Modulated Beam Photodetector
For
Mobile Robots

Peter Hiscocks

Department of Electrical and Computer Engineering
Ryerson University

phiscock@ee.ryerson.ca

November 13, 2002

Contents

1 Introduction

This paper is an exploration of using a modulated light beam method as the basis for an object detector useful in a mobile robot.

Mobile robots use their object detectors to, well, detect objects. These can be objects that are to be avoided (such as walls and bigger robots) or objects that are to be pursued (such as another sumo robot, or an object to be cooperatively pushed to a destination). How the detection is used is up to the path management software.

Various object detection devices and schemes are used at present. Some of the main candidates are:

Unmodulated Infared Illumination Kube [1] shows a simple approach to an unmodulated infared beam illuminator and detector. His arrangement claims a range of 10 inches, but it is not clear whether this is under conditions of strong ambient illumination.

Triangulation Beam Detector This device, Sharp GP2D02, uses special optics to generate a distance reading over the range of 10 to 80 centimetres. It requires a microprocessor for interfacing. The device is relatively expensive at about $20 CDN each.

Sonar Sonar echo-ranging, as described in [3], is a popular method of object detection. The circuitry is relatively complex, generally requires a microprocessor for interfacing. It also has some ambiguity problems with echoes and transmitter sidelobes.

Bumpers and Whiskers These mechanical devices are generally reliable and couldn't be simpler to interface. However, they are limited in range and can present problems in placement. A whisker switch is a fairly bulky device.

The object of this exercise was to develop an object detection method that

- was simple to implement
- could be used without a microprocessor
- had the maximum possible range

Notwithstanding the common use of infared emitters and detectors for simple optical ranging applications, the author prefers to use visible light. Being able to see where the light is going is an invaluable aid to debugging. Infared is misleading under some circumstances: surfaces that we would expect (based on our experience of visible light) to be reflective are not, and vice versa. There is no such ambiguity with visible light.

2 Photodetectors

Illumination may be detected with a Cadmium Sulphide CdS photoresistive cell, a visible-light photodiode or a phototransistor.

Cadmium Sulphide CdS Photoresistive Cell The CdS photoresistor is sensitive to visible light but responds relatively slowly (20 milliseconds response time is typical) to changes in illumination. As a result, it integrates ambient 120Hz illumination into its average level. Generally, there is no optical system associated with the CdS cell, and so it integrates the illumination over a wide receiving angle. This makes it useful for detecting overall illumination level, but difficult to use for object detection at a distance. The CdS cell is effective in a line-following application, but it has to be very close to the working surface [4].

It would be interesting to use the CdS cell with an optical focussing system. This might well make it useable for detection at large distances.

Visible Light Photodiode A light emitting diode will function as a visible light photodiode, in which case the reverse-biased current is proportional to illumination. The device is low sensitivity, requiring that it be used with a high-gain (eg, $5M\Omega$) transresistance amplifier. The response is very fast, up into the MHz region. To take advantage of this, the transresistance amplifier use a fast operational amplifier. A high intensity LED is equipped with a narrow beamwidth lens system. When it is used as a photodiode detector, this helps with discrimination against optical noise.

Phototransistor The device used in this investigation was the 2N5777, which is a darlington NPN phototransistor with a relatively wide beamwidth (70% relative response) of 45 degrees. A device like the BPW38 photodarlington, with a beam width of 5 degrees, might be a better choice.

The switching speed is in the order of hundreds of microseconds. This is problematic when the device is used to drive a resistor directly, because the Miller effect severely limits the speed of response. However, when it drives a transresistance amplifier, the voltage across the transistor does not change, and the response is relatively much faster. In this configuration, a switching frequency of 2KHz is quite feasible, and more limited by the frequency response of a low-cost operational amplifier. The required transresistance gain is much more modest than required when using a photodiode, about 10kΩ.

3 Rejecting Ambient Illumination

The photodetector should have some mechanism for ignoring ambient illumination. The interfering illumination takes two forms: a zero frequency (DC) component due to sources like sunlight, and an component at 120Hz due to the AC lighting current in artificial illumination sources, such as incandescent lamps. These sources may be treated as noise that worsens the signal-noise (S/N) ratio at the detector.

The object detector is assumed to have a local source of illumination, typically a high-intensity red LED or infared LED. So the S/N at the detector, which should be as large as possible, depends on the ratio of the total illumination into the detector vs the component generated by the local source.

Optical Methods

Two optical measures can be taken to improve the signal-noise ratio.

Restricted Beamwidth If the beamwidth of the detector is restricted, then the S/N is improved. This requires that the detector be able to see the illumination generated by the local source.

The beamwidth can be decreased by using a lens with a larger aperture than the detector or by restricting the field-of-view of the detector with a tube or pinhole (an *aperture stop*). Since a lens gathers light over a larger area, it will generate a larger signal at the detector. However, if there is adequate signal but the S/N is poor, an aperture stop will be sufficient. Put another way, a lens will extend the sensitivity of detection because it collects more energy. An aperture stop will improve the signal-noise ratio, but not extend the sensitivity of the detector. The aperture stop will work where there is a strong signal but poor S/N. The lens will work where the signal is weak.

Colour Filtering A narrowband colour filter restricts the bandwidth of the light entering the detector to that of the local source. By rejecting ambient light that is not at the same wavelength as the local source, it improves the S/N. A comprehensive examination of contrast enhancement by filtering is in [5].

Zero Frequency Ambient Light

The object detector is assumed to have a local source of illumination, typically a high-intensity red LED or infared LED.

There are two methods of rejecting the zero frequency (steady-state) component of background illumination.

Subtraction of Background The simplest method of rejecting the zero frequency component of ambient illumination is to take a reading of the background level with the illuminating LED switched off, and then subtract this from a reading with the illuminating LED switched on. If there is an object in the field of view, the second reading will be larger. The difference value is proportional to the reflectivity and distance of the target.

This method is most easily performed with a microprocessor which senses the detector level via an A/D converter. It can also be done analog circuitry with two sample-hold circuits that directs save the two readings to storage capacitors. The voltages of the capacitors are then subtracted to get the difference reading.

The method only works where the light due to the illuminating LED exceeds the background level. Obviously, this also works better if the field-of-view of the detector matches that of the illuminator LED.

High Frequency Switching If the illuminating LED is switched at some frequency, the detected signal may be AC coupled to remove the zero frequency component of ambient illumination **providing** that the ambient illumination is not so strong as to saturate the detector. This requirement usually results in a compromize between sensitivity of the detector and immunity to ambient light.

120Hz Ambient Light

Where the light source is operated from the 60Hz line, there will be a strong component of illumination at 120Hz. This component can be removed by averaging over one complete AC cycle or by high-pass filtering.

Averaging Over One Cycle The average value of an AC waveform is zero. Consequently, if a series of detector readings are taken over one complete cycle and then averaged, the result is the average or DC component of the waveform. In fact, two readings will suffice if they are taken $T_{120}/2$ seconds appart, where T_{120} is the period of a 120Hz waveform (8.3 milliseconds). This is most easily done with a microprocessor. Consequently, if you were using the subtraction method to detect the presence of a reflection, a good strategy would be

- take readings 8.3 milliseconds appart with the local source turned off. This gives you the DC component of ambient light.
- take readings 8.3 milliseconds appart with the local source turned on. This gives you the DC component of ambient light plus the local illumination source.

High Pass Filtering If the local source is chopped (switched) at some frequency f_s such as 2400Hz, then it may be separated from the 120Hz component f_n by high pass filtering. In practice, the 120Hz AC component can be much larger than the desired signal, so fairly severe (ie, multi-pole) filtering is required.

For example, I found that there was 80 millivolts of AC noise at 120Hz (f_n) in the presence of 10 millivolts of signal at f_s. We'd like to amplify the component at f_c, so that means the AC component at f_n must be substantially attenuated. In this case, if the noise component is to be reduced to the same amplitude as the signal, the noise would be attenuated by a factor of 8, or 18db. A two pole filter has an attenuation of 20db per decade, so a 2-pole Sallen-Key high pass filter with a cutoff frequency f_c of 1200 Hz would work. (See [6] for design of the Sallen-Key filter.)

The response of the high-pass filter is down 3db at the cutoff frequency and the phase shift is plus 45 degrees (ie, the output signal leads the input by 45 degrees), so it's desireable to have the signal frequency f_s at twice the filter cutoff frequency, where the attenuation and phase shift are much less.

Alternatively, one or more bandpass filters, centred on the signal frequency, may be used [7]. Whatever filter is used should be simulated to check that it's actual response is as required.

4 Baseline Correction

A *baseline correction* circuit is a feedback circuit that maintains the output of the sensor stage at some average value. It does this by adding to or subtracting from the sensor bias current, at such a frequency that it has no effect on the signal frequency. The basic sensor transimpedance amplifier circuit is shown in figure 1.

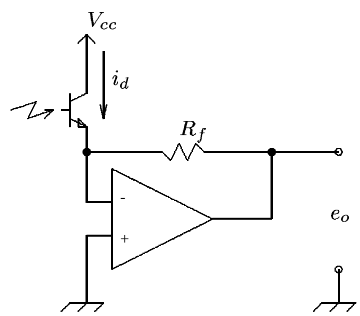

Figure 1: Transimpedance Amplifier

The output of this amplifier is simply

$$e_o = -i_d R_f$$

where i_d is the detector current induced by illumination on the phototransistor. Notice that the voltage of the emitter of the phototransistor is a virtual ground point, so the op-amp holds it at close to zero volts. The collector voltage of the transistor is V_{cc}, so the voltage across the transistor is fixed, which allows it to switch quickly.

The value of feedback resistor R_f is a compromise. On the one hand, it should be made as large as possible, to make the sensor stage output as large as possible. However, if there is a high level of ambient light, this will drive an excessive current through the phototransistor, which will drive the output of the op-amp to it's lower bound. A smaller value of R_f reduces the likelyhood of this happening.

The baseline correction circuit is shown in figure 2.

An slowly-responding integrator is connected to the output of the transimpedance amplifier so that it generates a signal proportional to the average value of the output. If the output e_o is non-zero, the integrator will drive current into the summing junction to restore e_o back to zero volts.

For example, if the ambient light level changes, then the value of e_o would go negative. The integrator output e_f gradually ramps negative, drawing more current out of the summing junction until e_o is restored back to zero volts.

Since the integrator responds slowly, it ignores high frequency signals such as the chopped light signal which appears at the output as previously. So the steady component of illumination is removed from the signal but the high frequency component remains.

A more complete circuit is shown in figure 3.

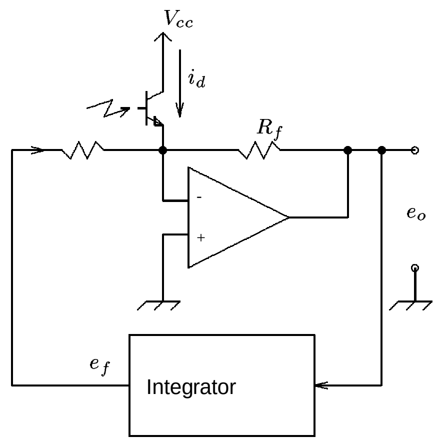

Figure 2: Baseline Corrector Circuit

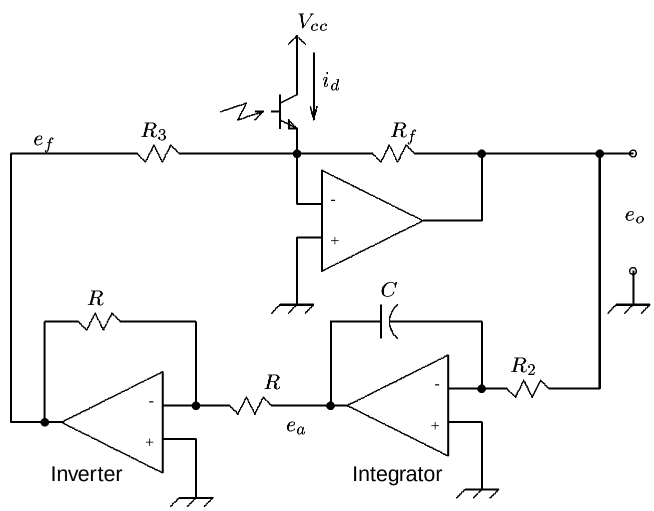

Figure 3: Baseline Corrector Circuit

An analysis of this circuit shows that the transfer function is

$$\frac{e_o}{i_d} = -\frac{sRC}{1 + s\frac{R}{R_f}C}$$

where $R = R_2R_3$ and $s = j\omega$. At high frequencies (where $s\frac{R}{R_f}C$ is large compared to 1), this reduces to $e_o/i_d = -R_f$ as before. However, R_f can be made larger because the baseline corrector will prevent saturation as the ambient light level increases.

This transfer function is a high-pass filter with cutoff frequency at

$$f_{cutoff} = \frac{R_f}{2\pi R_2 R_3 C} \text{ Hz}$$

This cutoff frequency should be positioned well below the chopped signal frequency. In the prototype, we used $R_f = 10k\Omega$, $R_2 = 1M\Omega$, $R_3 = 100\Omega$, $C = 100nF$ giving a cutoff frequency of 159Hz, well below the signal frequency f_s of 2400Hz. If the cutoff frequency was moved to something like 1000Hz, it would have helped reject the 120Hz ambient illumination. However, a different technique was used for that, as described in the next section.

Similar circuits are shown in [8] and [9].

5 Differential Current Technique

A differential techique can be used for cancelling out background illumination, whether zero frequency or 120Hz. Two phototransistors are used. One (the *reference* device) detects the background illumination. The other (the *detector* device) measures the background illumination plus the desired signal. The difference between these two measurements is the desired signal.

When using a transresistance amplifier, adding the reference device is as simple as adding a second phototransistor, as shown in figure 4.

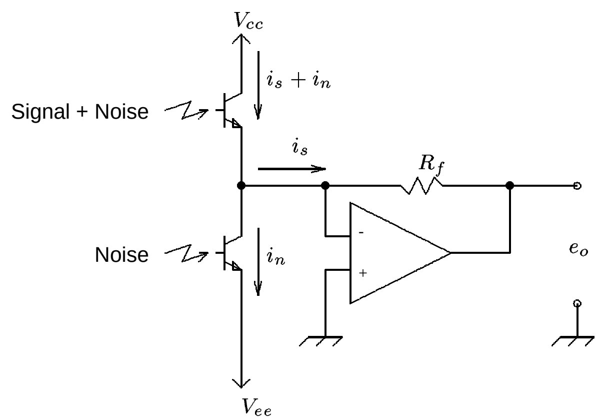

Figure 4: Differential Technique

When the reference terminal of the op-amp is grounded, the lower transistor must be biased by a negative supply V_{ee}.

In this arrangement, both transistors conduct a current proportional to the ambient illumination, but the feedback resistor conducts only the difference (signal) current, so it can be made larger than if it had to conduct current due to the ambient illumination.

6 Complete Circuit

A complete photodetector system, using the preceeding concepts, is shown in figure 5.

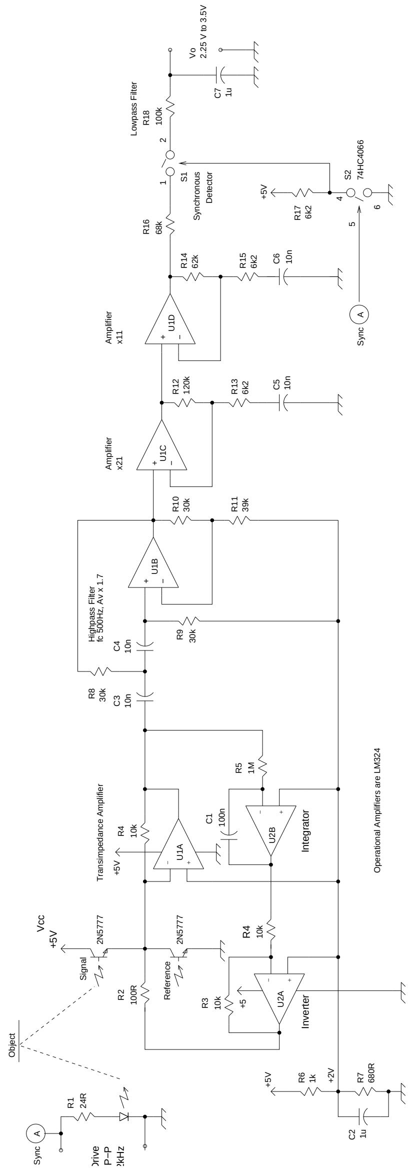

Figure 5: Photodetector System

The circuit is operated from a single +5 volt Vcc supply, so a 'pseudoground' voltage of 2 volts is generated by the voltage divider R6-R7 to act as a reference level between ground and Vcc. This establishes the level at the op-amp non-inverting terminals, so this becomes the reference for the various circuits. (The pseudoground is less than 2.5 volts because the maximum output of the LM324 amplifier is about 4 volts. The pseudoground is set to half this.)

The basic detector uses a baseline-corrected transimpedance detector with two phototransistors in a differential current configuration. This makes the detector stage relatively impervious to ambient light, though it can be disabled by an incandescent lamp with a foot or two of the sensors.

The signal is then filtered through a Sallen-Key highpass filter, designed using the information in [6]. The filter removes any remaining 120Hz component. It is biased in such a way that its output is established at a DC value of 2 volts.

The signal is further amplified in two AC coupled amplifiers, for a total gain through the highpass filter and two amplifier stages of approximately 400 volts/volt. The gain bandwidth of the LM324 is such that it has about 40db (x100v/v) of open-loop gain at 2KHz. The closed loop gain must be much less than this, so the gain must be spread over at least two stages of amplification. The amplifiers have a DC gain of unity, so that the DC level established at the output of the highpass filter propagates through to the output of U1D.

The amplified signal output is synchronously detected by an electronic switch S1 which is synchronized to the original 2KHz drive signal. The following lowpass filter R18-C7 averages the output of the switch with a time-constant of about 0.2 seconds (5 Hz cutoff frequency). The effect of the synchronous switch is to translate this 5 Hz filter to the 2KHz signal frequency and remove components below 1995Hz and above 2005Hz. Consequently, the signal-noise ratio is greatly improved.

Alternatively, one can view the switch as a modulator. In that case, the signal and the sync frequency are multiplied together, producing the sum frequency at $2 \times f_s$ and a difference frequency at zero Hz (DC). So if there is a signal frequency present, then the output will have a DC component.

One can also view this as a simple correlator circuit, in which the incoming signal is correlated against itself. If there is a match, then the result is a non-zero average during the switching interval, and a DC output will occur. If there is no correlation, the average voltage during the switching interval will be zero.

The net result is that any frequencies other than the signal frequency are essentially ignored, and the signal frequency is converted into a DC value proportional to its amplitude.

A spare switch S2 in the 74HC4066 package is used as an inverter to control the phase of the correlation switch S2.

Notice that the phase shift through the high pass filter must be minimal so that the signal waveform is in sync with the reference waveform that operates S1.

7 Results

The output level varied from approximately 2.25 volts for no object detected to 3.5 volts for an object close to the emitter and detector.

Without much attempt at optimization, the circuit can reliably detect a white square of paper, about 2 inches on a side, at a distance of 15 inches or so. Somewhat arbitrarily, the threshold for distance detection was taken as a change of approximately 70 millivolts in the output level.

The system works in the presence ambient illumination, which was tested by a bench mounted Luxo lamp. There is a small shift in output DC level (about 20mV typically) when the lamp is turned on and off. If the lamp is moved closer than within 12 inches of the detectors, the system stops operating. This is without any mechanical shielding or optical focussing of the detectors, both of which would undoubtably improve performance.

The system is sensitive to signal frequency on the power supply line. Consequently, if you attempt to drive the emitter LED (which requires substantial current) from the same supply as the detector circuit, it will malfunction. The emitter needs its own supply, or at least its own regulator.

The summing junction of the transimpedance amplifier is extremely sensitive to noise. It may be necessary to place a metal area under that point and connect the area to ground potential.

The resultant circuit is rather large, though it is made up of simple building blocks. Undoubtably, techniques that require less circuitry will be preferred in many applications. However, it the circuit illustrates several useful concepts and points the way to achieve long-range object detection using visible light.

References

[1] *A Minimal Infared Obstacle Detection Scheme*
C. Ronald Kube,
The Robotics Practitioner, Spring 1996, pp15-20.

[2] *Distance Measurement using Infared*
Elector Electronics, 7-8/2000, pp 88-90

[3] *Rotating Sonar: Modelling and Application*
Hong Zhang et al
1997 International Conference on Field and Service Robotics

[4] *eebot Technical Description*
Peter Hiscocks, 2002
`http:\\www.ee.ryerson.ca\˜phiscock`
A complete technical description of the *eebot* mobile robot used in a microprocessor laboratory course.

[5] *Contrast Enhancement Techniques*
Hewlett Packard Application Note 964, 1975

[6] *Active Filter Cookbook*
Don Lancaster
Howard W. Sams, 1975

[7] *Applications of Phototransistors in Electro-Optic Systems*
John Bliss
Motorola Application Note AN-508, 1970.

[8] *OPT201 Photodiode-Amplifier Rejects Ambient Light*
Mark Stitt and Wally Meinel
Burr-Brown Application Bulletin AB061, 1993

[9] *Wideband Photodiode Amplifier*
Michele Frantisek
Wireless World Magazine, Date unknown.

[10] *Small Aperture IR Optical Links Using LED Light Sources*
Lawrence E. Foltzer
QEX, August 1990, pp9-13