# Safe and Dependable Operation of a Large Industrial Autonomous Forklift

Ashley Tews

***Abstract*— For autonomous vehicles to operate in industrial environments, they must demonstrate safe, reliable, predictable, efficient and repeatable performance. To achieve this, two important high level factors are situational awareness and system dependability. The vehicle must be able to identify objects and predict the trajectories of dynamic objects in order to avoid unplanned interaction and to improve performance. In many environments, the vehicle is also required to operate for long periods of time over many days, weeks and months. Towards this goal, the vehicle needs to self-monitor its hardware and software systems, and have redundant primary systems. We have incorporated many of these requirements into our Autonomous Hot Metal Carrier which is a modified 20 tonne forklift used in aluminium smelters for carrying a 10 tonne payload between large sheds, in the presence of other vehicles and people. Our HMC has successfully conducted 100's of hours of autonomous operation in our industrial worksite. The main hardware and software systems will be discussed in this paper with particular focus on the redundant localisation and obstacle avoidance systems. Experiments are described to highlight the performance of the HMC systems in the presence of dynamic objects around a typical worksite.**

## I. Introduction

Vehicles operate constantly 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, Hot Metal Carriers (HMCs) perform the task of transporting molten aluminium from the smelter (where the aluminium is made) to the casting shed where it is turned into block products. The vehicles weigh approximately 20 tonnes unloaded and resemble forklifts except they have a dedicated hook for manipulating the load rather than fork tines (Figure 1). The molten aluminium is carried in large metal crucibles. The crucibles weigh approximately 2 tonnes and they can hold 8 tonnes of molten aluminium usually superheated above 700 degrees Celcius. Therefore, HMC operations are considered heavy, hot, and repetitive, with safety of operation a significant issue.

Our research is focused towards automating the operations of Hot Metal Carrier-like vehicles. There are many challenges in their operating environment considering they travel inside and outside of buildings. Inside, there is a vast amount of infrastructure, other mobile machines and people. In various areas, there are strong magnetic fields and high temperatures near the molten aluminium vats. Outside, their paths may be surrounded by infrastructure, fences, and their operation may be affected 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.

Ashley Tews is with the Commonwealth Scientific and Industrial Research Organization, Queensland, Australia Ashley.Tews@csiro.au

Fig. 1. A Hot Metal Carrier in the process of picking up the crucible.

At our worksite, we have fully automated a Hot Metal Carrier and have demonstrated typical operations of a production vehicle. Our vehicle is capable of autonomous start up, shutdown, navigation, obstacle management, and crucible pickup and drop off. It has conducted hundreds of hours of autonomous operations and demonstrated long periods of high reliability and repeatibility. The vehicle also has several safety systems incorporated into it to make its operations as safe as possible. The remainder of this paper outlines our research and results.

## II. Modules

To be fully capable of conducting all tasks of a manned vehicle, the autonomous HMC needs to address the issues of safety, reliability and repeatability. We have considered these issues when automating the HMC's hardware and software systems. A block diagram of the major hardware components is shown in Figure 2.

The major modules of the system are separated into high level and vehicle level. The high level modules provide commands for controlling the vehicle based on the requested tasks, vehicle state and observed state of the environment.

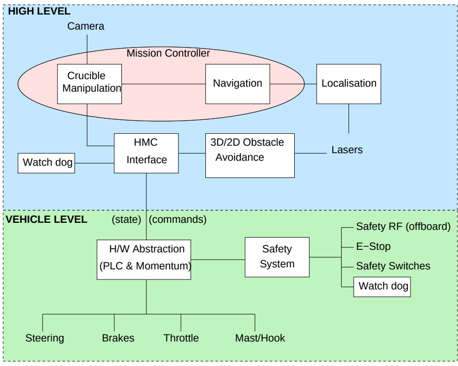


Fig. 2. Overview of the hardware architecture.

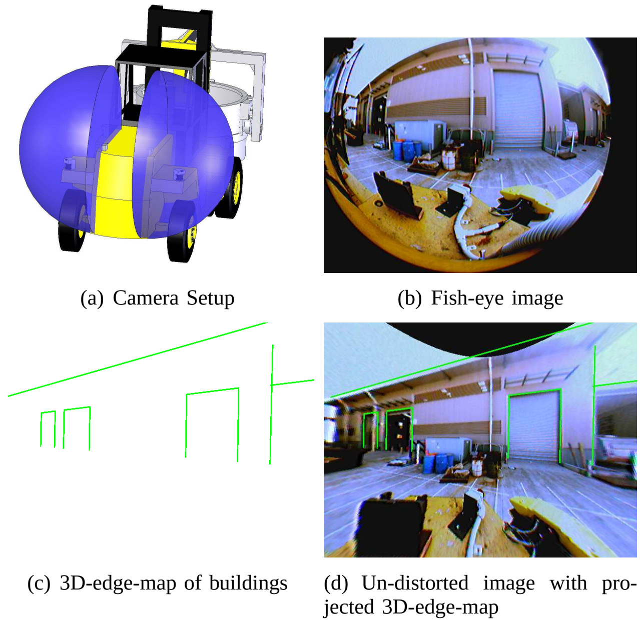

(a) Camera Setup

(b) Fish-eye image

(c) 3D-edge-map of buildings

(d) Un-distorted image with projected 3D-edge-map

Fig. 3. Examples of the vision-based localisation system. Two fish-eye cameras are placed at the front of the vehicle facing sideways (a and b). The blue hemispheres represent the field of view of the cameras. A surveyed edge map of the buildings (c) can be tracked in the images (d).

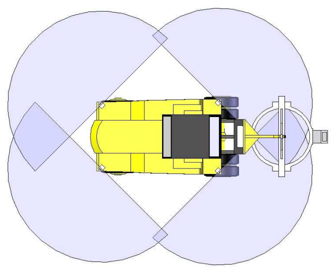

Fig. 4. The HMC's coverage from lasers mounted on the corners.

The vehicle level modules provide vehicle state information and act as the interface to the vehicle's control systems. They also take care of several low level safety interfaces including physical interlocks, heartbeat monitors between critical systems, and e-stop control.

The vehicle has a light stack on the hood to provide a basic visual indication of its state. The lights indicate whether an e-stop is active, user intervention is required, if the vehicle is operating in autonomous mode, as well as two programmable status indicators that can be used to indicate if a monitored system's parameters exceed or drop below a threshold value (e.g. pneumatic system pressure).

To allow the HMC to conduct autonomous operations safely without the requirement of a safety supervisor to be in the cabin, a RF safety remote is part of the low level interface. This allows the supervisor to be outside of the cabin to monitor operations. The unit has several programmable switches and an e-stop switch to stop the vehicle in an emergency.

The remainder of this section describes the main high level modules.

### A. *Redundant Localisation*

A fundamental requirement for any autonomous vehicle conducting reliable operations is localisation. It forms the basis of any high level navigation, path planning and obstacle avoidance systems. To achieve high reliability, single points of failure need to be reduced or removed completely. Many localisation system use a single type of sensor or fuse sensors into a single system. A hardware or sofware fault with these systems can render the localisation useless. Consequently, the vehicle may have little choice but to signal a fail and wait to be rescued. Using redundant hardware and software systems provides many benefits including the ability to continue with the complete failure of a system as well as the ability to cross-reference systems for bootstrapping, validity checking, and can also be used for offline data fusion.

The HMC's localisation consists of independent vision and laser-based systems. The vision-based localisation system is described in detail in [1]. It conists of a firewire fisheye camera mounted on each of the front mudguards that provide colour images back to an onboard computer (see Figure 3). The images are exposure compensated and edge-features extracted. Edges consist of the outline of major pieces of infrastructure such as sheds and doorways. The resulting edge-map is compared to an *a-priori* map generated offline from surveyed coordinates. The matches are determined probabilistically using a particle filter. The laser-based system uses the four outer lasers on the HMC (see Figure 4) and retro-reflective tape that forms artificial beacons highly visible on the lasers' intensity channel. The beacons have been placed at irregular intervals around the worksite with a maximum separation of 30m. Their locations have been surveyed and recorded in a database that is stored on the HMC. The system compares a sensed beacon constellation with the database to triangulate the vehicle position. A particle filter is also used for this purpose. The accuracy of this system is dependent on the density of sensed beacons and around our site, it is sufficient to allow the large HMC to navigate accurately through narrow doorways and roadways

Fig. 5. Entry to the storage shed where the crucible gets dropped off. Note the clearance between the vehicle and the doorway sides is less than 20cm. The vehicle successfully traverses through the doorway using waypoints which demonstrates the accuracy and repeatability of the localisation, navigation and control systems.

(e.g. Figure 5), some of which have a clearance of 20cm.

The combined localisation system works by use of an arbitration mechanism that compares the output and confidence of the vision and laser localisers. If the primary system has a low confidence or fails, the arbitrator promotes the secondary system to the primary and continues to monitor both for failure and recovery. The localisation output from the arbitrator is sent to the navigation module so any single system failure is transparent to vehicle operations. More details of this system are described in [2].

### *B. Obstacle Detection*

The obstacle detection systems consist of one of the most important safety aspects for any autonomous vehicle. We define an 'obstacle' as a significantly sized object that comes close to, or intersects the vehicle's volumetric trajectory. The volumetric trajectory consists of the bounding volume of the vehicle projected along its planned path. This includes overhangs such as the top of a shed door opening, side obstructions, and objects above a certain size on the ground. It is very difficult or expensive to outfit a vehicle such that it is entirely shrouded by a protective sensor curtain that can detect any object approaching or too close to the vehicle. As a result, the HMC uses 2D and 3D obstacle detection systems. These are supplementary systems that run in parallel and affect the vehicle's operations in different ways. These systems are described next.

*1) 2D Obstacle Detection:* The role of the 2D Obstacle Detection system is to provide a reactive protective envelope around the entire vehicle such that the vehicle will reduce speed and stop as an object approaches. This system is implemented using scanning laser rangefinders located at each corner of the vehicle, mounted approximately 1.4 m from the ground as shown in Figure 4.

These lasers are mounted with a slight downward tilt so they intersect the ground at around 25-30m. This module interacts directly with the hardware interface layer module (HMC Interface) to override any control commands and reduce the vehicle's velocity depending on the range of the object. It has two modes of operation depending on whether the crucible is on or not. When the crucible is on, it is detected in the rear laser scans and consequently, a shaped detection envelope is used instead. In this mode, the vehicle has a blind spot behind the crucible. In typical operations with the crucible on, the vehicle will only reverse when it is dropping off which is less frequent than other operations. However, we are addressing the blind spot issue as part of future work.

A second issue with using planar laser scans is that objects are only detected within the laser plane. Any obstacle above or below the scan is not detected. As a result, the main purpose of this system is to detect people close to the vehicle or nearby infrastructure (e.g. buildings, bollards or parked vehicles). In operation, the vehicle slows when it approaches the obstacle, or the obstacle approaches it until either the object is close enough to warrant the vehicle to halt or it passes. If the object is too close (approximately 50cm), the vehicle will remain stationary until the operator intervenes to remove the object, or drive the HMC around it manually.

*2) 3D Obstacle Detection:* The 3D obstacle detection system's primary purpose is to provide a more thorough analysis of the path in front of the vehicle. It consists of a system using a laser mounted above the cabin. The laser has a horizontal scan plane that intercepts the ground approximately 25m in front of the vehicle. This allows approximately eight seconds for the vehicle to come to a halt if travelling at high velocities around 3.0 m/s. An obstacle is determined as an object higher than approximately 5cm that lies in the path of the vehicle. The path is determined from the vehicle's current position past the next waypoint. The system works by accumulating scans as the vehicle travels. The ground plane is extracted from these scans and any object projecting from it identified as traversable or not. If it is not, the system sends a signal to the hardware interface to stop the vehicle and signal that an obstacle has been encountered. This signal consists of a flashing light on the vehicle's status light stack and sending a message through the software system. The vehicle remains halted until the object is removed and the status cleared by the operator via the safety remote. Manual, rather than automatic clearing of the status is a safety issue since in general, the vehicle's path should be clear and any unexpected object detected may indicate a problem in that area of the worksite.

### *C. Mission Controller*

The high level mission controller directs the navigation, tasking and path planning components as shown in Figure 6. The Mission Controller is responsible for switching between tasks and monitoring their performance. A task may be "drive along a section of road", "drop off the crucible", "start up the engine" or even "blow the horn". Currently 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 sub-task for that task. For example, a missed

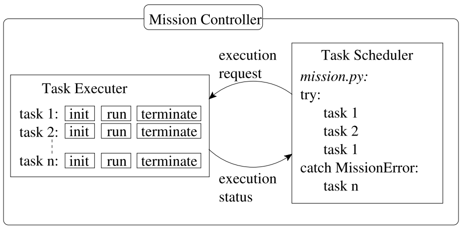


Fig. 6. Overview of the mission control architecture.

Fig. 7. The visual fiducials used to uniquely identify the crucible in the environment.

crucible pick up will trigger a “missed approach” signal and the HMC will move away from the crucible and retry the approach manoeuvre.

The mission controller is a generic component of our system: only the task implementations are specific to the HMC. For this reason, it is currently used on several of our platforms, including an autonomous submarine [3].

### D. Human-Robot Interface

The ultimate goal of an autonomous industrial vehicle is for it to be dependable enough to conduct tasks out of sight of an operator. To allow this, the vehicle needs to have some level of offboard control and the ability to report status and sensor data to a safety supervisor who may be monitoring several vehicles simultaneously.

The most basic level of offboard control consists of a remote e-stop that can be manually or automatically triggered. At more advanced levels, the vehicle may be fully controlled offboard by either a computer or physical interface (manual control panel and joystick), with full sensor displays, allowing immersive tele-operation.

Our system consists of a small remote RF portable control unit that has an e-stop, several programmable function switches and a range of approximately 150m. The unit sends a heartbeat signal out periodically which is received by the onboard RF receiver which is hardwired into the e-stop safety PLC circuitry on the vehicle. If a signal is not received within several milliseconds, an e-stop is initiated on the vehicle. We have programmed the switches to perform the functions of halting the vehicle, sounding the horn, and resetting from a ’detected obstacle’ event. The halt function forces the vehicle to stop moving and freezes all controls. Upon release, the vehicle will continue from that state. This function is particularly useful when testing.

The vehicle also outputs data from its internal sensors (e.g. engine parameters, mast information, brakes etc.) and external sensors (lasers and cameras) for external viewing. Visualisation software allows the safety supervisor to monitor all systems on the vehicle.

### E. Object Detection

Object detection in the system consists of detecting the crucible for pickup operations, and offboard detection and classification of dynamic objects in the environment. The pickup system is based on visually recognising the crucible in the environment [4]. Due to the similarity of the crucible’s round profile with other objects in the worksite, such as drums, the crucible is uniquely marked with self-similar landmarks as shown in Figure 7. Cameras are mounted on the mast of the vehicle looking rearwards for crucible detection.

The system has different modes of operation depending on whether the crucible’s location is known or not. If it is, which would be the case if the location was recorded when it was dropped off, the cameras are directed to locate the markers on the handle. Once positively identified, the relative location of the crucible is calculated with respect to the hook on the HMC. The vehicle then visually servos to the pickup point on the crucible where the remainder of the pickup procedure is managed as a task in the mission controller. If the location is known only approximately within a 20 by 20 m area, the system will execute a distributed search plan for the cameras to locate it. Once they have, a normal visual servo ensues. This is known as a ’long range’ pickup.

The offboard system is in its preliminary stages at present and consists of a static webcam monitoring one of the common areas for HMC operations. The purpose of this system is to track and classify objects in the scene to provide the HMC with greater situational awareness and offboard localisation ability. The system is based on [5], with enhancements to the classification part of the system. Basically, the system consists of:
1. Determining the background image
2. Performing background subtraction to highlight moving parts of the image
3. Merging proximally close moving parts into single blobs and tracking the blobs
4. Classifying the blobs as either ’vehicle’ or ’person/group’.

The system is capable of tracking and identifying multiple various dynamic objects in a scene, in sunlight and rain. It can handle objects being temporarily occluded or objects crossing paths. Examples of classification are shown below in Figure 8.

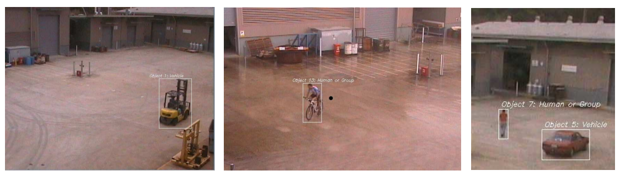

Fig. 8. An example of the various types of dynamic objects tracked. From left to right - forklift, cyclist in the rain, and a person after egressing a car.

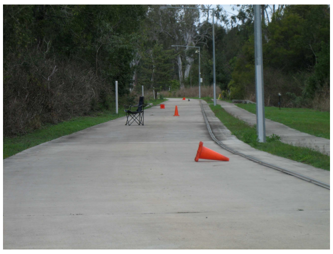

Fig. 9. Traffic cones and chairs used to test the 3D obstacle detection system.

### F. Other Modules

There are many other modules that complete the HMC's systems. These include the hardware interface, various safety systems including physical interlocks and heartbeat checks, navigation and crucible manipulation. The operation of the low level interface and safety modules are beyond the scope of this paper. The navigation module is based on waypoint traverses through pre-programmed path segments. The segments are stored in a mission database file and selected as part of the mission script. The inter-operation of the mission controller and navigation system is basic but effective. Picking up and dropping off the crucible are also basic programmed operations that do not vary once the parameters encoding the vehicle position versus the mast motion are tuned.

## III. Experiments

### A. Obstacle Detection

This consisted of testing the 2D and 3D obstacle detection systems.

*2D Obstacle Detection:* The 2D tests consist of placing a tall object in the path of the vehicle during forward and reverse manoeuvres, for each corner laser. In each case, the vehicle would slow to a stop as it approached the object. The system was also tested with people walking towards the HMC from various peripheral locations. The HMC perfomed as expected by slowing to a halt as the person approached.

*3D Obstacle Detection:* The 3D obstacle detection system was tested with a variety of obstacle shapes and sizes along different trajectories of the HMC. Example objects are shown in Figure 9. These objects were placed in the HMC's path during a prolonged experiment. The obstacle detection system correctly determined that each object was an obstacle which would then halt the vehicle when it was within approximately 15 m. The safety supervisor removed the object and reset the 'obstacle detected' system via a switch on the safety remote. The vehicle continued until then next object was found. A screenshot visualising the data on detection of a non-traversable object is shown in Figure 10.

Tests were also conducted with smaller objects consisting of chunks of concrete which were considered traversable by a human operator. In these cases, the vehicle would continue over them.

Fig. 10. A 2D visualisation of an 'obstacle detected' event in the 3D obstacle detection system. The HMC is the yellow object with the grey crucible attached behind it (left). The HMC's path is shown as the black line projected to the right. Environment features are shown in black and pink with the groundplane as the green dots. Along the vehicle's projected path is a red object (traffic cone). Since this object occurs within the width of the vehicle along its path, it is considered an obstacle.

### B. Redundant Localisation

The redundant localisation system was tested around our main workarea as shown by the blue square in Figure 11. The area is surrounded by buildings which is well-suited to the vision-based localisation method described previously. The main experiment involved simulating a power failure in the primary laser-based localisation system which reduced its confidence values ([1]). Upon detecting this, the arbitrator switched the primary localisation source to the vision-based localiser and the vehicle continued operations. The laser-beacon localisation system was then brought back online and since it produces slightly higher accuracy and therefore is considered as the primary localisation source, the arbitrator switched back to using its outputs.

### C. Long Duration Experiments

Three significant long duration experiments have been undertaken in the project to date. They consist of a two, five, and eight hour trial with the HMC conducting typical operations.


Fig. 11. The path (yellow) of the 2 hour experiment. The crucible pick up and drop off occurred in the open area at the end of the path on the left and the in-shed operations were conducted in the large shed on the upper right. The 5 hour experiment was conducted in the large area surrounded by buildings annotated by the blue box.

*Five Hour Trial:* the purpose of this trial was to test the integrity of all hardware and software systems continuously operating over five hours. The experiment was conducted in the area indicated by the blue square shown in Figure 11.

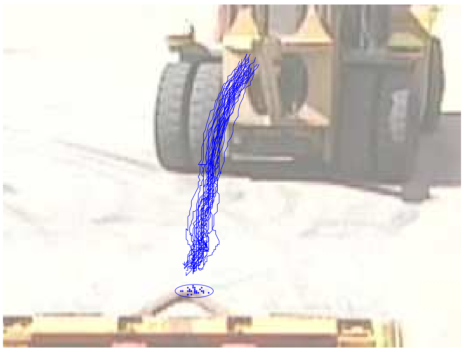

Fig. 12. Transposition of hook path for 29 crucible pickups undertaken at one of the pickup locations during the five hour trial. For reference, the width of the pickup point on the handle of the crucible is approximately 20cm.

The HMC's task was to pick up and drop off the crucible at opposite ends of that area with navigation loops in between tasks. The vision-based crucible detection system was used for locating and servoing to the crucible during pickups. The vehicle undertook the five hour test with the only halt being when the battery on the safety remote had to be replaced. This triggered an e-stop on the vehicle which was then reset and it continued on from where it stopped in the mission. Statistics from this test are shown in Table I. While it is difficult to determine the accuracy of the vehicle and crucible localisation systems due to the lack of a reliable ground truth (GPS is ineffective around built environments mainly due to multi-pathing), upon analysing the log files recorded during the test showed a maximum path spread of 0.3m over all paths with the average being less than 0.2m. The accuracy of the crucible pickups occurring at one end of the test area, which represent the accuracy of the vision-based crucible recognition system is shown in Figure 12.

*Two Hour Trial:* from the success of the systems tested in the five hour experiment, a trial was conducted with a longer traverse path along a narrow road and a crucible dropoff point inside a shed with a narrow entry and filled with equipment. This main path is shown in yellow in Figure 11 and the shed entry in Figure 5.

Three techniques for locating the crucible were tested. Two were vision-based as described in Section II-E. The third was based on servoing to the dropoff location of the crucible recorded from the laser-beacon localisation system. This provided a test of the accuracy of the laser-beacon localiser since any error in location would result in the HMC trying to pick the crucible up from the wrong location.

The mission script required the HMC to autonomously start up, traverse to the crucible scan location and conduct a 'long range' visual pickup. It would then traverse to the storage shed, drop the crucible off inside, drive out and conduct a laser-beacon localiser pickup in the shed. It would then traverse back to the start location, drop the crucible off, navigate around the area to a point where it would conduct a normal vision-based crucible pickup. This cycle to and from the storage shed constituted the remainder of the mission until the last cycle where the crucible was placed in its 'home' location and the HMC parked in its shed and shut down. All phases of the trial were conducted successfully. More details about the five and two hour experiments can be found in [4].

TABLE I

KEY STATISTICS FROM THE 5 AND 2 HOUR EXPERIMENTS

| Experiment | Total Dist. | Cycle Dist. | Velocity Range | Cruc. Ops. |
|---|---|---|---|---|
| 5 hr | 8.5 km | 0.3 km | -1.1 : 1.6 m/s | 58 |
| 2 hr | 6.5 km | 0.93 km | -1.4 : 3.0 m/s | 14 |

*Eight Hour Trial:* The purpose of this trial was to test automated door control, 3D obstacle detection and vehicle scheduling over a shift of normal vehicle operations. The mission was written such that every hour, the vehicle would signal an operator to enable the physical safety interlocks to allow it to conduct a task sequence. The sequence consisted of starting up in its shed and requesting the shed door to open via wireless communication to a receiver on the door built specifically for the purpose. Once the door signalled it was open, the HMC would move out, request the door to close and conduct the crucible pickup - navigation - dropoff cycle described in the two hour trial. Upon completing the 20 minute cycle, it would request the door to open, drive in and park with a final request to close the door. All operations were conducted successfully during the eight hours.

## IV. DISCUSSION

It is important for autonomous vehicles operating in environments with large amounts of infrastructure and in the presence of dynamic objects to be able to conduct repeatable, safe, predictable and reliable operations. Dynamic objects can manifest as people or vehicles moving about the environment, sometimes within close proximity to the robot. To provide the required dependable operations, the vehicle should have redundant self-monitoring systems that are fault-tolerant and where possible have redundant backups. Outside the vehicle, it needs to be 'situationally aware' of its surroundings with respect to its task. Local observations taken from environment sensors such as lasers may be insufficient to determine potential collisions with unseen dynamic objects. Offboard systems such as webcams mounted to infrastructure, or even the perception from other mobile bases can be used to augment this extra sensing.

We are in the process of providing these functionalities with the autonomous Hot Metal Carrier project. Many of the systems described in this paper have been designed to accommodate these requirements. In particular, the localisation, obstacle detection and object recognition systems. While the object recognition system is currently offboard the vehicle, it is capable of tracking and localising dynamic objects to report back to the HMC. We are currently undertaking experiments to demonstrate this utility. While the HMC consists of several basic systems, it has been successfully conducting autonomous operations over hundreds of hours of demonstrations and tests. The fundamental systems have proven reliable, but need to facilitate the redundancy and situational awareness capbilities mentioned above.

*Acknowledgements*

This work was funded in part by CSIRO's Light Metals Flagship project and by the CSIRO ICT Centre's Dependable Field Robotics projects. The author gratefully acknowledge the key contributions of the Hot Metal Carrier project team; in particular: Cedric Pradalier, Robert Zlot, Polly Alexander, Paul Flick, and Jonathan Roberts. The author would also like to thank the Engineering Support team for their maintenance efforts.

## REFERENCES

[1] S. Nuske, J. Roberts, and G. Wyeth, "Outdoor visual localisation in industrial building environments," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Pasadena, U.S.A, May 2008.

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

[3] A. Negre, C. Pradalier, and M. Dunbabbin, "Robust vision-based underwater target identification & homing using self-similar landmarks," in *Proceedings of International Conference on Field and Service Robotics*, Chamony, France, 2007.

[4] 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, April-May 2008.

[5] O. Javed and M. Shah, *Computer Vision - ECCV 2002*, ser. Lecture Notes in Computer Science. Spinger Berlin/Heidelberg, 2002, vol. 2353/2002, ch. Tracking And Object Classification For Automated Surveillance, pp. 439–443.