Light Metals 2013 *Edited by: Barry Sadler*
TMS (The Minerals, Metals & Materials Society), 2013

Autonomous Vehicle and Smelter Technologies

Ashley Tews[1], Paulo Borges[1]

[1]Commonwealth Scientific and Industrial Research Organisation (CSIRO)
Queensland Centre for Advanced Technology; Pullenvale, Queensland, Australia

Keywords: automation, vehicle, safety

Abstract

In 2005, the CSIRO automated a forklift-based Hot Metal Carrier to be capable of typical metal transfer operations around a smelter. The project was highly successful and the vehicle has demonstrated hundreds of hours of live autonomous operations to thousands of people from industry and the public. As a result of the exposure of the project, we have expanded its focus to demonstrate how various technology components can be utilised to gain benefits in other areas of smelter automation. These include asset tracking, vehicle usage analysis, pedestrian detection and infrastructure profiling. Each of these components was derived from the core autonomous vehicle technology suite and has shown a high potential for improving safety and efficiency of vehicle operations, metal tracking from pot to furnace, or improved diagnostic processes for measuring deformations of bakes and furnaces. The autonomous vehicle project and its extended technologies are described in this paper.

Introduction

Vehicles are in constant operation around industrial worksites. In many applications, they perform repetitive homogeneous tasks such as moving loads from one warehouse location to another. In the aluminium industry, the primary task of Hot Metal Carriers (HMCs) is to perform the task of transporting molten aluminium from the smelter to the casting shed. Vehicles can weigh 20 tonnes unloaded and are typically either converted forklifts (Figure 1) or articulated trucks. The molten aluminium is carried in large metal crucibles. The crucibles weigh over 2 tonnes and they can hold 8 or more tonnes of molten aluminium at 700 degrees Celsius. Therefore, HMC operations are considered heavy, hot, and hazardous, with safety of operation a significant issue.

Figure 1: A Hot Metal Carrier picking up the crucible.

Our primary research is focused towards automating the operations of material transport vehicles such as HMCs and is equally relevant to solid metal transfer vehicles and anode exchange vehicles. There are many challenges in their operating environment considering they travel inside and outside buildings. Inside, there is a vast amount of infrastructure, other mobile machines and people. In various areas, there are large magnetic fields and high temperatures near the pots and furnaces. Outside, their paths may be surrounded by infrastructure, fences, and their operation may be effected by the environmental conditions: rain, fog, snow, and heat. Research into automating these vehicles and their operations needs to consider the variability in operating conditions to produce repeatable and reliable performance of the task.

At our worksite, we have fully automated a forklift-style HMC and have demonstrated typical transport operations and movements of a production vehicle. Our vehicle is capable of autonomous start up, shutdown, navigation, obstacle management, tasking, crucible pickup and drop off. It has conducted hundreds of hours of autonomous operations and demonstrated substantial continuous periods of high reliability and repeatability.

As part of automation, we have developed many hardware and software modules to allow control and monitoring of the vehicle as well as outfitting sensors to allow it to monitor the surrounding area. We have also developed technology to allow external sensing systems such as site cameras to be used in improving the sensing range of the vehicle. The result can be considered as a relatively complete package for automating a smelter vehicle. However, there is additional utility in various components that allow them to be used independently. Examples include vehicle and asset tracking systems, infrastructure monitoring and pedestrian detection systems.

The remainder of this paper will overview the main modules in our autonomous HMC and how some of them can be used as stand-alone units for additional applications. We will finish the paper with relevant conclusions.

Our Approach

Our HMC has been automated to the level it can carry out all the operations of a conventionally operated vehicle with a driver on-board. However, whereas the driver of a conventional HMC is responsible for the efficiency, safety and sensing for the operations, the autonomous HMC has hardware systems to take this role. Apart from the obvious internal sensors that provide information about the state of the vehicle (e.g. temperature, oil pressure, odometry, hook height, mast tilt, etc.), the vehicle has external environment sensors to assist with navigation, obstacle management and crucible tasks. Four scanning laser range finders are positioned around the vehicle (Figure 2) and are tilted down to provide 360 degrees of coverage to a distance of approximately 30m, with the blindspots apparent in the figure. These lasers are used to provide beacon-based localisation and obstacle detection. Two Pan-Tilt-Zoom (PTZ) webcams are attached to the mast as the primary sensor for locating the crucible via markers on its handle [1] and three cameras around the front of the vehicle for pedestrian detection and loclisation (Figure 3.

The autonomous HMC's safety system consists of a number of physical interlocks, Emergency Stops (E-Stops), obstacle management, on- and off-board RF remote failsafe and software watchdogs. The E-Stops are located around the vehicle, inside and on the portable remote RF device. Activating an E-Stop brings the vehicle to a quick halt and shuts down the engine. Hydraulic controls are frozen at this point. Door interlocks are also included in the E-Stop loop to prevent access to the vehicle whilst running autonomously. The software safety systems consist of high-level velocity control when objects are detected close to the vehicle and low-level watchdog checks between interface level software and the low-level control software. A watchdog timeout initiates an E-Stop.

Figure 4 provides a high-level view of the software and hardware architecture of the autonomous HMC's systems. Low-level components such as throttle, brakes, steering, hook and mast controls are controlled through Programmable Logic Controllers (PLCs). The critical safety components, such as the E-stop buttons and the watchdog monitor, are controlled through higher grade failsafe PLCs. These PLCs provide redundancy checks of relay connections and continuously monitor the input and output state of hardware connections.

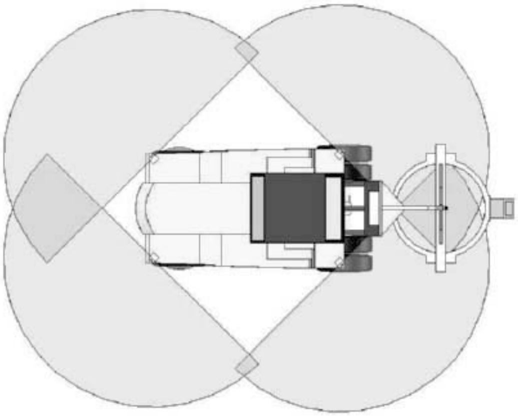

Figure 2: The HMC's lasers are located at each corner of the vehicle and offer overlapping coverage out to approximately 30m (indicated by purple sectors).

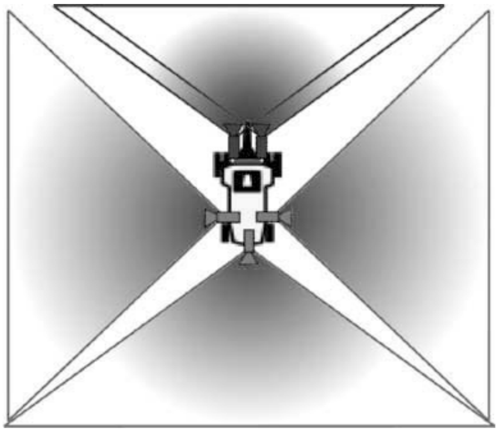

Figure 3: Camera locations and their fields of view shown on the autonomous HMC. Red cameras are used for pedestrian detection and vision-based localisation. Blue PTZ cameras are used for crucible identification.

The H/W Abstraction program converts the internal vehicle state sensors to human-readable signals and manages the vehicle demands in an opposite manner. High Level programs work directly with the external sensors and vehicle state to control the vehicle. Vehicle Level programs control and monitor the vehicle hardware systems. The main High Level programs are described next.

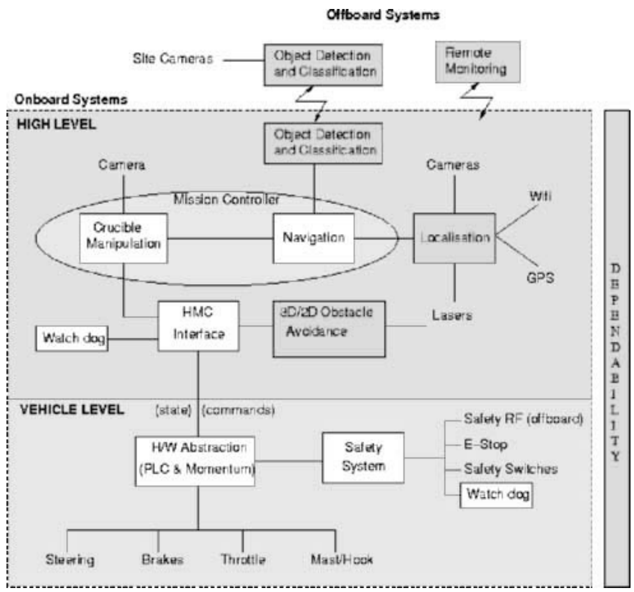

Figure 4: The HMC system architecture. The program blocks are shown in boxes or ellipses with leaves representing physical parts of the system. Dependability is common to all modules.

Localisation

The localisation systems are responsible for determining the position of the vehicle in site coordinates, whether it is in sheds or outside.

The choice of localisation system depends on the required accuracy for the operations and the sensors available. For autonomous vehicle control, high precision is required in certain parts of the environment such as around infrastructure, people and during payload manipulation, whereas traversing along wide roads can use a less precise method. The main environment sensors on the autonomous HMC are the scanning laser rangefinders (Figure 2) as they have been tested over many years in our industrial projects and proven robust and consistent. The camera suite act as a secondary sensor for localisation. We have developed localisers (described below) for each sensor, that are appropriate for different applications. The laser-based localisers are simpler and more robust than the camera-based since they are not affected by lighting variations or absence and only consist of 1D range and intensity data.

Laser Beacon-based Localisation The main localisation method used by our autonomous HMC is beacon-based. This consists of placing reflective tape (beacons) at regular intervals on permanent pieces of infrastructure. The tape locations are either surveyed or determined by other means, and stored in a reference database. The vehicle's onboard scanning laser rangefinders can identify the tape by its high reflectivity and therefore, can map the sensed beacon 'constellation' back to the database to estimate the vehicle's pose.

Laser-based SLAM The drawback with using artificial landmarks such as beacons is the initial setup required (selecting and surveying locations) and maintenance of existing and new locations. A more self-contained method is termed 'Simultaneous Localisation and Mapping' (SLAM) which automatically identifies key features in the environment and creates a reference map on-the-fly. As part of this process, the location of the vehicle is determined using its reference map and new areas added as identified. There are many challenges with implementing a successful SLAM system such as determining key features in the environment, determining the vehicle's motion, and recognising when the same area has been mapped which can occur sometime after it was last visited. In our implementation [2], the system is robust to people and vehicles moving around and can assimilate changes to the environment into the reference database. This system is useful for tracking, as opposed to controlling vehicles around a worksite.

Vision-based Localisation An important issue with developing a reliable autonomous vehicle is to ensure there is redundancy in primary systems wherever possible and practical. This is especially relevant with the localisation system and to provide redundancy, a different sensing modality on a different power and communication path is required, that can provide similar accuracy and performance to the primary. Our approach is to use colour cameras as they can be applied to multiple applications and are inexpensive. As with other localisation systems, a feature database is required as a reference. Our system uses the strong edge features of major infrastructure. These features are derived either by surveying, or by an automated approach we have developed [3] using a 3D SLAM system [4]. The vision processing system receives streaming images from the onboard cameras and processes it into an edge representation of the surrounding environment to map to the feature database [5]. We have conducted experiments to demonstrate the utility of using this system as a backup to the beacon-based localiser with successful results during daytime operations [6].

Navigation

The autonomous HMC's navigation system uses waypoints derived automatically by driving the required route of operations. Waypoints are recorded after a certain change in distance since the last waypoint or a certain change in vehicle heading. Each waypoint also contains a velocity so ramping speeds can be utilised for smoother navigation. The resulting waypoint list is split into task segments with each segment being a homogeneous action such as a forwards traverse (used for normal navigation) or backwards traverse (used for crucible manipulation tasks). Within a segment, the navigation system switches to the next waypoint in the list when it is close to the current waypoint.

Laser-based Object Detection

Object detection consists of determining where obstacles and the payload are. When the proximity of nearby objects such as people or infrastructure gets closer to the HMC, it slows and will eventually stop if they get too close. This is disabled for the hook-end when the HMC is approaching or carrying the crucible. The lasers are also used for detecting the location of crucibles in the environment [7] which is used for pickup operations.

Crucible Operations

The key functionality of a forklift-type Hot Metal Carrier is its ability to handle the crucible. Two main operational phases can be distinguished: crucible pickup and crucible drop off. Drop off is an easy maneuver from an automation point-of-view. No sensing is required and a simple ballistic manoeuvre is sufficient.The pickup manoeuvre is more difficult. It can be divided into two steps: first, an approach step where the hook is visually guided toward to the pickup point in the middle of the crucible handle, then the actual pickup. The latter is an easy manoeuvre, again a ballistic movement, similar to a drop off.

The approach part is more complex. It is principally based on the onboard mast-located PTZ webcams detecting the crucible from about 20m. We use specially designed barcodes placed on the bale arm to identify the crucible visually to the HMC [1]. As with most outdoor computer vision applications, it requires proper management of sensitivity to lighting conditions.

Mission Planning (Tasking) and Recovery

The Mission Controller is responsible for switching between tasks and monitoring their performance. A mission is a sequence of tasks with each task returning its status during execution. Once a task has finished, the Mission Controller selects the next task. Contingencies occurring during task execution cause the Mission Controller to select the contingency subtask for that task. For example, a missed crucible pickup will trigger a "missed approach" signal and the HMC will move away from the crucible and retry to pick it up. During a 5 hour experiment of continuous operations, the only halt in operations occurred after approximately 4 hours when the remote Safety RF unit's battery went flat. This triggered an E-Stop on the vehicle. When the battery was replaced, the vehicle was restarted and as it was about to pickup the crucible at the time of the E-stop, it executed a missed pickup and successfully continued its operations. In a fully autonomous system, missions can be allocated by schedulers.

Applications of Component Technologies

The systems discussed in the preceding section form the main technology components for the autonomous HMC. Their joint utility is for creating a reliable and functional autonomous unit. However, some provide utility as standalone units or have spawned new technologies useful to the overall goal of improving smelter operations through automation. A selection of these are described next.

Anti-collision systems

As discussed in the Object Detection section, the lasers effectively provide proximity detection zone in a relatively horizontal plane around the vehicle. Any object impinging on this plane will result in the vehicle slowing to a stop depending on the range. This simple approach can be implemented on manned vehicles that can reduce the speed or stop the vehicle in the same manner.

Pedestrian Detection (onboard and offboard)

Situational awareness for industrial vehicles is crucial to ensure safety of personnel and equipment. We have implemented camera based pedestrian detection systems on the autonomous HMC, which can help reducing the risk of collision between the vehicle and people. Two different approaches have been developed, using (i) onboard cameras and (ii) offboard cameras.

Onboard cameras provide autonomous pedestrian detection from the perspective of the vehicle, where cameras can be mounted facing forward, sideways or to the back, as illustrated in Figure 3. Any pedestrian close to the HMC triggers a warning signal to the driver.

Offboard cameras cover areas that are not necessarily visible from the vehicle. While human drivers and onboard sensors are able to detect pedestrians within line-of-sight, in complex environments obscured dynamic objects can unpredictably enter the path of a vehicle. The safety system we have developed [8] integrates a vision-based offboard pedestrian tracking subsystem with the onboard localisation and navigation subsystems discussed earlier in this paper. This combination enables warnings to be communicated and effectively extends the vehicle operator's field of view to include areas that would otherwise be blind spots. A simple flashing light interface in the vehicle cabin (illustrated in Figure 5) provides a clear and intuitive interface to alert drivers of potential collisions. The pedestrian detection is carried out with efficient computer vision pattern recognition techniques developed by our research group [9, 10].

Hundreds of hours of tests under real operations have illustrated the applicability of the system, with positive feedback from industrial drivers.

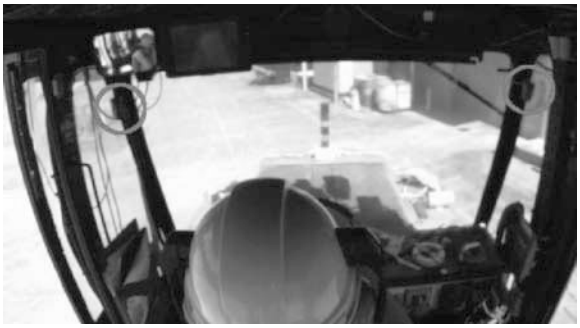

Figure 5: Lights mounted inside the cabin, highlighted by the orange circles in this figure. During manual operation, these lights alert the driver of a potential collision route between the vehicle and a nearby pedestrian.

More recently, we have extended the system to perform joint pedestrian detection [11], combining information from on-board and off-board cameras simultaneously. This novel approach is currently being trialled. Preliminary positive results indicate the method can further improve the pedestrian detection system reliability.

Vehicle Utilisation

Localisation can be used as a standalone component requiring only lasers and an Inertial Measurement Unit connected to an onboard computer, in the case of beacon-based or SLAM systems. This allows a wide variety of applications and analyses to be undertaken in efforts to reduce maintenance, improve safety or improve efficiency:

Tracking
The data from the localiser consists of timestamped position, from which, the vehicle velocity can be determined. This information can be logged for holistic analysis of vehicle usage statistics associated with areas on site. E.g, over the course of a shift, all transit times from payload pickup to delivery at set locations can be recorded and automatic reports generated. Live vehicle locations can also be displayed on an online site map via a web interface, including whether they are carrying a payload.

Event Analysis Adding a simple G-force sensor (IMU) to the localisation information stream allows for monitoring of high G events such as payload contact, collisions, fast cornering and braking, and even potentially identifying where road surfaces have deteriorated causing bumps.

Automatic Speed Reduction Since vehicle position and velocity data is available from the localisation system, it can be used to impose speed restrictions on the vehicle aligned with those required in different areas. This would require interfacing with the vehicle's control system (brakes, throttle) and storing the different area boundaries with their speed limits in the onboard PC. When the vehicle enters the area as determined by the localiser, the local speed limit can be referred to and enforced.

Asset Tracking

A problem that can arise with manned payload delivery is with wrongful delivery. This can be a particularly costly operation for smelters where the wrong metal put in to a furnace can render it unusable. We have developed a system [12] that uses site cameras to detect and read identification labels on stationary payloads (Figure 6). The cameras can be mounted in key areas where verification checks are required, such as at a potline, weighbridge or furnace. The system is highly reliable and uses check digit verification at the end of the label to ensure the label identification is accurate. The information can then be used online to track the payload.

Infrastructure Profiling

Apart from being used for extracting environment features, the 3D SLAM system can create relatively accurate models of environments or infrastructure. We have applied the technology to mapping caves, underground mines, building interiors, and worksites. The purposes have ranged from visualisation, deformation analysis over time, and dimensioning of features/objects. In the smelter-related applications, we have used it to map sites to create 3D edge models for vision-based localisation and Carbon Bake Furnace (CBF) pit scans. The

(a) Initial detection (b) Zoomed in detection

(c) Text extraction of ID

Figure 6: Identifying a crucible label in the environment without initial knowledge of its location.

CBF application can also be applied to cathode autopsies, temporal pot deformation analysis and other infrastructure monitoring applications. An example of an interior shed mapping application is shown in Figure 7.

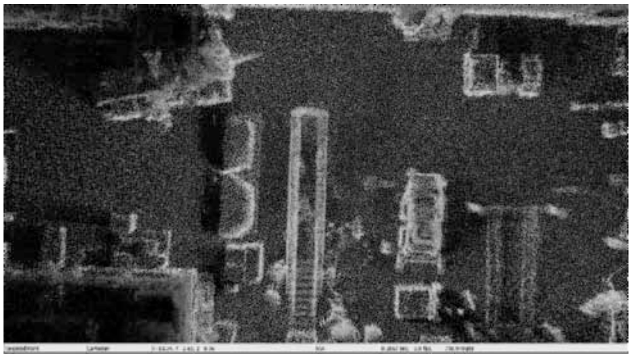

Figure 7: 3D view of the processed data from scanning our robotics bay. The colours are heights from blue (low) to red (high). Note the maintenance pit in the centre.

Summary

This paper has overviewed CSIRO's autonomous Hot Metal Carrier and the application of its main automation components. The components can be be considered as part of an automation solution necessary for either retrofitting or developing a new material transport vehicles such as HMCs, ingot transfer, or anode exchange vehicles. These components can also be considered independently for their application in non-autonomous vehicle roles. By considering them from both perspectives, an incremental roadmap can be rolled out for integrating automation into a smelter's operations around material transport. The independent components offer a value return along the integration path.

References

[1] C. Pradalier, A. Tews, and J. Roberts, "Vision-based operations of a large industrial vehicle: Autonomous hot metal carrier." *Journal of Field Robotics*, vol. 25, no. 4-5, pp. 243–267, 2008.

[2] M. Bosse and R. Zlot, "Map matching and data association for large-scale 2D laser scan-based SLAM," *International Journal of Robotics Research*, vol. 27, no. 6, pp. 667–692, June 2008.

[3] P. Borges, R. Zlot, M. Bosse, S. Nuske, and A. Tews, "Vision-based localization using an edge map extracted from 3d laser range data," in *International Conference on Robotics and Automation*, May 2010, pp. 4902 –4909.

[4] M. Bosse and R. Zlot, "Continuous 3D scan-matching with a spinning 2D laser," in *International Conference on Robotics and Automation*, May 2009, pp. 4312–4319.

[5] S. Nuske, J. Roberts, and G. Wyeth, "Visual localisation in outdoor industrial building environments," in *IEEE International Conference on Robotics and Automation*, May 2008, pp. 544–550.

[6] J. Roberts, A. Tews, and S. Nuske, "Redundant sensing for localisation in outdoor industrial environments," in *6th IARP/IEEE-RAS/EURON Workshop on Technical Challenges for Dependable Robots in Human Environments*, May 2008.

[7] A. Tews, "3d payload detection from 2d range scans," in *Intelligent Robots and Systems*, Sep 2011, pp. 834 –840.

[8] P. Borges, R. Zlot, and A. Tews, "Pedestrian detection for driver assist and autonomous vehicle operation using off-board and on-board sensing," in *Australasian Conference on Robotics and Automation*. ACRA, Dec 2010.

[9] P. V. K. Borges, "Blob motion statistics for pedestrian detection," in *2011 International Conference on Digital Image Computing: Techniques and Applications*. IEEE, Dec 2011.

[10] S. Johnsen and A. Tews, "Real-time object tracking and classification using a static camera," in *IEEE International Conference on Robotics and Automation - Workshop on People Detection and Tracking*, May 2009.

[11] J. Miseikis and P. V. K. Borges, "Joint human detection from on-board and off-board cameras," in *Workshop on Robots and Sensors Integration, IEEE International Conference Intelligent Robots and Systems*, Oct 2012.

[12] E. Markey and A. Tews, "A system for reliable text recognition in an industrial environment," in *Australiasian Conference on Robotics and Automation*, Melbourne, Australia, Dec 2011.