the closest path in the set of pre-recorded paths in global coordinates (line 10). It transforms this path to an ego-centric coordinate system (line 11) and decides if the starting point of the path is in front of the vehicle (line 12) or behind. If so, which corresponds to case (2) above, then the path generation algorithm is used to generate a path to the starting point (line 13); otherwise the remaining subset of the path still in front of the vehicle is used.

In addition to the path that defines how the vehicle should steer we also need to compute the desirable speed. The path generation algorithm is designed to guarantee safe distances to the tunnel walls, and the speed is hence limited based only on user input and the path curvature, as computed by `safePathSpeed` (line 18), and the maximum visible distance of the traversable area in the direction of motion.

Both the path and the computed speed are fed to the existing path tracking and rate controllers, that navigate the vehicle along the calculated path and run at 25 Hz, synchronous with the top-level control loop. The path and speed information, together with the detected tunnels, are also sent to the OPS to be presented to the operator.

V. Experimental Results

A. Execution

Five subjects were enlisted. These were operators from Atlas Copco's development and service organizations, with strong experience in driving the LHD both in teleoperation and on-board. All subjects had also driven the vehicle during a six month field test of Atlas Copco's teleoperation system.

Each subject drove the test path 10 times for each mode. Learning effects were reduced by allowing the subjects to practice driving in all the modes for up to half an hour before starting the test. To further limit learning and exhaustion effects the order in which the test cases were run was different for all subjects.

During the experiment the run times and the significant events were recorded to be able to asses productivity as well as the impact on the wear and maintenance requirements. The three main events that were tracked were: (1) slamming the bucket into the wall when driving backwards, (2) collisions, and (3) driving into the simulated ditch. Of these events the first two are not uncommon in normal operation of LHD vehicles, while a vehicle that drives into a ditch often needs to be towed and sometimes requires extensive repairs.

Events (1) happen easily since the front side tip of the bucket sweeps outside the perimeter of the wheel tracks when turning. This seldom has any direct negative effect on the machine, but often results in pieces of rock falling out of the bucket onto the roadbed, which can cause damage to the tires during the following runs if not removed. These events therefore lead to either reduced productivity (if the road bed is cleaned) or to increased maintenance costs (if not).

For events (2), both frontal collisions and scraping the side of the vehicle were counted. Both events (1) and (2) were extracted from audio feed back. As the operators station was located in a small cabin inside the test area and collisions give a significant audio impulse, relevant collision events were easy to notice.

A ditch event (3) was defined as when a tire crossed the painted line between the cones. The lines were visually inspected after each run when the visual feedback of the OPS indicated that the vehicle had been close to a ditch.

B. Quantitative Results

The measured times t_{Tunnel} and t_{Turn} for all runs of the test subjects are displayed as box plots in Figure 8. Each box has lines at the lower quartile, median, and upper quartile values. Dashed lines extend to the maximum and minimum values of the data. No data could be collected for the local autonomy mode from the fifth test subject because the test vehicle required unscheduled maintenance.

No learning or fatigue effects could be identified in any of the subject and driving mode data sets.

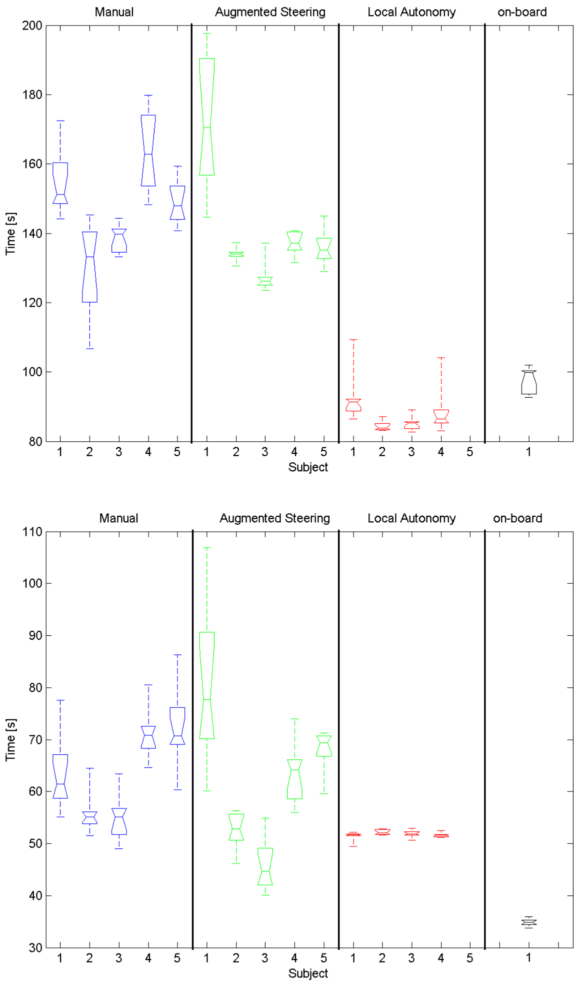

Fig. 8. t_{Tunnel} (top) and t_{Turn} (bottom) for all the 10 runs of each subject and mode. The result from the Manual runs are displayed to the left (blue), Augmented Steering in the middle (green) and Local Autonomy runs to the right (red). The rightmost column shows the result of the 5 runs in on-board mode (black).

For comparison we also let an ordinary LHD operator drive the test path manually on-board the test vehicle. In on-board manual mode the path was driven five times.

1) Manual mode vs Local autonomy mode: From Figure 8 it is clear that the subjects are able to teleoperate the vehicle significantly faster in local autonomy mode than in the traditional manual mode. This is confirmed by ANOVA analysis of the result of the first four subjects for both t_{Tunnel} ($F_{(1,78)} = 472.1,\ p < .0001$) and t_{Turn} ($F_{(1,78)} = 52.8,\ p < .0001$). We can thus with a high degree of confidence reject the first and second null hypotheses, i.e. it is unlikely that an operator can drive the vehicle as fast in manual mode as in local autonomy mode.

Analysis of the number of wear events as displayed in Table I, show that we can reject the third null hypothesis with a high degree of confidence ($F_{(1,6)} = 29.4,\ p < .002$). It is thus unlikely that the number of wear events would be as low when driving in manual as in local autonomy mode.

From Figure 8 it is clear that local autonomy even outperform an on-board operator when driving inside tunnels.

TABLE I

WEAR EVENTS THAT OCCURRED DURING THE EXPERIMENT.

Mode	Man					AS					LA
Subject	**1**	**2**	**3**	**4**	**5**	**1**	**2**	**3**	**4**	**5**	**1-4**
Ditch	1	1	2	2	1	2	0	2	1	1	0
Collision	2	0	0	1	0	2	0	1	0	0	0
Slam bucket	1	2	0	2	0	0	1	3	1	0	0

When turning at intersections, however, the on-board manual operator is significantly faster. The main reason for this is that the top-speed when playing back the pre-recorded paths at intersection was limited to 7 km/h, which is much slower than how an on-board operator drive through the intersections. In retrospective, we now believe that this was a too conservative choice, and that path generation and path following at intersection could safely be run at higher speeds.

2) Manual mode versus Augmented steering: Figure 8 indicates that most of the subjects are able to drive the vehicle faster in augmented steering mode compared to manual mode. This is confirmed when analyzing the result with respect to both mode and subject, where there is a significant difference for t_{Tunnel} ($F_{(1,90)} = 10.31,\ p < .002$). We can thus reject the first null hypothesis, concluding that it is likely that most operators can drive the vehicle faster in tunnels in Augmented steering mode than in Manual mode.

From a maintenance and wear point of view there is no difference between the two modes since the number of wear events are almost identical (Man: 15, AS: 14).

C. Qualitative Results

In order to analyses their subjective perception, the drivers were asked to provide answers ranging from 0 ("I totally disagree") to 4 ("I completely agree") on the questionnaires.

From these answers we have a strong indication that the subjects prefer the local autonomy mode compared to manual mode in both the driving situations $Tunnel$ (mean $\mu : 3.75$, STD $\sigma : 0.50$) and $Turn$ ($\mu : 3.75, \sigma : 0.50$).

On the question if the subjects believed that they would be able to drive the machine for a full eight hour shift at the speed they achieved in the experiment, the answers differ significantly ($F_{(1,6)} = 9.8,\ p < .03$) between manual mode ($\mu : 1.20, \sigma : 1.10$, where 1 means "Disagree") and local autonomy mode ($\mu : 3.25, \sigma : 0.50$, where 3 means "Agree").

The questionnaires show that the subjects are pretty much undecided if they prefer Augmented steering mode compared to manual mode in both driving situations $Tunnel$ ($\mu : 2.60, \sigma : 1.14$) and $Turn$ ($\mu : 2.40, \sigma : 1.52$). Neither is there any relevant difference between the two modes regarding if the subjects believe to be able to keep up the speed they achieved in the experiment for a full shift: ($\mu : 1.20, \sigma : 1.10$) for manual, ($\mu : 1.60, \sigma : 0.89$) for augmented steering.

D. Discussion

Based on the result above, the productivity improvement of the tramming part of the work cycle is close to 50% when comparing manual and local autonomy mode. If the loading, dumping and driving in and out of the load drift is approximated to take 90 s, the productivity improvement of a teleoperated LHD in a scenario corresponding to the one in the user study would still be 30%. Combined with the fewer wear events this leads to the conclusion that local autonomy mode is superior to manual teleoperation. When it comes to augmented steering compared to manual mode all figures speak mildly in favor of augmented steering, and there is a significant difference in the speed achieved in the *Tunnel* sections of the route.

Another thing to notice is that the standard deviation for t_{Tunnel} and t_{Turn} is much smaller for local autonomy than for manual mode. This indicates that the impact of operator performance and attitude is smaller in the local autonomy mode, which may make the productivity during a specific shift easier to predict.

Interestingly, there were more ditch events than collision events, even though driving into a ditch has more severe consequences than colliding with a wall. We speculate that the test subjects did not take the simulated ditches seriously and thus drove more recklessly and perhaps faster than they would have done if there had been real ditches in the test scenario.

Several of the drivers also attempted to drive with third and fourth gear enabled during the manual baseline runs, but concluded that they were not able to handle the vehicle at speeds exceeding 10-12 km/h even in straight tunnels. The industry practice of limiting the transmission to second gear in manual teleoperation thus seem justified and motivate the use of local autonomy.

VI. Conclusions

The main contribution of this paper is a user study evaluation that indicates that local autonomy functionality in mining LHD vehicles can lead to significant productivity improvements and reduced maintenance costs as compared to normal teleoperation. This result is meaningful insofar it is feasible to implement a system which is able to generate local paths similar to the ones used in our test implementation, while lifting the assumptions made in Section IV. Given the current state of the art in path planning, we consider this as possible. Our result shows that such a system would enable mines to improve LHD operator safety, while maintaining high productivity and low maintenance costs. Local autonomy is very interesting from a commercial point of view since it opens up a market with hundreds of potential customers as opposed to dozens of potential customers of fully autonomous systems.

As an additional contribution, we have proposed a new method, *Augmented Steering*, to improve the performance of teleoperated LHD vehicles. This method is simple and it can be applied to any teleoperated articulated vehicle since it only needs the input from an articulation angle sensor. The user study indicates that this technique can improve performance in teleoperation of articulated vehicles, which should be contrasted with the almost negligible cost of its implementation.

References

[1] C. Swart, "Automation test goes live," *Mining and Construction*, no. 1, 2008.

[2] K. Casteel, "Underground haulage equipment trends," *Engineering and Mining Journal*, 2008.

[3] "AutoMine," http://www.miningandconstruction.sandvik.com.

[4] L. Conway, R. Volz, and M. Walker, "Teleautonomous systems: projecting and coordinating intelligent action at a distance," *IEEE Trans. on Robotics and Automation*, vol. 6, no. 2, pp. 146–158, 1990.

[5] T. Fong, C. E. Thorpe, and C. Baur, "Advanced interfaces for vehicle teleoperation: Collaborative control, sensor fusion displays, and remote driving tools," *Autonomous Robots*, vol. 11, no. 1, pp. 77–85, 2001.

[6] E. Slawinski and V. Mut, "Control using prediction for teleoperation of mobile robots," in *Int Conf on Mechatronics and Automation. ICMA*, 2007, pp. 1782–1787.

[7] A. Steinfeld, "Interface lessons for fully and semi-autonomous mobile robots," in *Proc of the IEEE Int Conf on Robotics and Automation. ICRA*, 2004, pp. 2752–2757.

[8] R. Wegner and J. Anderson, "Agent-based support for balancing teleoperation and autonomy in urban search and rescue," *Int Journal of Robotics and Automation*, vol. 21, no. 2, pp. 120–128, 2006.

[9] E. S. Duff, J. M. Roberts, and P. I. Corke, "Automation of an underground mining vehicle using reactive navigation and opportunistic localization," in *Australasian Conf on Robotics and Automation*, Auckland, NZ, 2002, pp. 151–156.

[10] F. Ogaki and K. Suzuki, "Adaptive teleoperation of a mobile robot under communication time delay," in *Int Workshop on Robotic and Sensors Environments. ROSE*, Ottawa, Canada, 2007, pp. 1–6.

[11] K.-B. Sim, K.-S. Byun, and F. Harashima, "Internet-based teleoperation of an intelligent robot with optimal two-layer fuzzy controller," *IEEE Tran on Industrial Electronics*, vol. 53, no. 4, pp. 1362–1372, 2006.

[12] N. Mitsou, S. Velanas, and C. Tzafestas, "Visuo-haptic interface for teleoperation of mobile robot exploration tasks," in *15th IEEE Int Symp on Robot and Human Interactive Communication. ROMAN*, 2006, pp. 157–163.

[13] J. Babić, M. Budišić, and I. Petrović, "Dynamic window based force reflection for safe teleoperation of a mobile robot via internet," in *IEEE/ASME Int Conf on Advanced Intelligent Mechatronics*, Zürich, Switzerland, 2007, pp. 1–6.

[14] S. Kim, C. Roh, S. Kang, and M. Park, "Outdoor navigation of a mobile robot using differential gps and curb detection," in *Proc of the IEEE Int Conf on Robotics and Automation*, 2007, pp. 3414–3419.

[15] M. Wang and J. N. K. Liu, "Interactive control for internet-based mobile robot teleoperation," *Robotics and Autonomous Systems*, no. 52, pp. 160–179, 2005.

[16] J. Jacko and A. Sears, *The human-computer interaction handbook: fundamentals, evolving technologies, and emerging applications*. Lawrence Erlbaum Associates, 2003.

[17] C. Wohlin, P. Runeson, M. Hst, M. C. Ohlsson, B. Regnell, and A. Wessln, *Experimentation in software engineering: an introduction*. Kluwer Academic Publishers, 2002.

[18] F. G. Shinskey, *Feedback Controllers for the Process Industries*. McGraw-Hill Professional, 1994.

[19] J. Marshall, T. Barfoot, and J. Larsson, "Autonomous underground tramming for center-articulated vehicles," *Journal of Field Robotics*, vol. 25, no. 6-7, pp. 400–421, 2008.

[20] J. Larsson, "Reactive navigation of an autonomous vehicle in underground mines," Licentiate thesis, Örebro University. Örebro, Sweden, 2007.