An Intelligent Active Range Sensor for Vehicle Guidance: System Overview

N. E. Pears
Department of Engineering
Cambridge University
Cambridge, CB1 1PZ
email: nep@eng.cam.ac.uk

Abstract

The mechatronic design and the two stage calibration procedure for an eyesafe laser rangefinder, based on the lateral-effect photodiode (LEP), are presented. The sensor acquires two-dimensional range data, and is active in the sense that it can change the orientation of its field of view in order to track useful range features. An analysis of LEP operation shows that image position measurement repeatability, normalised with respect to the detector half length, is equal to the signal current to noise current ratio. This result allows accurate estimation of the variance of individual range measurements, making the sensor particularly amenable to statistically based range feature detection. Range data acquisition, range feature extraction and control of the active head behaviour are all implemented on a local network of six transputers. This parallel structure is described and it is shown how the sensor constitutes an intelligent agent in a balanced sensor suite for the guidance of close range mobile robot manœuvres.

1 Introduction

Mobile robot research aims to achieve vehicle navigation in increasingly unstructured environments with increasing levels of autonomy, and with a strong emphasis on robustness within that autonomy. It has become evident that this requires the coupling of a broad range of real-time sensor based activities which operate over a variety time scales; for example, continuous map updating, occasional path planning and vigilant, reactive obstacle avoidance. This requires a multiprocessor architecture that is real-time, modular, expandable and loosely coupled. Transputer technology is well suited to these requirements, and underpins the Locally Intelligent Control Agent (LICA) architecture developed in the Oxford Autonomous Guided Vehicle (AGV) project [3]. Sensor processing within the LICA architecture is configured so that one or more processing units are dedicated to each data stream entering the system so that only the salient features and associated uncertainties of that data need to be presented to the rest of the robotic system. Some of the sensors creating that data streams are active, in the sense that sensor orientation can be changed in order to extract the maximum amount of information relevant to the current task. Since the information extracted locally can guide orientation changes for a particular sensor, that sensor and its associated processing units constitute a modular device termed an *intelligent active sensor.*

The aim of this paper is to present a system overview, and some more specific ideas on calibration and range variance estimation, for an intelligent active range sensor, which is used to guide close range mobile robot manœuvres such as parking and docking. In section 2, a number of range measurement techniques are reviewed and, in the following section, the sensor design and performance are summarised. Section 4 describes the two stage calibration procedure, whilst section 5 derives the relationship between image position measurement repeatability and the signal current to noise current ratio on the detector. This is required to estimate the variance associated with a range measurement. Section 6 describes how the rangefinder operates as a conceptual intelligent sensor, and outlines the algorithms that are implemented within its local processing structure. The final section before the conclusions provides a discussion of the operation of the sensor and its role in a balanced on-board sensor suite.

2 A review of range measurement techniques

Passive vision provides the most comprehensive source of sensory information from a single robotic sensing device; but the projection of a 3-D scene into a 2-D image engenders ambiguity, making range recovery difficult. Several different cues have been exploited to recover range such as stereo disparity [9], and visual motion [4]. However, the use of passive vision is often precluded, since the large bandwidth of raw data requires fast, potentially expensive processing hardware to extract the required range information. Sensors designed specifically for range measurement can often provide task-specific range information at a sufficiently high rate for the real-time performance requirements more cheaply. Here, we will compare optical radar, sonar, and optical triangulation. Comprehensive reviews of ranging techniques can be found in Jarvis [12] and Everett [5].

In optical radar, the phase relationship between an amplitude modulated light beam (laser, or collimated LED) and its reflection is used to calculate range [2]. In sonar, the speed of sound is sufficiently low for the round trip of pulses to be timed [10]. Optical triangulation is a geometric means of ranging and involves projecting a light source onto the scene and observing the image position of that projection with a lens/detector combination offset from the axis of projection. Scene coverage can be achieved either by scanning a spot or line stripe [13], or by projecting a pattern of dots or lines [1] on to the area of interest.

Sonar has a number of disadvantages when compared to optical methods. Firstly, the low speed of sound means that range measurements can only be made at a low rate (10-100Hz, depending on range). Secondly, its large wavelength means that it suffers from specularities. This means that the spatial density of measurements is low, since only certain parts of the scene can be detected, such as cylinders, corners, and planar surfaces which are oriented within the effective beamwidth of the ultrasonic pulse. Thirdly, specularity in combination with a large beamwidth engenders ambiguity in the angle of the range measurement. For these reasons, we believe that sonar sensing is unsuitable for guiding close range vehicle manœuvres.

Of the optical methods of active ranging mentioned, optical radar is more practical than optical triangulation in outdoor environments where, at large ranges, triangulation becomes inaccurate. In addition, the missing parts problem, which is a characteristic of any stereo/triangulation system, is eliminated, since the transmitted and received beams are co-axial. However, over short ranges, such systems are generally of lower accuracy when compared with active triangulation systems. For this reason, we elected to design a triangulation system. General discussions of the tradeoffs involved when designing optical triangulation range sensors can be found in Pears and Probert [7].

3 Summary of the sensor design

In optical triangulation schemes, it is apparent that the angle between the projection axis and the optical axis of the lens is one of the parameters that define the accuracy and depth of field in ranging, which is problematic if the light source is to be scanned across the scene. However, if the detection lens is scanned in synchronism with the projected light, so that their angular separation remains constant, then the performance in ranging varies very little with scan angle [8]. The implication for close range vehicle manœuvres, such as obstacle avoidance and docking, is that a large scan angle (large field of view) can be built into the sensor design without compromising depth of field, and vigilance during vehicle movement can be maintained.

An optical means of scanning the laser and lens in exact synchronism is shown in plan view in fig. 1, which is an adaptation of Livingstone and Rioux's [11] configuration. Its physical dimensions are 21cm wide, 16cm deep, and 4.2cm high. Referring to this figure, a collimated and modulated laser beam is projected onto a small mirror (a), which deflects the beam onto the front face of the scanning mirror (b). The

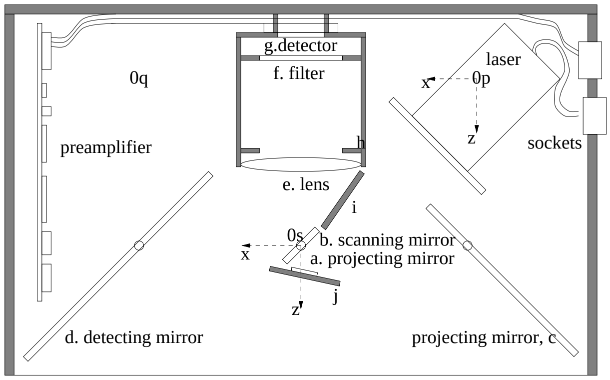

Figure 1: Plan view of the sensor head

scanned beam is then projected into the scene by a larger mirror (c). With this arrangement, the laser is scanned over twice the angle over which the scanning mirror deflects, and the centre of scanning is at the virtual image point O_p in the sensor. (Note that this point defines the origin of the *projection frame* and measurements referenced to this frame are denoted with a superscript p. This frame is displaced from the *sensor frame*, in which calibration measurements and range feature extraction takes place. The origin of this frame, denoted O_s on fig. 1, is coincident with the axis of rotation of both the sensor head and scanning mirror, and measurements referenced to this frame are denoted with a superscript s).

Laser light, scattered by objects in the scene, is collected by the large detection mirror (d) and is deflected onto the rear of the scanning mirror. Light leaving the scanning mirror is focussed by the lens (e) and passes through an optical filter (f), matched to the laser wavelength, before forming an image of the projected spot on the image position measurement device (g). With the geometry described above, the lens is effectively scanned, in exact synchronism with the laser, around virtual image point Oq in the sensor on an arc with radius equal to the separation between the scanning mirror and the lens.

The synchronised scanning optics in fig 1 is termed the *sensor head*. This sensor head is mounted on a servo driven platform which can rotate the field of view of the sensor head between +90 and -90 degrees relative to the forward looking direction. This head drive allows the scanning field of view to be centered at the optimal position for the current manœuvre. For example the sensor may centre on navigable freespace or may turn so that the obstacle being avoided does not leave the scanning field of view.

3.1 Image position measurement

The geometric means of range measurement described above requires the one-dimensional measurement of image position. An analogue means of measurement using the lateral-effect photodiode (LEP) was found to suit our application in preference to the more commonly used linear CCD because of close range accuracy requirements. The device generates a

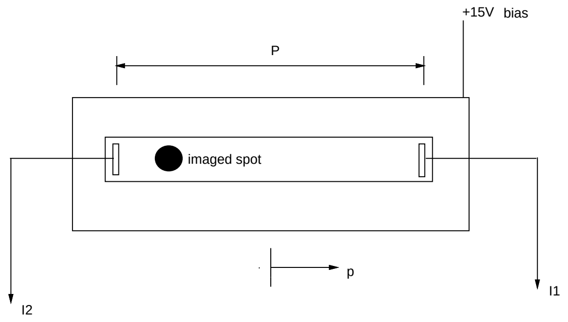

Figure 2: The lateral-effect photodiode

photocurrent at the site of the imaged spot and, since it has uniform resistivity along its length, it acts as a current divider so that the position of the light centroid relative to the centre of the device (see fig. 2) is

$$p = \frac{I_1 - I_2}{I_1 + I_2}\Big(\frac{P}{2}\Big), \qquad \Big\{-\frac{P}{2} \leq p \leq +\frac{P}{2}\Big\} \qquad (1)$$

and the detector current, I_0, is the sum of the terminal currents, I_1 and I_2.

3.2 Sensor performance

Table 1 summarises the specifications of the sensor head and sensor head drive.

Table 1. Sensor specifications

Specification	value
field of view	40 degrees
depth of field	2.1m
stand of distance (min. range)	0.4m
maximum range	2.5m
range measurements per scan	256
range measurement frequency	2.5kHz
bandwidth of detector	1kHz
scan frequency	9.8Hz
laser power	0.9mW
laser wavelength	670nm
laser class	II (max 1mW)
sensor head position resolution	0.36 degree
head response time (90 deg. step)	0.5s

In order to determine ranging repeatability, 1000 readings of image position, measured range, and detector current were taken, at target ranges equally spaced between 0.75m and 2.5m. Standard deviations for all target ranges, and their values as a percentage of the target range, are shown in table 2. In addition, the standard deviation of image position and the average detector current in nanoamps are shown. Note that repeatability, as defined by one standard deviation, is 1mm or better up to a range of 0.89m and is 1cm or better up 1.57m. These results are relatively immune to the effects of ambient lighting as laser modulation and lock-in detection are used in the sensor design [8].

Table 2. Ranging Results

z(m)	σ_z(cm)	% rep.	I_o (nA)	$\Delta p(\mu m)$
0.75	0.05	0.067	48.0	3.81
1.0	0.148	0.148	26.0	6.615
1.25	0.463	0.37	16.0	11.3
1.5	0.798	0.532	11.0	15.745
1.75	1.568	1.1	8.3	22.095
2.0	2.547	1.27	6.4	29.265
2.25	3.713	1.65	5.0	35.28
2.5	7.662	3.06	4.3	45.07

4 Two Stage Calibration

4.1 Depth (z) calibration

In the first stage of the calibration procedure, the AGV is positioned normal to a laboratory wall, so that the origin of the sensor frame is 2.5m from the wall. It then uses its bar code scanner guidance system to move incrementally towards the wall with a predefined step size, taking a scan between each move. Spatial filtering is applied in the direction of the scan angle, θ, by application of the one dimensional convolution operator [1,4,6,4,1]. This approach causes minimal distortion of the look-up table as a model of the ranging process, since image position data is much denser in the θ axis than the z axis (256 data points compared to around 40), and the non-linearity is less severe. After N moves, a two dimensional look up table of dimension 256 angles by N range values has been built. Subsequently, range can be calculated as

$$z^s = T_z(p, \theta) \qquad (2)$$

where T_z represents an interpolation in the calibration look up table.

4.2 Determination of projection vergence and projection origin

If the scan angle, θ, is zero when the laser points along the z^p axis then, from fig. 1, the x^s coordinate of the range measurement can be calculated from the interpolated z^s measurement and the scan angle as

$$x^s = x^s_{0p} + (z^s - z^s_{0p}) \tan\theta \qquad (3)$$

where $[x^s, z^s]^T_{0p}$ is the position of the origin of the projection frame in the sensor frame. The angle of the laser projection is made up of the projection angle due to the scanning mirror orientation, which is known accurately from the scanning mirror drive feedback signal, and a small, unknown vergence angle. This vergence angle represents the fact that the scanning field of view is not be centered on the z^p axis, but is inclined slightly towards the z^s axis, depending on the orientation of the projecting mirrors (a) and (j) in figure 1. Thus, a second stage of calibration is required to determine the projection vergence angle γ and the position of the origin of the projection frame in the sensor frame, $[x^s, z^s]^T_{0p}$, so that equation 3 can be employed to determine x^s.

A calibration target with broad black and white stripes of equal and known arbitrary spacing is placed at an arbitrary range, perpendicular to the z^s axis, as in fig. 3. For the three intensity changes at $i = 0, 1, 2$ on fig. 3, we can subtract pairs of equations of the form 3 above to give

$$x^s_{i+1} - x^s_i = D = (z^s - z^s_{0p})(\tan\theta_{i+1} - \tan\theta_i), \quad \{i = 0, 1\} \tag{4}$$

Equating the right hand side of the pair of equations in 4 gives

$$\tan\theta_2 - 2\tan\theta_1 + \tan\theta_0 = 0 \tag{5}$$

We can then write the absolute angles θ_i as the sum of the scan angle θ_{s_i}, as measured by the scanning mirror feedback signal, and an unknown projected vergence angle, γ, and expand the tangent of a sum in the above equation. Multiplying out and rearranging yields an equation implicit in the projected vergence angle.

$$\alpha\tan^3\gamma + \beta\tan^2\gamma + +\alpha\tan\gamma + \beta = 0 \tag{6}$$

where

$$\alpha = \tan\theta_{s_0}\tan\theta_{s_1} - 2\tan\theta_{s_0}\tan\theta_{s_2} + \tan\theta_{s_1}\tan\theta_{s_2} \tag{7}$$

$$\beta = \tan\theta_{s_0} - 2\tan\theta_{s_1} + \tan\theta_{s_2} \tag{8}$$

Since both β and $\tan^3\gamma$ are small, an estimate of projection vergence angle, accurate to around 0.5 % in simulations ($\gamma = 5$ degrees), is given by

$$\gamma = \tan^{-1}\left(-\frac{\beta}{\alpha}\right) \tag{9}$$

If greater accuracy is required, the above equation can provide the initial estimate in a Newton-Raphson numerical solution to equation 6.

The positioning of the target, shown in fig. 3 is such that the intensity change associated with $i = 0$ is positioned at $x^s = 0$. Firstly, this prevents equation 5 from being poorly conditioned because of too much symmetry (the worst case being when $\theta_1 = 0$ in fig. 3), and secondly it provides a means of determining the origin of the scan axis, since the intensity changes are at known positions in the sensor frame (z^s is known from the sensor itself as it has been calibrated in that dimension, and x^s is known since the stripes are of known width).

The projection angles at which the intensity changes occur are extracted by a standard difference of Gaussian edge detector. Subsequently, the standard least squares solution to the intersection of the five laser projection lines shown in fig. 3 generates the scan origin coordinates,

$$\hat{\mathbf{x}}^s_{0p} = [\hat{x}^s, \hat{z}^s]^T_{0p} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{y} \tag{10}$$

where $\mathbf{H}$ is the stacked measurement matrix

$$\mathbf{H} = \begin{bmatrix} 1 & -\tan\theta_{-2} \\ \vdots & \vdots \\ 1 & -\tan\theta_2 \end{bmatrix} \tag{11}$$

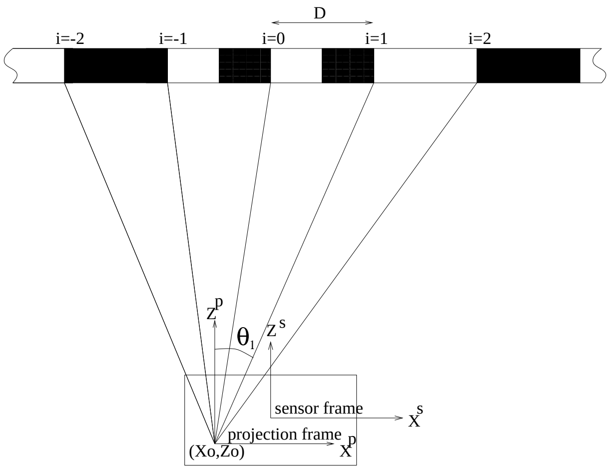

Figure 3: Calibration rig for x_0, z_0, γ

and $\mathbf{y}$ is given by

$$\mathbf{y} = \begin{bmatrix} -2D - z^s_t\tan\theta_{-2} \\ -D - z^s_t\tan\theta_{-1} \\ -z^s_t\tan\theta_0 \\ D - z^s_t\tan\theta_1 \\ 2D - z^s_t\tan\theta_2 \end{bmatrix} \tag{12}$$

where z^s_t is the range at which the calibration target is placed in the sensor frame.

5 Range variance estimation

If image position variance can be computed, it can be scaled by the triangulation gain, which is the magnitude of the local gradient, $|\frac{\partial z}{\partial p}|_{z,\theta}$ in the calibration table T_z, to give an estimate of range variance. This variance information is essential to allow robust algorithms to be applied to the raw range data; in particular, it is used in the EKF algorithm for line segment extraction, outlined in section 6.2. This section establishes the relationship between the variance associated with an image position measurement and the detector current for that measurement.

Rewriting equation 1, the image position on a LEP can be normalised with respect to the detector half-length so that

$$p_n = \frac{p}{\left(\frac{P}{2}\right)} = \frac{I_1 - I_2}{I_1 + I_2}, \qquad \{-1 \le p_n \le +1\} \tag{13}$$

and the uncertainty in normalised position, Δp_n, is given by:

$$\Delta p_n = \left|\frac{\partial p_n}{\partial I_1}\right|\Delta I_1 + \left|\frac{\partial p_n}{\partial I_2}\right|\Delta I_2 \tag{14}$$

where

$$\left|\frac{\partial p_n}{\partial I_1}\right| = \frac{2I_2}{(I_1+I_2)^2} = \frac{2I_2}{I_0^2} \quad (15)$$

and

$$\left|\frac{\partial p_n}{\partial I_2}\right| = \frac{2I_1}{(I_1+I_2)^2} = \frac{2I_1}{I_0^2} \quad (16)$$

Now it can be shown that thermal noise in the one-dimensional LEP and the shot noise due to the dark current are significantly greater than than the shot noise due to the signal current. This latter, relatively minor noise, divides between the LEP terminals according to image position. The dominant noise sources, however, divide equally between the LEP terminals. Thus, ignoring shot noise due to signal current, and denoting the remaining rms noise current by I_n:

$$\Delta I_1 = \Delta I_2 = \frac{I_n}{2} \quad (17)$$

Substituting equations 15, 16 and 17 into equation 14 gives

$$\Delta p_n = \frac{\Delta p}{\left(\frac{P}{2}\right)} = \frac{I_n}{I_0} \quad (18)$$

Thus the image position repeatability is inversely proportional to the signal current to noise current ratio. In order to estimate the constant noise current, I_n, the above equation can be linearised by taking logarithms and the results in table 2 can be used to estimate the constant noise current, I_n, by standard least squares methods.

6 Integration of the sensor into the LICA architecture

6.1 System overview

In addition to underpinning the LICA architecture, transputer technology provides simple interfacing to external sensing and actuating devices via readily available commercial link adaptors such as the CO11. This is essential given that the physical manifestation of robotic systems is an assembly of sensing and actuating devices with which it interacts with the real world.

Fig. 4 shows the integration of the intelligent sensor into the AGV's LICA architecture. The system hardware consists of the on-board intelligent sensor, which is connected to rest of the on-board LICA based mobile robot architecture through a single transputer link. It can also be connected to a remote debugging system through a pair of RS422 differential links.

The hardware termed the intelligent sensor consists of the sensor head and sensor head drive, a number of sensor interface cards (system timing and control, lock in detection of LEP signals, galvo scanner drive and sensor head drive), and four LICA motherboards which house six processing transputer modules (TRAMs), a dumb CO11 based analogue to digital converter TRAM, and a dumb CO11 based digital IO TRAM.

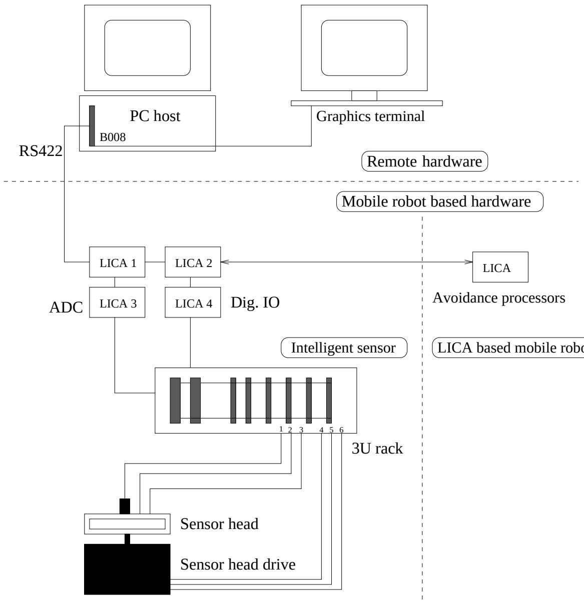

Figure 4: System outline

The PC host houses a transputer motherboard holding the system root TRAM, which configures the system and provides debugging facilities, and a graphics TRAM, which provides a real-time display of the range scans. If debugging and graphics is not required the RS422 umbilical cord can be disconnected and the remaining transputers can be booted from EPROM.

Figure 5 shows the detailed connectivity of the processing units which form part of the locally intelligent sensor. The four TRAMs indicated on left of the figure are the T805 processing TRAMs that do the main body of the processing, which includes extraction of line segments and control of the sensor head orientation. The other two processing TRAMs are T2s which handle the data flow to and from the CO11 based digital IO and ADC devices. It can be seen that the transputers are arranged in a pipeline, which is tailored to the point by point basis on which data arrives as the sensor sweeps across a scene. The purpose of the first three processing transputers after the ADC TRAM is to acquire range data continuously and extract features in sequences of sensor scans. The anchor point of the local intelligent system is a single transputer called the local sensor planner. This decides where to move the sensor head on the basis of extracted features. Subsequent transputers handle the control of the head (the gaze controller), and the interfaces to the system digital IO.

The following section briefly outlines the algorithms implemented within the intelligent sensor's transputer pipeline.

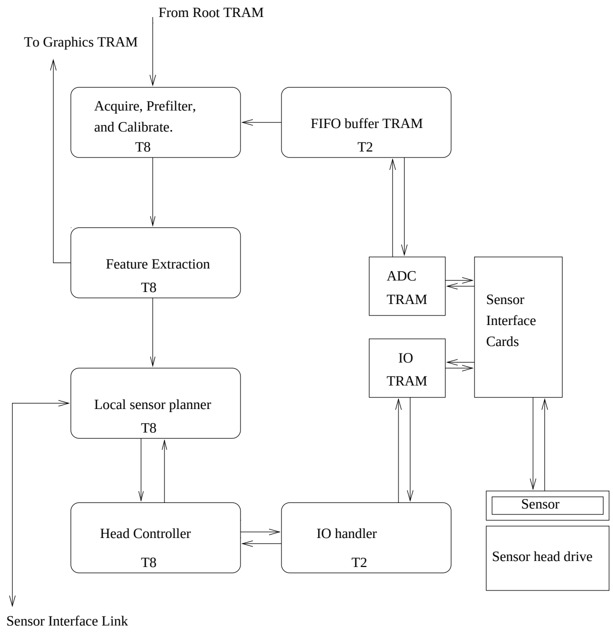

Figure 5: Transputer pipeline

6.2 The transputer pipeline

Data acquisition and buffering. Sensor data enters the TRAM network synchronously at 2.5kHz, by externally triggering A to D conversion of the amplified and demodulated LEP terminal signals on the sensor's system clock. The ADC TRAM sends the data along a transputer link to a FIFO buffering TRAM as it is converted, which removes synchrony at the front end of the pipeline, and so relaxes timing constraints. This means that the maximum scan rate, as defined by pipeline processing limitations, depends on the average time to process a point, and not the maximum time to process a point.

Filtering and range mapping. The subsequent TRAM requests data from the FIFO buffer and calculates normalised image position, with which there is an associated scan angle, detector current and sensor head angle. Any image positions registered outside the calibration range for a given angle are ignored. Also, range positions registered with low detector current are rejected in accordance with a dynamic threshold. This threshold requires a predefined minimum detector current at the defined minimum range, $I_{0_{min}}^{z_{min}}$, and readings are rejected for which

$$I_0 < \left\{ I_{0_{min}} = I_{0_{min}}^{z_{min}} (\frac{z_{min}}{z})^2 \right\} \tag{19}$$

Effectively, this dynamic signal threshold filters out areas of poor reflectivity, since it is the same form as the inverse square law of image intensity with range. In particular, the dynamic threshold is effective in filtering out those noisy range readings at edges, which are associated with the (near) zero detector current of missing parts. Such areas, which can be thought of as having zero reflectivity, can be rejected at close ranges without blanking out measurements which are of low detector current due to their larger range. For the remaining points, the sensor calibration described in section 4 provides the mapping

$$(p, \theta)^T \rightarrow (z, x)^T \tag{20}$$

Also, before each data point is passed to the feature detection TRAM, an associated range variance is calculated, as described in section 5.

Range discontinuity detection. It is assumed that the range data set for a given scan can be associated with a piecewise linear model of the world. If the parameters of such a model can be extracted from a scan, they can be used to guide both sensor head movements and vehicle movements in a purposeful, task-oriented manner. For many real environments in which autonomous vehicles operate (containing, for example, walls, pillars and boxes), a piecewise linear assumption is not unrealistic. Parts of the environment where the assumption does not hold can be identified because the discontinuities extracted will not exhibit predictable behaviour when the vehicle and sensor head move.

The feature extraction algorithm must estimate the parameters of the line segments, the position of range discontinuities, and their associated uncertainties. The standard least squares estimator is inappropriate, since the world coordinates x^s and z^s are not independent, but are related by equation 3. Also, the algorithm must cater for sensor head movements, since they can be significant in the time it takes for a single scan. The Extended Kalman Filter (EKF) satisfies these requirements and provides a computational framework in which sensor head movements are catered for by the evolution of an appropriate state [6]. This algorithm is recursive and so maps very well onto a point by point pipelined transputer architecture. The advantage of this approach is that edges are found with minimal latency as the sensor scans across the scene. Typically this latency is 2 to 3 range sample periods, or around 1ms, which compares favourably with a batch processing approach which must incur delays of more than 256 range sample periods, or around 100ms.

The edge position information is employed both by the intelligent sensor itself to control the sensor field of view and at the tactical level of the robot navigation system for obstacle avoidance and docking.

Local sensor planner. The local sensor planner implements a finite state machine, where each state is a particular sensor head behaviour. State sequencing is dependent on the type of manœuvre, the current state, and the local sensor observations provided by the line segment extraction algorithm. Currently, these states include:

- (1) Slaving the head orientation to the tangent of the vehicle path.
- (2) Centering on the nearest line segment endpoint in the vehicle frame.

- (3) Moving the head so that it is parallel to an extracted line segment (i.e the sensor x axis and the line segment are the same orientation in the vehicle frame).
- (4) Hold the head at a defined position.
- (5) Centering on freespace.

The first of these is the default state, reflecting the fact that the sensor should 'look' in the direction that the vehicle is moving. The others operate when a potential collision is detected. Centering the head on the nearest edge endpoint has been used to maintain observation of a reference point on the obstacle throughout an avoidance manœuvre. This behaviour may be integrated into a number of different levels of vehicle control; for example, a reactive obstacle avoidance capability has been employed where the vehicle controller orientates the vehicle by monitoring the sensor head angle, until the sensor head angle is driven to 90 degrees. Another approach would be to dynamically replan a local path around the obstacle using the increased certainty of the object position. (Maintaining an obstacle in the field of view facilitates the temporal integration of edge position data). Fig. 6 shows a sequence of range scans as the vehicle approaches and avoids the corner of a concrete pillar. Note that the corner of the pillar remains at the centre of the field of view as a result of the sensor head control. Moving the head so that it is parallel to a line segment has been used to align the vehicle to the side of a cardboard box obstacle. Here the parameters of the line segment in the vehicle frame (the one extracted with most confidence if there are more than one) are used to generate the steering control signals. Centering on freespace has provided a means of directing attention to possible routes around an obstacle, once an obstacle has been detected.

State sequencing between the above behaviours is dependent on the type of manœuvre implemented. To date, only ad-hoc sequencing has been applied to specific manœuvres, within specific environments. A simple example is illustrated in figure 7, where the vehicle is directed towards the centre of a cardboard box, which is in an arbitrary rotational position. The task is to side-step the obstacle and subsequently continue moving in the same direction. In position *a* the vehicle is moving forward with the sensor slaved to the path (state 1). As the obstacle comes into view, the vehicle decelerates so that it is stopped at a predefined distance from the box. The sensor then fixates on the box corner if it is visible (state 2), otherwise it aligns with the visible box facet (state 3). The vehicle rotates on the spot (this axis of rotation is coincident with the axis of the sensor head rotation) until it is parallel with the obstacle facet AB (position *b*). During this manœuvre the sensor head angle changes as the sensor fixates on a feature and, if it is driven to 90 degrees relative to the forward looking position, it is held there (state 4). The vehicle then moves forward steering so as to keep a fixed distance from the obstacle (position *c*) and, if the sensor is not already in state 4, it will be driven to that state since it is still

Figure 6: Centering a corner by sensor head control

fixating on the corner. This manœuvre continues until corner B is no longer within the sensor's field of view. The vehicle then rotates bringing the box corner back into view, which initiates the sensor to fixate on corner B (state 2), although no head motion will occur until the corner moves past the sensor's field of view centre. The vehicle continues rotating until segment AB has an observed vehicle frame orientation which is approximately equal to that when the manœuvre was started (position *d*). Finally the sensor transits to state (1) and the vehicle moves off in its original direction (position *e*).

Of course this is a contrived situation with no generality (the manœuvre must be made in an uncluttered area!). However, it has shown that the sensor is fast, accurate and robust enough to guide real-time vehicle manœuvres. The real-time decision making required for effective active sensor control in a more general sense is not addressed here, though the active sensor provides a useful tool for implementing and evaluating such schemes.

Sensor head position controller. Demand sensor head positions are passed to the head position controller from the sensor planner TRAM at the scan rate,

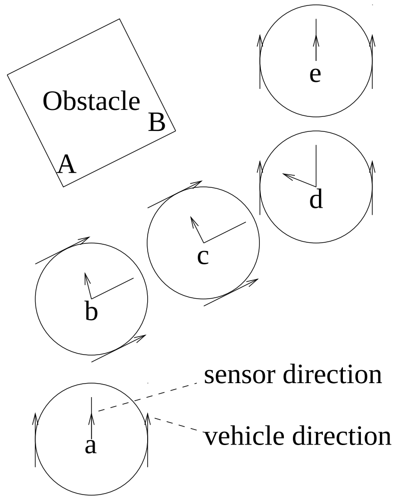

Figure 7: A simple avoid manœuvre

9.8Hz. This controller is a synchronous process in which a sensor scan time interval is divided into ten control time intervals. An appropriate demand velocity is computed and output to the head motor control board at each of these control time intervals.

7 Discussion

The type of task oriented active sensing implemented by the intelligent sensor optimises the use the sensor's finite resources such as its bandwidth. The essential point is that, for robotic tasks in which targets are localised, the sensor's field of view should be of an appropriate size and oriented towards that region of interest to allow all of the bandwidth to be used in accomplishing the robotic task. For docking and manœuvering around obstacles, this is not possible when using an unselective 360 degree range scan and much of the sensors bandwidth will provide information that is useless for the current vehicle manœuvre. (360 degree range scans are, however, appropriate for localisation and for providing vigilance to moving obstacles.)

Figure 6 indicates that the repeatability of the sensor improves dramatically at close ranges. This close range repeatability is particularly important, since it allows the safety margins during vehicle manœuvres to be smaller. This, in turn, means that manœuvres around obstacles can be made which would have otherwise been deemed as unpassable. At larger ranges (above 2m), the sensor only needs to detect the fact that an obstacle exists, to initiate deceleration. Thus, as the vehicle approaches obstacles, features must be extracted much more accurately, and this behaviour is inherent in the LEP sensor's operation.

8 Conclusions

The sensor has proved to be suitable for guiding close range real time mobile robot manœuvres because it is optically based (and therefore can provide a high spatial density of range readings), fast, accurate, it has a wide field of view, and is active and locally autonomous. The sensor is a locally autonomous agent which is a useful platform with which to investigate the interaction between active sensor control (sensor planning) and vehicle control for close range mobile robot manœuvres.

References

[1] Lo H.R. Blake A., McCowen D. and Konash D. Epipolar geometry for trinocular active range sensors. In *Proc. British Machine Vision Conf.*, 1990.

[2] Miller G.L. and Wagner E.R. An optical rangefinder for autonomous robot cart navigation. In *SPIE 852*, pages 122–134, 1987.

[3] Hu H. Sensor-based control architecture. In S. Cameron and P. Probert, editors, *Advanced Guided Vehicles*, Series in Robotics and Automated Systems, chapter 3, pages 17–35. World Scientific Publishing, 1994.

[4] C. Harris. Geometry from visual motion. In A. Blake and A. Yuille, editors, *Active Vision*. MIT Press, 1992.

[5] Everett H.R. Survey of collision avoidance and ranging sensors for mobile robots. *Robotics and Autonomus Systems*, 5:5–67, 1989.

[6] Pears N.E. An intelligent range sensor: Design, calibration and processing for mobile robot guidance. *Cambridge University Engineering Department Technical Report, CUED/F-INFENG/TR 262*, 1996.

[7] Pears N.E. and Probert P.J. Active triangulation rangefinder design for mobile robots. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 2047–2052, 1992.

[8] Pears N.E. and Probert P.J. An optical range sensor for mobile robot guidance. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 1993.

[9] S.B. Pollard, J.E.W. Mayhew, and J.P. Frisby. PMF:A Stereo Correspondence Algorithm Using A Disparity Gradient. *Perception*, 14:449–470, 1985.

[10] Kuc R. A spatial sampling criterion for sonar obstacle detection. *IEEE Trans. Pattern Analysis and Machine Intell.*, 12(7):686–690, 1990.

[11] Livingstone F. R. and Rioux M. Development of a large field of view 3-d vision system. In *SPIE 665*, pages 188–194, 1986.

[12] Jarvis R.A. A perspective on range finding techniques for computer vision. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 5:122–139, 1983.

[13] Reid G. T. and Rixon R. C. Stripe scanning for engineering. *Sensor Review*, 8(2):67–71, 1988.