MVA2005 IAPR Conference on Machine VIsion Applications, May 16-18, 2005 Tsukuba Science City, Japan
3-16

A Heterodyning Range Imager

M. J. Cree, A. A. Dorrington and D. A. Carnegie

Dept. Physics and Electronic Engineering
University of Waikato, Hamilton, New Zealand
E-mail: `cree@waikato.ac.nz`

Abstract

We describe a full-field range imager based on a heterodyning configuration. The scene is illuminated with a modulated light source. The received reflected light is mixed in a high-speed shutter with a frequency shifted version of the light source modulation. This results in a low frequency signal, with the range encoded as phase, that is easily measured with standard CCD machine vision camera technology. We report tests on repeatability and accuracy of measuring range in a small local region of pixels for an object placed in a 1–5 m range. We achieve 1 to 3 cm (standard-deviation) accuracy for measuring range when the modulation frequency is at 10 MHz.

1 Introduction

Simultaneous full-field active image ranging has been described in a number of papers [2, 5, 6, 7], either using an image-intensifier as a high-speed shuttering element or by implementation of specialised IC light sensor technology that includes a high-speed shutter as part of the IC design. The use of new IC designs is currently restricted to low resolution (160 × 124 pixels is the largest we are aware of [6]). Those that use image intensifier technology typically use a homodyne configuration (the high speed shutter is modulated or pulsed with the same signal as the light source) which then encodes range into the intensity of the signal received at the CCD camera. This method is severely limited by the dynamic range of current camera technology and range accuracies of centimetres is typically reported for ranging over distances of a few metres.

We describe here a simultaneous full-field active ranger that is based on the principle of heterodyning [4]. The hardware used is very similar to that described in the homodyne configurations. The major difference is that the high-speed shutter is driven with a frequency shifted version of the light modulation signal so that a low frequency signal, with the range encoded as phase, is received at the camera.

In this paper we report on the accuracy and repeatability with which we can range to an object at a range of 1 m to 5 m. Our current hardware can only be considered proof-of-concept as a number of components we use are sub-optimal. Nevertheless, we report ranging accuracy comparable to many other methods despite the fact that they are using more sophisticated hardware than our current prototype. With a number of easily identifiable hardware upgrades we predict the Waikato system is capable of an order of magnitude improvement in ranging accuracy than reported here [1].

Also our image intensifier has a burnt in image on the photocathode from previous studies and this seems to add extra significant phase offsets with a spatial distribution corresponding to the burnt in image. As we have not attempted to calibrate for this extra problem—a new image intensifier is on order!—we therefore do not present any full-field range reconstructions in this paper.[1] Only the reliability and repeatability with which range can be measured in a small local patch of pixels in the camera view is reported here.

2 Review of Known Technology

Simultaneous full-field image rangers acquire a full-field image in one measurement process without mechanical scanning. They normally use pulsing or modulation methods to encode range into some other signal parameter that is more straight forward to measure. There are many possibilities here, each having pros and cons.

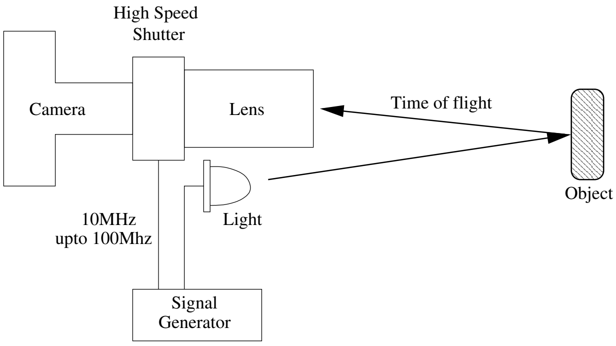

Figure 1. Schematic of a Full-Field Ranger

A generalised hardware for modulated ranging is shown in figure 1. A light source, illuminating the scene, is continuously modulated at a frequency f. The light arriving back at the ranger from the scene is delayed by a phase angle ϕ given by

$$\phi = 2\pi f\tau = \frac{4\pi f d}{c} \qquad (1)$$

---

[1] We hope to have some full field range reconstructions ready for presentation at the conference itself.

where $\tau = 2d/c$ is the time-delay of the signal, d is the range to the point in the scene under consideration and c is the speed of light. For full-field imaging ϕ should be considered a function of 2D position. The light impinges on a high speed shutter that is also modulated, and the two signals mix to produce a low-frequency signal with the phase delay ϕ encoded in it. This means that a standard CCD camera sensor, which typically can only sample light signals thirty times a second, can be used for the light sensing element, despite the fact that the light may be modulated at radio frequencies (typically at 100 MHz).

Standard radar/lidar techniques are amenable to image rangers, and phase-shift-keying (PSK) and frequency stepped continuous wave (FSCW) have been reported in the literature [5]. Both these methods cope well with different received signal intensity due to varying scene conditions (colour, texture, etc.), and FSCW has the added benefit of being able to distinguish multipath effects. However, both require very linear electronics and optics: any harmonics introduced ruin the range reconstruction. This is a fundamental problem and the use of these modulation methods have only been reported by researchers developing specialised light sensitive IC arrays that include the high-speed shutter process in the IC design. As the image-intensifier, which we use as a high-speed shutter, is non-linear, these modulation techniques are not suitable with our hardware.

Scott [7] and Christie *et al.* [2] report a homodyne configuration, in which the high speed shutter is modulated with exactly the same signal as the light source. In this case the phase delay is encoded into intensity of light received. This is complicated by the differing intensities of light reflected due to the scene colour and texture, and normally three or four images are captured with 90° or 180° phase shifts introduced to cancel out background illumination and intensity variations inherent in the scene. The major problem with this method is one of dynamic range, as CCD cameras typically have 8–12 bit resolution, and taking off an estimated 3–4 bits for natural scene intensity variations, one is left with range encoded with a limited number of bits of accuracy. Over a measurement range of 10 m one can really only hope to achieve measurement precision of the order of centimetres at best. Christie *et al.* [2] report 2 cm measurement uncertainty when measuring over a range of 0.8–1.8 m.

What has not been considered before is the possibility of a heterodyne configuration for a full field ranger. Heterodyning is a well known technique in lidar and laser applications to measure range. However, these applications have always used light sensors capable of measuring high frequency modulation (the mixing is done electronically) or have been restricted to a single beam of light. It is also interesting to note that heterodyning has been used in full-field imaging to measure fluorescence lifetime decay times (of the order of nanoseconds) in fluorescence lifetime imaging microscopy (FLIM) [3, 8]. Indeed, FLIM uses a hardware configuration similar in many details to that we describe below for our range imager. What has not been realised before now, is that this technique can be used for high accuracy ranging in a full-field ranger.

In the heterodyne configuration, the high speed shutter is modulated at a frequency f_2 that has a small offset to the modulation frequency f_1 of of the light source, namely $f_2 = f_1 + \delta f$. The two signals mix in the high speed shutter to produce a low frequency signal of frequency δf in which the phase of the incoming signal (from the scene) is preserved in the low frequency signal after mixing. Thus to range it suffices to recover the phase of the low frequency signal that is easily sampled by a machine vision camera. This has the advantages that the phase delay of light propagation is encoded in a continuous variable (phase) that is not subject to dynamic range limitations by analogue-to-digital conversions, and unlike the standard radar methods, is robust against harmonic distortion in the measured signal.

3 The Waikato Ranger

Our hardware implementation follows that described by Scott [7] and Christie *et al.* [2]. A bank of LEDs is used to illuminate the scene. These are easily modulated at radio frequency. A Photek 25 mm diameter, 1 MCP, image intensifier is used as the high speed shutter as it can be modulated up to 30 MHz. A Nikon (manual focus) 80–200 mm focal length camera lens forms an image of the scene on the photo-cathode of the image intensifier. A Pulnix TM9701 CCD camera views the phosphor screen of the image intensifier via a 25 mm fixed focal length lens with a 5 mm spacer between the camera and lens. This enables a working distance of 80–90 mm between camera and image intensifier, with the image of the image intensifier just touching the top and bottom of the CCD field-of-view. Electronics currently limit the Waikato ranger to operation at 10 MHz, giving an unambiguous ranging distance of 15 m.

3.1 Signal Processing

There are a large number of algorithms to detect the phase of a sinusoidal signal [9]. Many, such as the phase-stepping algorithms used in profile measurement with fringe patterns, are only suitable for pure uncontaminated signals. As components in the Waikato Ranger, most notably the image intensifier, are both noisy and non-linear, we require an algorithm that can detect the phase of a single frequency in the presence of harmonics and noise. Fourier based methods are commonly used in such situations and high accuracy phase measurements have been reported in the literature [9].

A straight-forward discrete Fourier transform (DFT) of the received heterodyne beat signal is used to isolate the fundamental frequency of the beat. The angle of the complex quantity of the bin corresponding to the fundamental frequency gives the phase shift of the signal, and hence, via eqn. 1, the range to the point of the object that reflected the illumination. The Waikato ranger produces a video sequence sampled at 29.97 frames per second (fps) with a resolution of 768 × 484 pixels. The image intensifier phosphor screen fills an approximately 512 pixel diameter circular region of the camera view. The DFT is performed down the time axis of the video sequence for each pixel in the camera view. A range value is thereby obtained independently for each pixel.

4 Testing Methodology

Currently the camera start of acquisition cycle cannot be synchronised to the start of the heterodyne beat frequency, thus any captured sequence has an arbitrary constant phase added across the whole camera view. We therefore always put a reference object (a small flat surface) in the field-of-view at a fixed known distance. This is used to determine the unknown constant phase and subtract it off the whole reconstructed view.

The characteristics of the ranger have been measured over a range of 1 m to 5 m in three staged experiments as the complete range could not be measured in one experiment due to the limited depth-of-field of the camera lens. Range was measured in one experiment over 1–1.75 m, in another over 1.55–2.75 m, and in the third over 2.55–5 m. For each of these experiments the camera focus was set to the middle of the range under study. All experiments were performed with a modulation frequency of 10 MHz and with heterodyning beat frequencies of 1 Hz, 2 Hz and 5 Hz.

In the 1–1.75 m study a reference block was placed at 0.800 m range adjacent to the optic axis of the camera view, and a moveable block was placed at distances between 1.000 m and 1.750 m in 0.250 m increments. The moveable block was placed as close as possible to the camera optical axis. Three separate video sequence of 300 samples (i.e. 10 s long) were taken for the moveable block at each of its distances.

The same procedure was repeated for the other two studies, except with the reference block at 1.500 m for the 1.55–2.75 m study, and with the reference block at 2.500 m for the 2.55–5 m study. The moveable block was moved forward from the furtherest position in 0.250 m increments except for the last step in which it was moved by 0.200 m, hence giving the nearest 1.550 m and 2.550 m ranges.

Straightforward Fourier analysis of the signal in each pixel was used to reconstruct range images. A region of interest (ROI) of between 1600 and 2500 pixels was used for the reference block for each of the three studies. This ROI was kept the same within a study as the reference block always appears in the same place in each image. The mean and standard deviation of range was calculated over this ROI, and a constant range offset was subtracted off the range image to shift the mean range of the ROI of the reference block to the correct range (0.800 m, 1.500 m or 2.500 m). Another ROI of 100 pixels was assigned to the moveable block. This ROI was kept the same for each video sequence of the moveable block over a study. The mean and standard deviation of measured range over the ROI is calculated for the moveable block for each range reconstruction. We then add an extra constant offset to all the results of a single study to exactly align the common measured range between two studies. That is, for example, for the first (1.00–1.75 m) study, we add a constant offset to all distances measured so that the range measured for the 1.75 m point is identical to the 1.75 m measurement of the second (1.55–2.75 m) study.

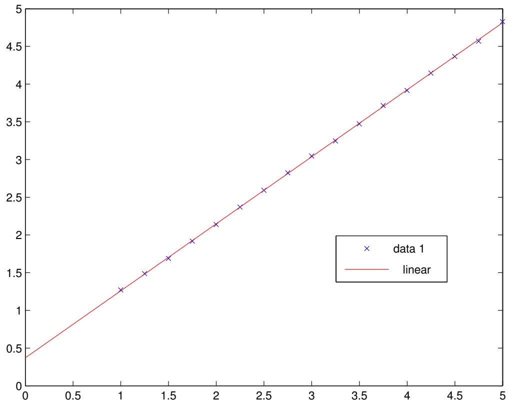

Figure 2. Measured range to a moveable target versus actual range

5 Results

The results of ranging to a small ROI on a flat object moved between 1 m and 5 m range is plotted against known range in figure 2. A line, obtained by standard linear regression, is plotted through all data points. This line has a slope of 0.89, indicating a miscalibration, between measured range and actual range. We repeated this study three times, and also have performed similar experiments with a 2 Hz beat frequency and a 5 Hz beat frequency. The same anomalous slope (varying between 0.89 and 0.93 for the various studies) is present. We have moved the position of the ROI in the image for analysis to be 50 pixels off the optical axis and the anomalous 0.9 slope is still present (this eliminates the possibility of off-axis distortions due to the increased path length the light must travel). These results are quite repeatable and ranges are measured to an accuracy of 1–3 cm standard deviation error depending on the intensity of the signal received.

6 Discussion

The results demonstrate good repeatability and reasonable accuracy of ranging (1–3 cm standard error), but are marred by an anomalous slope of 0.9 when measured range is plotted against actual range. In principle this anomaly could be calibrated for and removed from the results since the results are quite repeatable under a number of conditions, however a thorough understanding of the anomaly is a better approach. Thorough checking has confirmed that the measurement and signal generation equipment used is well calibrated, that errors due to changing path length of the light signal as one moves off the optical axis of the ranger are far too small to be the explanation, and that the reconstruction algorithm gives a correct slope of 1.0 for simulated data.

A theoretical analysis of heterodyning shows that the delayed phase of the signal is dependent on only two things: the range to the object and the modulation frequency. We conclude that contamination of the range signal being received at the image intensifier is occurring. Multipath effects can be excluded (light being reflected off some other object, such as a wall, onto the scene) as these increase the measured range. Experiments have been repeated with extra light blocking material over nearby reflecting objects and in between the ranger and the LED light source and the same anomalous 0.9 slope is measured. It is now believed that the problem is direct electromagnetic inductive coupling between the LED light source and the image intensifier. Theoretical estimations of how measured range is affected by a partial direct coupling of the light source signal to the image intensifier, that we are carrying out at present, confirms this hypothesis. Better electromagnetic shielding of signals and equipment should therefore fix the anomaly.

The Waikato ranger described here can be improved in a number of ways, some of which will make substantial improvements to ranging resolution. This is very significant, since with off-the-shelf prototype equipment and no attempt at refinement, we have achieved ranging precision similar to that reported elsewhere in the literature where much more sophisticated hardware has been used. Improvements include increasing the modulation frequency to 100 MHz which will give an order of magnitude improvement in precision (i.e. into the millimetre range). The coupling between the image intensifier and the camera can be improved, either by reducing the working distance with a higher quality lens, or by using direct fibre coupling to the CCD element of the camera. Currently between 4 to 5 bits of resolution in the received light signal intensity is being obtained and it is known that such low bit resolution has an adverse effect on measuring phase [10]. By using low noise and more sensitive camera technology with a 12 bit ADC, better quality signals would be obtained, allowing significantly better phase estimations. We are currently working on building a high precision ranger to image in the 1–10 m range in a lit room, that will achieve 1 mm precision in range as well as x-y position.

7 Conclusion

All full-field rangers based on image intensifiers have in the past used homodyne modulation. The range derived from such an approach is severely limited by the dynamic range of the digitising camera. We described an image intensified ranger which was run in a heterodyne configuration. Experimentation with modulation at 10 MHz and with measured signals of 4 to 5 bits in intensity achieved 1 to 3 cm range resolution, depending on the heterodyne beat frequency, over distances of 1 to 5 m. A systematic anomaly in the measured range compared to the actual range in the results is believed to be due to electromagnetic coupling between the LED light source and the image intensifier. With a number of relatively easy improvements to the hardware it is believed that a ranging resolution of 1 mm, with near real-time acquisition, is achievable.

References

[1] D. A. Carnegie, M. J. Cree, and A. A. Dorrington. An accurate real-time range finding system. *Rev. Sci. Inst.*, 2005. Submitted for publication.

[2] S. Christie, S. L. Hill, B. Bury, J. O. Gray, and K. M. Booth. Design and development of a multi-detecting two-dimensional ranging sensor. *Meas. Sci. Tech.*, 6:1301–1308, 1995.

[3] R. Cubeddu, D. Comelli, C. D'Andrea, P. Taroni, and G. Valentini. Time-resolved fluorescence imaging in biology and medicine. *J. Phys. D*, 35:R61–R76, 2002.

[4] A. A. Dorrington. Range sensing system, 2003. New Zealand Patent 525241.

[5] P. Gulden, M. Vossiek, P. Heide, and R. Schwarte. Novel opportunities for optical level gauging and 3-d-imaging with the photoelectonic mixing device. *IEEE Trans. Instr. Meas.*, 51:679–684, 2002.

[6] T. Oggier, R. Kaufmann, M. Lehmann, P. Metzler, M. Schweizer G. Lang, M. Richter, B. Büttgen, N. Blanc K. Griesbach, B. Uhlmann, K.-H. Stegemann, and C. Ellmers. 3D-imaging in real-time with miniaturized optical range camera. In *OPTO 2004 Conference*, Nürnberg, Germany, May 2004.

[7] M. W. Scott. Range imaging laser radar, 1990. US Patent 4,935,616.

[8] E. M. Sevick-Muraca and D. Y. Paithankar. Fluorescence imaging system and method, 1999. US Patent 5,865,754.

[9] P. Shea. Phase measurement. In J. G. Webster, editor, *Electrical Measurement, Signal Processing and Displays*, chapter 5. CRC Press, Boca Raton, Fl, 2003.

[10] B. Zhao and Y. Surrel. Effect of quantisation error on the computed phase of phase-shifting measurements. *App. Opts.*, 36:2070ff, 1997.