Redundant Sensing for Localisation in Outdoor Industrial Environments

Jonathan Roberts, Ashley Tews and Stephen Nuske
Autonomous Systems Laboratory
CSIRO ICT Centre
PO Box 883, Kenmore, Qld 4069, AUSTRALIA
Email: firstname.lastname@csiro.au

***Abstract*— We describe our experiences with automating a large fork-lift type vehicle that operates outdoors and in all weather. In particular, we focus on the use of independent and robust localisation systems for reliable navigation around the worksite. Two localisation systems are briefly described. The first is based on laser range finders and retro-reflective beacons, and the second uses a two camera vision system to estimate the vehicle's pose relative to a known model of the surrounding buildings. We show the results from an experiment where the 20 tonne experimental vehicle, an autonomous Hot Metal Carrier, was conducting autonomous operations and one of the localisation systems was deliberately made to fail.**

I. Introduction

Heavy industries that use large ground vehicles for material transport are beginning to explore the use of automation and have recently begun to automate some of their vehicles. Steelworks and Aluminium smelters typically contain small fleets of large vehicles that are used to move bulk products around large work sites (typically hundreds of metres at a time, sometimes kilometres). Such Automated Ground Vehicles (AGVs) must be highly reliable, both in a mechanical and performance sense. It is clear that AGVs operating in the heavy industrial applications described above must be dependable and ideally should be capable of continued operation in the presence of a partial failure of a localisation system.

The first AGVs to appear in these application areas have been confined to areas of the operation where people and other vehicles are prohibited. The vehicles move slowly and the routes contain significant infrastructure for guidance such as buried wires in the concrete, painted lines, etc. These vehicles can be thought of a trains without tracks. The next stage of AGV development for these applications is for vehicles that can travel at higher speeds, in and out of buildings, along roadways and potentially operate with other vehicles. Operational constraints will require that these vehicles may have to operate in areas where personnel are working, or at least transiting (in vehicles or on foot).

Over the past four years, our team has been developing robust localisation techniques for a class of heavy vehicles used in the aluminium smelting industry. Hot Metal Carriers (HMCs) are large vehicles used to transport molten aluminium from the smelter (where the aluminium is made) to the casting shed where it is turned into block products. They operate 24 hours a day, seven days a week. HMCs are large (approximately 20 tonnes unloaded) forklift type vehicles except they have a dedicated hook for manipulating the load rather than fork tines (Figure 1). The aluminium is carried in large metal crucibles which weigh approximately 2 tonnes and hold 8 tonnes of molten aluminium (usually super-heated to 700 degrees Celsius).

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

The operating environment of an HMC presents many challenges. The vehicles must travel inside and outside buildings, day and night and in all weather. Inside, there is a vast amount of infrastructure, other mobile machines and people. In the area immediately around the smelter cells, there are large magnetic fields and high temperatures. Outside, the vehicle's path may be surrounded by infrastructure such as buildings and fences, and their operation may be effected by the weather: 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.

This paper will describe our experiences with two independent localisation systems, one using a laser scanner and the other using cameras, both developed for the target HMC application. A navigation system is developed to take input from the multiple localisation systems, and compare and arbitrate the input information to decide the appropriate navigation actions. The system is designed to combine these independent and unrelated localisation systems to give redundancy, which provides improved levels

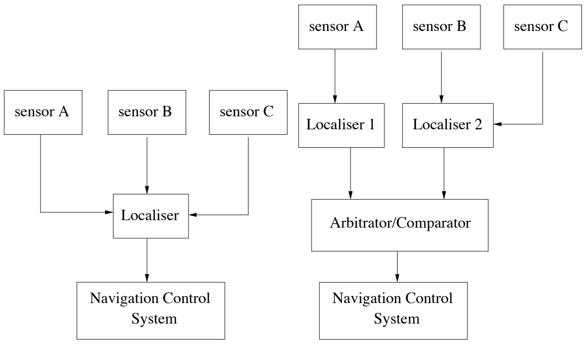

(a) Multi-Sensor Data Fusion (b) Redundant Localisation Systems

Fig. 2. Diagrams depicting the difference between navigation using multi-sensor data fusion and navigation using independent localisation systems.

of dependability in operation. Results from experiments demonstrate continued performance of the vehicle when one localisation system degrades.

II. Related work

The area of dependability in outdoor (terrestrial) field robots did not gain significant attention until the early-to-mid 2000's when the DARPA Grand Challenge events were held. See the Journal of Field Robotics Special Issue on the DARPA Grand Challenge ([1] and [2]) for a comprehensive set of papers by some of the successful and unsuccessful teams. The more recent DARPA Urban Challenge again focused teams into the area of dependability. However, the research results from this event have yet to be published and it is unclear as to whether any teams had redundant localisation systems. The first Grand Challenges have relied heavily on the use of GPS — which is something that cannot be utilised much of the time in our application of interest. In our environments, which are often indoors or in so-called urban canyons, there is little satellite coverage and GPS signal is not received on-board the vehicle.

The use of multiple sensors for localisation has been well researched and has been widely applied in the area of field robotics. For the most part, multiple sensor information is fused to form a single localisation system. This approach can improve the situation where the sensors individually cannot provide enough information for continuous and/or reliable localisation. In multi-sensor data fusion (Figure 2(a)), the aim is to provide a single localisation system a more complete set of input sensor data by fusing all available sensor information. However, when sensors fail, provide erroneous readings or have a limited view of the world, the accuracy and confidence of the localisation estimates degrade. Hence, the data fusion process is not focused on providing redundancy. Examples of sensor fusion in the literature are Majumder et al. [3] who fuse sonar and camera information for an underwater vehicle, Miura et al. [4] fuse laser and stereo camera data into an obstacle map and Arras and Tomatis [5] fuse tracked features extracted from laser and camera data into a single EKF.

In our work we use lasers and cameras in outdoor environments, however most previous laser and camera systems were developed for indoor environments. Newman et al. [6] is one of the only examples of outdoor localisation using both laser and a camera. They use these sensors in a single localisation system, whereas our work presents two individual and unrelated localisation systems.

Examples of redundant sensing are high-integrity inertial sensing with pairs of inertial sensors to achieve high levels of reliability [7]. In this case the sensors are duplicated and the sensor readings themselves compared (i.e. they are not completely unrelated sensors). Scheding et al. [8] use multiple redundant sensors, a laser and a gyro to identify system faults. They assert that the probability of identical sensor fault modes is much lower using sensors with different physical principals, as opposed to using multiple of the same sensor. In their work the only sensor that can perform localisation is the laser, the gyro is just measuring motion and detecting faults. A similar technique is used in standard GPS processing engines that use more than the required minimum number of satellites to obtain a reliable position estimate.

The use of multiple, and often independent sensing and control systems has been widely used by spacecraft engineers since the beginning of human spaceflight. [9] describes the Saturn V guidance and control system that used complete subsystem duplication in many of its operations to achieve the required reliability. Similarly, the Space Shuttle exploits four primary computers at the heart of its fly-by-wire control system [10], [11]. We believe that it will require similar practices to achieve the required reliability for certain field robotics applications - especially those of heavy machinery operating in human populated environments.

III. Proposed System

The system presented in this paper uses multiple sensors in an alternative and more dependable manner. The unrelated sensors are used by independent localisation systems, which provide redundancy to the navigation system. To the authors' knowledge, the use of multiple sensors for multiple-independent-localisation systems has rarely been investigated in the area of field robotics research. Figure 2 shows the fundamental difference in this approach. A system using independent localisation systems (Figure 2(b)), uses an additional process - an arbitrator or comparator - to monitor the pose estimates from the multiple localisation systems and cross-checks them for consistency. It is only in very recent times that field roboticists have had the ability to compare pose estimates from independent localisation systems as until now it has been difficult to deploy more than one working localisation system on a field robot.

We have now developed two high reliability localisation systems that are optimized to work in large outdoor industrial environments. One is a system based on the use of multiple 2D laser scanners and reflective beacons. The

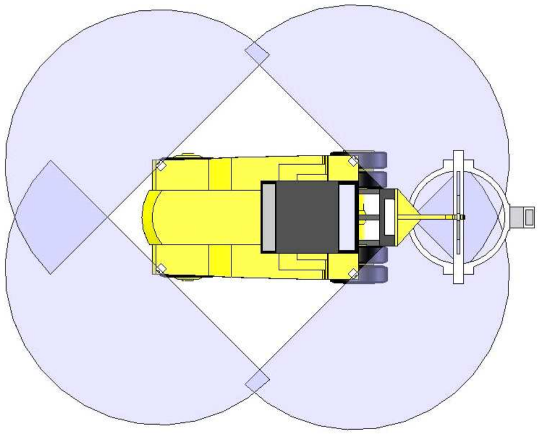

(a) Laser setup and coverage

(b) The industrial environment

Fig. 3. Laser localisation system. Four lasers are placed at each corner of the vehicle and (a) demonstrates the coverage of the laser scans. (b) is an image of the industrial setting, reflective strips on the posts and walls of the site are detected by the laser scanners.

other uses a vision system to estimate the vehicle's pose based on an *a priori* edge map of the buildings in the environment. Both these systems have been operating on the autonomous HMC and both can be used to guide the HMC around our test site. The remainder of this paper describes the two localisation systems and shows the results from experiments where one of the localisation systems (the laser-based system) was disabled.

A. Laser Localisation

Our laser localisation system, previously published in [12], is comprised of four laser rangefinders placed on the four corners of the vehicle (Figure 3(a)). The lasers detect reflective beacons that are placed around the environment on the posts and walls at surveyed locations (Figure 3(b)). The beacons' locations are used to triangulate the vehicle's position to a site-referenced (global) coordinate system when detected.

B. Camera Localisation

Our vision-based localisation system, appearing in [13] uses two fish-eye cameras mounted sideways on the vehicle (Figure 4(a)). A sparse 3D-edge-map of the building environment (Figure 4(c)) can be tracked in the camera images

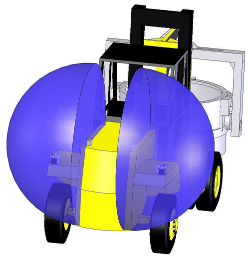

(a) Camera Setup

(b) Fish-eye image

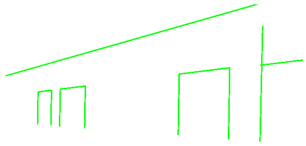

(c) 3D-edge-map of buildings

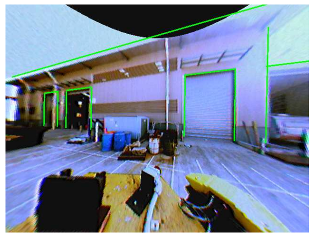

(d) Un-distorted image with projected 3D-edge-map

Fig. 4. Examples of the vision-based localisation system. Two fish-eye cameras are placed at the front of the vehicle facing sideways (a). 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).

giving the pose of the vehicle. The 3D-edge-map tracking is facilitated in a particle filter and processed on a standard GPU (Graphics Processing Unit). The incoming fish-eye images (Figure 4(b)) are first corrected for distortion (Figure 4(d)) and then passed through an edge filter. The 3D-edge-map can then be projected onto the undistorted edge images for direct comparison.

The comparison score is calculated as the alignment between the 3D-edge-map and the camera edge-image and is computed for every particle by the GPU. This comparison score gives an indication of the likelihood a particle is at or near the correct pose estimate and is used by the filter to re-sample the particles each iteration.

A confidence measure of whether the particle filter is still correctly tracking the buildings is calculated as the mean alignment score of the best 5% (with the highest likelihood) of the particles. This confidence is used by the vehicle's navigation system for decisions regarding when and how to use the vision-based localisation.

C. Navigation System

The vision and laser localisers are two independent systems that are each able to provide the inputs for navigation of the vehicle. However when combined together there is redundancy in localisation. An independent process is used - an arbitrator or comparator - which accepts these two inputs, evaluates a confidence in each system and determines the appropriate pose estimate for the navigation system (Figure 2(b)). We propose four modes of vehicle operation post-failure of the localisation system:

1) Termination of operation to an immediate safe state (fail-safe behaviour).

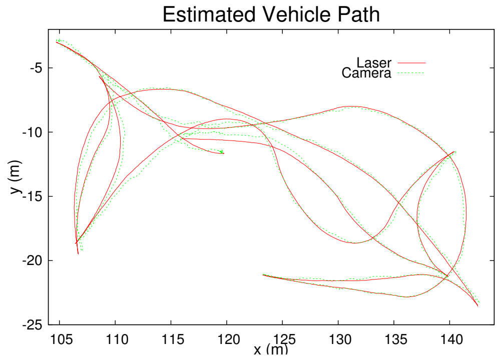

(a) Path Discrepancy

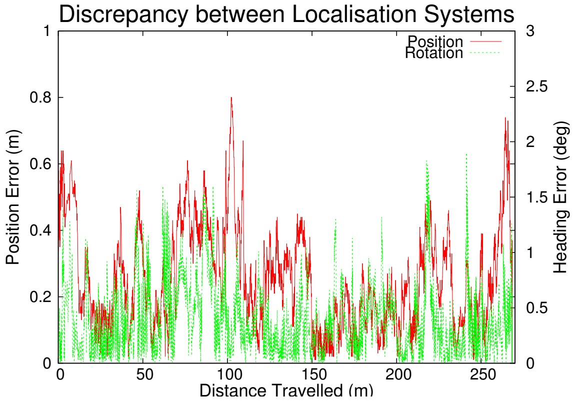

(b) Position and Rotation Discrepancy

Fig. 5. Discrepancy between the laser and camera localisation systems, recorded while the vehicle drives around the site.

2) Termination of operation where the vehicle defaults to "limp home" type navigation after which it can be investigated and repaired.
3) Continued operation with a degradation in operational performance (e.g. slower speed operation).
4) Continued operation with no performance loss.

Ultimately, Mode 4, is the target of the research outlined in this paper, where vehicles can continue to operate, even after one localisation system fails. The system failure would then be repaired by a maintenance crew at the next available opportunity. However, even the development of Mode 1 is a challenge as this requires that the AGV system correctly detects the localisation system failure. Apart from a partial or complete sensor failure, it can be difficult detecting when a single localisation system becomes inaccurate. In particular cases where a localisation system's pose estimate slowly drifts from the correct solution, a second, and independent localisation system is required for comparison. This sort of functionality is much better performed with multiple localisation systems.

D. Localisation Arbitration

The autonomous HMC's primary global localiser is the laser-based system. It provides accuracies within 100mm for navigation and crucible operations - docking and drop-off. The performance of the vision system has been compared with the laser system as seen in Figure 5. The figure

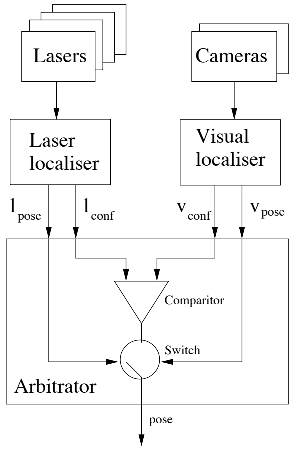

Fig. 6. A very simple arbitrator used for the experiment.

shows the two separate systems have similar outputs which would each be a suitable basis for navigation. Both systems provide internal estimates of their confidence of operations which we have determined to be reasonable metrics that reflect the system's accuracy. To provide redundancy and reliability in localisation, an arbitration module is used to provide the most accurate pose estimate by comparing the systems and making decisions about which is providing the highest confidence and most accurate estimates to pass through to the navigation system. Furthermore, the arbitrator also passes through the confidence value which the navigation system can use as a dynamic guide for setting the upper limits for velocity control - if the confidence in localisation is low, then the vehicle's maximum forward and reverse speeds should be reduced (Mode 3 in Section III-C).

The input parameters for the arbitrator are the vision and laser pose estimates (v_{pose} and l_{pose}), and their confidence measures (v_{conf} and l_{conf}). Currently, the arbitrator will always choose l_{pose} and l_{conf} as output values unless either of the following cases occur:

1) l_{conf} and v_{conf} are low
2) l_{conf} is low and v_{conf} is high

The choice of low and high thresholds for these evaluations are currently empirically determined based on previous testing of individual systems. In case 1, the navigation system will slow the vehicle to a stop since it has assumed inaccurate localisation from all available sources (Mode 1 in Section III-C). In case 2, the arbitrator will switch to the visual localiser and use its confidence and pose estimate as an output (Mode 4 in Section III-C).

IV. Results

To test our idea we devised an experimental trial in which the HMC was tasked to perform a normal crucible pickup, transit and drop off. The trials were run outside in a compound area - surrounded by large industrial sheds as shown in Figure 7(a). The mission of the HMC was to:

1) From a parked position, drive to the crucible position (known from a previous autonomous mission).

2) Pick up the crucible.
3) While carrying the crucible, complete a circuit of the compound area.
4) Return to the crucible pick-up position and drop off the crucible.
5) Return to the parking position.

The failure mode that was tested was a loss of the laser-based localiser triggered by a simulated power failure to the lasers. The failure was timed to occur during the transit phase of the HMC (just after the crucible pick-up). A simple arbitrator was created (Figure 6) that took input from the two localisation systems and output the pose of the system that it trusted most. The output value was then used by the navigation system. It did this by continuously monitoring the confidence values of the localisation systems (l_{conf} and v_{conf}). For this experiment, the arbitrator was programmed to trust the laser localiser more than the vision localiser as long as the laser localiser's confidence was greater than 0.4 (on a scale of 0.0 to 1.0). If the laser localiser's confidence dropped below this threshold then the arbitrator used the vision localiser's output and continued the mission. It should be noted that as each system is independent, then so are the confidence values. Both systems report confidence in the 0.0 to 1.0 range but confidence estimates are not calibrated. The speed of the vehicle changes depending on the confidence value of the arbitrator as follows:

$$\begin{array}{rll} 1.0 <= & aconf >= 0.75, & speed = 100\% \\ 0.75 < & aconf >= 0.5, & speed = 75\% \\ 0.5 < & aconf >= 0.3, & speed = 50\% \\ 0.3 < & aconf >= 0.0, & speed = 0\% \end{array}$$

Figure 7(b) shows the confidence values plotted against time. The initial high confidence values in the figure are derived from the laser localiser. The simulated laser failure occurred at approximately the 80 second point in the figure. At this point the arbitrator switched to the vision localiser. The confidence values from that point onwards are from the vision localiser. It is clear from Figure 7(b) that when the vision localiser takes over, it reports relatively low confidence values between the 80 and 160 second marks. During this phase the HMC was performing a full 360 degree loop around the compound. This was the most difficult section for the vision-system to track the buildings due to the rapid change in vehicle orientation. While the confidence reported by the vision system was low, the navigation system moderated the speed of the vehicle as per the rules indicated above. The HMC had completed the turn at the 160 second point, and from then until the end of the mission the vision localiser reported a high confidence. As a result the HMC was able to complete its mission at full speed, dropping off the crucible and returning to the parking position. Figure 7(c) shows the HMC's path and indicates the location of the parking position, crucible pickup and drop-off point and the location of the simulated laser power failure.

(a) The compound area navigated during the experiment

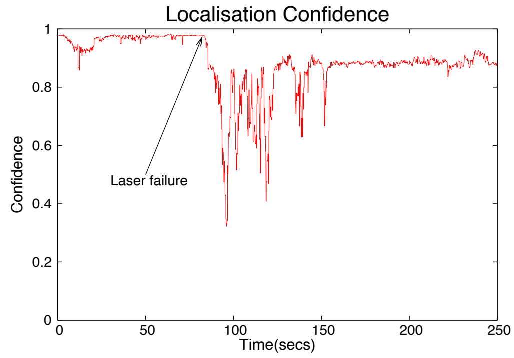

(b) Localiser confidence values during the experiment

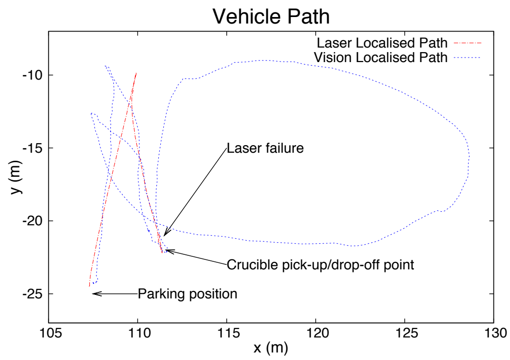

(c) The path of the HMC during the experiment. The vehicle uses the laser-based system until it fails then the remainder of the path is the vision-based system.

Fig. 7. Figures showing the experiment area and experimental results

V. Conclusion

Over the past five years, we have been developing an autonomous navigation system for a heavy duty industrial transport application - that of the movement of molten aluminium around a smelter. We have now developed and implemented a number of independent localisation systems. The idea of sensor fusion in field robotics has been widely exploited over the past decade. However, the motivation for this sensor fusion has often been to achieve the reliable operation of a single localisation system. Algorithms, sensors and computing hardware has now reached a point where it is possible to deploy multiple localisation systems that can work throughout a mobile robot's environment. This finally allows us the opportunity to compare the estimates from these systems and start to investigate the best methods of choosing the most reliable, the most trusted or developing methods to optimally combine them to achieve a more dependable outcome.

VI. Further Work

While the basic localisation arbitrator described in this paper is a relatively simple mechanism for system switching on failure, it represents a fundamental change in the HMC's architecture which has been successfully utilised for hundreds of hours of autonomous operation. We are currently in the process of developing a far more sophisticated arbitrator that is capable of monitoring many sources of localisation data. In the near term we will have pose estimates from seven localisation sources:

- laser-beacon localiser
- vision localiser
- laser scan matching from a SLAM derived site map
- GPS (where it works)
- wheel encoder-based odometry
- laser (scan matching) odometry
- vision-based odometry

The first four forms of localisation information are absolute (and given in the world co-ordinate frame), whereas the last three forms of localisation data are relative in nature (i.e. they drift over time). Note that not all of the above seven sources of localisation data are independent in that some use the same sensor (all the laser scanner-based systems) and some of the absolute system require data from the relative systems to function. The research problem to be addressed is how these estimates are compared in a reliable way. It is hypothesised that the higher the correlation between multiple inputs, the higher the confidence of each systems' performance which can be reflected back on their own performance estimates. This allows a more robust solution for evaluating the different independent inputs. Much work in the area of track-to-track correlation/association has been carried out over the past three decades [14]–[18]. This research has developed ways to correlate aircraft tracks from multiple radar tracking installations. Here, the system must determine which tracks belong to the same aircraft and which are from separate aircraft. The problem also manifests itself in the arena of tracking for missile defence. Here the tracks are analogous to the trajectories of our robots and we investigate that many of the same techniques can be applied to determine how well trajectories from the independent localisers correlate.

Other interesting issues for this application relate to specific areas of the site where one localiser outperforms another. For example, laser-based systems are good when there is infrastructure close by (buildings, etc) and do not perform well in open areas. GPS on the other hand performs well in the open but extremely poorly, or not at all close to or inside buildings. GPS also has the additional problem of reporting high confidence values (GDOP) when it is clearly inaccurate - it believes that it is good when it is not. This problem may exist in other localisation systems and so any arbitrator must contain some sort of 'trust' measure on the individual localisers that it is monitoring and comparing. This may involve some sort of machine learning. It may also include some 'teaching' by a site expert to train the arbitrator where certain localisers can and can not be trusted.

Finally, we will be further developing and testing ideas on how to allow a vehicle with limited confidence from its localisation system to safely and reliably 'limp home' (Mode 2 in Section III-C).

Acknowledgment

This work was funded in part by CSIRO's Light Metals Flagship project and by the CSIRO ICT Centre's Autonomous Ground Vehicles project. The authors gratefully acknowledge the contribution of the rest of the Autonomous Systems Lab's team and in particular: Paul Flick, Polly Alexander, Leslie Overs, Clinton Roy, Mike Bosse, Cedric Pradalier and John Whitham, who have all contributed to the project.

References

[1] K. Iagnemma and M. Buehler, "Special issue on the DARPA Grand Challenge (Part 1)," *Journal of Field Robotics*, vol. 23, no. 8, August 2006.

[2] ——, "Special issue on the DARPA Grand Challenge (Part 2)," *Journal of Field Robotics*, vol. 23, no. 9, September 2006.

[3] S. Majumder, S. Sheding, and H. Durrant-Whyte, "Multisensor data fusion for underwater navigation," *Robotics and Autonomous Systems*, vol. 35, pp. 97–108, 2001.

[4] J. Miura, Y. Negishi, and Y. Shirai, "Mobile robot map generation by integrating omnidirectional stereo and laser range finder," in *International Conference on Intelligent Robots and Systems*, 2002.

[5] K. O. Arras and N. Tomatis, "Improving robustness and precision in mobile robot localization by using laser range finding and monocular vision," in *Third European Workshop on Advanced Mobile Robots*, 1999.

[6] P. Newman, D. Cole, and K. Ho, "Outdoor slam using visual appearance and laser ranging," in *In Proc. International Conference on Robotics and Automation*, 2006.

[7] M. A. Sturza, "Navigation system integrity monitoring using redundant measurement," *Journal of The Institute of Navigation*, vol. 35, no. 4, pp. 69–87, 1988.

[8] S. Scheding, E. Nebot, and H. Durrant-Whyte, "High-integrity navigation: A frequency-domain approach," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 4, pp. 676–694, July 2000.

[9] F. B. More and J. B. White, "Application of redundancy in the saturn v guidance and control system," NASA Marshall Space Flight Centre, Tech. Rep. NASA-TM-X-73352, 1967.

[10] C. E. Price, "Fault tolerant avionics for the space shuttle fault tolerant avionics for the space shuttle," in *IEEE/AIAA Digital Avionics Systems Conference*, Los Angeles, CA, October 1991, pp. 203–206.

[11] J. R. Sklaroff, "Redundancy management technique for space shuttle computers," *IBM Journal of Research and Development*, vol. 20, no. 1, pp. 20–28, 1976.

[12] A. Tews, C. Pradalier, and J. Roberts, "Autonomous hot metal carrier," in *Proceedings of IEEE International Conference on Robotics and Automation*, Rome, Italy, Apr. 2007, pp. 1176–1182.

[13] 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 (to appear)*, Pasadena, U.S.A, May 2008.

[14] R. Singer and A. J. Kanyuck, "Computer control of multiple site track correlation," *Automatica*, vol. 7, pp. 455–464, 1971.

[15] Y. Bar-Shalom, "On the track-to-track correlation problem," *IEEE Transactions on Automatic Control*, vol. 26, no. 2, pp. 571–572, 1981.

[16] F. R. Castella, "Theoretical performance of a multisensor track-to-track correlation technqiue," *IEE Proceedings Radar, Sonar and Navigation*, vol. 142, no. 6, pp. 281–285, December 1995.

[17] B. F. La Scala and A. Farina, "Choosing a track association method," *Information Fusion*, vol. 3, pp. 119–133, 2002.

[18] D. E. Maurer, "Information handover for track-to-track correlation," *Information Fusion*, vol. 4, pp. 281–295, 2003.