Lidar Design, Use, and Calibration Concepts for Correct Environmental Detection

Martin D. Adams, *Member, IEEE*

***Abstract*—The useful environmental interaction of a mobile robot, is completely dependent on the reliable extraction of information from its immediate surroundings. A particular class of sensors often now applied to this problem is the lidar (light detection and ranging) system. The aim of this article is to examine the performance limits and sources of error in these sensors at their design and calibration stages and during their general use. A framework, aimed directly at optimizing the quality of the output information, for use in mobile robot navigational algorithms, is given. The design concepts for producing correct range estimates in the presence of a large dynamic range (>120 dB) of surface albedo is addressed. The performance limits, which can be expected in terms of systematic and random range errors, are theoretically analyzed and modeled to provide a correct calibration procedure. During this derivation, it will be shown that the naive determination of the sensor to target distance as a function of any lidar's output signal, in general provides a false calibration.**

The possible scanning speed and data sampling rates are derived as functions of a lidar's geometrical and electronic temporal design specifications. Finally the issue of temporally averaging of several range values is demonstrated and it will be shown that under certain quantified conditions, range variance reduction is possible.

The text addresses the use of amplitude modulated continuous wave (AMCW) and time of flight lidars in general, but makes several references to a particular lidar design example, giving results and conclusions from an actual engineering AMCW lidar implementation.

***Index Terms*—AMCW range estimation, lidar, sensing.**

I. Introduction

THE correct interpretation of the data produced by any sensor, scanning within indoor environments should begin with an analysis of its hardware design. Indeed the sources of electronic noise, nonlinear behavior, signal saturation, and even erroneous signals within a sensor should, at least, be known and understood during its use or, ideally, minimized at the sensor's design stage.

Many research institutions are now utilizing *lidar* sensors for robot navigational experiments [1]–[3] and the motivation for this work is from the experience gained in robot navigational experiments, using commercially available sensors [4], [5]. The first aim of this article is to show how the design of the individual modules within a lidar sensor can affect the information extracted from such a device. This information is often sampled from the sensor's outputs, the sensor being assumed to be a sealed black box, and used to attain some form of reliable environmental mapping. A view inside the black box is taken to show how various design factors influence the data interpretation or indeed its misinterpretation. An overview of the crucial components of an AMCW lidar sensor is given in Section II, the design of which are used to optimize a real engineering implementation and resulting use of such a device. The intended application is mobile robot navigation, which is reflected within the design considerations, in terms of size, scanning speed and reachable area. The lidar, which was designed and tested in this article, is a coaxial sensor, transmitting collimated, amplitude modulated light into the environment via a single scanning mirror. This can be continuously rotated about a vertical axis, and simultaneously swept with respect to a horizontal axis thus giving limited 3-D coverage [6]. Since the device is coaxial, detection takes place in a direction retro to that of transmission, thus eliminating the disparity problem associated with triangulation methods [3]. Range estimation results from the phase difference between the transmitted and received signals, this being defined modulo half the AMCW wavelength (15 m here).

Section III takes an in depth view of the critical design factors by defining 'worst-case' sensor to target parameters under which lidar range estimation should just be possible. A theoretical relationship between the minimum detectable photocurrent, the lidar's transmitter power and the detector aperture is then derived.

Section IV presents the theoretical performance limits, resulting from various noise sources, which can be estimated before electronic construction takes place. Section V explores the causes of systematic range errors, which can cause problems in mobile robot feature detection and matching algorithms. This section presents remedies for effectively reducing the dynamic range of the processed signals, and hence the nonlinear dependence of the range estimate upon the received signal amplitude.

Most feature detection and map building algorithms rely on statistically based algorithms, which make optimal state estimates, given noisy, uncertain sensor measurements. These require knowledge of the sensor uncertainty and therefore here, the use of the amplitude of a lidar's received signal is analyzed and related to the range variance. Provided both the range estimate and the signal amplitude are available, the uncertainty of each range reading can be determined, and made available for such algorithms [8]. Therefore, the second aim of this article is to provide a correct calibration procedure for lidars in Section VI.

For real-time navigation, the speed at which an optical beam can be scanned and hence independent range samples recorded must be addressed and is the subject of Section VII. By mod-

Manuscript received October 13, 1999; revised July 13, 2000. This paper was recommended for publication by Associate Editor G. Hager and Editor S. Hutchinson upon evaluation of the reviewers' comments. This work was supported by Prof. G. Schweitzer, Institute of Robotics, ETH Zurich, Switzerland.

The author is with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798.

Publisher Item Identifier S 1042-296X(00)11645-3.

eling the phase locked loop (PLL), a tool often used for phase measurement, as a closed loop control system, an upper limit for the scanning speed and temporally uncorrelated range data acquisition rate is derived.

Finally, Section VIII explores the averaging of several range estimates, recorded at high speed, for range estimation improvement, without loss of angular resolution.

II. AMCW Lidar Modules

A brief overview of the modules used in lidars is now given, and for a more detailed analysis, the reader is referred to [6]. When suggesting methods for improving the quality of range estimates, specific references will be made to these modules.

Most of the literature to date emphasises the importance of a noise analysis before design takes place [9], [5], [10], [11]. Indeed in general lidar design, whether a pulse must be detected and its amplitude measured, or a continuous signal reproduced, the primary design objective should be low noise and not the exact reproduction of the actual signal [12].

As a result of this literature survey, and the experiences of the author with lidar design principles, the implementation of a usable AMCW ranging system can be considered to consist of six major components, namely the: *transmitter module, receiver module, signal compression unit, relative phase discriminator, output filter stage*, and *received signal envelope detector*, the interconnection of which is shown in [6].

III. Critical Lidar Design Factors

In order to gain an understanding for the systematic and random errors, a review of the physics of reflection and signal reception is necessary. When incident upon an opaque surface, light undergoes both *specular* and *diffuse* reflection simultaneously, and it is the diffuse component which dominates the range estimate, and which is of interest in lidar design[1] [15]. If the transmitter produces an RMS radiant power P_T incident upon a surface at an angle θ relative to the local surface normal (Fig. 1), the reflected power per steradian, in the direction retro to the incident beam is $I_R = P_T \rho \cos\theta/\pi$, where ρ is the diffuse reflectivity, which is, in general, a function of the transmission wavelength.

If the receiver aperture has an area A_R and is situated a distance r from the illuminated spot (Fig. 1), then it subtends a solid angle $\alpha = A_R/r^2$. The total power then received is P_R where

$$P_R = \eta \frac{A_R P_T \rho \cos\theta}{\pi r^2} \tag{1}$$

[1]Horn suggests that Lambert's cosine law is a reasonable approximation for many materials such as paper, snow and matte paint, when illuminated by a point source, as in this case [13]. Different lighting such as uniform radiance or hemispherical "skies" yield different brightness to surface orientation relationships. Effects such as *scatter* and *speckle*, caused by the random interference of coherent light waves from a diffusely reflecting object, also lead to deviations from Lambert's cosine law [14]. Experiments in indoor environments with the particular lidar's range/amplitude data presented here, and the use of Lambert's cosine law, have indeed produced albedo images similar in quality to black and white photographs, suggesting that Lambertian reflection is a reasonable approximation for typical indoor surfaces [8], [6].

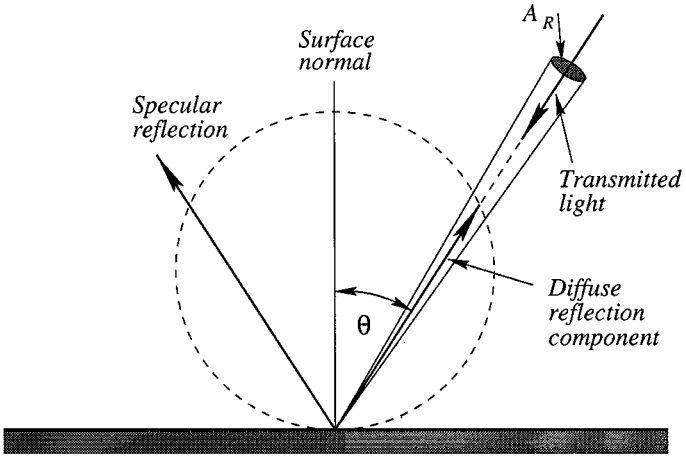

Fig. 1. The variables which affect diffuse reflection according to Lambert's cosine law.

where η is the receiver's efficiency. By entering 'worst case' values of θ, ρ, and r into (1), the minimum current which must be faithfully detected can be derived, which, in the design specifications here, is 4.0 nA. It can also be seen from (1) that the received power is proportional to $\rho\cos\theta/r^2$. In indoor environments, diffuse reflectance ratios can vary between approximately 0.02 for dark objects and almost 1.0 for white surfaces [15]. As an example, if objects are to be visible to the sensor at incidence angles $0° < \theta < 80°$ (i.e., near tangential reflection) and for ranges $0.2 < r < 15.0$ m, the received signal can have a dynamic range of $1.620 \times 10^6 : 1$ or 124 dB[2]. This issue is considered further in Section V-A-1 where electronic compression circuits will be addressed, to cope with the actual possible dynamic range of the received signal.

IV. Performance Limits—Noise

To ensure that the above criterion can be met, the sources of noise and their possible reduction within the sensor's receiver must be addressed. The total noise current is primarily caused by the following four effects [6]:

1) a *shot noise* component as a result of the photoreceiver's *dark current*, $\hat{i}_{\mathrm{dark-shot}}$;
2) a noise current source due to *avalanche multiplication* (if an APD is used), $\hat{i}_{\mathrm{apd}}$;
3) a *shot noise* component due to back ground illumination, $\hat{i}_{\mathrm{bg-shot}}$;
4) a *shot noise* component due to the induced signal current itself, $\hat{i}_{\mathrm{rec-shot}}$.

The total RMS noise current is then

$$\hat{i}_{\mathrm{total\,(RMS)}} = \sqrt{\hat{i}^2_{\mathrm{dark-shot}} + \hat{i}^2_{\mathrm{apd}} + \hat{i}^2_{\mathrm{bg-shot}} + \hat{i}^2_{\mathrm{rec-shot}}} \tag{2}$$

It now remains to determine the minimum signal current amplitude which needs to be detected and selected from the APD, and ensure that this is much larger than the RMS total noise current defined in (2). This gives rise to a further

[2]These were the design specifications for the construction of the lidar used in this article.

question: "How high does the signal-to-noise ratio (SNR) need to be for reliable range estimation?" By estimating the nature of the probability distribution of the phase (and hence range) estimate, Brownlow derived an expression for the probability that the error in a given range measurement is less than a predefined value [11]. As would be expected, this probability value increases dramatically with increasing SNR, and indeed for a 10-MHz modulation index, it can be shown that to achieve 99% confidence that all range measurements are within a tolerance of 1% of the maximum range, a minimum SNR of 30 dB is necessary [6].

Reference [6] shows that substituting each individual noise current estimate into (2) gives the result

$$\hat{i}_{\text{total (RMS)}} = \sqrt{KB + 2qI_{\text{rec}}B} \tag{3}$$

where K is the mean square noise current per Hertz due to the dark current, back ground illumination and avalanche multiplication and is 1.485^{-24} A^2/Hz in this case. For an SNR of 30 dB, the mean received current must be at least 32 times larger than $\hat{i}_{\text{total (RMS)}}$ so that

$$I_{rec} \geq 32\sqrt{KB + 2qI_{rec}B}. \tag{4}$$

The maximum quadratic solution for I_{rec}, which turns the above inequality into an exact equality, corresponds to the minimum received current necessary to attain the above defined confidence in the range estimate. It is clear from (3) that the noise current is significantly reduced by minimizing the receiver's bandwidth B. An ideal choice of receiver is therefore a simple tuned resonant circuit as used in AM radios, with a resonant peak at 10 MHz [11]. This value must coincide with the minimum necessary detectable current from the design specifications (4.0 nA here). Hence, it is therefore necessary to proceed with the receiver analysis by

1) adjusting the design parameters (receiver aperture size and sensitivity, transmitter power), or sensor specifications (maximum range requirement), such that the minimum detectable current (4.0 nA) is greater than I_{rec} this being the solution to inequality (4);
2) constructing a low-bandwidth receiver capable of selecting this signal [11].

V. Causes of and Remedies for, Range Errors

A. Systematic Range Errors

In most lidar systems, systematic range errors are reported to be of greater concern than random errors [16]. This is demonstrated in Fig. 2. The left plan shows a simple line model of the environment surrounding the sensor (located at the centre of the triangle). The right scan was recorded from a commercially available AMCW lidar sensor [5]. In order to demonstrate the effect of the large dynamic range of the received signal, a dark green piece of paper was mounted on the upper pillar between B and C. The remaining parts of the pillar between A and B and between C and D were white. Due to the differing amplitudes of the received signals from each part of the pillar, a clear systematic range error has occurred. This indicates the necessity for some form of reliable amplitude control of the received signal, to ensure linear range estimation throughout the entire specified dynamic range of the received signal, which was shown in Section III to be higher than 120 dB in the design specifications presented.

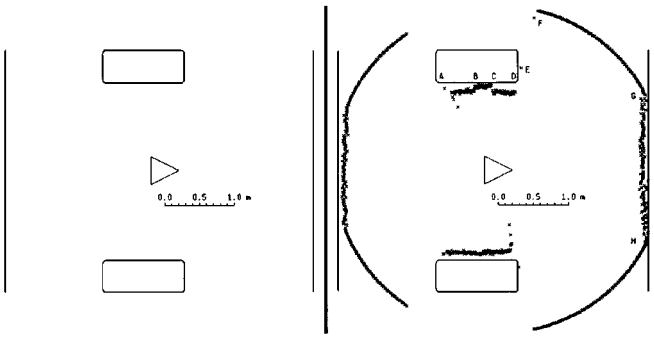

Fig. 2. The left plan shows a simple map of a laboratory environment surrounding a lidar at the centre of the triangle shown. The right plot shows a single 360° scan. Only the unadjusted range data is shown, and each data point is represented as a cross. Curved regions such as FG correspond to out of range depth readings, and are shown as crosses at 2.5 m from the centre of the mobile robot.

Spurious ranges present themselves somewhat differently at each pillar edge, a problem which can cause range edge detectors to fail [16], [4]. This occurs because the range estimate, at the sensor's output, results from a combination of the reflectivities, beam to target angles of incidence and ranges from *all* objects intersecting the projected optical footprint (known as multiple path effects [4], [17]), and any "ghost" or internal leakage[3] path within the sensor. Depending on the particular optical footprint position, when the range output is sampled, the resulting range estimate can vary tremendously, as [6] shows that the received signal amplitude is considerably weakened. If the sample was taken when one of the artifacts such as D or F dominated, a range estimate geometrically between these artifacts (such as E) results. Alternatively, if the ghost signal is higher than that produced by any artifact within the footprint, at the time of sampling, points closer in range, such as those on the right-hand side of the lower pillar can result.

1) Dynamic Range Compression: The gain of the receiver stage should be set such that the largest received signal, with which the sensor is to function under its design specifications, is linearly amplified meaning that no unwanted phase shifts are produced due to saturation of its output signal. At the output of this stage, weak signals can still be too small for use in the mixing and phase discrimination stages of the sensor.

Techniques used in radar technology include the application of log-limiting amplifiers which guarantee minimal phase shift between input and output over a very large input dynamic range [17]. This form of amplitude control is demonstrated in the two graphs of Fig. 3 where it can be seen that weak signals are linearly amplified by the cascade of amplifiers, whereas strong signals are effectively clipped, whilst preserving the phase information. The SL-531 log limiting amplifier for example, quotes a maximum phase shift between input and output of 1.1° over

[3]This results due to either direct electronic cross talk between the transmitter and receiver or an optical path which exists directly between the transmitter and receiver. The detection and removal of these points is covered in [6], [16].

its entire logarithmic range [18]. Over a maximum range of 15 m,

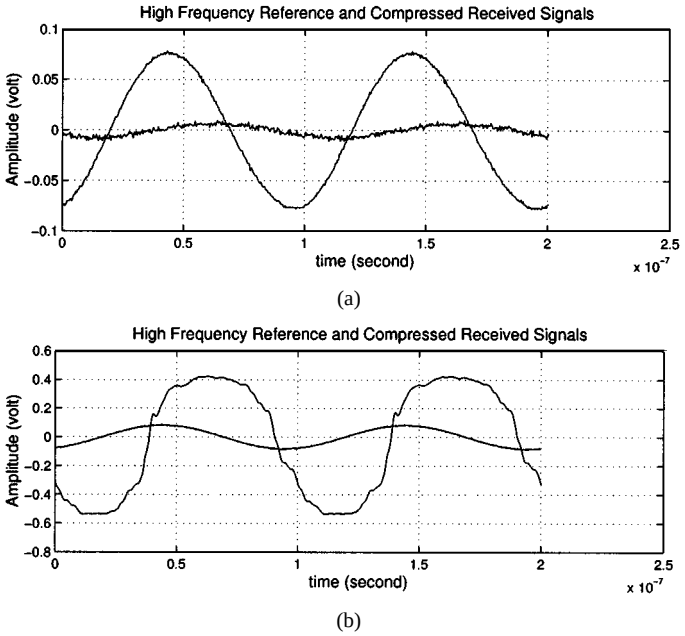

Fig. 3. (a) Weak HF received signal and (b) strong HF received signal, after amplification by the log-limiting amplifier cascade. The reference signal is also shown in each case [(the larger sine wave in (a)]. Note that the peaks of the larger amplified received signal (b) are rounded.

this corresponds to a range drift of 4.5 cm.[4] The important information which remains in the lower graph, is the relative phase between the reference and compressed received signal, which is within 1.0° of that in the upper graph.

What ever form of dynamic compression is used, the most important factor is that the phase information between the transmitted reference and the received light signals is affected minimally by the amplitude of the received signal, or is affected in a predictable and repeatable way. This issue will be addressed further in Section VI, where a correct calibration procedure for lidar sensors is explored.

B. Random Range Errors

Although the receiver noise sources were presented in Section IV, the *propagation* of noise through to the range estimate, was not quantified. This is now necessary for lidar calibration purposes, allowing the sensor to provide quantified range uncertainty information for robot state estimation based navigational algorithms.

The photodiode acts as a current source which produces a time varying current at the frequency of the modulating signal. However, as shown in Section IV, a noise current $\hat{i}_{\text{total(RMS)}}$ is also produced. The noise in the amplitude of the received signal is not directly of interest in AMCW measurement systems, since the range estimate arises from the relative phase between the

[4]In time of flight (TOF) lidars, the above problem is less pronounced, but still arises due to the finite rise time of the received pulse, as a function of the received signal amplitude. The time at which the pulse is recorded depends on the exact point during its transition at which it was considered detected, and requires the use of a *constant fraction timing discriminator* [19].

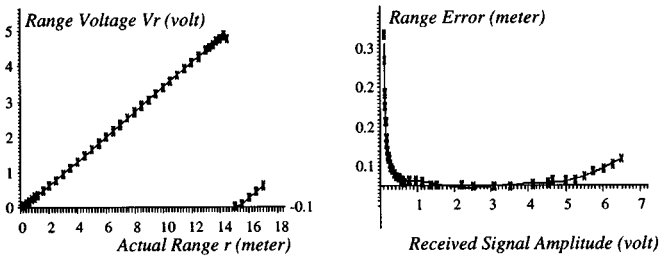

Fig. 4. Calibration curves 1 and 2 for correct lidar calibration. The left graph shows the range-voltage to actual range relationship (1 V = 3.00 m) provided the target used produces a constant received signal amplitude. Note the ambiguity interval of 15 m. The right graph shows that the sensor is extremely linear for received signal amplitudes (which have been logarithmically compressed) between 1.0–4.5 V (very small range error). It also provides the necessary range correction for all possible received signal amplitudes.

zero crossings on the ωt axis of the received and transmitted signals. Based upon a simple noise triangle analysis, the resulting range variance σ_r^2 (m^2), varies with the received signal amplitude V_r (volts) according to the following equation [6], [5]:

$$\sigma_r^2 \approx \left(\frac{\lambda \sigma_n}{4\pi}\right)^2 \left(\frac{1}{V_r}\right)^2 + \sigma_e^2 \tag{5}$$

where λ is the modulation wavelength (30 m in the design example here), σ_n^2 (volts^2) is the combined constant variance of the electronic noise sources (quantified in Section IV), and σ_e^2 (m^2) is the additive electronic noise range variance which results *after* the amplification, mixing, and phase comparison stages.

VI. Correct Calibration Procedures

This section considers the necessary procedures for determining the *three* relationships necessary to provide a full calibration of an AMCW lidar, namely, Calibration 1, the output range voltage versus actual range; Calibration 2, the erroneously, internally induced, electronic phase (and hence range) shift versus received signal amplitude,[5] and Calibration 3, the range variance versus received signal amplitude.

Calibration 1: To eliminate the varying effect of Calibration 2, it is *essential* that when initially calibrating range-voltage versus range, the returned signal strength is held constant, by using, for example, different colored targets. This crucial factor is often over-looked by mobile robotics researchers, employing lidar sensors. The left graph in Fig. 4 shows an initial calibration of range-voltage versus actual sensor to target distance. This curve will only approach linearity, if optical and electronic leakage between the transmitter and receiver is minimized [6]. This graph offers a correct range-voltage versus actual range calibration for a target at any range from the sensor, provided it returns the signal amplitude at which the calibration took place. Initially, this calibration seems useless as, in general, targets will of course return varying received signal amplitudes within the lidar. This is the reason that Calibration 2 is necessary.

Calibration 2: The data points in the right graph of Fig. 4 show the error caused by the amplifiers within the receiver circuit relative to the initial calibration in the left graph. Various

[5]Although this effect is reduced by dynamic range compression (Section V-A.1), it will not be completely eliminated, meaning that Calibration 2 is in general still necessary.

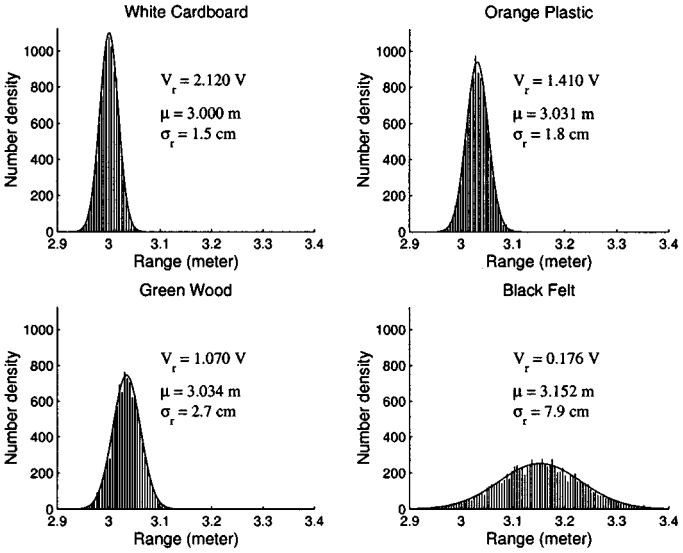

Fig. 5. Histograms showsing the effect of different colored targets at a given range. All targets were at a true range of 3.00 m from the sensor. The signal strength values V_r, sample means μ and range standard deviations σ_r are shown with each graph. The continuous curves show calculated Gaussian distributions with the same mean and variance as the discrete data.

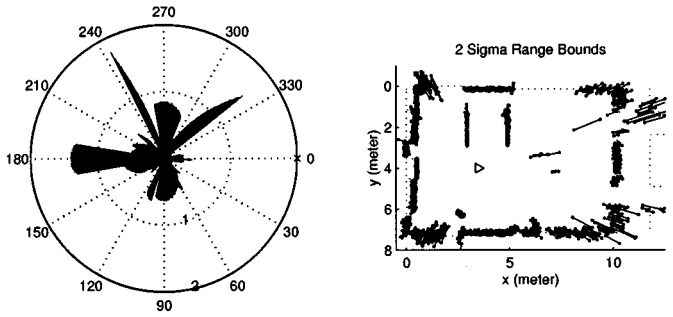

Fig. 6. Left: Signal amplitude (radial axis [volts]) versus sensor scanning angle (degrees). Right: Lines of length $2\sigma_r$ centred on the range estimates. The triangle in the right graph shows the position of the mobile robot. The dotted line represents a plan view of the actual environment.

combinations of target reflectance; orientation of target normal relative to the emitted light beam, and sensor to target distance will affect the returned signal strength [20]. *Experiment shows that the factors which affect the returned signal strength are irrelevant as far as calibrating the sensor is concerned and it is only the returned signal strength itself which is of importance.* An analytical model for the right curve is not necessary here and would provide no *general* insight into the problem, as similar sensors exist which use other circuits before phase detection [20], [5], [21], [15]. This calibration is essential for correcting the range output of a lidar, and to our knowledge is often omitted by mobile robotics researchers, which causes large mapping errors when weak or strong received signals are encountered from the scanned environment.

Calibration 3: To establish the range variance as a function of the received signal amplitude, 10 000 independent range measurements were made of fixed targets with the sensor stationary. The histograms in Fig. 5 have horizontal axes showing the measured range r, produced from the left calibration curve of Fig. 4, and vertical axes showing the number density. Note that the distributions are normalized, since the sum of the heights of all the range measurements is constant (10 000 in this case). All of the histograms in Fig. 5 were produced from different targets at a fixed range (3.0 m) from the sensor. As expected, different signal strength values correspond to different variances within the range values. Note that the distributions are approximately Gaussian.[6] Fig. 5 also shows the changes in the sample mean of the ranges for different signal strengths, which must be compensated for by Calibration 2.

These results can be used to determine the unknown constants σ_n and σ_e in (5) and hence the numerical relationship between σ_r^2 and V_r[7]. It should be noted that the range variance cannot increase without limit, since the phase measurement of an AMCW range finder is defined modulo 2π. Therefore the range is defined modulo $\lambda/2$ (15 m here) which is the ambiguity interval of an AMCW lidar (left graph, Fig. 4).

The use of Calibration 3 is demonstrated in Fig. 6 where the left scan shows the amplitude of the received signal in polar coordinates (radial axis in volts) as a function of the scanning angle (angular coordinate in degrees), and the right scan shows the standard deviation in range (plotted in Cartesian form). The right scan shows lines of length $2\sigma_r$, calculated from the received amplitude in the left scan and (5), centred on the actual range estimates themselves. The dotted line represents a plan view of the actual environment. The length of the line segments in Fig. 6 gives a quantitative assessment of the uncertainty associated with each range estimate, particularly useful for weighting the influence of each range value for feature detection purposes [6].

VII. Possible Scanning Speed

This section provides the lidar user with a tool for deriving the speed at which range information can be faithfully sampled from a lidar. In any lidar system, this derivation requires a model of the range estimation electronics, and again here, reference will be made to the *phase estimation* electronics within the AMCW process.

A reliable tool for producing square waves locked in phase to almost any noisy periodic input signal is the PLL (see Fig. 7). To derive the range sampling constraints, the PLL must be examined. A PLL can be modeled as a phase comparator, which produces a digital signal representing the phase difference between its two input signals. This signal is then low-pass filtered, and used as an input to a *voltage controlled oscillator* (VCO) which in turn can produce a square wave with frequency proportional to the input voltage [23]. This can be used in a closed

[6]For very weak received signals, it can be shown that the distributions are Rayleigh in form [22], [6].

[7]Note that a TOF lidar will also produce randomly distributed range estimates but the analysis should be based upon the finite rise time of the received pulse as a function of the received signal intensity. In general TOF lidars suffer more than their AMCW counterparts if the received signal is weak as false detection, or no detection at all can result. An AMCW lidar will produce a noisy but consistent range estimate, assuming correct calibration [19].

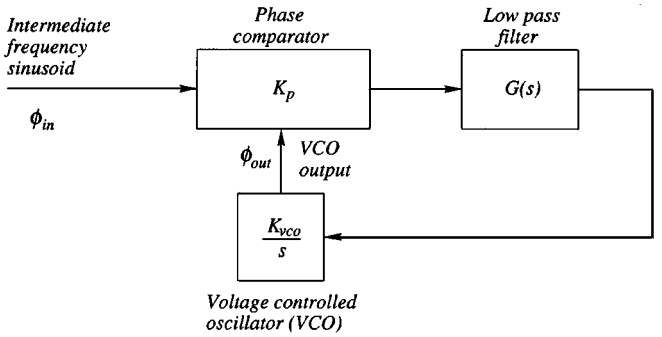

Fig. 7. Block diagram representation of the phase locked loop.

loop system (as shown in Fig. 7) to produce a clean, locally produced signal, with the same frequency as the noisy input signal (derived from the sensor's amplified received signal) and a constant relative phase relationship.

To ensure that the PLL is able to track the dynamic phase variations of the received signal as the sensor scans, it is necessary to analyze the low-pass filter used in conjunction with the phase detector and VCO. A schematic block diagram of the PLL is shown in Fig. 7. The stability is improved if a 'lead-lag' low-pass filter is used meaning that $G(s)$ in Fig. 7 has the form

$$G(s) = \frac{1 + sT_2}{1 + s(T_1 + T_2)} \tag{6}$$

where s is the Laplace frequency variable and T_1 and T_2 are time constants, dependent on the components used in the lead-lag filter. If the gain of the phase comparator is K_p and that of the VCO is $K_{\rm vco}/s$ as shown in Fig. 7, the overall closed loop transfer function between the phase of the input sinusoid $\phi_{\rm in}$ and the output square wave $\phi_{\rm out}$ is given by:

$$\frac{\phi_{\rm out}}{\phi_{\rm in}} = \frac{K_{\rm vco}K_p(1 + sT_2)}{s^2(T_1 + T_2) + s(1 + T_2K_{\rm vco}K_p) + K_{\rm vco}K_p} \tag{7}$$

which produces a classical second-order response to changes in the input phase ϕ_{in}, caused by actual range changes. The denominator of (7) can be arranged to have the standard form $s^2 + 2\zeta\omega_n s + \omega_n^2$ where ζ is the damping factor and ω_n the natural frequency of the response. With knowledge of the possible speed at which the input phase can change with respect to time (i.e., rate of range change) and the desired settling time for the locally produced VCO output square wave, values for ζ and ω_n can be calculated and implemented by choosing the correct components in the lead-lag low-pass filter. The highest frequency changes in range which need to be recorded correspond to a change in full range (15 m) divided by the time necessary for the scanning mirror to rotate through the effective beam width of the light spot. The estimated beam-width of the lidar considered here, is 0.0067 rad, calculated by measuring the optical foot print diameter at maximum range. With a maximum mirror scanning speed of 2.5 revs/s for example, the time between independent range measurements is 0.43 ms. Within this time interval, it is necessary that all transient effects of the second-order transfer function of (7) have reached an acceptable level. Brownlow defined this "acceptable level" as the time period $t = 3\pi/\omega_n$ after which any overshoot has reduced to less

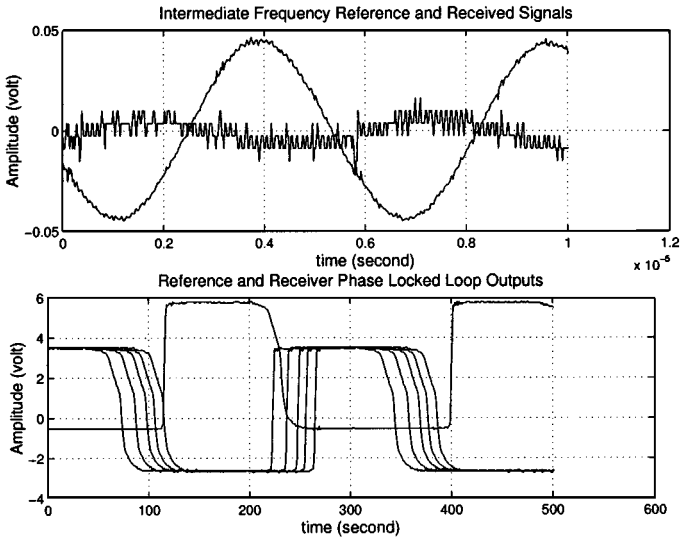

Fig. 8. Reduced frequency reference and received signals (top graph) and their corresponding VCO outputs (bottom graph) for a weakly reflecting target at 7.0 m used in earlier experiments. In the lower graph, the output of the received signal's VCO was recorded at five different times.

than 1-cm range error [11]. Applying this criterion, ω_n is approximately 22 000 rad/s, meaning that a bandwidth of almost 3.5 kHz is necessary.

The range estimate, as a result of the frequency lock-in detection capability of the PLL, is demonstrated in Fig. 8 where the top graph shows the received (low amplitude, noisy wave) and reference signal from a target at 7.0 m. The lower graph shows the square wave outputs from the two VCOs running on separate PLLs. The reference 'square' wave has been shifted vertically by 2.0 V to clarify them. The subsequent processing necessary to produce an analogue output proportional to range, simply requires a suitable phase detection circuit with both of these square waves as inputs. To demonstrate the effect of the received noise, Fig. 8 shows the results recorded from the VCO outputs at five different time intervals, these being superimposed upon each other in the lower graph. The smaller, highly noise-corrupted sine wave in the top graph is the received signal. It can be seen in this graph that the time axis crossing of the received signal is ill defined due to its low SNR. This effect reproduces itself at the outputs of the VCOs in the form of phase noise as shown in the lower graph of Fig. 8. Since this is a very weak signal it can be seen that a large phase uncertainty, and hence ultimately, range uncertainty results.

A "good" reflector placed 7 m away from the sensor, was used for the same experiment in Fig. 9. This time, the larger signal is the received signal in the upper graph, and once again the received signal's VCO output was recorded at five instants in time relative to the reference VCO signal. It can be seen that the time axis crossings are more clearly defined and the phase noise is greatly reduced.

An estimate of the range simply results from low-pass filtering the output of the phase comparator.[8] Clearly this final block in the range estimation circuitry requires an optimally set cut-off frequency to yield the necessary measurement band-

[8]Detailed methods for phase discrimination can be found in [6].

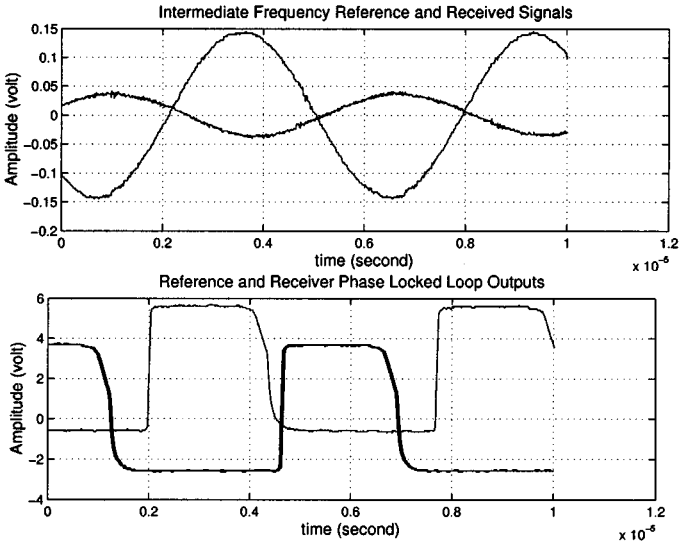

Fig. 9. Reduced frequency reference and received signals (top graph) and their corresponding VCO outputs (bottom graph) for a reflective target at 7.0 m used in earlier experiments. In the lower graph, the output of the received signal's VCO was recorded at five different times.

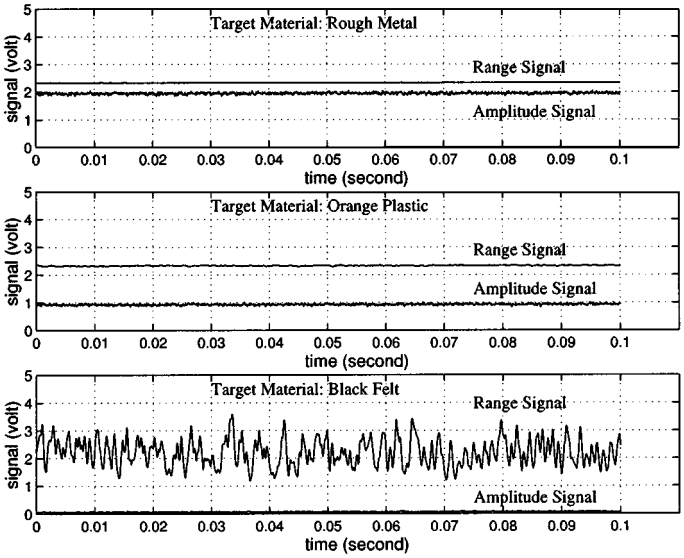

Fig. 10. Range and amplitude outputs for differing target reflectivities at a fixed range (7.0 m). The maximum range (15 m) corresponds to a signal output of 5.0 V.

width, but at the same time minimize the high frequency phase noise at the output. Applying the same argument as above, the bandwidth of the low-pass filter should be chosen to allow for the maximum possible range changes at the maximum mirror scanning speed (2.5 rev/s here). The results of low-pass filtering the output of the phase comparator for various targets, again situated at a range of 7.0 m from the sensor, are shown in Fig. 10. The lower curve in each graph shows the amplitude estimate of the signal concerned. The graphs are scaled such that full range in each case corresponds to 0–5 V, which for the range signals corresponds to 0.0–15.0 m, and for the amplitude output corresponds to the weakest and strongest received signals. The top graph shows the results obtained from a relatively good diffuse (white) reflector at 7.0 m. It can be seen that the amplitude measurement is only at 2.0 V, however the 0–5.0-V scale allows for very large signals produced by specular reflectors such as mirrors or shiny metal, which can still produce valid range estimates. Note that the range estimate, situated at 2.33 V appears almost noise free and, due to the relatively strong signal amplitude gives a reliable, low noise range estimate. In the middle graph, the target is substituted with a darker material, again situated 7.0 m from the sensor. The detected signal amplitude is approximately halved, whilst the range estimate remains at an average value of 2.33 V, although slightly more noise is evident. The lower graph in Fig. 10 demonstrates the case of an extremely weak reflector (black felt), again situated 7.0 m from the sensor. This time the amplitude of the received signal is only 1/20 of that of the good reflector, even after semi-logarithmic amplification by the dynamic range compression module. The range estimate is almost random within the 0–5.0-V output boundary, although it can be noted that given enough time to average the range output, a reasonable range estimate would still result.

VIII. Averaging of Range Data

Section VI quantified the range variance of a *single* range sample. The lower graph in Fig. 10 suggests that, given enough time, *several* range signal samples could be averaged to form a single range estimate with lower variance. To be of use in mobile robotic applications however, a lidar is usually used in continuous scanning mode, meaning that each range sample could result from a different environmental object. It was mentioned in Section V-A however, that all lidars produce a finite optical footprint, rather than an infinitesimal point, on the object(s) being sensed. The sampled range data can therefore be considered to be the output of a dynamic system, as no matter what range changes actually occur, the range signal cannot change instantly because of two constraints.

1) **Electronic Constraint:** The output range signal is the result of low-pass filtering the phase estimate and is therefore band-limited, thus limiting the speed at which *new* range information can be sampled.
2) **Geometrical Constraint:** Due to the finite size of the optical footprint, the range signal results from the convolution of all reflected light signals within the footprint, thus lowering the angular resolution [6].

The aim of this section is to manipulate these two effects to optimize the sampling and averaging of the range output. Constraint 1 will be quantified to give the number of consecutive samples to be averaged, resulting in a single point with lower range variance than the individual samples. Constraint 2 will then be quantified to ensure that no loss of angular resolution results in the new, lower density scan. This is important in applications such as feature extraction, where the location of edges needs to be known, with some precision, for mobile robot localization [6], [24].

Quantification of Constraint 1: It can be shown that if the sampling time interval ΔT is much less than the low-pass filter time constant T_f (high correlation between successive samples)

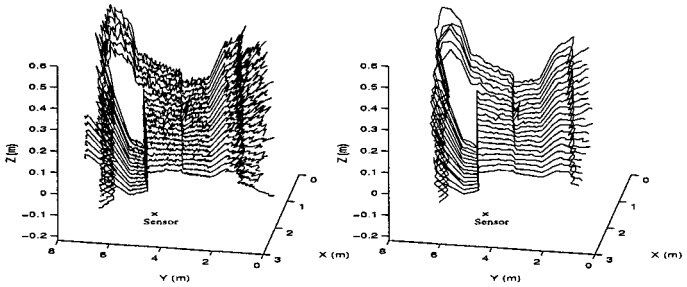

Fig. 11. 3-D range data showing the corner of a room from a single 3-D scan. The left scan shows all recorded data points from the section under consideration, and the right scan shows only one sample averaged from every four horizontally scanned range points.

then the standard deviation $\overline{\sigma}_r$ of the average of n samples taken at time intervals ΔT is

$$\overline{\sigma}_r = \frac{\sigma_r}{\sqrt{n_{\text{eff}}}} \tag{8}$$

where σ_r is the standard deviation of each individual sample. For a single pole, low-pass filter, if $n\Delta T \ll T_f$, then $n_{\text{eff}} \approx 1$ [15]. If however

$$n\Delta T \gg T_f \tag{9}$$

$n_{\text{eff}} \approx (n\Delta T/2T_f)$. Note that this result is only true for $\Delta T \ll T_f$ and in any case, n_{eff} can never be larger than n, the number of samples recorded. Hence if a target can be sampled such that the product $n\Delta T$ is greater than $2T_f$ an improvement in the confidence in the range estimate results, since $\overline{\sigma}_r$ will be lower than σ_r.

The above criterion was used to reduce the range error in Fig. 11 where two 3-D scans are shown after systematic range error compensation (Section VI). The left scan shows a corner of an environment containing cupboards and a chest of drawers, each sample taken every 0.5° of sensor head rotation (every 0.7 ms).[9] Every four of these were averaged to form a single data point in the right range map. The improvement in the range variance is evident. In this case $4\Delta T = 2.8$ ms, which is about ten times larger than T_f, which in the design example is 0.28 ms, corresponding to a filter cut-off frequency of 3.5 kHz.

Quantification of Constraint 2: To produce these results with no loss of angular resolution, the time necessary to record one full 2-D section of a scan, T_{scan}, has a lower limit. This is because (for no resolution loss) the averaging should be restricted to a scanned area smaller than, or equal to, the optical footprint, i.e., expressed as a time constraint: $n\Delta T \leq$ optical footprint traversal time, $T_{\text{footprint}}$. From geometrical considerations [6]

$$T_{\text{footprint}} = \frac{bT_{\text{scan}}}{\pi R_{\max}} \tag{10}$$

where b = optical footprint radius at maximum range $R_{\max}$. Substituting for $n\Delta T$ in inequality (9) and resolving for the 2-D scan time T_{scan} gives

$$T_{\text{scan}} \gg \frac{\pi R_{\max} T_F}{b}. \tag{11}$$

[9]Since the sensor head was scanning at 2 rev/s about its vertical axis.

In the design example, $R_{\max}$ = 15.0 m, $T_F = 0.28$ ms, $b = 0.05$ m (beam radius at 15-m range), meaning that the lower limit for $T_{\text{scan}} = 0.26$ s. The 2-D scanning rate of 2 revs/s just satisfies this so that no loss in resolution is observed in the right scan of Fig. 11.

IX. Summary

Research with feature detection, localization and map building algorithms in mobile robotics, has inspired a thorough investigation into a commonly used sensor data source—the lidar. Before the construction (or purchase) of any such sensor takes place, it was demonstrated that the minimum detectable current induced in the receiver, for a worst case target, can be derived. The limiting factors, in the form of noise, attributed to the receiver electronics, which oppose the reliable detection of this current must be addressed. A theoretical quantification relating the power of the transmitter (and hence eye safety); the sensitivity of the receiver (and hence sensor cost); the receiver aperture area (and hence sensor size) and the receiver bandwidth (and hence possible range sampling and scanning speed) can give the user an informed choice of necessary sensor for a given application.

A solution for minimizing systematic range distortion as a function of the received signal intensity was given. For reliable sensor data manipulation in general, an invaluable quantity is the range variance. By considering the physics of noise propagation from the receiver to the output range estimate, it is possible to provide a model which produces a unique, provably correct range variance with each range value. An important point often overlooked by robotics researchers is that, in general, the naive determination of the output range voltage from AMCW or TOF lidars, as a function of the actual sensor to target range, provides a *false* calibration.

A simple second-order system which models the phase response of PLL systems to actual range changes between the sensor and its target was used to derive the speed constraints of the sensor as a function of the transmitted optical beam-width.

Finally, an interesting potential results if it is possible to densely sample the range output, and average several range values to produce a single estimate. In this way, it is possible to improve the range variance of a single range sample, producing a less noisy, lower density range scan. By manipulating the geometry of a lidar's optical footprint, it is possible to maintain the same angular resolution within the lower density scan, provided the geometrical and temporal constraints of Section VIII are adhered to.

Acknowledgment

The author would like to thank M. Brownlow for invaluable advice that he provided and the initial motivation for the electronic designs used in this work.

References

[1] M. Buchberger, K. W. Joerg, and E. Puttkamer, "Laser radar and sonar based world modeling and motion control for fast obstacle avoidance of the autonomous mobile robot MOBOT-IV," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1993, pp. 534–540.

[2] I. Kweon, R. Hoffman, and E. Krotkov, "Experimental characterization of the perceptron laser rangefinder," Carnegie Mellon Univ., Robotics Institute, Tech. Rep., 1991.
[3] J. Borenstein, H. R. Everett, and L. Feng, *Navigating Mobile Robots: Sensors and Techniques*. Wellesley, MA: Peters, 1995.
[4] M. D. Adams and P. J. Probert, "The interpretation of phase and intensity data from amcw light detection sensors for reliable ranging," *Int. J. Robot. Res.*, vol. 15, no. 5, pp. 441–458, 1996.
[5] G. L. Miller and E. R. Wagner, "An optical rangefinder for autonomous robot cart navigation," AT&T Bell Labs., Princeton, NJ, Tech. Rep., 1987.
[6] M. D. Adams, *Sensor Modeling, Design and Data Processing for Autonomous Navigation*, Singapore: World Scientific, 1999.
[7] Sick Optic Electronic, Sick Aktuell 2/97, Erwin Sick AG, CH-6370 Stans, Switzerland, 1997.
[8] M. D. Adams and A. Kerstens, "Tracking naturally occuring indoor features in 2-D and 3-D with lidar range/amplitude data," *Int. J. Robot. Res.*, vol. 17, no. 9, pp. 907–923, 1998.
[9] R. Jarvis, "A laser time-of-flight range scanner for robotic vision," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-5, pp. 505–512, 1983.
[10] F. K. Hopkins and J. T. Boyd, "Dark current analysis of insb photodiodes," *Infrared Physics*, vol. 24, no. 4, pp. 391–395, 1984.
[11] M. J. Brownlow, "A Time of Flight Optical Range Sensor for Mobile Robot Navigation," Ph.D. dissertation, Univ. Oxford, Oxford, U.K., 1993.
[12] K. Hyyppa, "On a Laser Anglemeter for Mobile Robot Navigation," Ph.D. dissertation, Lulea Univ. of Technol., Sweden, 1993.
[13] B. K. P. Horn, *Robot Vision*, 12th ed. Cambridge, MA: MIT Press, 1988.
[14] J. C. Dainty, *Laser Speckle and Related Phenomena*. Berlin, Germany: Springer Verlag, 1975.
[15] D. Nitzan, A. E. Brain, and R. O. Duda, "The measurement and use of registered reflectance data in scene analysis," *Proc. IEEE*, vol. 65, pp. 206–220, 1977.
[16] M. Hebert and E. Krotkov, "3-D measurements from imaging laser radars: How good are they?," in *Proc. Int. Conf. Intelligent Robots and Systems*, 1991, pp. 359–364.
[17] M. I. Skolnik, *Introduction to Radar Systems*. New York: McGraw-Hill, 1962.
[18] *GEC Plessey Semiconductors*, SL531 250 mHz True Log IF Amplifier, Tech. Rep., 1993.
[19] P. Vuylsteke, C. B. Price, and A. Oossterlinck, "Image sensors for real-time 3D acquisition, part 1," in *Traditional and Non-Traditional Robotic Sensors*, T. C. Henderson, Ed. New York: Springer Verlag, 1990, pp. 187–210.
[20] G. Allègre and H. Clergeot, "A two dimensional laser rangefinder for robot visiion and autonomy," in *Proc. Int. Conf. Intelligent Robots and Systems*, 1991, pp. 371–376.
[21] D. E. Maxwell, "A 5 to 50 MHz direct-reading phase meter with hundredth-degree precisioin," *IEEE Trans. Instrum. Meas.*, vol. 15, pp. 304–310, 1966.
[22] F. R. Connor, *Noise*. London, U.K.: Edward Arnold, 1982.
[23] F. R. Horowitz and F. R. Hill, *The Art of Electronics*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1993.
[24] J. A. Castellanos and J. D. Tardos, *Mobile Robot Localization and Map Building: A Multisensor Fusion Approach*, 1st ed. Norwell, MA: Kluwer, 2000.

Martin Adams received the first degree in engineering science at the University of Oxford, U.K, in 1988, and continued to study for the D.Phil degree at the Robotics Research Group, University of Oxford, which he received in 1992.

He then moved to Switzerland, where he continued his research as a Postdoctoral Research Assistant at the Institute of Robotics, Swiss Federal Institute of Technology, Zurich. He was a Guest Professor and taught control theory in Buchs, St. Gallen (Switzerland) from 1994 to 1995. From September 1996 to August 2000, the author was a Research Scientist in robotics and control at the European Semiconductor Equipment Centre (ESEC), Cham, Switzerland. He is currently an Assistant Professor at the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. His interests include mobile robot navigation, sensor design, data interpretation, vehicle modeling, and adaptive control.