Chapter 6

Optical Triangulation Range Sensors

N. E. Pears[1]
Department of Engineering Science, University of Oxford.

6.1 Introduction

The close range manœuvres of an autonomous vehicle such as obstacle avoidance, parking and docking require range sensing which is robust, high bandwidth and accurate at short ranges. Range sensors based on active optical triangulation are currently among the best means of realising these requirements. Although the geometric principle of sensing is simple and proven, it is not obvious, for a given application, what sensor configuration will give the best performance. This is because there are a large number of performance trade-offs to be dealt with which interact over the boundaries of optical, mechanical and electronic sensor design.

This chapter is mainly practical in content and is divided into two main parts: The first part (sections 6.2 to 6.4) describes, in general terms, the design choices and trade-offs which must be confronted when developing an optical triangulation range sensor. The design options for optical source, image detector, and means of scene coverage are detailed in the context of the performance requirements of a given application.

The second part of the chapter (sections 6.4 to 6.8) is more specific and deals with a short range design (maximum range 2.5m) for close range vehicle manœuvres which has been implemented within the Oxford AGV project. The optical, electronic and mechanical design of the sensor is described and subsequently its performance is characterised using least squares on logarithmic plots to estimate intrinsic sensor parameters. This is necessary both to assess the success of the design and to apply

[1]From 1998, author at Department of Computer Science, University of York, UK.

meaningful algorithms to the raw range scans. Finally results are presented which demonstrate the high short range accuracy of the scanning sensor.

Further detail on sensor design, characterisation and the processing of range data can be found in Pears and Probert [5], [6] and [7] respectively. Note that this document contains scanned figures, please refer to the book for better figures.

6.2 Optical triangulation geometries

Scene coverage can be achieved either by scanning a spot or line stripe, or by projecting a pattern of dots or lines onto the area of interest. With the latter approach, the bandwidth of the sensor is increased and the need for potentially expensive scanning is obviated at the expense of reducing the depth of field, imposing a fixed spatial resolution on the scene and introducing a correspondence problem. Some systems have tackled this correspondence problem or node labelling problem using coding of line thickness, space or colour. Others have utilised two or more different viewpoints of the pattern or have constructed trinocular systems consisting of one camera and two projectors [1]. The characteristics of projected pattern techniques suggest that they are most suited to fast inspection applications where there is a known and fixed depth of field over which ranging takes place. In such cases the designer can choose the thickness of the lines to prevent them breaking up and hampering the correspondence process at range discontinuities.

In scanning systems, either a light spot, or a light stripe expanded from a spot using a cylindrical lens, is scanned. Using such techniques, there is no correspondence problem, since only one portion of the scene is examined at any time. A light stripe is attractive for applications in which the scanned beam is accessible to the human eye, since it presents a much smaller eye hazard for a given amount of laser power output. Stripe projection, however, can only be used with 2-D segmented detectors such as the CCD. In addition, stripe systems have poor immunity to ambient light because the signal power is dispersed over a large image. This means that a longer integration time is required to obtain the same signal level as a spot system and, as a result, the noise due to ambient light will be increased.

The remainder of this section concentrates on geometries for scanned systems since these are more suitable for the larger depths of field required for autonomous vehicle applications.

6.2.1 Conventional scanning geometries

The basic geometry of scanning optical triangulation is shown in fig. 6.1. It can be seen that the axis of detection is determined by object range along the projection axis, and the position of the principal point of the light collecting lens. If we assume that the range of the object is large compared with the focal length of the collecting

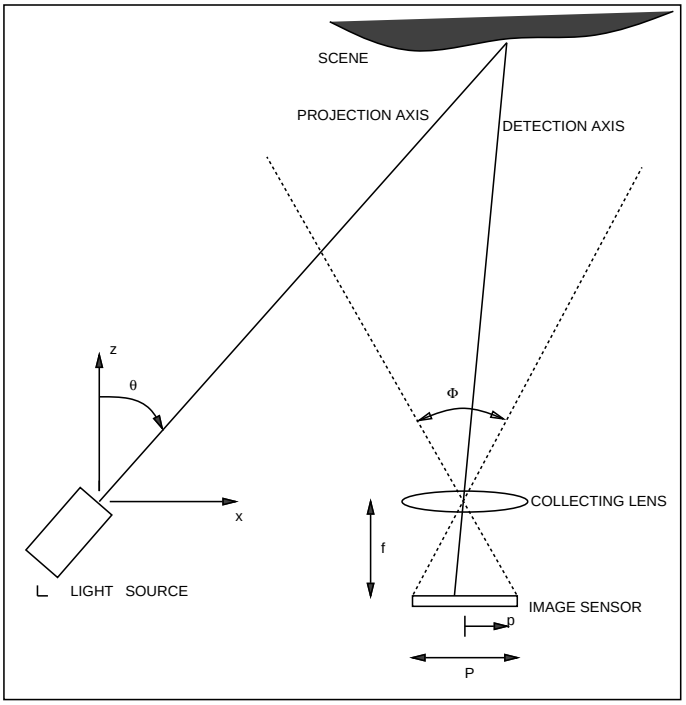

Figure 6.1: Active triangulation geometry.

lens, then the focal plane is at a distance f from the principal point of the lens and simple analysis shows that:

$$z = \frac{fd}{p + ftan\theta} \tag{6.1}$$

$$x = ztan\theta \tag{6.2}$$

In the limit, the ratio of image resolution to range resolution is defined as the triangulation gain (G_p) and, from eq. 6.1, is given by

$$\frac{\partial p}{\partial z} = G_p = \frac{fd}{z^2} \tag{6.3}$$

This shows that range resolution, for a given image resolution, is proportional to source/detector separation and focal length, and decreases with the inverse square of range. In a scanned ranging system, there is an additional effect on ranging accuracy developed by the error in measurement of the projection angle. From eq. 6.1 we have

$$\frac{\partial \theta}{\partial z} = G_\theta = \frac{dcos^2\theta}{z^2} \tag{6.4}$$

The effect on ranging of each of the geometrical/optical parameters of fig. 6.1 can be summarised as

- *Baseline:* A small baseline gives a compact sensor and good immunity to the missing parts problem. Missing parts refers to portions of the scene in which no measurements can be made either because the projected light can not access them or because the projected light is obscured from the detector by other parts of the scene. A large baseline exacerbates the latter problem but improves range resolution.

- *Detector length and focal length:* A larger detector length can provide either a larger field of view or an improved range resolution or a limited degree of both, depending on the choice of focal length. (A short focal length gives a large field of view at the expense of accuracy and vice-versa.) However, the geometric benefit from an increased detector length must be offset with increased sensor head size and a possible degradation of electrical characteristics.

6.2.2 Synchronised scanning geometries

In the geometry described above, the angle between the projection axis and the optical axis of the imaging lens changes as the sensor scans across the scene. This angle defines, along with other intrinsic sensor parameters, most of the important performance measures such as minimum range, maximum range, depth of field for ranging, and accuracy.

Geometries based on synchronised scanning are employed in multi-dimensional ranging in order to maintain a more uniform triangulation geometry as the sensor scans over the scene. In addition to being a valuable property in itself, this uniformity reduces the trade-offs that have to be made in the scanning geometry of the previous section. Oomen and Verbeek [2] introduced a lateral synchronization scheme in which the optical source scanned in a direction perpendicular to the detector field of view and Rioux [4] described a longitudinal synchronisation scheme in which the scan direction is parallel to the detector field of view. The latter geometry is schematically illustrated in fig. 6.2 and it can be seen that the scanned beam is tracked by an imaging mirror so that the angular separation between the projection axis and the optical axis of the imaging lens remains constant. [2] This means that the field of view of the detector rotates with the projection angle and that, to a first approximation, the scanning motion is subtracted from the image of the projected beam. Thus, the position of the image on the detector is much more dependent on range than on projection angle and nearly the whole of the detector length can be used to monitor range.

Simple geometric analysis applied to this scheme yields

$$z = \frac{fd}{p} cos^2\theta + \frac{d}{2} sin2\theta \tag{6.5}$$

and x is again given by eq. 6.2. The effects of image resolution and scan angle resolution on range accuracy can be described by the equations

$$G_p = \frac{fdcos^2\theta}{z^2} \tag{6.6}$$

$$G_\theta = \frac{1}{dcos2\theta - 2ztan\theta} \tag{6.7}$$

Benefits of this geometry over a conventional scanning geometry include

- Reduced susceptibility to angular error in scanning (eq. 6.7) as compared to a conventional geometry (eq. 6.4).
- Focal length can be adjusted to trade off resolution and depth of field (z range), as in a conventional geometry, but now there is no reduction in field of view (x range), since this is dependent mainly on the angle over which scanning takes place. This means that there is no limitation in the aspect ratio of the viewing area.
- Ambient light immunity is improved, because the field of view rotates with the projection angle.

[2] Note that it is more difficult to keep the triangulation baseline constant over the scanning range for a number of reasons which include a finite imaging mirror, imaging lens separation.

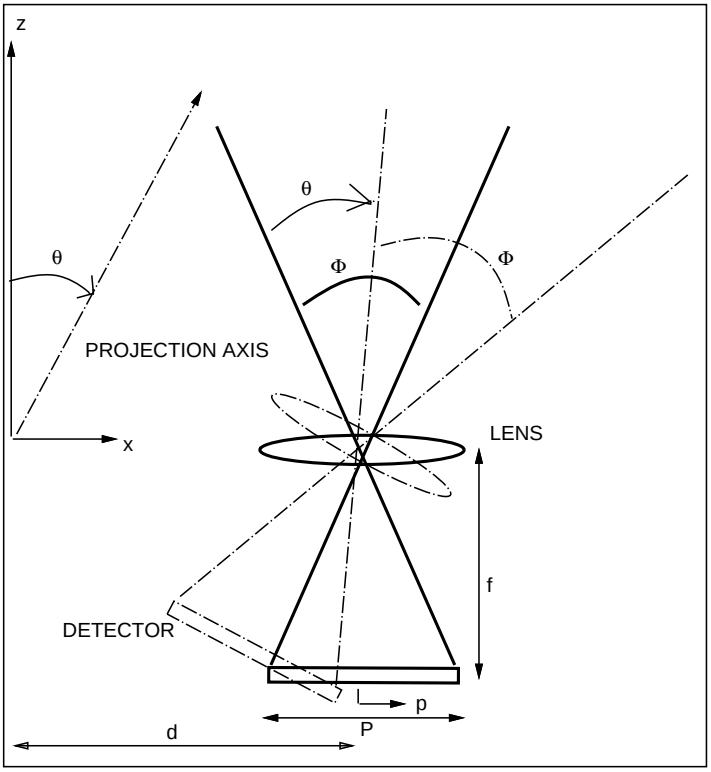

Figure 6.2: Synchronised scanning geometry.

6.3 Image position sensing devices

In this section a comprehensive coverage of image position detection devices is not attempted. Instead the two most common devices for image position sensing within optical triangulation systems are described and compared in detail. These devices are the charge-coupled device and the lateral-effect photodiode.

6.3.1 The charge-coupled device (CCD)

A charge-coupled device consists of a layer of semiconductor which is insulated from an array of electrodes by a thin layer of silicon dioxide. The operation of a CCD, for optical imaging and its many other applications, rests on the principles of charge storage and charge transfer which it is assumed that the reader is familiar with.

In optical applications, light impinging on the CCD generates electron-hole pairs in the semiconductor. The electrons are trapped in the potential well whereas the holes disappear into the substrate. The number of minority carriers (electrons for n-channel) collected under a given electrode within a given period of time (integration period) is proportional to the local light intensity. Readout of the optical image is achieved by transferring the packets of charge to adjacent CCDs which are shielded from the incident light. (If the image is read out directly from the imaging CCD, appreciable image smear can result since light is still impinging on the device during the readout period.) Clocking the charge packets out of the device is done by grouping the electrodes into three or four sets which are called phases; a different phase is applied to each set of electrodes to move the potential wells and, therefore, the charge packets across the device.

6.3.2 The lateral-effect photodiode (LEP)

LEPs, also known as position sensing detectors, are less well known in the robotics community than CCD sensors.

The LEP is constructed from a slice of silicon with P and N doped layers forming a PN junction. Charge carriers are generated by light impinging on the device, are separated in the depletion region, and are distributed to the electrodes at either end of the device. Since the device has uniform resistivity in the implanted layers, photocurrent is divided according to the position of the centroid of light intensity. Thus the position of the light centroid relative to the centre of the device (see fig. 6.3[3]) is

$$p = \frac{I_1 - I_2}{I_1 + I_2}(\frac{P}{2}), \qquad \{-\frac{P}{2} \leq p \leq +\frac{P}{2}\} \tag{6.8}$$

[3]Apologies for the poor quality of the figures in this PDF document: these are scanned images - see the book itself for better figures.

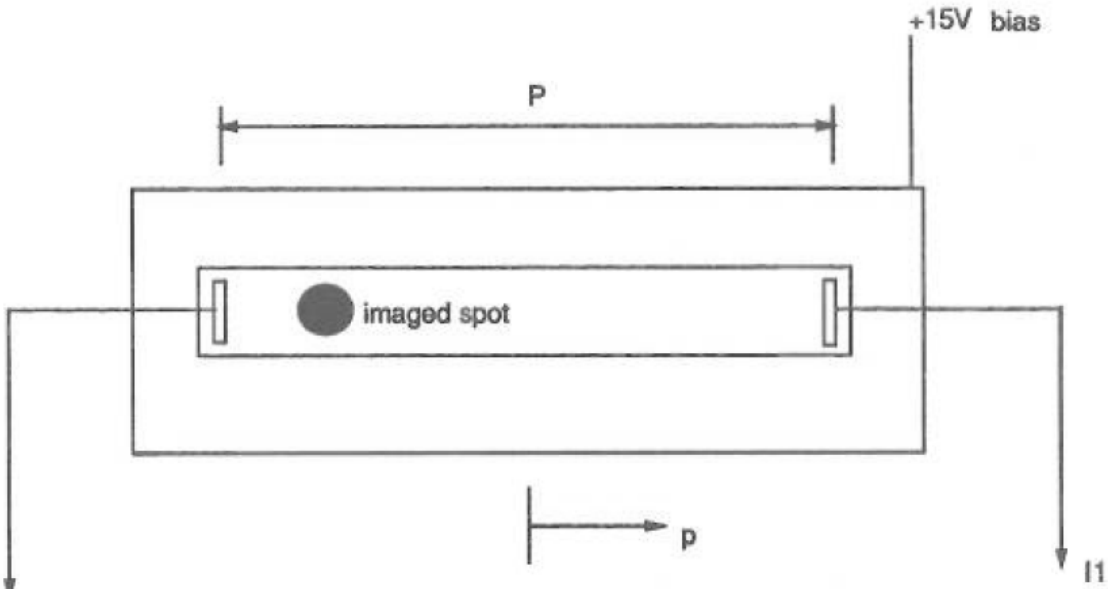

Figure 6.3: The lateral-effect photodiode

and the signal current is simply

$$I_s = I_1 + I_2 \tag{6.9}$$

6.3.3 CCD/LEP comparison

In a CCD, image resolution is dependent on the pixel density. If pixel size is small compared to the image size, sub-pixel accuracy may be obtained by fitting a curve through a set of pixel intensity values. This, however, increases the cost and complexity of analysing the data. In contrast, an LEP is a continuum, and image position is derived from the normalised difference in current signals at the terminals of the device. The resolution of this image position is a function of the signal to noise ratio over the measurement bandwidth. This implies that at short ranges or with high power light sources, when the image irradiance is high, an LEP will outperform a CCD. At longer ranges and lower power light sources, a CCD will resolve image position more accurately. For a given application, the range must be determined above which a CCD can improve on LEP resolution in order to make a sensible choice of image sensor.

In addition to a comparison of image resolution, the rate at which range measurements can be made must be considered (although the two are related through temporal averaging). In a 2-D CCD, measurement is restricted to n times the frame rate, where n is the number of pixels along the length of the imaged stripe (typical values of 512 pixels and 25Hz frame rate gives a measurement rate of 12.8kHz). For

a 1-D CCD, the measurement rate is the pixel clock out rate divided by n giving measurement rates of a similar order of magnitude to 2-D CCDs. The measurement bandwidth of an LEP, however, is only limited by the capacitance and resistance of the device. For a 10x2mm 1-D LEP, this is typically around 2MHz although, in practice, measurement bandwidth in the signal processing circuitry must be restricted in the interests of signal to noise ratio.

In addition to a comparison of image position resolution, the following points must be considered when comparing the two devices

- The segmented structure of a 2-D CCD provides us with the option of scanning a stripe of light to make n range readings per frame on a profile perpendicular to the sweep direction. LEPs (1-D and 2-D), however, can only employ flying-spot coverage.

- The LEP is much more insensitive to badly collimated light sources and defocusing with range variations as compared with the CCD, since it integrates light intensity over its surface to provide image position. A CCD based system can combat these problems with peak detection circuitry and Scheimflug geometries.

6.4 Optical sources

6.4.1 Laser diodes and LEDs

The coherence of the laser allows collimation into a very narrow, parallel beam of light, bringing the advantages of specific localised probing of the environment, compact scanning optics, and a small image size. Also, their spectral purity allows the use of very narrow band optical filters to reject ambient light and their very high temporal bandwidth allows high frequency square wave modulation without introducing phase shifts and signal distortion. The great problem of a laser, when projected into an open environment, is the eye safety problem, which sets an upper limit on the power output.

A light emitting diode (LED) is an incoherent light source which emits over a range of wavelengths and has a limited temporal bandwidth. Although the refractive index of glass varies with wavelength, the main problem in collimating an LED lies with its incoherence. Collimation is particularly difficult with high power LEDs since they emit over a large area. The great advantage of the LED is that it is unconditionally eye safe. This means that in open environments high power LEDs may be employed, which eliminates the signal to noise problems encountered with either noisy or insensitive detectors.

From the above discussion, it seems that lasers are preferable to LEDs in all respects except for the laser eye-safety problem which may limit power output. If high

powers are required in an open environment, LEDs must be used for 1-D (unscanned systems). For scanned systems, however, higher power laser sources may be used if eye-safety is guaranteed through a laser shutdown mechanism on scanner failure and a minimum accessible distance to the laser projection.

6.4.2 Spectral frequency of projection: visible, UV, or IR?

A visible beam is of great help in the alignment of any optical system and has the advantage of increased eye-safety for collimated output powers up to 1mW (class II limit) as it has the ability to stimulate the blink reflex. For beams above this power, the blink reflex is not fast enough to prevent damage to the eye. Whereas visible light is safer and more practical than infra-red, the responsivity curves of most optical sensing devices show that they are more sensitive to infra-red. Typically, an increase of sensitivity of around 50% can be obtained by moving the projected wavelength from visible light to infra-red. Infra-red is useful indoors because man made objects reflect infra-red energy well. Outdoors, however, radiation emitted by roads and other bodies limits the use of infra-red projection, particularly if the source is unmodulated. Here, UV is more useful since the ozone layer blocks sunlight which is less than 0.3 microns [3]. These factors can present a difficult trade-off since, in the interests of ambient light filtering, we need to select a light source which peaks at a spectral frequency away from that of ambient light and yet we require the detector to be responsive to that light source.

6.5 Sensor design

The preceding sections have described the general design trade-offs that need to be confronted when building optical triangulation range sensors. This section puts these principles into practice and describes a specific implementation, designed for close range autonomous vehicle manœuvres.

6.5.1 Image sensor and optical source

Since high accuracy was required at short ranges (0.5m-1m) and the maximum range was limited by the specifications embodied in the geometric design to 2.5m, the LEP was selected as the image position measurement device. The LEP is often cheaper and simpler to use than a CCD and can offer a greatly improved image position resolution at shorter ranges. The range at which a CCD system begins to outperform an LEP system depends on many optical, geometric, and electrical parameters. Our system requires ranges over 1m before a 1024 pixel CCD (one-pixel resolution) becomes more accurate than a simple LEP.

The optical source is a 670nm laser diode which is modulated at 10kHz and projects an average power of 0.9mW into the scene. This power complies with the

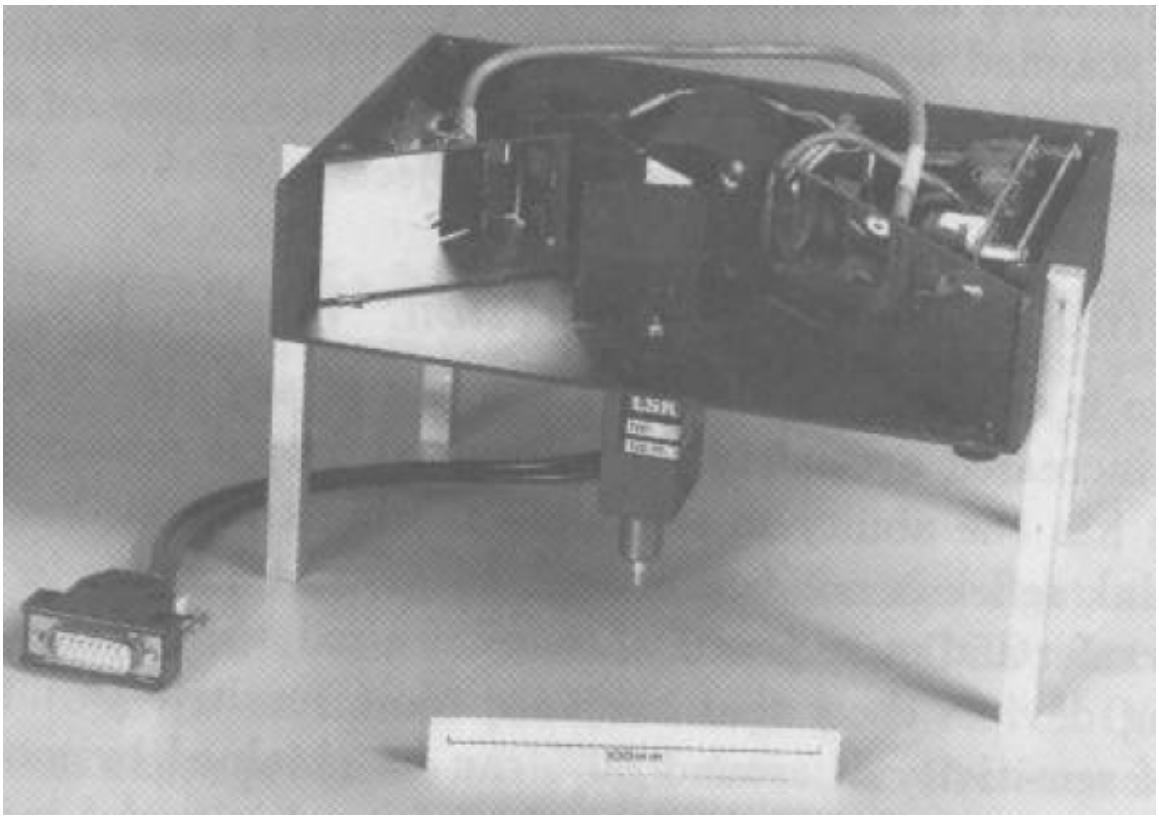

Figure 6.4: The Oxford LEP sensor

class II limit (1mW) for visible lasers and is both practical and safe to work with. The sensor is shown in Figure 6.4.

6.5.2 Electronic design

Careful circuit design for low noise performance is required for the efficient use of an LEP. The 10kHz modulation on the laser diode allows the use of synchronous detection in the receiver to maximise signal to noise ratio and eliminate dc offsets.

The key elements of the electronic design are shown in fig. 6.5. The first stage, a transimpedance amplifier, is designed with three specifications in mind: low noise density (required for good image position resolution), very high gain (required to detect signals of a few nanoamps) and a wide bandwidth (required to include the first few harmonics of the 10kHz modulation). The requirement of a high bandwidth, which is a consequence of using square wave modulation in synchronous detection, is in direct conflict with the need for a high transimpedance ($180M\Omega$). A typical stray capacitance of 3pf in parallel with a feedback resistor of $180M\Omega$ rolls the preamplifier response off at around 300Hz, whereas amplification of the first four harmonics of the 10kHz modulation requires a bandwidth greater than 70kHz. To allow this, the

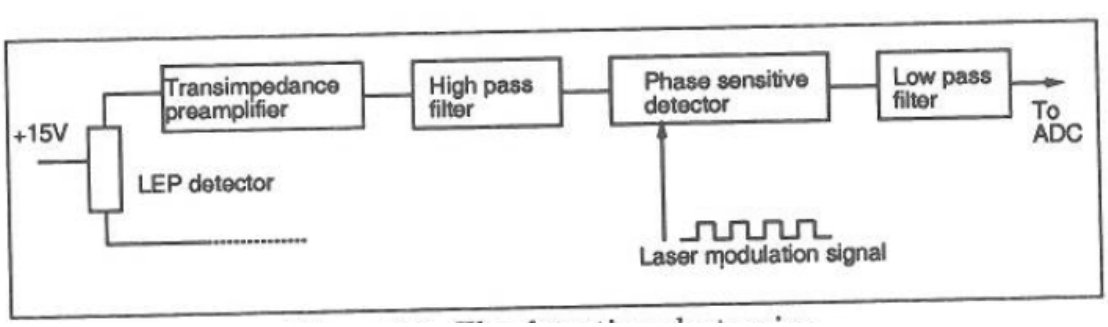

Figure 6.5: The detection electronics

preamplification is divided into two stages: a transimpedance amplification, implemented by a ultra low noise FET op amp, with gain 620kΩ, and a voltage amplification stage, implemented by a wide bandwidth FET op amp, with gain 150. A further factor of 1.9 is provided by later processing stages. The distribution of gain between the transimpedance and voltage amplification stages is such as to make their bandwidths equal. In this way, the total transimpedance is maximised for the chosen bandwidth (80kHz).

The function of the synchronous detection is to centre the signal content around the harmonics of the modulation, thus avoiding the effects of ambient light, dc offsets and drifts in the preamplification stages, and flicker ($\frac{1}{f}$) noise. This allows us to approach the theoretical limit of performance as defined by the white noise density over the measurement bandwidth. A high pass prefilter precedes the phase sensitive detector in order to allow for its non-ideal behaviour which is comprised of uneven mark-space ratio in the modulation, and imperfect matching of the two gains (inverting and non-inverting) in the phase sensitive detector.

The low pass filter shown in fig. 6.5 defines the measurement bandwidth of the system. This must be set considerably below the modulation frequency, at a level which is an appropriate compromise between signal to noise ratio and speed of response. Our filter is a fourth order filter with gain 1.9 and 3db cut off at 1kHz. Since this stage dominates the dynamic response of the system, a Bessel response was chosen to avoid any ringing at edges in the range scan.

6.5.3 Geometric design

To provide a depth of field of 0.4m to 2.5m, whilst observing the tradeoffs described in section 6.2, a baseline $d = 9.5cm$, a focal length $f = 5cm$, and a detector length $P = 1cm$ were used. A synchronised scanning head adapted from Livingstone and Rioux [8] was designed. Fig. 6.6 shows a plan view of the sensor, which is 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 use of this small mirror prevents the laser body from

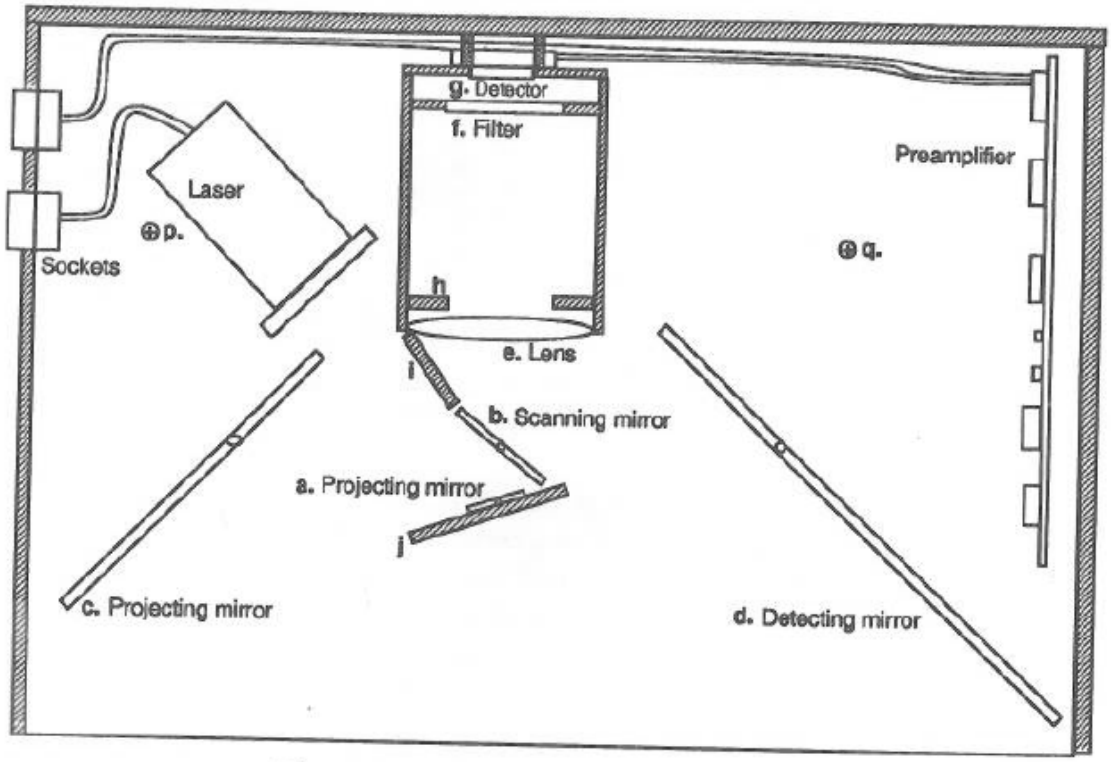

Figure 6.6: The obstacle avoidance sensor

blocking the sensor aperture. The scanned beam is then projected into the scene by a larger mirror (c), which can be pivoted to set the direction of the projection angle when the scanning mirror is at rest in its zero position. 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 p in the sensor.

Laser light, scattered by objects in the scene, is collected by the large adjustable 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 lateral-effect photodiode (g). To minimise noise, the detector signals are amplified inside the camera head before being passed to the synchronous detector in the sensor interface rack.

With the geometry described above, the lens is effectively scanned around virtual image point q in the sensor on an arc with radius equal to the separation between the scanning mirror and the lens. In our design, this separation is kept as small as possible to minimise variations in triangulation baseline over the scanning range. The dimensions and positioning of the detection mirror (d) are critical and ensure that the full sensor aperture (the full surface of the scanning mirror) is accessible over all

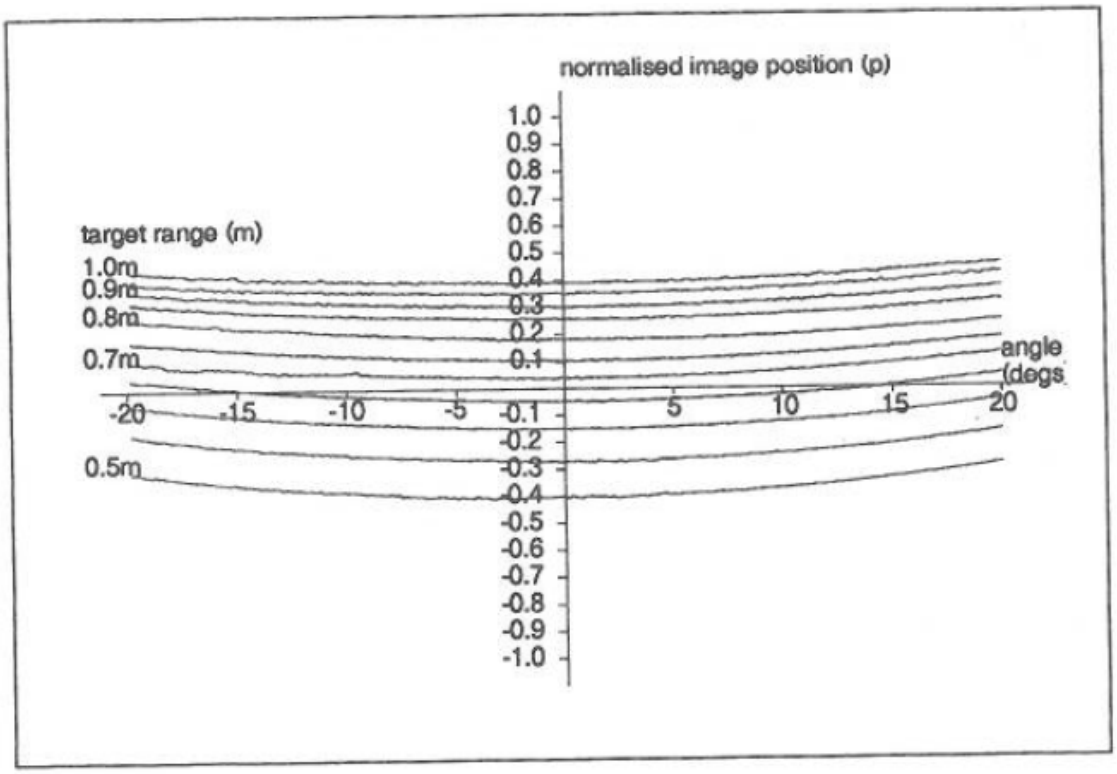

Figure 6.7: Calibration curves: image position against scan angle

combinations of scan angle and target range.

The sensor is designed with as large an aperture as possible that is consistent with our scanning requirements because of the LEP's dependence on a good signal to noise ratio. Most of the aperture derives from the depth of the scanning mirror (4cm) rather than its width (2.2cm) in order to limit rotational inertia. Direct optical paths, in which laser light is focussed directly from the scene onto the LEP, have been prevented with the use of a 'cats-eye' aperture stop (h) behind the lens and shielding plates (i) and (j).

Fig. 6.7 shows the recorded image position as the sensor scanned across a planar target at ranges of 0.5 to 1.0m. Since the lens is scanned in synchronism with the laser, there is only a small variation in image position over the scanning range. The fall in triangulation gain with target range is shown by the decreasing separation of the image position curves.

The final specifications of the sensor are: field of view (scan range) 40 degrees, depth of field (z range) 0.4m to 2.5m, samples per horizontal scan 256, scan frequency 10Hz, sample frequency 2.56kHz, projected laser power 0.9mW at 670nm (class II), detector bandwidth 1kHz (set by fourth order Bessel filter).

6.6 Sensor characterisation

A comparison of the actual performance of the sensor and the performance predicted from the specifications of the LEP, preamplification electronics and sensor geometry is essential to determine how close the sensor is to the physical limits of ranging accuracy. In the following analysis, only the effect of image position resolution on ranging resolution is considered. However, in a scanning sensor, errors in scan angle measurement will also contribute to ranging error. The specifications of our scanner quote a mirror angle repeatability of 0.1mrad which makes that the effect of scan angle resolution on accuracy negligible.

6.6.1 Experimental conditions

The sensor was set up with the laser pointing along the z-axis and image position, p, was calibrated against range, z. A target, consisting of an off-white cardboard box was set up at ranges between 0.75m and 2.5m in steps of 0.25m. At each range, 1000 measurements of signal current, I_s, and normalised image position, p, were made. For each image position measurement, range, z, was interpolated from the calibration look up table.

6.6.2 System noise estimation

Using logarithmic plots of averaged signal current against range and image resolution (the standard deviation of image position) against range, the effective noise density in the system can be estimated.

If it is assumed that the projected spot is small enough and distant enough to be treated as a point source and that purely lambertian scattering occurs in the scene, then the total signal current I_s at range z, for laser power P_l, can be approximated by

$$I_s = \frac{S}{z^2} \tag{6.10}$$

The signal current constant, S, is given by

$$S = \frac{\rho A cos\theta_d T_1 T_2 R_\phi P_l}{2\pi} \tag{6.11}$$

where ρ is scene reflectivity, A is aperture (m^2), θ_d is the orientation of the aperture with respect to the target surface, T_1 is the reflectivity of the projection optics, T_2 is the reflectivity of the detection optics, R_ϕ is the responsivity of the LEP (A/W) and P_l is the power projected by the laser (W).

Rather than measuring or estimating these individual parameters, in which case the errors in each value would contribute to the error in the estimation of the sensor

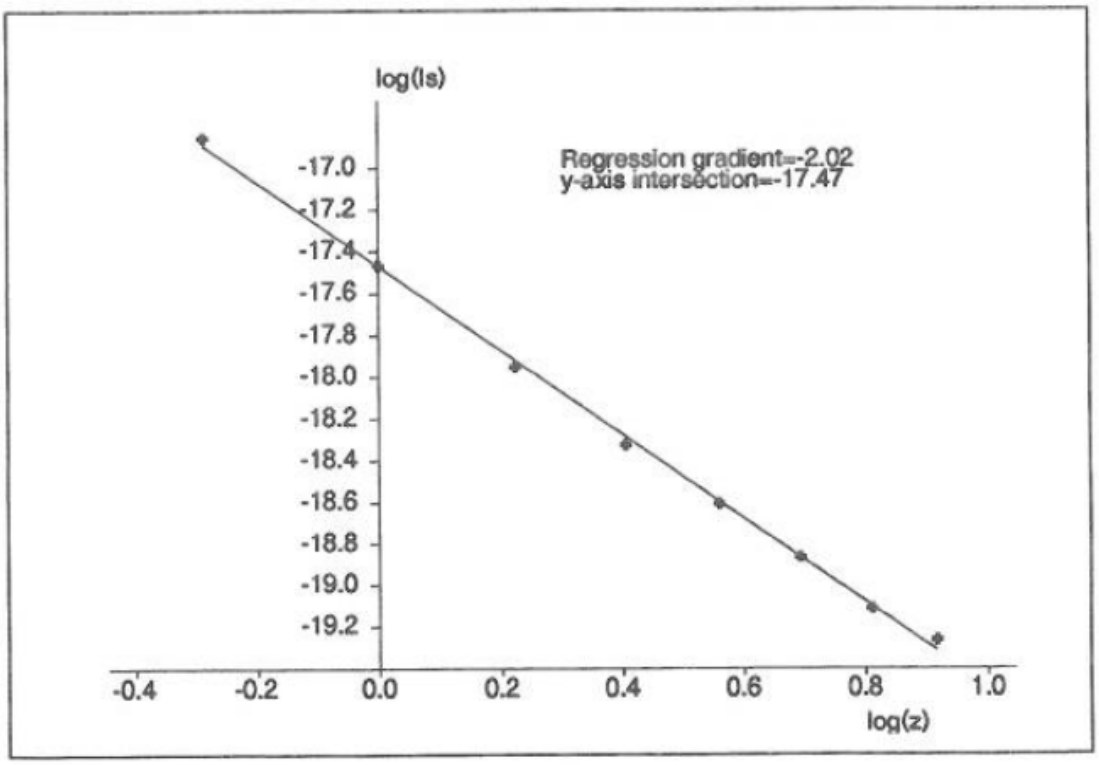

Figure 6.8: Logarithmic plot of signal current against range

performance, the composite parameter, S, can be estimated accurately from a logarithmic plot of signal current against range. This plot is given in fig. 6.8 where the crosses represent the actual measurements and the solid line is a least squares fit. The intersection of the regression at the y-axis gives

$$S = 2.59 \times 10^{-8} Am^2 \tag{6.12}$$

Now, image resolution for an LEP is

$$\Delta p = \frac{P}{2\frac{I_s}{I_n}} \tag{6.13}$$

where P is the detector length and $\frac{I_s}{I_n}$ is the signal current to noise current ratio of the detection. Since image position measurements are normalised to lie in the region $\{-1 \leq p \leq 1\}$, then

$$\Delta p_{norm} = \frac{I_n}{I_s} = \frac{I_n z^2}{S} \tag{6.14}$$

where I_n is the total noise current over the measurement bandwidth. Thus a logarithmic plot of image resolution against range, given in fig. 6.9, gives the constant $\frac{I_n}{S}$ and thus the noise current I_n

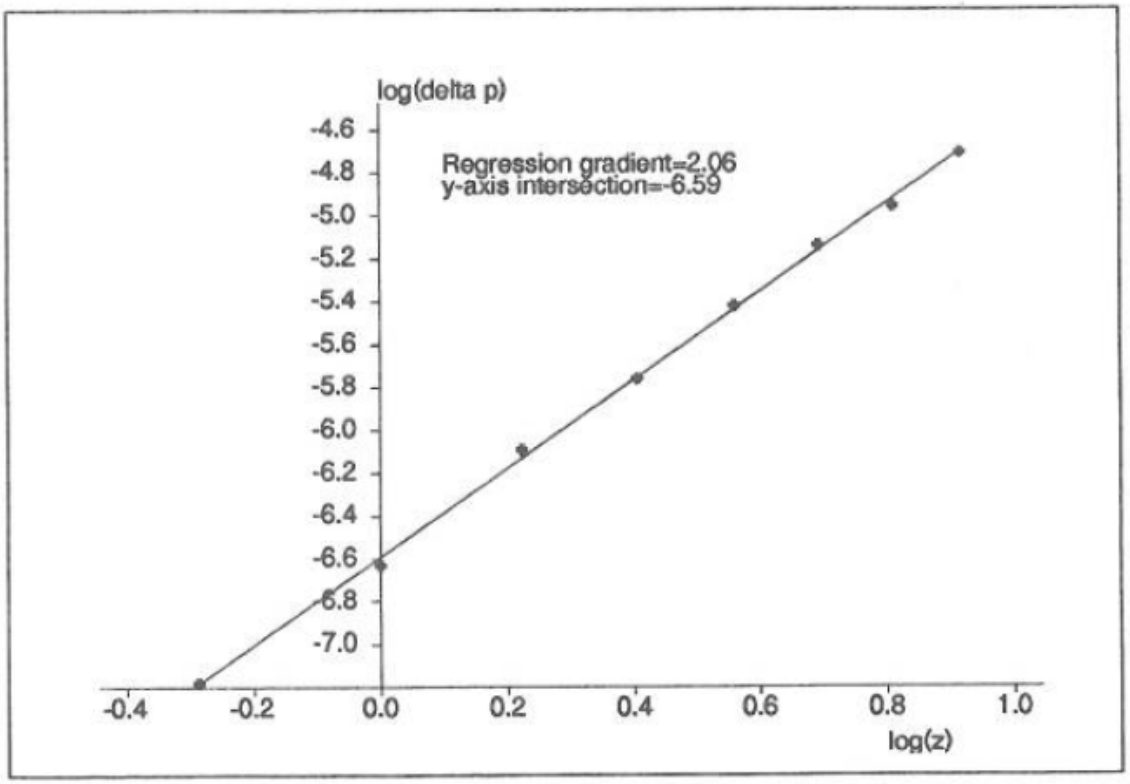

Figure 6.9: Logarithmic plot of image resolution against range

$$I_n = 36.5pA \tag{6.15}$$

In these tests, the measurement bandwidth was set by a fourth order low pass Bessel filter with cut off at 1kHz. Thus the approximate current noise density, assuming an ideal filter response, is

$$i_n = 1.15pAHz^{-\frac{1}{2}} \tag{6.16}$$

6.6.3 Theoretical noise calculation

In calculating the theoretical noise density, it is assumed that shot noise due to the signal current, and shot noise due to any current generated by ambient light passing through the optical filter, are negligible.

The total detector noise density is the rms summation of shot noise from the dark current and thermal noise. For the LEP used in the sensor, dark current $I_d = 100nA$, $R_s = 50k\Omega$, and assuming $T = 298K$:

$$i_{nd} = \sqrt{2eI_d + \frac{4kT}{R_s}} = 0.6pAHz^{-\frac{1}{2}} \tag{6.17}$$

Preamplifier noise, including feedback resistor noise is given as:

$$i_{na} = \sqrt{i_{n_a}^2 + \frac{e_{n_a}^2}{R_s} + \frac{4kT}{R_f}} = 0.2pAHz^{-\frac{1}{2}} \tag{6.18}$$

Calculating the rms of eq. 6.17 and eq. 6.18 gives the total current noise density (I_n) for the detector-preamplifier combination as $0.632pAHz^{-\frac{1}{2}}$. The noise density calculated from measurements ($1.15pAHz^{-\frac{1}{2}}$) is only 82% above this value which shows that the use of synchronous detection has given a reasonably close agreement between the theoretical performance limit of the sensor and its actual ranging performance.

6.7 Sensor performance

6.7.1 Measured accuracy

Fig. 6.10 shows a histogram of range measurements when a target was placed at 1m from the sensor, and 1000 readings were taken. In this figure, the crosses show the frequency with which a measurement fell within a particular range interval. The solid line shows the equivalent Gaussian, generated using mean and standard deviation of the batch of 1000 measurements. From fig. 6.10 it is evident that, in any processing of the raw sensor data, the assumption of of a Gaussian form is reasonable. Standard deviations for all target ranges, and their values as a percentage of the target range, are shown in table 1.

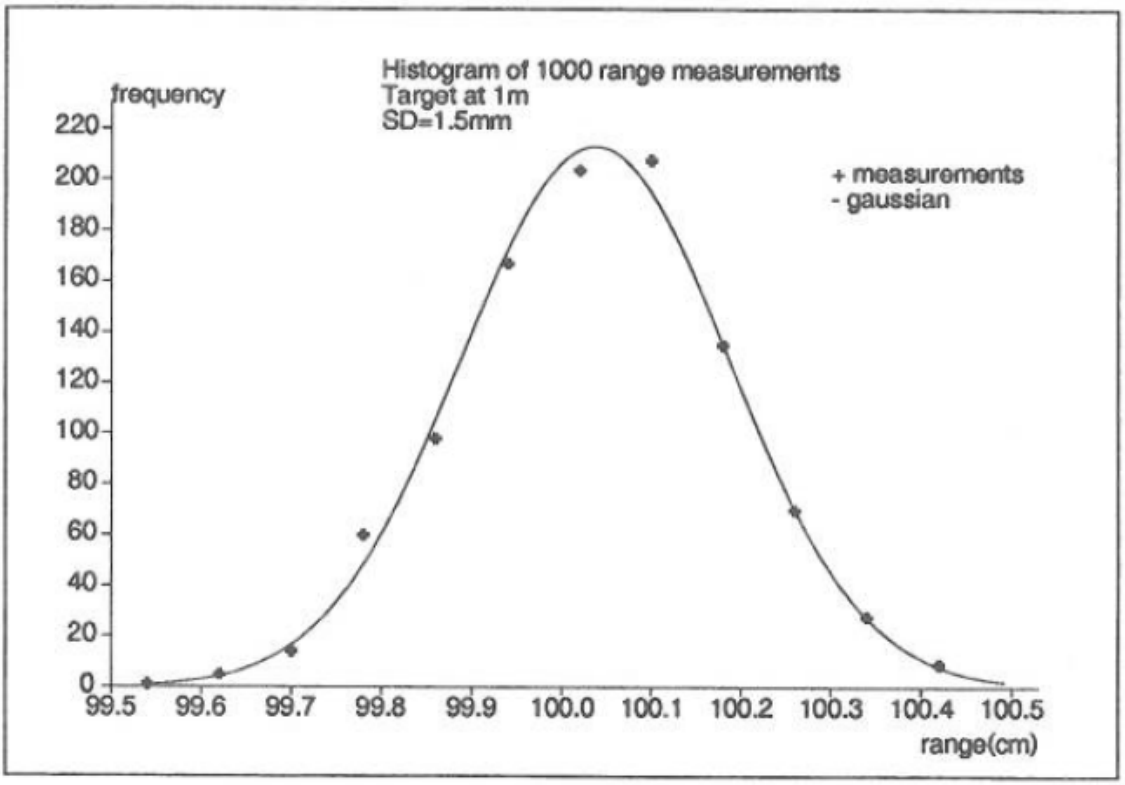

Figure 6.10: The distribution of range measurements at 1m

Table 1. Ranging accuracy

Range(m)	Standard deviation (cm)	SD as percentage of range
0.75	0.05	0.067
1.0	0.148	0.148
1.25	0.463	0.37
1.5	0.798	0.532
1.75	1.568	1.1
2.0	2.547	1.27
2.25	3.713	1.65
2.5	7.662	3.06

6.7.2 Predicted accuracy

In order to calculate range resolution from image position resolution the triangulation gain constant, fd, must be estimated. Combining eq. 6.6 and eq. 6.14) for $cos\theta = 1$ gives

$$\Delta z = Tz^4 \tag{6.19}$$

where

$$T = \frac{PI_n}{2fdS} \tag{6.20}$$

Thus, a logarithmic plot of Δz against z will give the constant T and, since I_n and S are known from previous logarithmic plots, fd can be obtained. The intersection of the regression line with the y-axis in such a plot gives $fd = 0.0044m$.

Range resolution against range can now be plotted for the measured noise density ($1.15\ pAHz^{-\frac{1}{2}}$) and the theoretical noise density ($0.632\ pAHz^{-\frac{1}{2}}$) using eq. 6.19 and eq. 6.20. These two plots, and the points at which ranging accuracy was measured (from table 1), are shown in fig. 6.11.

When scanning, ranging accuracy will be poorer at the edges of the scan than at the centre since, from eq. 6.6, triangulation gain drops with $cos^2\theta$ and, to compound this, signal strength will be reduced because of the apparent reduction in aperture with $cos\theta$. The total effect is to reduce ranging accuracy at the edges of a scan to around 83% of its value at the centre. A typical plot in fig. 6.12, shows the excellent response over abrupt edges and indicates that the variation in ranging accuracy with scan angle is small compared to the variation with range.

6.8 Use of the sensor within the Oxford AGV project

The sensor is being used within the Oxford AGV project to investigate

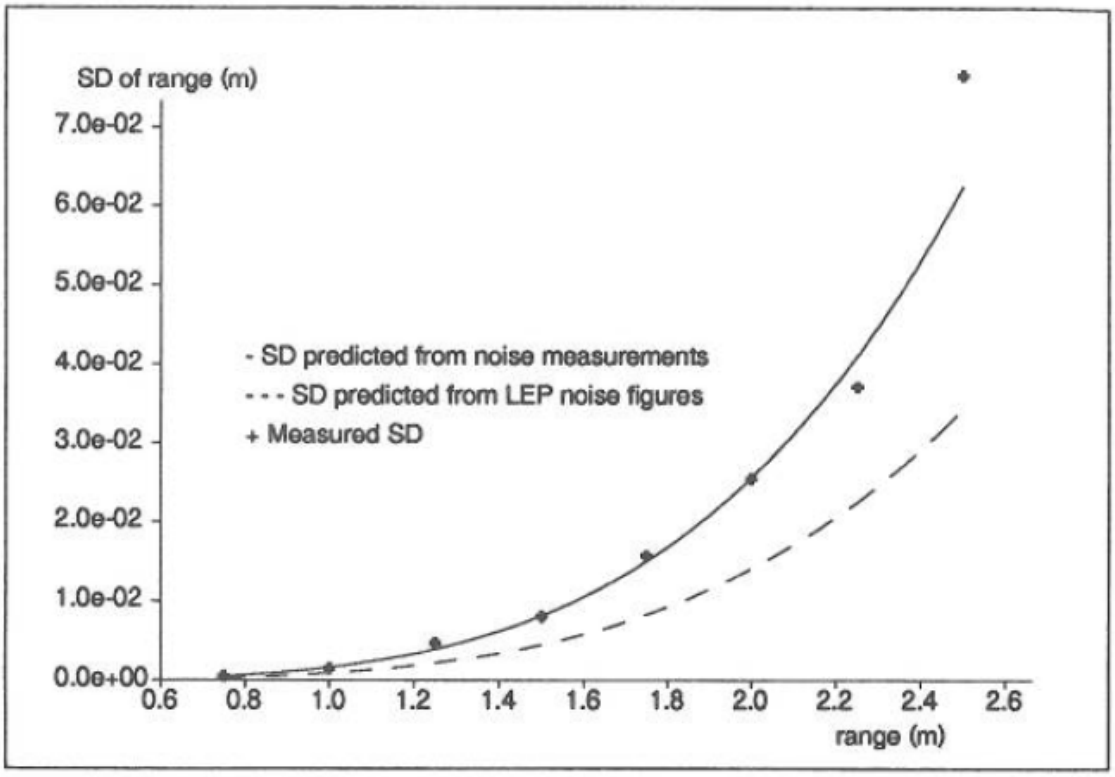

Figure 6.11: Plot of predicted and actual SD of range against range

- How intensity data can be used in conjunction with range data in order to identify and locate range features.

- How active sensing strategies can be used to facilitate robust close range autonomous vehicle manœuvres.

6.8.1 The use of intensity data

An attractive feature of the LEP is the Gaussian noise characteristic of image position measurement and the direct relationship between its variance and the intensity signal. These facts make the sensor particularly amenable to statistically based feature detection, tracking and data fusion algorithms such as the Kalman filter.

Experimentation with the sensor has shown that intensity data may be used to detect certain qualities of, or classes of, feature. These include

- Features from intensity signatures: It has been noted that signal amplitude data may be useful for identification and localisation of certain types of feature. Fig. 6.13b illustrates that, in the case of a cylindrical object, a distinct peak is found at the centre of the feature.

- Range discontinuities: Intensity data can be very useful in detecting small range discontinuities since it is zero when the projected spot is occluded by the discontinuity (missing parts). More generally, intensity can be used to form an accurate validation gate for range discontinuity detection within the Kalman filter algorithm.

- Feature reflectivity: If the sensor is calibrated with a target of known reflectivity, feature reflectivity can be estimated from intensity and range data.

6.8.2 Active sensing

The full implementation of the sensor described in this chapter is active, both in the sense that it uses structured light and in the sense that it can actively look for and track features, which are predicted to be observable, rather than waiting for them to enter a static field of view. When operating actively, it can rotate its whole body in a horizontal plane between +90 and -90 degrees and can vary the scan of the laser over any angular range between -20 and +20 degrees.

The sensor has been developed as a locally autonomous sensor agent in which the software for calibration, filtering, feature extraction, feature tracking, sensor planning and head control are mounted on a small network of transputers dedicated to the sensor. The local intelligence of the sensor is embodied in the local sensor planner which receives a state vector describing the vehicle motion over a transputer link connected to the vehicle controller. Using this information and the sensor's own information

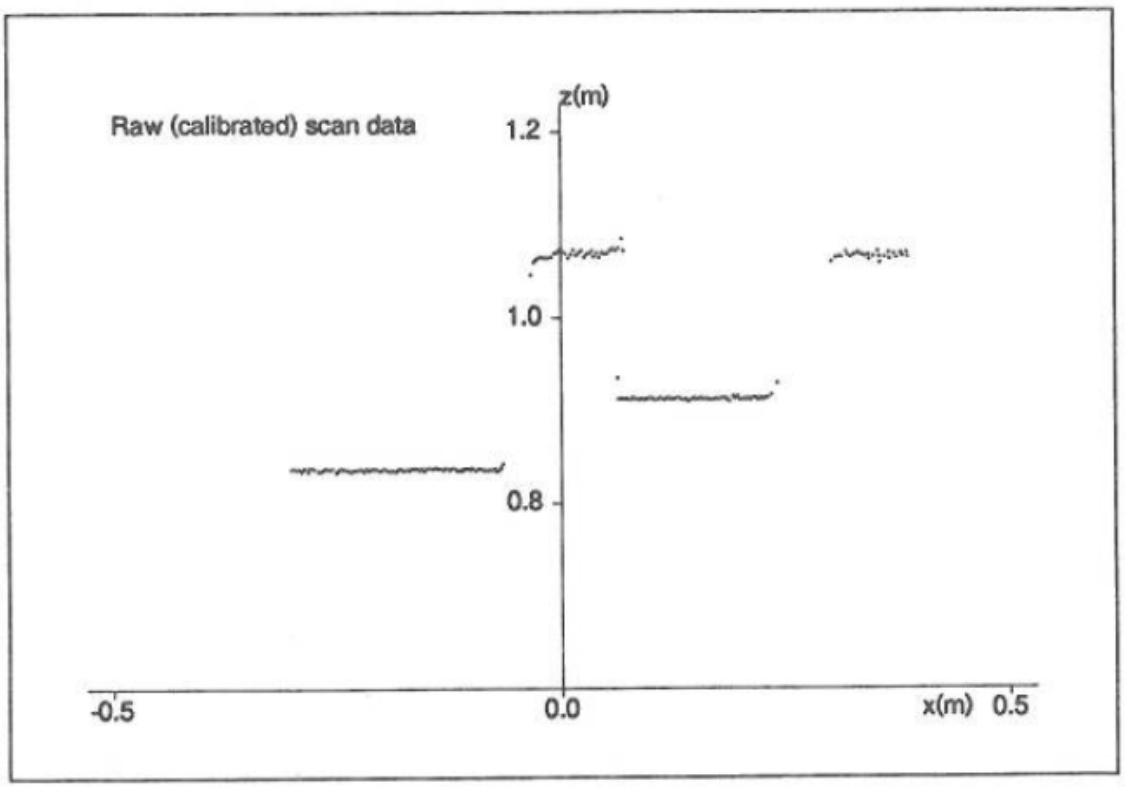

Figure 6.12: Range image as the sensor scans over a scene incorporating several cardboard boxes.

concerning features and their movements, it can decide which feature tracks to maintain based on the costs and utility of active sensing actions. A vector describing the local environment is passed back from the sensor network to the vehicle's path planner.

This active realisation is being used to demonstrate close range manœuvres such as obstacle avoidance on a Robosoft mobile platform.

6.9 Conclusions

Optical triangulation range sensor design has been discussed with a view to implementing a short range sensor for close range autonomous vehicle manœuvres such as obstacle avoidance, parking and docking. This discussion has shown the implications of choosing a certain sensor configuration (optical source, image sensor and means scene coverage) in terms of performance criteria such as accuracy, bandwidth and useful range.

A wide field of view range sensor for short range mobile robots manœuvres has been described. Measurements have shown that ranging is accurate to 0.15% at

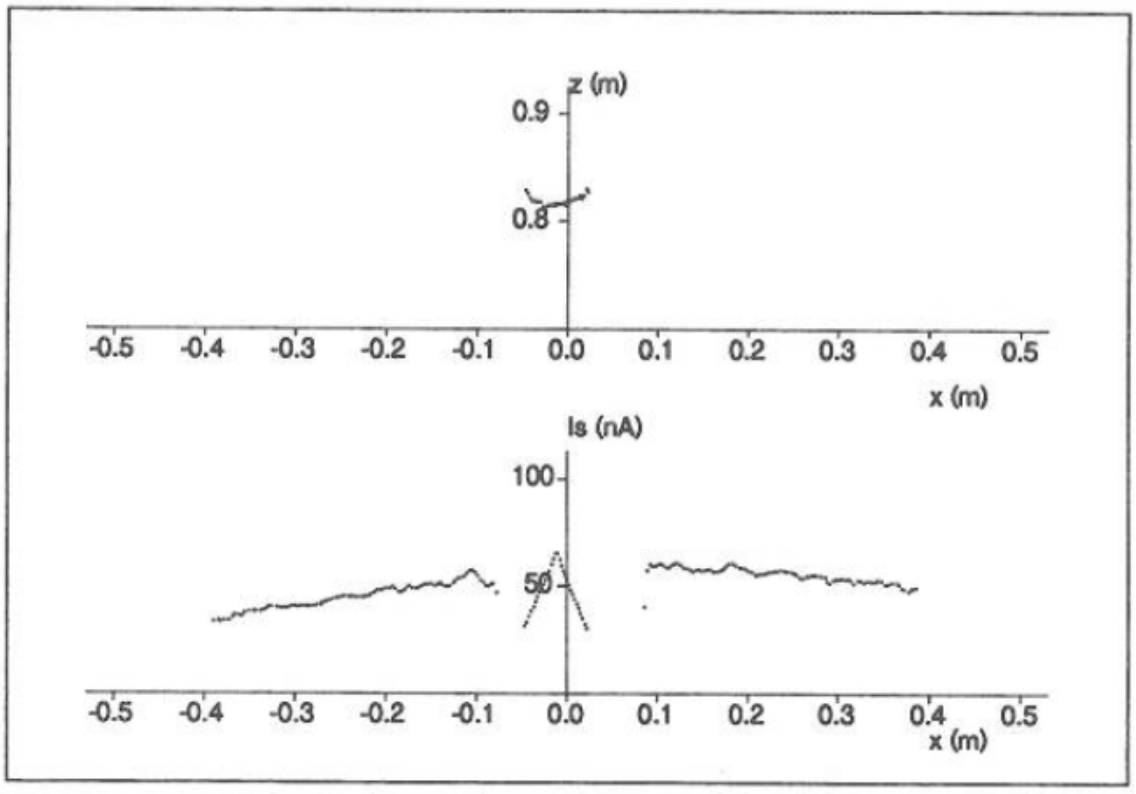

Figure 6.13: Range image (top) and intensity image (bottom) as the sensor scans over a scene incorporating a cylinder.

1m, 1.3% at 2m, and 3% at 2.5m for a measurement bandwidth of 1kHz, which is suitable for real-time planning, and an average projected power of 0.9mW, which is eye-safe. These performance figures compare well with other rangefinders available at comparable cost. In addition, we have extended the application range of LEPs, which are normally used with non eye-safe lasers over much smaller depths of field.

A sensor characterisation showed that the measured range resolution is only 82% above the theoretical performance limit as defined by the detector and preamplifier noise specifications. This good performance is the result of careful preamplifier design and the use of synchronous detection. Also, the sensor characterisation allows a good estimate of range variance to be made by measuring the received photocurrent. This variance estimation, along with an approximate Gaussian form on the ranging errors, makes the sensor particularly amenable to statistically based feature detection, tracking and data fusion algorithms.

Bibliography

[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] Oomen G.L. and Verbeek W.J.P.A. A real-time optical profile sensor for robotic arc welding. In *SPIE 449*, pages 62–71, 1984.

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

[4] Rioux M. Laser rangefinder based on synchronised scanners. *Applied Optics*, 23(21):3837–3844, 1984.

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

[6] N.E.Pears and P.J.Probert. A lateral effect photodiode based scanning sensor for autonomous vehicle obstacle avoidance. In *1st IFAC Int. Workshop on Intelligent Autonomous Vehicles*, 1993.

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

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