Table 2

Data bus	Command bus	Operation
03	02	Load '03' into high byte register.
25	10	Load Position Data Bus into Θ_1 adder.

The switches S1 and S2 control overall operation of the interface. The data path from the Position Data Bus can be broken by closing S1, which causes the robot arm to halt at its current position. In addition, the adders can be cleared by closing S2, which functionally restores the robot controller to its original state, without the adders.

4. Sensitive skin and sensor interface

4.1. Sensor interface

The purpose of the sensor interface circuit is to allow computer access to the sensor. Its two major components are the analog to digital converter, and the one-shots that control the addressing of a particular sensor. The overall sensor interface block diagram is shown in *Fig. 7*, and the circuit that implements the sensor interface is shown in *Figs. 8* and *9*. Points labeled 'A', 'B', and 'C' on *Fig. 8* are connected to the corresponding points in *Fig. 9*.

Communication to the sensor system from XPDCS 2 is handled by one link adaptor, which is connected with twisted pairs for noise immu-

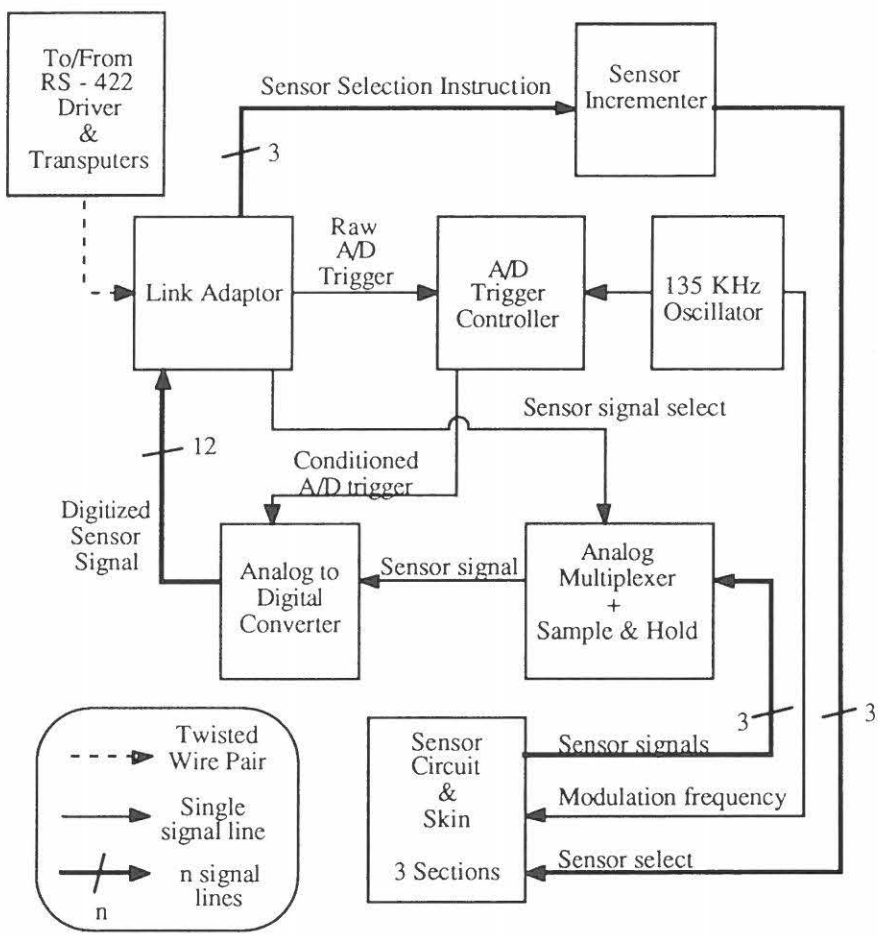

Fig. 7. Sensor interface block diagram.

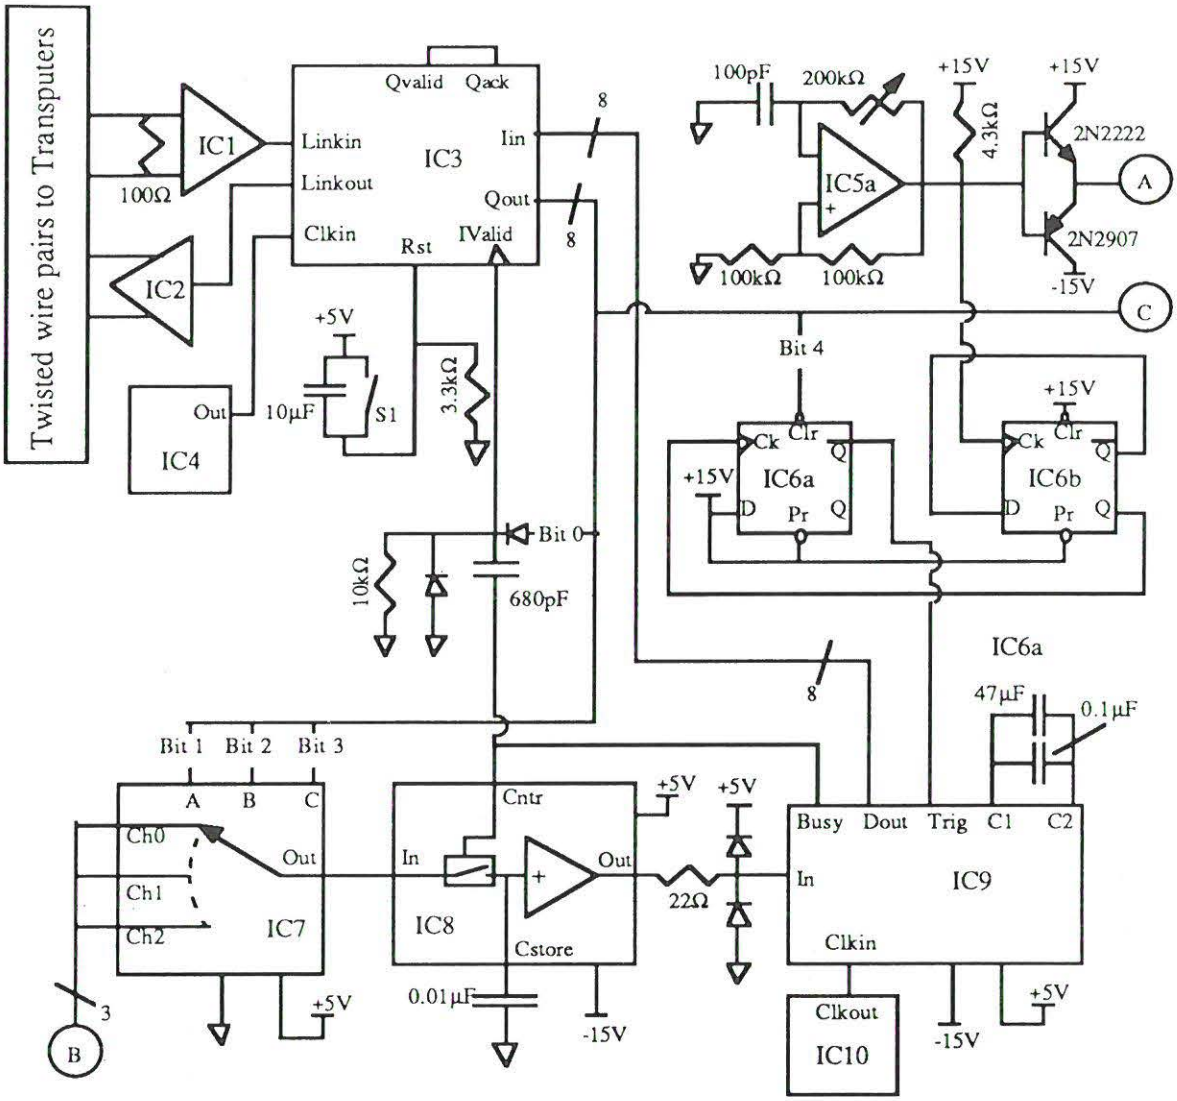

Fig. 8. Sensor interface circuit, part 1. Points A, B, C refer to the corresponding points on *Fig. 9.*

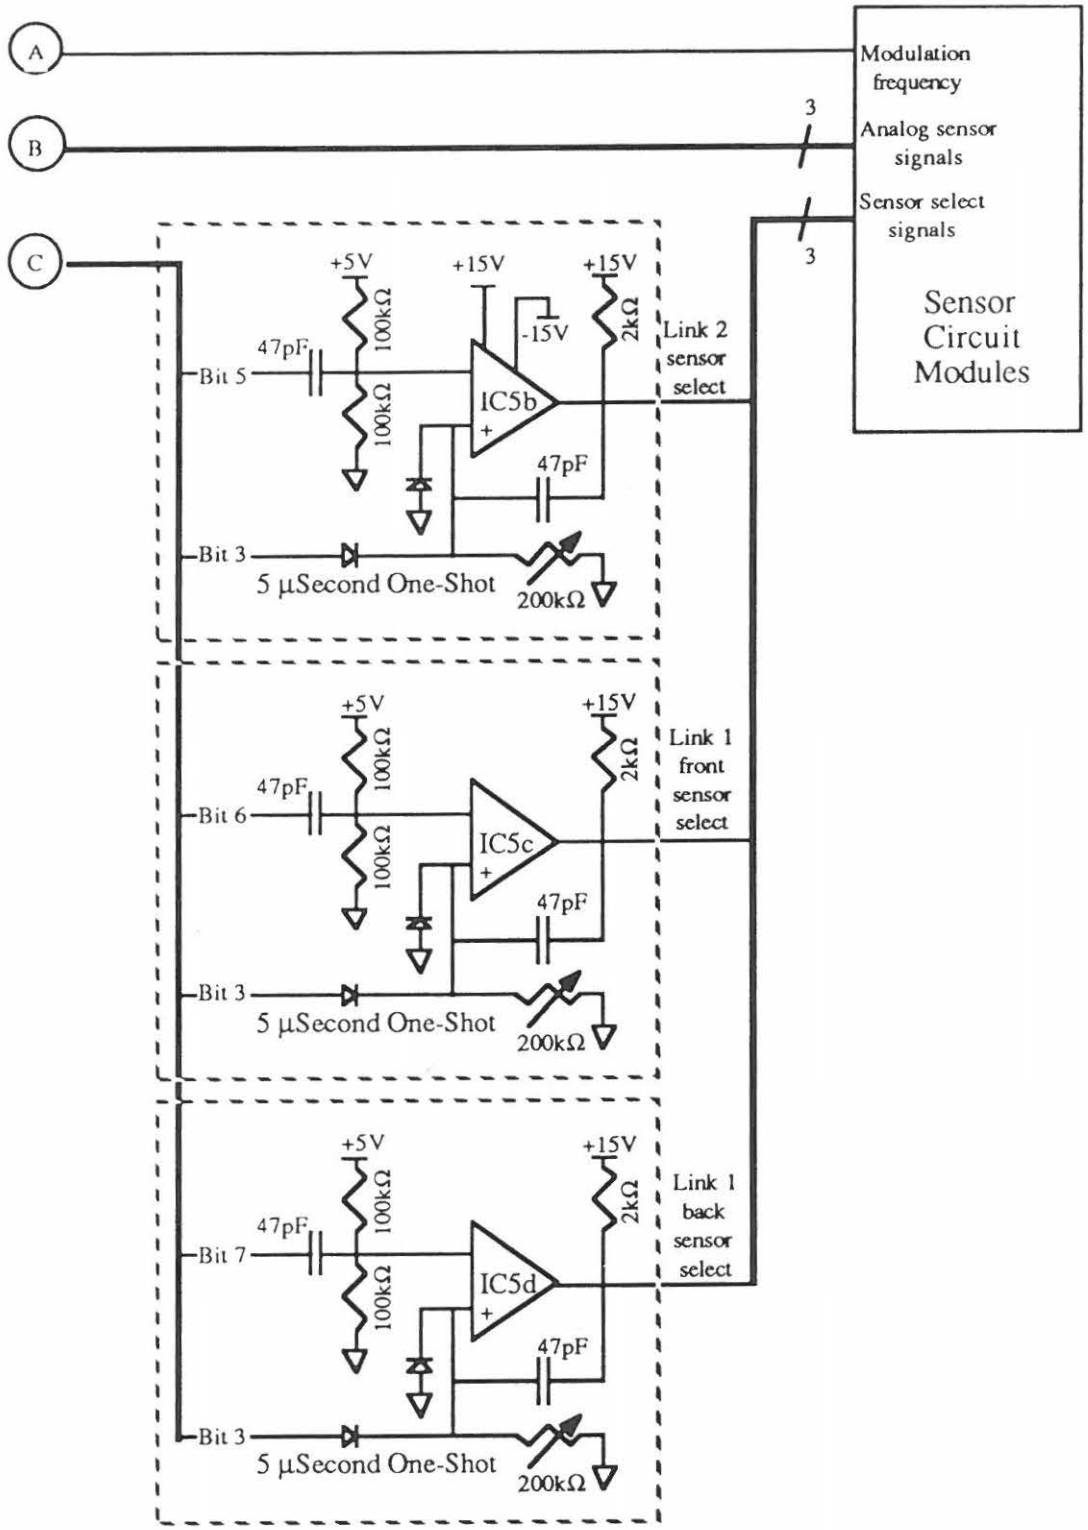

Fig. 9. Sensor interface circuit, part 2.

Table 3
Bit functions of the command byte (bit 0 is the least significant bit)

Bit	Function
0	Link adaptor send
1	Analog channel select
2	Analog channel select
3	Analog channel select
4	Analog to digital converter trigger
5	Increment l_3 sensor
6	Increment l_2 back sensor
7	Increment l_2 front sensor

nity (*Fig. 8*). The output byte of the link adaptor at 'Qout' is used as the command byte. Similar to the robot interface command byte, asserting a particular bit of the command byte triggers an action in the sensor interface, which in turn can cause an action in the sensor circuit module. The function of each bit of the command byte is shown in Table 3.

The frequency of the transmitted signal by the sensitive skin is generated on the sensor interface circuit board by op-amp IC5a. This signal is distributed to the three sensor processing circuit modules, which are mounted on the surface of the arm, underneath the sensitive skin. This signal is also distributed to other parts of the sensor interface circuit.

The analog to digital converter (ADC) converts the analog signal at its input 'In' to a digital number available at the 'Dout' output, triggered by a high to low transition on its 'Trig' input. During the conversion, which takes about 5 μsec, the 'Busy' output remains low. This signal controls the sample and hold (IC8), and the 'IValid' signal on IC3. This latter data path allows IC9 to trigger IC3 directly as soon as the conversion is finished, and eliminates the need for the transputer to constantly poll the interface. Note that IC9 is a 12 bit ADC, the eight least significant bits of the digitized data are sent as soon as the conversion is finished. The most significant four bits are multiplexed on the eight bit data bus via a buffer internal to IC9. This data is placed on the output data bus 'Dout' when the 'Trig' input is high.

The analog sensor signal that is digitized by the ADC is switched by the analog multiplexer IC7 [8]. Bits 1 through 3 of the command byte select which output of the three sensor circuit modules is digitized. The binary number represented by the previously mentioned bits corresponds to the number of the channel that is selected, with bit 1 presenting the least significant bit of the channel number. For example, if bit 1 = 0, bit 2 = 1, and bit 3 = 0, then channel 'Ch2' is selected.

The triggering of the ADC is initiated by bit 4 of the command byte. In principle, this signal can be connected directly to the 'Trig' input of IC9, but the trigger signal is preconditioned by IC6 to reduce the noise level of the digitized signal. The origin of this noise is as follows. The sensor interface circuit is located a few feet from the skin, at the base of the robot. Some of the signal from the 135 kHz oscillator formed by IC5a is capacitively coupled to the analog sensor signal

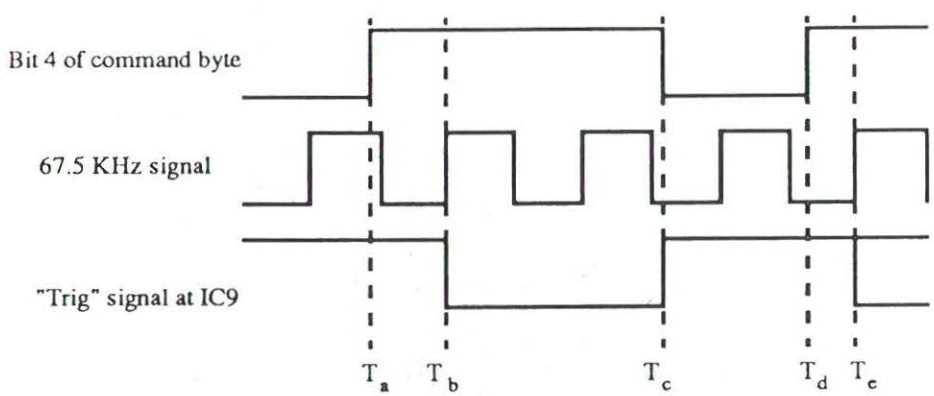

Fig. 10. Triggering the ADC in phase with the 67.5 KHz signal.

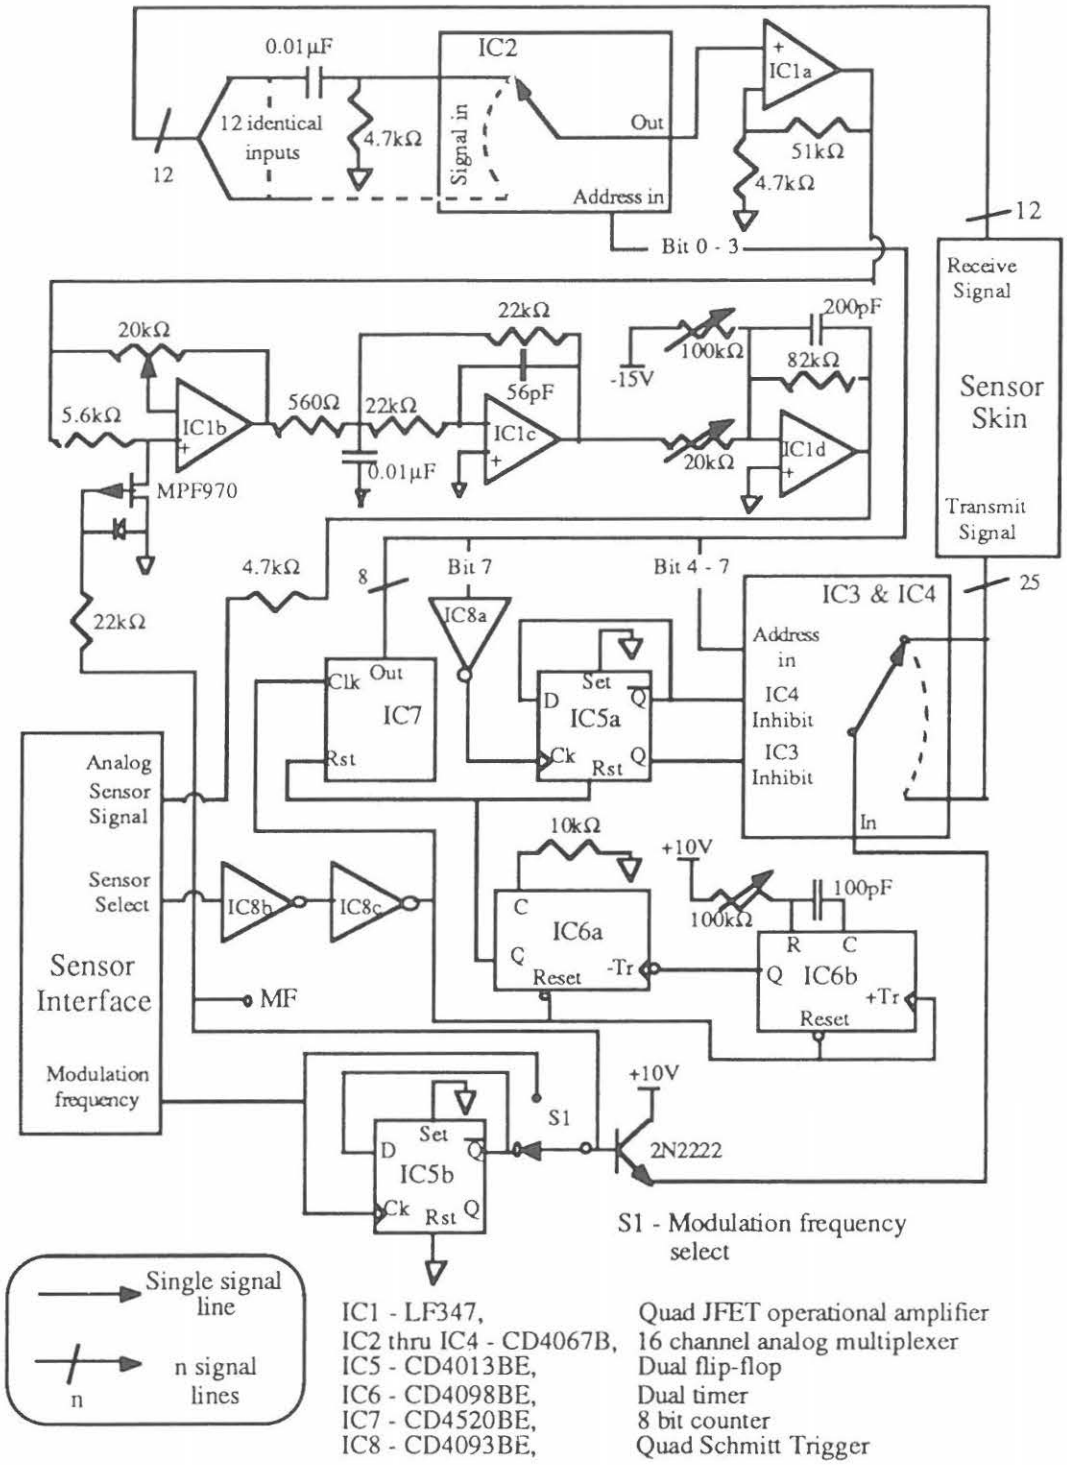

Fig. 11. Sensor circuit module.

lines. In addition, the 135 kHz signal, and another signal at 67.5 kHz are used for the modulation and demodulation of signals on the sensor circuit module. The result is that there exists a small amount of the modulation frequencies' signal on the analog sensor signal lines. This noise could have been removed by using a high-pass filter, but it is rejected in a different way. First, IC6b divides the frequency of the 135 kHz signal by two to 67.5 KHz, this signal appears at the chip's 'Q' output. The trigger for the ADC is synchronized by IC6a to a rising edge of the 67.5 kHz signal. Thus when bit 4 of the command bus is asserted by the transputer, the ADC is not triggered until the next rising edge of the 67.5 kHz signal. By triggering the ADC at the same point with regard to the phase of the 67.5 kHz signal, signals that are integer multiples of the 67.5 kHz signal are filtered from the digitized data.

Referring to *Fig. 10*, at point T_a bit 4 of the command byte is brought high, but the ADC is not triggered until at point T_b. When bit 4 returns low at T_c, the trigger signal of IC9 is reset for another sample, which occurs again at points T_d and T_e.

A slight drawback of this triggering scheme is that the trigger command will be delayed by at most 14.8 μsec, however this short amount of time can be considered negligible in the application described.

The sensors are addressed in a serial fashion by a short pulse on the Sensor Select input s of the sensor circuit module, and are reset by a long pulse on this same signal line. The short pulse is delivered by the One-Shots formed by IC5b, IC5c, and IC5d. Each sensor section is pulsed by its own One-Shot as indicated on *Fig. 9*. The exact pulse length is adjustable by the 200 kΩ variable resistor. The long high pulse required for the reset function is provided by bit 3 of the command byte. The Sensor Select signal remains high as long as bit 3 is a '1'. By asserting bit 3, the entire skin is reset together, synchronizing the address counters on the three sensor circuit modules.

4.2. Sensor circuit modules

A sensor circuit module contains the majority of the electronics necessary to implement the sensing function. There are three such modules in the current version of the system, each connected to a different part of the skin. The module is fastened to the arm using Velcro fasteners, and the sensitive skin is then wrapped over it.

The sensor circuit module can be divided into two parts, see *Fig. 11*. The first part is the sensor addressing circuit, which decodes which sensor is being addressed currently, the second part is the sensor detection circuit which amplifies and filters the signals from the light detectors.

To increase the reliability of the sensor system, it is desirable to minimize the amount of interconnecting wires that run across the joints of the robot. There is a conflicting requirement, however: it is also desirable to be able to address each sensor individually, so that the position of obstacles with respect to the arm can be determined. If a parallel addressing scheme is used, similar to how computer memory is typically accessed, $\log_2 n$ lines will be needed, where n is the total number of sensors. By using a serial addressing scheme, the sensors can be addressed using just one signal line.

The addressing scheme is organized as follows. Each sensor circuit module has a counter which keeps track of which sensor is currently being addressed. The counter is incremented on every rising edge (low to high transition) of the serial clock line, which causes a new sensor to be selected. Thus pulses of any length, including short ones, will increment the counter. The counter is reset to zero with a long pulse by using a pulse discriminator on the sensor circuit module. In the implemented system, pulses that are 'high' for longer than 10 μsec are considered reset pulses; pulses shorter than 10 μsec increment the counter. This addressing scheme allows just one signal line to address a potentially infinite number of sensors.

An obvious drawback of this scheme is that random addressing is not possible. All sensors with addresses lower than a desired sensor will have to be selected when picking a particular sensor. By keeping the pulses as short as possible, one can quickly address a particular sensor, and minimize the amount of wasted time. Note, however, that in general in such a sensor system there will be a need to poll *all* sensors periodically. In this case the order of addressing is immaterial, and hence the advantages of the serial addressing scheme outweigh its disadvantages.

Referring to *Fig. 11*, the Sensor Select signal from the Sensor Interface is first 'cleaned up' by Schmitt triggers IC8b and IC8c, and is then connected to the 'Clk' input of the eight bit counter, which keeps track of which sensor is currently being selected. The pulse discriminator, which decides when the counter should be reset is composed of the dual one-shot IC6.

The pulse discriminator operates as follows. While the Sensor Select line is low, the one-shots' outputs 'Q' are low, and the eight bit counter is not reset. On the arrival of a pulse on the Sensor Select line at T_a, one-shot IC6a is triggered by the positive edge, and its output 'Q' goes high. If the Sensor Select line stays high longer than 10 μsec, it will allow IC6a to time out, causing its output to drop low at T_b. This triggers IC6b, and its output 'Q' goes high, which resets the counter. Thus pulses longer than 10 μsec cause the counter IC7 to be reset. On the other hand, if the Sensor Select signal goes low before IC6a times out, such as at T_d and T_e (*Fig. 12*), no reset pulse will be generated, and the counter increments normally.

The light transmitted by the Infrared Emitting Diode (IRED) is amplitude modulated, and then synchronously detected to increase the immunity to other light sources. It also allows operation on several 'channels', since light transmitted by one sensor can be rejected by another detector which is demodulating at a different frequency. This will be explained in Section 4.3 below.

The output byte of IC7 at its output 'Out' controls the analog multiplexers that switch the appropriate optical components to the sensor circuit. The least significant four bits are connected to the analog multiplexer (mux) IC2, which selects among the 12 signals from the preamplifiers on the sensitive skin. The signal from the analog switch is first high-pass filtered by IC1 a to remove components due to room lighting from the received signal [9]. It is then connected to the synchronous detector ICb, which demodulates the transmitted signal. Operation of the detector is explained below. After demodulation, the signal is then low-pass filtered by the three pole Butterworth filter composed of IC1c and IC1d. The cut-off frequency of the Butterworth filter is at 10 KHz. The output of IC1d is then connected to one of the input channels on the Sensor Interface Board via a 4.7 kΩ resistor which provides short circuit protection for IC1d's output stage.

The settling time of the Butterworth filter, which determines overall sensor response time, is approximately 0.25 msec. A higher bandwidth filter would settle in less time, but would also make the circuit more noise prone. Since the links of the arm are polled in parallel, the limiting factor in the sensor update rate is the polling of the larger of the two skin sections, which is link l_3. In the current implementation, the sensor polling rate is such that the entire skin is polled once every sixteenth of a second.

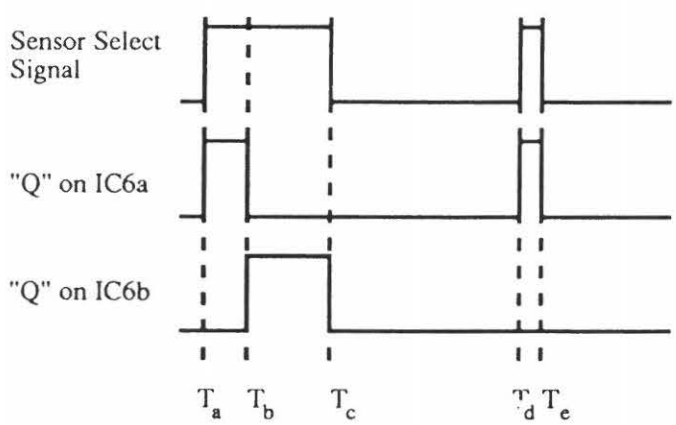

Fig. 12. Pulse discriminator wave forms.

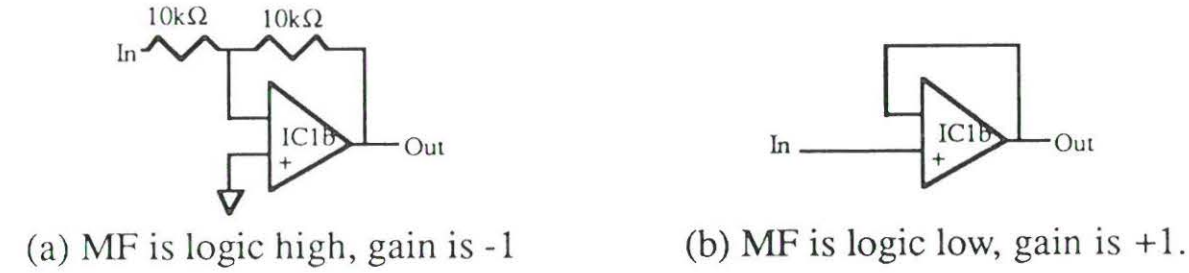

Fig. 13. Approximate circuits for the demodulator depending on the state of signal 'MF'.

The switch S1 selects the operating frequency of a particular Sensor Circuit Module. By selecting the output of IC5b, the circuit operates at 67.5 KHz, or one half the frequency of the signal 'Modulation frequency', supplied by the Sensor Interface Circuit. The lower frequency (67.5 KHz) is used for sensors on link l_2, and the higher frequency (135 KHz) is used for sensors on link l_3.

4.3. Operation of the demodulator

Instead of using a full-fledged synchronous demodulator, which requires an analog multiplier, a more simple way of demodulating the signal is used. The demodulator IC1b operates by alternately amplifying the signal at the output of IC1a by $+1$ and -1 in step with the digital signal at the gate of the MPF970 transistor. This signal is also the same one that is transmitted by the selected IRED, thus the flashing of the IRED is in step with the operation of IC1b. The demodulator operates as follows. Assuming that the 20 kΩ variable resistor is at its center, the circuit can be broken into two versions. One version, shown in *Fig. 13a*, is the approx. equiv. circuit when the signal 'MF' is high; the other, shown in *Fig. 13b*, is when that signal is low [10]. The gain of the demodulator stage thus alternates between $+1$ and -1 depending on the logic level of the signal 'MF'.

To analyze the effect of the demodulator on its input signal, we model it as an analog multiplier with one input connected to a square wave of frequency ω_s and its output connected to a low-pass filter.

A true demodulator uses a sine wave as its demodulation signal. By using the described circuit, a square wave becomes the demodulation signal. The effect of the demodulator on its input signal is analyzed in two ways, in the frequency domain and in the time domain.

The operation of the demodulator is analyzed by applying a sine wave of known frequency ω_t to the input of the demodulator. The magnitude of the spectrum of the demodulation signal (square wave) is shown in *Fig. 15*.

The multiplication of the two signals at the input of the multiplier produces signals whose frequencies are the sum and difference of the frequencies of the original signals, and is equiva-

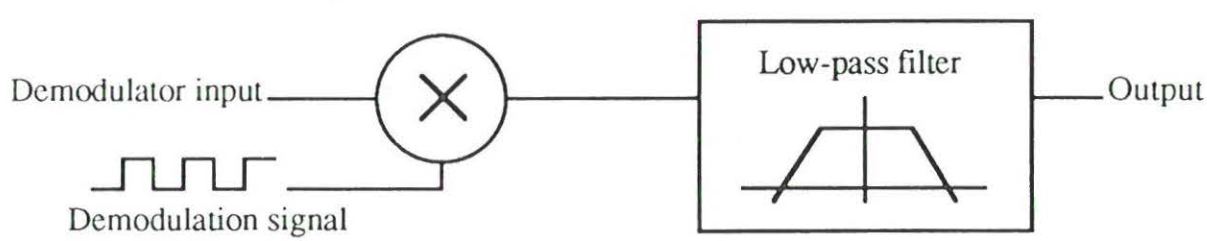

Fig. 14. Model of the demodulator and low-pass filter.

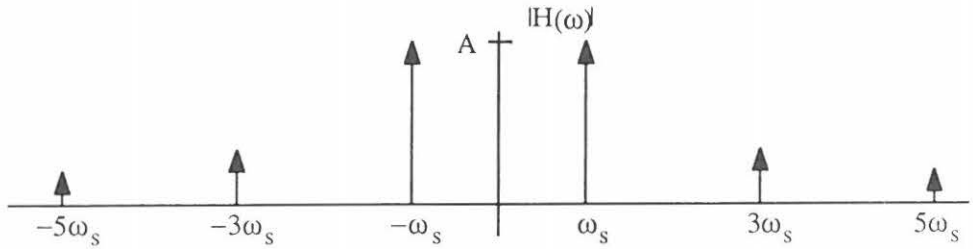

Fig. 15. Sketch of the magnitude of the spectrum of a square wave at angular frequency ω_s.

lent to translation of the spectrum of the input signal [11]. Thus the spectrum of the signal at the output of the demodulator is that at the input shifted by $\omega = \pm\omega_s$, $\pm 3\omega_s$, $\pm 5\omega_s, \ldots$, etc. The output of the demodulator is fed to a low-pass filter to remove unwanted components. If a component of the input signal is shifted into the pass band of the low-pass filter, it will be passed unattenuated. This will happen when the input signal has a component at $\omega_t = n\omega_s$, where $n = 1$, 3, 5... etc. More formally, the frequency response of the switching synchronous detector is the convolution of the spectrum of the square wave and the response of the low-pass filter (*Fig. 16*).

To gain a more intuitive understanding of the operation of the demodulator, let us examine its operation in the time domain. In *Fig. 17*, the input signals to the analog multiplier are shown on the left, and the output signals of the low-pass filter on the far right. The output signal of the demodulator is filtered by the low-pass filter, 14.

Let us consider three examples. First, assume that the demodulation signal and Signal 1 are in phase, and the signal swing is between ± 1 Volt (*Fig. 17a*). The resulting output signal is a constant 1 Volt because the inversion operation of the multiplier is in phase with the polarity of Signal 1.

In the second example, the input signal is Signal 2. The output of the demodulator, shown to the right of Signal 2, is filtered by the low-pass filter, *Fig. 17b*. Since this signal has an average (DC) value of 0 Volts, it is blocked by the low-pass

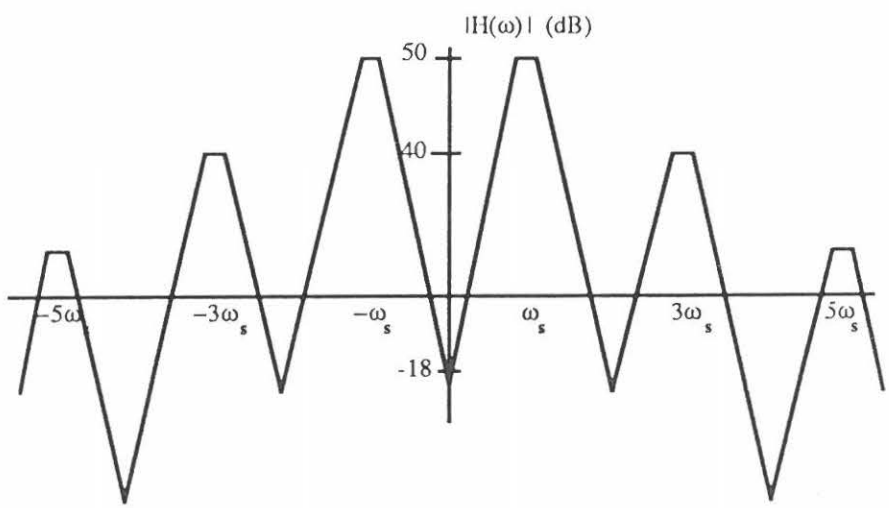

Fig. 16. Frequency response of the switching synchronous detector; ω_s is the frequency of the demodulating square wave.

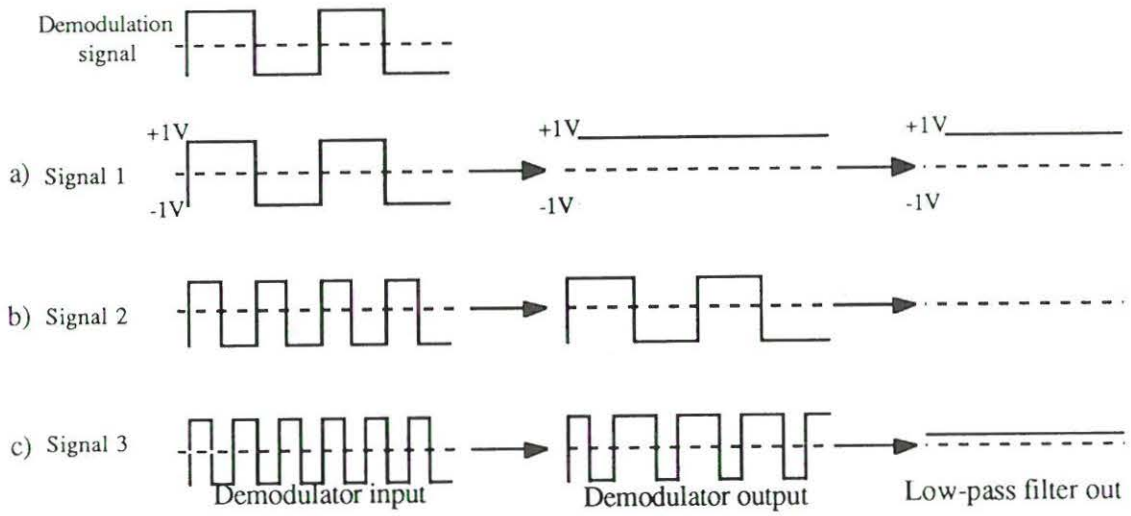

Fig. 17. Waveforms for the demodulator examples.

filter, and there is no signal at the output of the low-pass filter. Thus signals at twice the frequency of the demodulation are highly rejected by this system. It can similarly be shown that signals at one half the frequency of the demodulation frequency are also highly rejected. This response is the reason why the operating frequency of the sensors of link l_2 and l_3 are at a 2:1 ratio. Light emitted by one link's transmitter does not affect the operation of a receiver on the other link.

In the third example, the frequency of Signal 3 at the input of the demodulator is three times the demodulation frequency, and its phase is as shown in *Fig. 17c*. In this case, the output of the demodulator has a DC value of 0.3333 Volts. Thus the output of the low-pass filter is 0.3333 Volts. By continuing this example for signals of higher frequencies, it can be shown that signals at odd multiples of the demodulation frequency are passed. More precisely, if:

the voltage swing of the input signal
$= \pm A$ Volts
the frequency of the demodulation signal
$= \omega_s$,
the frequency of the input signal
$= \omega_t$, where $\omega_t = n\omega_s$, $n = 1, 3, 5, \ldots$ etc.,

then the DC value of the output of the low-pass filter $< A/n$ Volts.

The inequality in the last expression is due to the uncertainty of the phase between the input signal and the demodulation signal. The results of this time domain analysis agree with the results of the frequency domain analysis. Signals at odd multiples of the demodulation frequencies are passed, while signals at even multiples of the demodulation frequency are rejected.

The drawback of the simpler demodulation circuit used in the sensor system is its sensitivity to signals at the odd multiples of the demodulation signal frequency. It is unlikely, however, that there will be stray optical radiation modulated at such a high frequency in the robot environment. Furthermore, for the incident light to produce a false reading, it must have a strong component in the infrared region since the optical components are molded from a specially colored plastic. In addition, high frequency components are attenuated by the $1/n$ factor mentioned in the time domain analysis. Even flashing light sources such as cathode ray tubes operate in a much lower frequency band (< 16 KHz).

The advantage of our switching synchronous demodulator over a multiplying synchronous de-

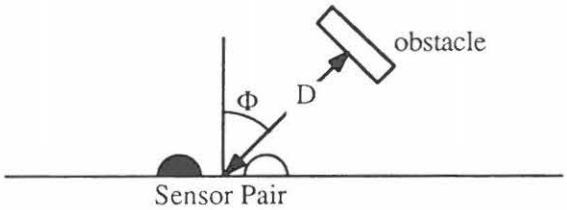

Fig. 18. Test setup for measurement of sensor response.

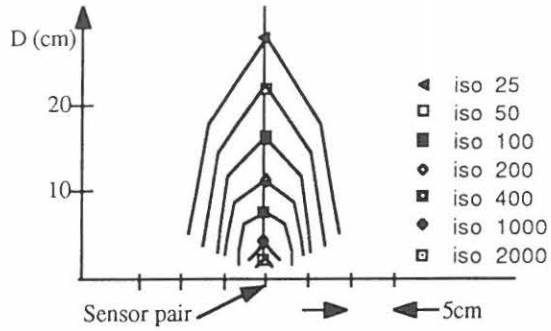

Fig. 20. Directional response of one sensor pair. Iso X refers to the line of constant sensor output of X.

modulator is its simplicity and compactness. By using the switching detector, the need for a pure sine wave and an analog multiplier is eliminated. Since the transmitted signal is a square wave, it can easily be used as the control signal for the switching detector. The detection function can be performed using a single op-amp, and the need for circuits that can fit in a small space can be important in these applications.

4.4. Performance of the sensor circuit

The sensor response was assessed experimentally using a randomly selected sensor pair and a $3'' \times 3''$ piece of white paper as the test obstacle. The response is measured as a function of the distance D from the skin to the test obstacle, and the angle Φ between the perpendicular to the test obstacle and the sensing axis, see *Fig. 18*.

The output of the sensor, where each unit represents 1.2 mVolts (thus 1000 = 1.2 Volt), is shown in *Fig. 19* for four values of the angle Φ, $\Phi = 0°$, 20°, 40°, and 60°. Twenty samples were taken for each position of the obstacle, and the results were then averaged. Note a decrease in response for non-zero Φ values due to the beam width of the IRED.

In *Fig. 20*, the sensor data is plotted to show the sensitivity cone for one sensor pair. Lines of constant sensor reading are shown; e.g., 'iso 25'

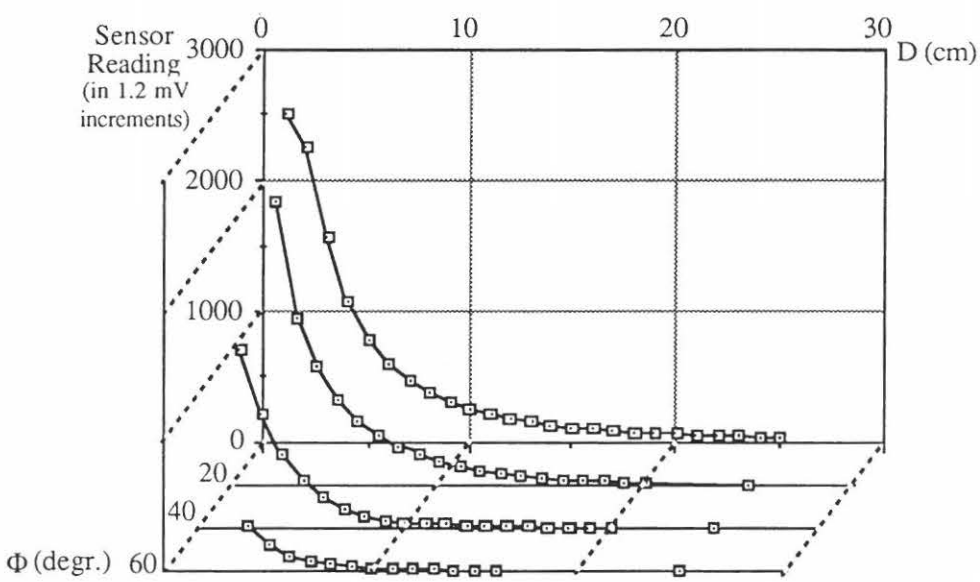

Fig. 19. Sensor output for various values of the distance D to the obstacle, and for four values of the angle Φ between the obstacle and the sensing axis.

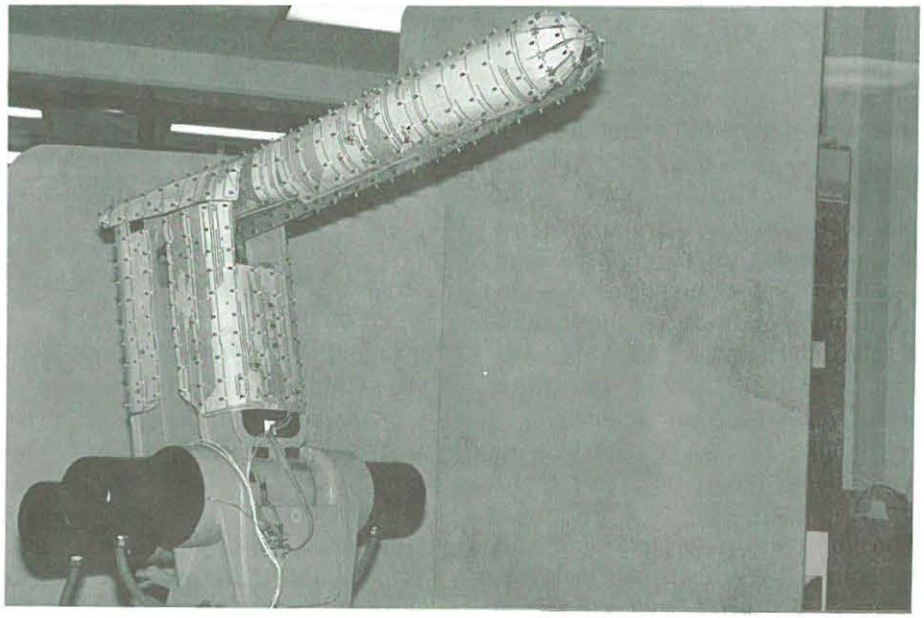

Fig. 21. Front view of the arm with installed sensitive skin. Note the skin dome covering the robot wrist.

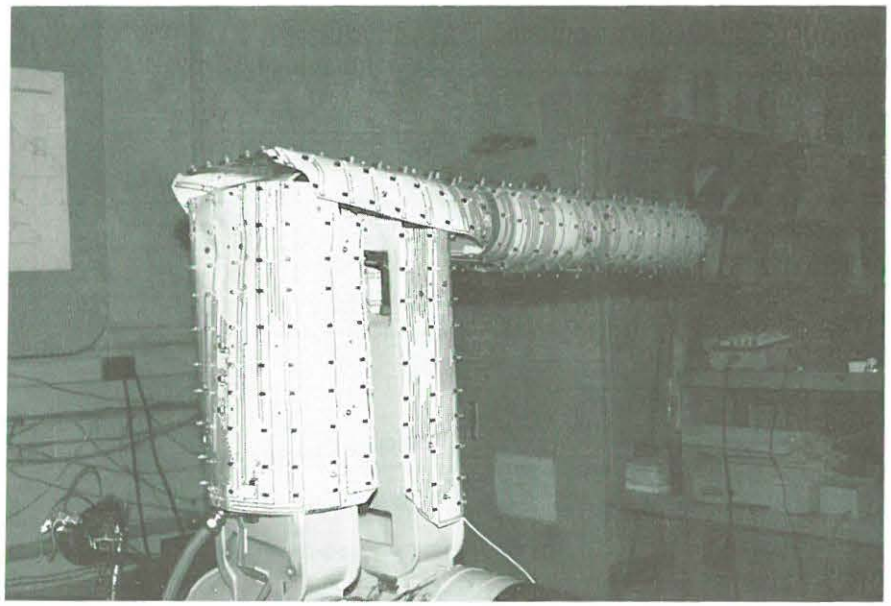

Fig. 22. Back view of the arm with sensitive skin. Note the sliding section of the skin around the joint j_3.

refers to the line of a constant sensor reading of 25. The tick marks on the horizontal axis indicate the locations of sensor pairs. Note that if the skin is placed flat, adjacent sensitivity cones overlap when an obstacle similar in color and size to the test obstacle is being detected. Where the skin is curved to fit the contour of the arm, the sensor spacing is reduced to ensure a similar overlap of the sensitivity cones, see *Fig. 26*.

Because of the sufficient density of sensor pairs along the sensitive skin (less than or equal to 5 cm between the neighboring sensor pairs), the neighboring sensor beams overlap, and so no obstacle that is at least as detectable in terms of its size and texture as the test obstacle can approach the arm undetected. In fact, if the test obstacle were the smallest possible obstacle in the work space of the arm, and the skin were placed flat, the minimum sensor spacing could be increased to 11 cm. However, if the sensor spacing is increased, or if the sensor skin is bent over a sharp corner, the resulting sensitive area around the arm will have a very irregular or 'bumpy' shape. This may lead to rough motion while following the contour of sharp obstacles, since the motion planning system servos around a constant sensor reading. Therefore, sharp corners in the skin material should be avoided and corners should be rounded gradually, see *Fig. 26*.

4.5. The sensitive skin

The skin base is manufactured of a Kapton based material, only 0.0085″ thick. Both sides of the material are copper cladded, resembling a regular printed circuit board. The material is etched and drilled using standard printed circuit techniques. After processing, the circuit board provides for both structural support and electrical interconnection for the optical components. Availability of this material has drastically reduced the problem of mounting the optical components. It allows the placement, with great precision, of the optical devices with little labor.

There were three major challenges faced when covering the arm with the sensitive skin. The first relates to covering link l_2 of the robot arm. As *Fig. 1* shows, link l_2 of our arm presents a trape-

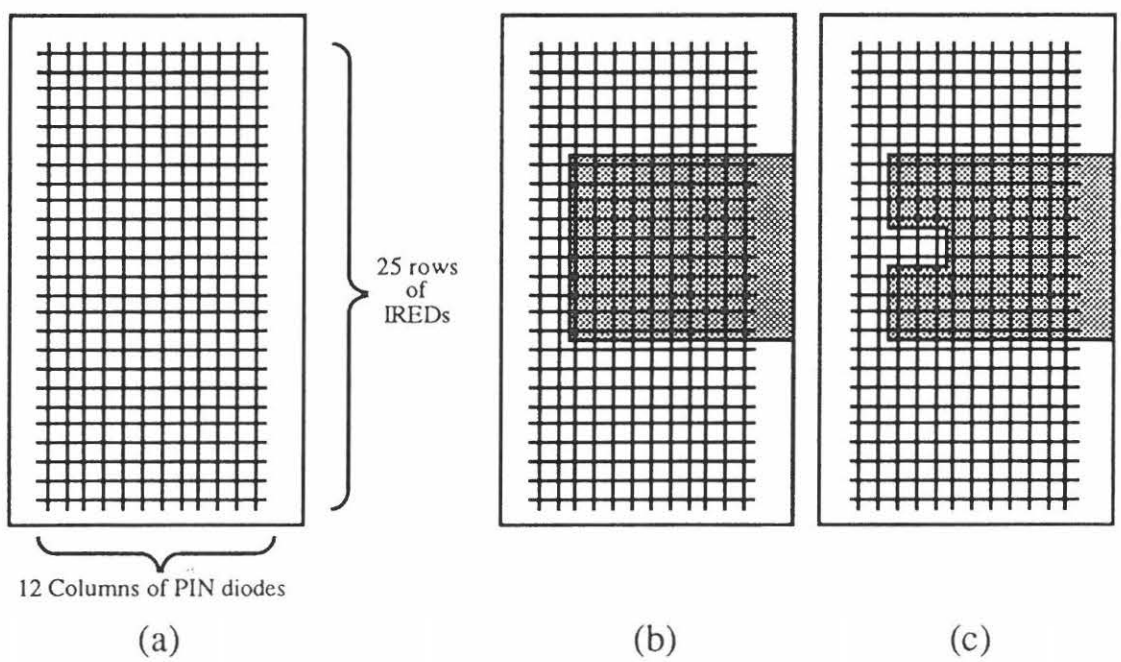

Fig. 23. Sheets of the sensitive skin. (a) The standard sheet as used for link l_3 (12×25) sensor pairs. (b) Piece of the sheet (shaded area) cut for the back part of link l_2; contains (10×10) sensor pairs. (c) Piece of the sheet cut for the front part of link l_2; contains ($10 \times 10 - 6$) sensor pairs.

zoidal closed kinematic chain whose width varies with the change in the joint value Θ_2. Human skin stretches and flexes while maintaining its sensing function, an ability that is difficult to duplicate in electronic devices. To solve this problem, each of the two sections of link l_2 is covered by its own skin sheet which moves independent of the other. The second challenge was covering the wrist of the arm. Although only the first three degrees of freedom of the arm are directly controlled, to provide collision-free gross motion, we also wish to maintain utility and safety of the wrist. The solution used here was to cover the wrist with a domed section, see *Fig. 21*, that has a slot, narrow enough to assure sensing beam overlap at neighboring sensor pairs but sufficiently wide to allow for the wrist motion. The third challenge relates to covering the joint j_3. To avoid wear and tear of the skin around the joint, the solution chosen was to have the section of sensitive skin covering link l_2 'slide' under the section covering link l_3, see *Fig. 22*.

Altogether, three sections of the skin are used to cover the robot arm, two sections for link l_2, and one section for link l_3. To save labor and production costs, a single large 'standard' skin section was designed. One copy of the standard section was used to cover link l_3; the two smaller skin sections covering link l_2 are cut out of another copy of the standard section, see *Fig. 23*. This design standardization necessitated some small compromises in the overall placement of sensor pairs. Alternatively, the circuitry on the Sensor Circuit Module could probably have been incoporated onto the sensor skin together with the optical components, but it would have necessitated interrupting the regular pattern of sensor pairs.

The overall size of the resulting section covering link l_3 is 61 cm × 122 cm (2′ × 4′), while the smaller sections covering link l_2 are 51 cm × 51 cm (20″ × 20″). One challenge during the production was the shear size of the artwork – several companies specializing in printed circuit

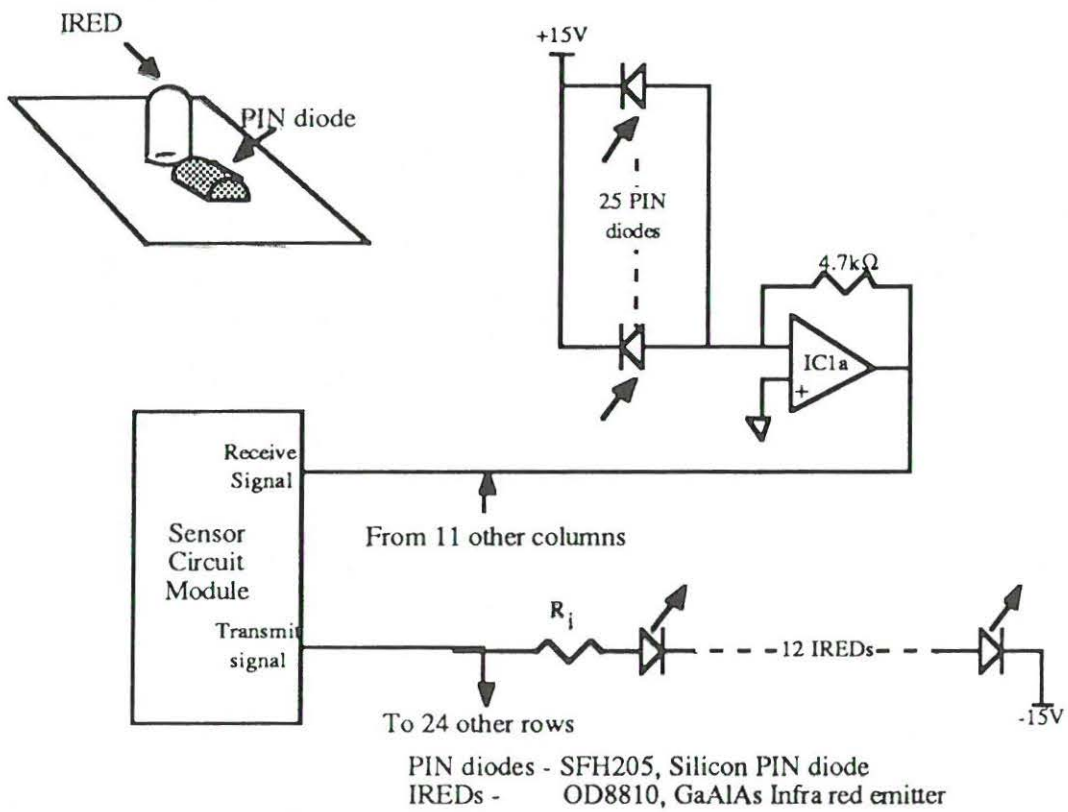

Fig. 24. Sensor skin circuit.

boards were not able to accommodate such a large pattern. The resolution of the circuit production equipment is also important because lines as narrow as 0.05″ have to be reproduced in some skin sections.

Out of the few copies of the standard sheet in the manufacturing batch, three were used for a 'dry run': the sensor sheets were cut, drilled, and mounted on the arm without optical components. This allowed final adjustments of the placement of mounting posts and mounting holes and, more importantly, allowed fine-tune planning of the wrist coverage. As mentioned before, the latter includes slits cut into the end of the artwork and joined into a dome shaped enclosure around the wrist. A large opening was then cut into the resulting dome, to allow the wrist motion, see *Fig. 21*. This necessitated transplanting several sensor pairs to new locations.

4.6. Skin circuitry

To reduce the amount of wiring on the arm, the optical components are arranged in an irregularly spaced row–column format. On the (standard) section of sensitive skin covering link l_3, there are 25 rows of IREDs, each containing 12 emitting diodes, see *Fig. 23*. The PIN diodes are arranged in 12 columns, each containing 25 light detectors. All the PIN diodes in a column are wired in parallel, and then placed in a reverse bias configuration. As a result of this wiring configuration, the output of all the PIN diodes are summed, and the signal from each of the 12 columns is preamplified by an op-amp operating in transresistance mode. The output of the op-amp is then connected via a connector to the Sensor Circuit Module (*Fig. 24*). The IREDs common to a row are wired in series. The resistor R_i limits the pulsing current to a 20 mAmp peak. This arrangement causes all the IREDs in a row to pulse in unison.

Two major criteria were considered during the selection of optical components: electrical and optical properties. Cost was less of a factor because in general the prices of these components do not vary widely. As far as electrical properties are concerned, the most efficient and powerful devices available were selected. Selecting the optical properties required more thought, however. In order to reduce the effect of obstacle size on the amount of reflected light (and thus sensor reading), it is desirable to reduce the beam width of the transmitted signal as much as possible. This comes from the fact that at distances at which the sensor operates the emitted light of the IRED illuminates only a spot on the obstacle, causing the rest of the obstacle to be invisible to the sensor. A compromise is needed, however: if the beam widths of the emitters are too narrow, small obstacles may be missed if they can fit between two sensors' beams. Ultimately, the limited choices available between the beam widths led to selecting an IRED with a half-power angle of 10° [12].

As far as the PIN diodes are concerned, to reduce the effect of room lighting, the material of its enclosure should double as an ambient light filter. The last element to select is the kind of lens provided with the PIN diode, which deter-

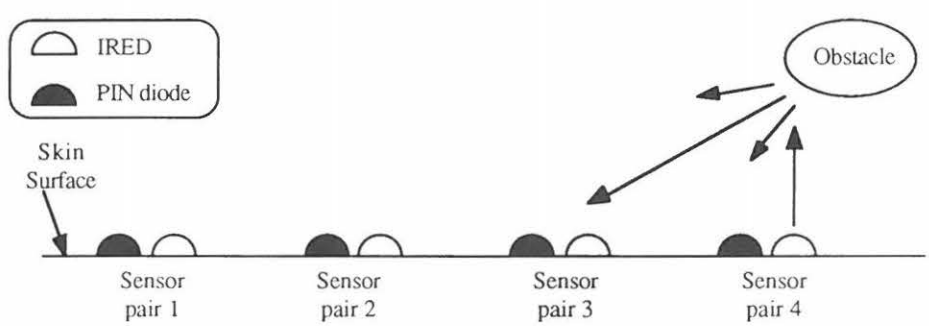

Fig. 25. Light from the emitter in one sensor pair can be received by the detector in another sensor pair.

mines its directional sensitivity. It turns out that by a proper choice of the lens the sensor insensitivity to the surface reflectance property of the obstacle can be improved. The wavelength of infrared light is in the range of hundreds of nanometers, scattering light equally in all directions, and tending obstacles to appear matte [6]. Many obstacles, however, do reflect light specularly to some degree, causing light to be reflected away from the transmitter. This problem can be partially solved using the fact that the outputs of all the PIN diodes in one column are summed. Thus, the light beam emitted from one sensor pair can be picked up by a PIN diode at another location along the arm. This effectively results in a receiver that is physically distributed over the arm. Light reflected to another sensor pair – for example from Sensor pair 4 to Sensor pair 3, in *Fig. 25* – will be hitting the PIN diode at a large angle Φ. For this light to be received, the PIN diode should have a concave lens, or no lens at all. Since the former type is not available, PIN diodes with no lenses were selected.

4.7. Mounting the skin on the arm

The skin is mounted on the robot arm using several fasteners distributed over the surface of the arm. A cross section view of link l_3 and the sensitive skin is shown in *Fig. 26*. Each fastener presents a bolt attached at an angle of 10° to the surface of an acrylic foot. The foot is a $1'' \times 1''$ piece of acrylic glued to the surface of the arm. The 10° angle accommodates the curvature of the skin at the mounting point. Two nuts threaded onto the bolt clamp the skin down, and allow continuous adjustment of the cross-section shape of the skin. The two skin sections covering link l_2 are fastened using a similar technique.

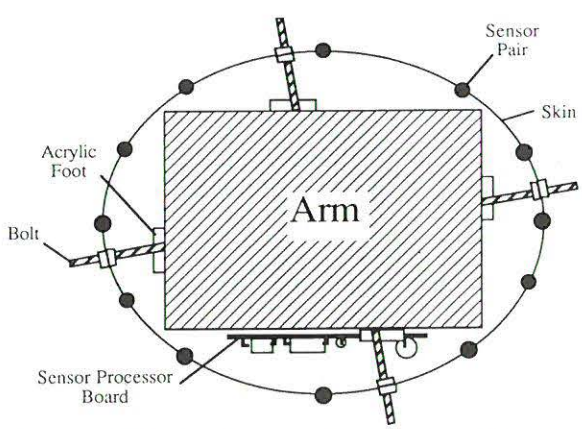

Fig. 26. Cross-sectional view of the arm with sensitive skin. Note the gradual curve in the skin and the variable sensor spacing – the distance between sensors is smaller at the areas of greater curvature of the skin.

Acknowledgments

The research work described in this paper has been supported in part by the National Science Foundation Grants DMC-8712357 and IRC-8805943.

References

[1] E. Cheung and V. Lumelsky, Motion planning for a whole-sensitive robot arm manipulator, *Proc. 1990 IEEE Conference on Robotics and Automation*, Cincinnati, OH (May 1990).
[2] E. Cheung and V. Lumelsky, Proximity sensing for robot manipulator motion planning: System and implementation issues, *IEEE Transactions on Robotics & Automation* (December 1989).
[3] M. Buehler, L. Whitcomb, F. Levin and D. Koditschek, A distributed message passing computational and I/O engine for real-time motion control, *Proc. American Control Conference*, American Control Society, Pittsburgh, PA (June 1989).
[4] *Transputer Development System* (Prentice-Hall, New York, 1988).
[5] C.A. Hoare, *OCCAM 2 Reference Manual* (Prentice-Hall, New York, 1988).
[6] D. Ballard and C. Brown, *Computer Vision* (Prentice-Hall, Englewood Cliffs, NJ, 1982).
[7] *The TTL Data Book* (Texas Instruments, 1981).
[8] *GE Solid State Data Book High Reliability CMOS ICs* (General Electric, Somerville, NJ, 1988).
[9] P. Horowitz and W. Hill, *The Art of Electronics* (Cambridge University Press, Cambridge, 1980).
[10] *Linear Data Book* (National Semiconductor Corporation, 1982).
[11] B. Lathi, *Signals, Systems and Communication* (Wiley, New York, 1965).
[12] *Opto-Diode Data Book* (Opto-Diode Corporation, 1988).