

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

A Prototype Tactile Sensor Array

James P. Christ
and
Arthur C. Sanderson
Department of Electrical Engineering
and The Robotics Institute
Carnegie-Mellon University
Pittsburgh, PA 15213

September 15, 1982

'k was supported In part by the National Science Foundation through a graduate fellowship to
hrist

629.892

C28r

82-14

cop. 3

Abstract

Although there are many potential applications for a tactile sensor array, very few practical implementations of such a sensor have been demonstrated. A practical tactile array sensor needs to be very durable, have a high resolution, have a small physical size, be relatively insensitive to noise, and have a compliant surface. In addition there is both physical and electrical coupling between elements of the array which should be eliminated or reduced as much as possible.

In order to investigate some of these problems, a prototype tactile sensor was constructed. The sensor was made of a sheet of conductive foam sandwiched between layers of conductors. When the foam is compressed at some point, the resistance through the foam decreases. By selecting the appropriate conductor on each side of the foam, the resistance at any one of 256 points could be measured. The spatial resolution of the sensor was 1/4 inch.

As a sample application, an object recognition system was implemented using the sensor. There are many questions involved in how to build a recognition system using a tactile sensor array. These include how to separate the object from the background, what features to use, and forms of preprocessing to perform on the tactile image. The system implemented was a first step at answering the first two questions. When an object was presented to the sensor array, the resulting data was thresholded in order to separate the background from the object. Three features were then computed from the object, the area and the second moments along the major and minor axes. These three features were then used in a suboptimal decision rule (the nearest mean normalized by the standard deviations) to classify the object. A total of 50 trials were performed using 5 objects. Two classification errors were made.

Table of Contents

Abstract	1
Introduction	3
1. Background Material	5
1.1. Sensor Technology	5
2. The Design of the Tactile Array	9
2.1. A Prototype Sensor	9
2.2. Construction of the Sensor	12
2.3. Performance of the Sensor	13
3. The Hardware	17
3.1. Measuring the Resistance	17
3.2. Connecting to the Tactile Array	18
3.3. Input/Output Circuitry	20
3.4. The Control Logic	22
A. The Recognition System	31
4.1. Segmentation of the Data	31
4.2. The Feature Set	32
4.3. The Decision Rule	33
4.4. Results	34
4.5. Sources of Error	34
Conclusion	45
Bibliography	47

List of Figures

Figure 2-1: The Missing Corner Problem	10
Figure 2-2: An Individual Sensor Element	11
Figure 2-3: Sensor Reading versus Compression	12
Figure 2-4: Image of a Right Angle Bend	13
Figure 2-5: Impulse Response	14
Figure 2-6: Response to Two Impulses	14
Figure 3-1: resistance to voltage conversion	18
Figure 3-2: Detailed View of Connection to the Sensor Array	20
Figure 3-3: Input Timing Cycle	21
Figure 3-4: Output Timing Cycle	21
Figure 3-5: State Diagram for the Finite State Machine	23
Figure 3-6: Four state finite state machine	24
Figure 3-7:	26
Figure 3-8:	27
Figure 3-9:	28
Figure 3-10:	29
Figure 4-1: Data for the meter and tape	35
Figure 4-2: Data for the can and battery	36
Figure 4-3: Data for the tube	37
Figure 4-4: Scatter diagram of $m_{2,0}$ versus Area for the training set	38
Figure 4-5: Scatter diagram of $m_{0,2}$ versus Area for the training set	38
Figure 4-6: Scatter diagram of $m_{0,2}$ versus $m_{2,0}$ for the training set	39
Figure 4-7: Sample output from the recognition routine	40
Figure 4-8: Summary of the trials of the recognition system	41
Figure 4-9: Scatter Plot of $m_{2,0}$ versus Area for test data	42
Figure 4-10: Scatter plot of $m_{0,2}$ versus Area for test data	42
Figure 4-11: Scatter plot of $m_{2,0}$ versus $m_{0,2}$ for test data	43
Figure 4-12: Discrete approximation to the area of an object	44

Introduction

Although vision systems for robots have been around for several years and have become very sophisticated, touch sensors are either very primitive or nonexistent. Force transducers can tell the controller when the gripper is pushing against something, but they can tell nothing about what the gripper is holding or what it is touching. This requires an array of sensors which can transduce information about surface contours.

There are many applications for touch sensors. One of the most basic is in handling an object. When a robot picks up an object, it needs to know what orientation the object is in, where it is grasping the object, and whether or not the object is slipping. Some research systems use vision to answer these questions. The camera determines the orientation of the object before it is picked up, allowing the robot to grasp it in some particular place. For most applications this works well enough, even though there is no feedback to ensure that the object was actually picked up correctly. With tactile feedback, the robot could determine the position and orientation of the object within the hand *after* the object has been picked up.

There are many other applications where tactile sensing can be used either to augment or replace vision. This includes such things as working in the dark and bin picking. Before any of this can be done, there is still a large amount of work which must be done on tactile sensors. The purpose of this project was to develop a prototype tactile sensor which could then be used to explore some of the issues involved in the use of tactile sensors.

Chapter 1

Background Material

A survey conducted by Harmon [6] provides a summary of desirable characteristics for tactile sensors. These include:

- Tactile sensors should be compliant and durable.
- Sensor arrays should be intelligent. Most of the information processing should occur *before* the data is received by the robot controller.
- Resolution should be on the order of $\frac{1}{10}$ inch, although some special applications may require greater or lesser accuracy.
- The sensor should be able to detect pressures as low as five to ten grams.
- A large dynamic range is desirable, on the order of 1000 to 1.
- Sensors should have a monotonic response.
- Sensors must *not* exhibit hysteresis.

1.1. Sensor Technology

The touch sensors marketed by Unimation are a good example of the state of the art in industry. One of their sensors is simply an on/off contact switch, which barely qualifies as a tactile sensor. Unimation also offers a slide probe which can be touched against an object, the travel of the probe is encoded digitally to yield the height of the surface. The most common touch sensors used in industry today are force transducers in the wrist of the robot. These sensors allow the controller to tell whether or not the gripper is touching a surface, how hard the robot is pressing against the surface, and at what angle.

In research laboratories we can find much more sophisticated sensors. Many attempts have been made to develop arrays of transducers using several different technologies. Sensors have been designed using spring-loaded switches, potentiometers, carbon fibers, conductive rubber, conductive sponge, strain gauges, piezo-electric materials, and piezo-resistive materials.

Some of the early arrays were built from spring-loaded switches. This method yields a binary (on/off) image of the object being touched. However, by varying the tension of the springs, the threshold of the image can be changed. Another early type of sensor was the push-rod. By connecting the rod to a digital encoder or a potentiometer, the displacement of the rod can be measured. This gives a continuous response, and is still used today (by Unimation, for instance).

Sensors have been constructed from various types of conductive rubbers and polymers. This approach has yielded low-cost, compliant arrays with a continuous response, but there are several drawbacks. These include noise, nonlinearity, hysteresis, fatigue, long time-constants, low sensitivity, and drift in some combination [5]. A paper by [14] discusses some of the properties of these materials. There have been some successful sensors constructed from these materials, for instance a pair of sensors which could recognize a set of three dimensional geometric shapes [4].

Most of the sensors built from conductive materials have measured the changes in the resistance of the material during compression. The sensors described in [13] and [8] measured resistance changes due to changes in contact area during compression. This led to improved sensitivity and a higher sensor density, but it is not clear how the other problems associated with conductive polymers were affected.

Another sensor which measures changes in contact resistance rather than changes in bulk resistance was constructed from carbon fibers [9]. In this sensor, two bundles of carbon fibers were laid perpendicular to each other. When an object pressed down at their intersection, more of the fibers in the two bundles would come into contact with each other, lowering the contact resistance. A low noise level was attributed to the large number of fibers involved in the contact. This sensor does not suffer from material fatigue, but has low compliance.

Research has also occurred in developing a tactile sensor from semiconductor devices (e.g. piezo-diodes or strain gauges). Sensors developed* from these devices do not suffer from many of the problems that conductive elastomer arrays do (e.g. sensitivity, hysteresis), but they do have their own problems such as a relatively high cost, large size, and lack of surface compliance. A paper by Bejczy [1] is one review of this type of sensor.

Sensors have been constructed from many other technologies, although they are not very common. A paper by Wang [16] includes descriptions of some of these sensors. One unusual¹ sensor was developed at SRI [7]. Pressure on the sensor causes a pin to protrude into the path between an LED and a phototransistor, decreasing the received light. This produces a fairly simple, linear sensor.

Some sensors have been constructed using magnetic or capacitive effects. These sensors have generally not been very successful due to their sensitivity to the material from which the object being touched is constructed.

Once a tactile sensor has been designed, there is still the problem of how to use the information obtained from it. Research in this area has not progressed very far, although a few sample systems have been demonstrated which can recognize three-dimensional objects by touch. An interesting approach to this problem was taken at JPL [15]. A tactile sensor was constructed by placing conductive rubber on the surface of a VLSI integrated circuit. The IC contained an array of contacts with a processing element per contact. This allows very fast initial processing of the data (e.g. filtering the data), although there is no reduction in the quantity of data produced.

The work done by Hillis at MIT [8] is a good example of the state of the art in tactile sensor research. This sensor was a 256 element array, in about a 1 inch square area. The sensors were made from anisotropically conductive silicone rubber, which will conduct current along only one axis of the rubber sheet. The rubber was laid against conductors on a printed circuit board running perpendicular to the axis of conduction; each of the intersections formed a sensor element. A nylon mesh was placed between the rubber and the circuit board to pull the two apart under no load, the resistance of the sensor varies as the contact area increases under a load. This sensor had a range of 1 to 100 grams of force.

A version of this sensor was mounted on a "finger" designed to resemble a human finger. This finger was then used in an object recognition system with the objects selected from a set of small fasteners, e.g. screws and pins. The object is felt to get a good view, and then pushed to see if it will roll. The features used in discriminating between objects were shape (long or round), whether it had bumps (up, down, or none), and whether it rolled when pushed between the finger and a flat surface. In processing the image, the data was first run through a simple filter, and then reduced to a two bit per pixel image by comparing to set thresholds. Everything below the first threshold was background, below the second threshold was a depression, and above the third threshold was a bump. Since all of the objects were small enough to fit into a single tactile view no attempt was made to merge multiple views of an object into a single representation, although this was one of the problems the Hillis felt should be looked into next.

Research in tactile sensors is not limited to developing sensors which will measure the contours of an object. Some of the other areas of research are detecting and preventing slip of an object in the gripper and measuring the texture of an object (is it rough, or is it smooth).

Chapter 2

The Design of the Tactile Array

When designing a tactile sensor for a robotic hand, there are several requirements to consider. The sensor needs to be fairly light, so that the dynamic performance of the robot does not suffer. The surface of the sensor needs to be compliant, so that it can conform to ridges and valleys in the object being touched. Although not necessary, it is desirable for the entire sensor to be flexible, so that it can conform to the "fingers" of the robot. Some other factors are ease of manufacture, durability, sensitivity to noise, and resolution.

Although it would seem to be desirable for the sensor to have a linear response, this is not at all necessary. Most applications only require that the response is monotonic. If a linear response is really needed, it is easy enough to convert the actual response to a linear response through a look-up table (or read-only memory) or by using the appropriate equation. Linearity was therefore not considered a major requirement in this project.

Another consideration is the resolution of the sensor array. Little can be said about this aspect without a specific application. More resolution is needed for handling small components than for handling turbine blades. Despite this some estimates of the needed resolution can be found, they generally tend to be about 100 sensors per square inch. This is not enough for some applications (sensing the leads on an integrated circuit, for example), and too many for others (e.g. handling turbine blades). Since resolution does not affect the algorithms needed to use the sensor in a practical environment, a goal of 16 sensors per square inch was accepted for this project.

2.1. A Prototype Sensor

A tactile array was designed and constructed using carbon impregnated foam as the sensing element. The basic idea behind this sensor is that when the foam is compressed at some point, the resistance through the foam decreases.* If a contact is placed on each face of a sheet of foam, the resistance between the contacts can be used to measure how much the foam has been compressed near the contacts. By placing many contacts on each face of the foam, an array of sensor elements can be constructed-

If a separate wire is used for each of the contacts any reasonably large sensor array will have too many wires and require too much circuitry to be practical- in order to avoid this, a

row of contacts can be made by a single conductor running along one face of the foam, if the conductors on one side of the foam are perpendicular to the conductors on the other side, the compression at a given point of the array can be measured by measuring the resistance between the two conductors which intersect at that point.

Although this seems to be a simple enough task, there are several things which complicate it. On each face of the foam, there are a number of parallel conductors. Since the conductors are in contact with a resistive sheet there is some finite resistance between them. In an $n \times n$ array with a distance d separating the conductors, each conductor will be of length nd . For n small, the resistance between two adjacent parallel conductors may be reasonably large; unfortunately the resistance decreases as $\frac{1}{n}$. The resistance between two adjacent conductors can quickly drop below the resistance we need to measure. In the limiting case of very long conductors spaced very closely together the resistance between adjacent parallel conductors is essentially zero, and the sensor array may be modelled as a layer of resistive foam between two conductive plates. This implies that for a large enough n , we cannot measure the resistance at a point in the array simply by measuring the resistance between the appropriate pair of perpendicular conductors. A solution needed to be found before a practical sensor array could be constructed and the solution will be discussed in section 3.2.

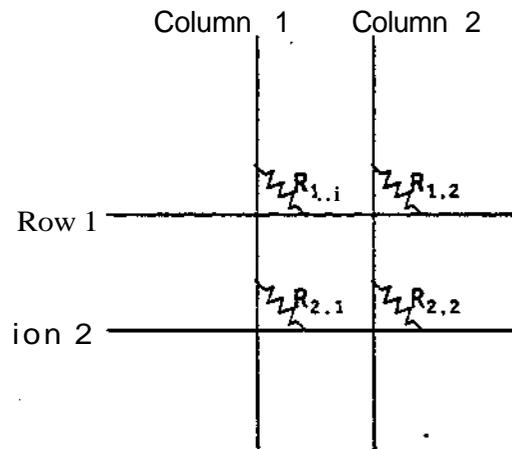


Figure 2-1: The Missing Corner Problem

A similar but somewhat more subtle problem is the "missing corner problem". If the pad is compressed at three corners of a rectangle we get the situation diagrammed in figure 2-1. Assuming that $R_{11} \ll R_{12} \ll R_{21} \ll R_{22}$ then the resistance seen by the outside world between row 1 and column 2 is $R_{11} + R_{21} + R_{22}$ in parallel with $R_{1,2}$. This can be approximated as $R_{11} + R_{21} + R_{22}$ which is much less than the actual resistance. The result is that all four locations are seen as having a low resistance; the missing corner has been filled in. This problem can be eliminated by masking the resistances R_{11} and R_{22} when measuring $R_{1,2}$. The method described in section 3.2 to eliminate the previous problem has the added benefit of eliminating the missing corner problem.

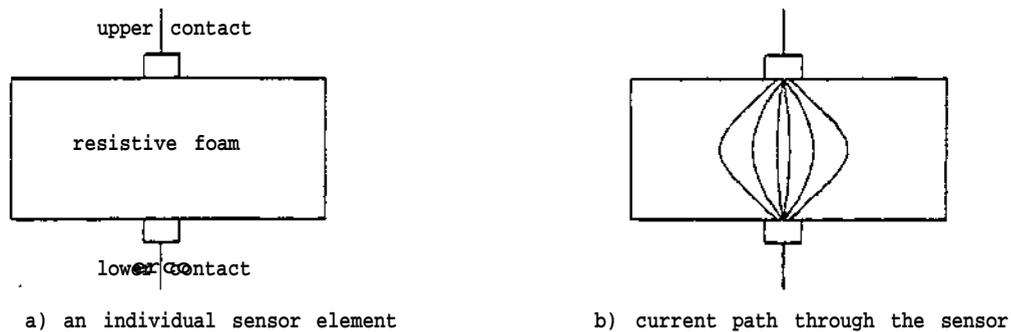


Figure 2-2: An Individual Sensor Element

The cross-section of a sensor in the array can be diagrammed as in figure 2-2a. The resistance between the two contacts is determined by the bulk resistance of the foam at each point, and by the path through the foam of the current used to measure the resistance. Although the contacts may be one-dimensional points, the foam is three-dimensional, so the current spreads as it passes through the foam. The path will look similar to the diagram in figure 2-2b.

This spreading can have a significant impact on the resolution of the sensor array. The resistance seen between the contacts can be thought of as the average resistance seen along each path through the foam, weighted by the fraction of the total current which travels that path. The current spreads through the entire pad, but the further from a straight line the path is the less current will travel along it. If contacts are spaced farther apart than the effective distance the current spreads, then the resolution is limited by the contact spacing. If the contacts are much closer than the distance the current spreads, the resolution is limited by how much the current spreads. Since the effective spreading radius of the current is determined by the thickness of the foam, the maximum resolution of the sensor array is determined by the thickness of the foam.

There are, of course, other factors which affect the resolution of the sensor array. One purely mechanical factor is the mechanical relaxation of the foam near a point under compression. If you compress the foam at a single point, the foam immediately around that point will also be compressed. As the distance from the point of compression increases the compression of the foam decreases, eventually becoming negligible. If the distance required for this expansion is too large, it may become the limiting factor in determining the resolution of the sensor.

In the construction of the sensor, conductors were run across the foam to make the contacts needed. If the elasticity of these conductors is low (which was indeed the case)* compressing the foam at some point along the conductor will push the conductor down. The conductor will then tend to compress the foam somewhat everywhere along the length of the conductor. The compression tends to spread preferentially in the direction of the conductors, with the effect becoming more pronounced as a larger portion of the conductor is depressed. The only way to reduce this problem is to increase the elasticity of the conductor.

One general concern about a tactile sensor is its linearity, or lack thereof. To investigate this aspect of a sensor constructed from a resistive foam, we need to investigate the mechanism through which the resistance of the sensor changes. If we consider the foam to be a solid block of material whose bulk resistivity does not vary with compression and model each sensor as a cylinder of effective radius r with height h , the resistance after a compression of distance x is proportional to $\frac{(h-x)}{r^2}$, which is linear in x . Unfortunately, this is not a very good model for the resistance of the foam during compression. The foam is a honeycomb of air bubbles, and any current must flow through the walls of the bubbles. As the foam is compressed, some of the bubbles are squeezed flat. This results in a shorter path for the current since it no longer has to flow along the walls of the bubble to get from the top to the bottom. There is no reason to expect this effect to be very linear, especially when the foam has been compressed almost to its limit and most of the holes have been squeezed shut.

2.2. Construction of the Sensor

Several sensors were constructed, all embodying the same concepts. The first sensor array was a 4×4 array with a spacing of about $\frac{1}{2}$ inch, the final sensor array was a 16×16 array with a spacing of $\frac{1}{4}$ inch, resulting in a 4 inch by 4 inch sensor array with 256 elements. The foam used was $\frac{1}{4}$ inch thick. The conductors were 32 gauge wire, chosen for their low impedance and easy availability.

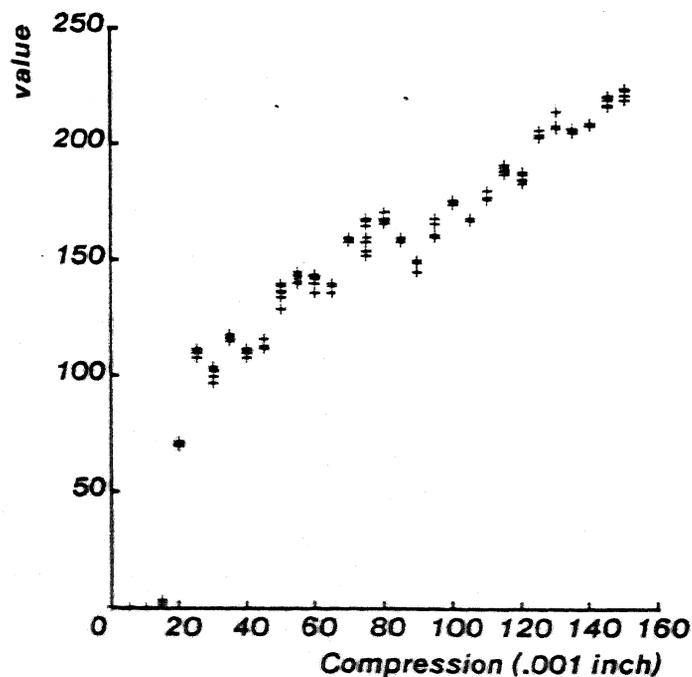


Figure 2-3: Sensor Reading versus Compression

```

0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 2f 75 67 ff 9b 7e 72 0 0 0 0
0 0 0 0 4 4b 69 b 3b 15 ac 9f 0 0 0 0
0 0 0 0 0 0 0 0 0 0 4a 44 0 0 0 0
0 0 0 0 0 0 0 0 0 0 45 3d 0 0 0 0
0 0 0 0 0 0 0 0 0 0 45 3c 0 5 0 0
0 0 0 0 0 0 0 0 0 0 19 13 0 0 0 0
0 0 0 0 0 0 0 0 0 0 3e 32 0 0 0 0
0 0 0 0 0 0 0 0 0 0 b 5 0 0 0 0
0 0 0 0 0 0 0 0 0 0 45 3d 46 0 0 0
0 0 0 0 0 0 0 0 0 0 40 37 0 0 0 0
0 0 0 0 0 0 0 0 0 0 54 4a 6a 0 0 0
0 0 0 0 0 0 0 0 0 0 8f 85 cd c 0 0
0 0 0 0 0 0 0 0 0 0 ff ef ff b 0 0
0 0 0 0 0 0 0 0 0 0 a0 95 bc 0 0 0
0 0 0 0 0 0 0 0 0 0 3c 33 19 0 0 0

```

opening allen4.
Mean: 96.936172, Std: 67.568848.
thresh: 29, area: 39

```

*****
** * **
**
**
**
**
**
**
**
**
**
**
**
**
**
**

```

Figure 2-4: Image of a Right Angle Bend

The response from the pad for varying distances of compression is graphed in figure 2-3. Although the number received from the interface circuit does not convert directly into resistance, for this trial a value of 0 corresponded to a resistance of 33.9 K Ω , a full scale reading of 255 corresponded to 14.8 K Ω . Figure 2-4 shows the raw data and a thresholded image of a right-angle bend (an allen wrench), and is included to give some idea of how well the "missing corner" problem and the resistive coupling between conductors have been dealt with. The two legs of the image should be the same width, but the differences are due more to how the wrench lined up with the rows of sensor elements in each direction than to a preferential spreading of the image due to the inelasticity of the conductors.

2.3. Performance of the Sensor

Overall, the sensor performed reasonably well. The major test of its performance was the construction of an object recognition system; this will be discussed in chapter 4. However, there were some problems with the sensor, some of which will be apparent in the discussion of the recognition system.

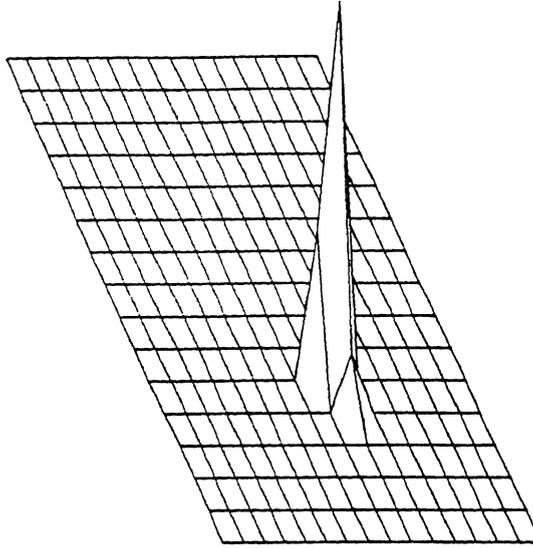


Figure 2-5: Impulse Response

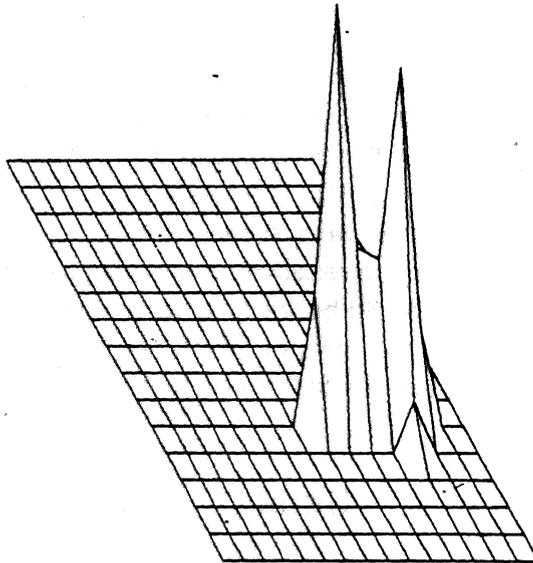


Figure 2-6: Response to Two Impulses

One problem with the sensor is noise in the unloaded state. The wires which form the contacts to the conductive foam are only resting on the surface when there is no object. This causes the resistance seen at some point of the array to vary substantially, presumably due to causes such as vibration in the table. This noise decreases substantially when an object is placed on the pad, pushing the wires into firm contact with the foam. Another factor which helps to eliminate the problem with noise in the unloaded state is that the resistance of the sensor varies by several orders of magnitude under compression while the circuit can encode only a range of 256 to 1. When the circuit is adjusted so that a moderate compression is full scale, elements without compression have a high enough resistance that they are encoded as a 0.

The problem of preferential spreading of the image was mentioned earlier, and is illustrated in figures 2-5 and 2-6. The first figure shows the response of the sensor array to a single point of pressure; there is limited spreading. However, in figure 2-6, a second point of pressure is added along the same row wire. The two points tend to depress the wire between them, causing the saddle between the two peaks to be much higher than the background, although it is still considerably lower than the peaks themselves.

The type of foam the sensor array is constructed out of can make a large difference. Variations in the resistance between different types of foam can be compensated for in the control circuitry, but physical differences cannot. The foam used for the sensor array was somewhat stiffer than desired. While this had little impact on the hardware and software used to drive the array, it did mean that an undesirably large pressure was needed to compress large areas of the array.

There were also problems with variations in the resistance of the uncompressed foam at different points of the array. Although this was not serious in the object recognition task performed, some form of calibration would have to be performed before much work could be done with three dimensional objects.

Chapter 3

The Hardware

Given the sensor, some form of interface between it and the LSI-11/23 was needed. Two modes of operation were envisioned for this interface. In the normal mode, the controller would select a location in the sensor array, convert the resistance seen at that location into a binary number, and then transmit the resulting number to the LSI-11/23. The controller would then move to the next location in the sensor, encode its resistance, and send that number to the LSI-11/23. This would enable the 11/23 to scan the entire tactile array by executing a sequence of input operations. In the second mode, the LSI-11/23 would output a pad location to the interface which would then transmit the resistance at that location to the 11/23. The 11/23 could then output a new location, and receive the resistance at that location. In the actual implementation the interface operates in a cross between these two modes. The interface will scan the tactile array sequentially until the LSI-11/23 outputs a new address; at that time the interface will begin a sequential scan starting at the new address. Both of the desired modes can be easily obtained from this.

In the following sections, the circuitry on the interface board is divided into four major functions. We will first cover the circuits needed to convert the resistance at some location of the tactile array to a binary number. Then we will discuss the circuitry needed to select a single location of the array. The third section details the communication with the 11/23, and the final section covers the design of the control logic which makes the interface board work. In many of the circuit diagrams used, there are clocked logic elements shown without any clock connection. Unless explicitly shown otherwise, all clocks are assumed to be connected directly to a common clock.

3.1. Measuring the Resistance

There are many ways to measure an unknown resistance. Two of the most common methods are to pass a current through it and measure the resulting voltage or to place a voltage across it and measure the current. Both of these methods were considered for measuring the resistance across the tactile array; the latter method was chosen for reasons which will be discussed later.

In order to place a known voltage across the sensor array, one column wire of the array was connected to + 5 volts. The voltage on one of the row wires was then forced to match a fixed

reference voltage (controlled by a potentiometer), thereby providing a constant voltage difference across the pad. In order to control the voltage on both sides of the sensor pad, the current through the pad needed to be adjusted, with the necessary current through the sensor proportional to the resistance of the pad. The current through the pad was controlled by connecting the row wire to the collector of a transistor and varying the base voltage (current). The emitter current was then passed through a resistor, producing a voltage which could be measured. The circuitry to do all this is shown in figure 3-1.

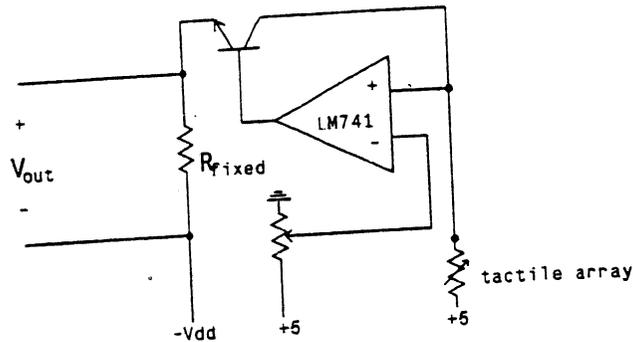


Figure 3-1: resistance to voltage conversion

Once the resistance of the sensor has been converted to a voltage, some form of an analog to digital converter is needed. The circuit used takes a digital to analog converter, and increments the input until the output voltage exceeds the voltage being measured. This is not a very efficient way to measure the voltage in terms of clock cycles needed; it requires an average of 128 cycles to do an 8 bit conversion, given random input. In this application it does much better than this, since the typical input voltage will be small except where there is an object on the pad. Even so, a successive approximation circuit is significantly faster, requiring 8 clock cycles to do the conversion (on any input data). However, the successive approximation circuit is more complicated, and the LSI-11/23 is not really fast enough to keep up with the circuit used. The logic which controls the analog to digital conversion will be discussed later; without this logic the conversion circuit is not very complicated so it is not included as a separate figure. For more detail, see figure 3-7.

3.2. Connecting to the Tactile Array

We have seen how the resistance of the sensor pad is measured, now we will see how a single location on the pad is selected. The basic idea is to use an analog switch to connect a wire on the bottom of the sensor to +5, and another analog switch to connect a wire on the top of the sensor to the measurement circuit. If all the other wires are then ignored, and if the pad itself is ignored except in the vicinity of the intersection of the two wires, the resistance between the two wires is essentially the resistance at the point of intersection.

This is fine if we can afford to ignore the unconnected wires and virtually all of the sensor array, but first we must look at why these can be ignored. We can ignore unconnected wires if the resistance between the unselected wires and the selected wires is large compared to the resistance between the selected wires. Looking at figure 2-1, it is clear that this difference can easily be as small as a 3 to 1 ratio (or smaller). This neglects the fact that two parallel wires along a resistive surface will tend to have a low resistance between them. The second assumption was that most of the pad could be ignored without affecting the measurement. This may be true, but it equally well may not be true; it depends on the thickness of the conductive foam compared to the distance between wires on the foam. The wires should be far enough apart that the resistance seen at one element of the array does not vary greatly with changes in the resistance at an adjacent element of the array. In summary, one of the assumptions we would like to make is *always* false, the accuracy of the other can be ensured during the construction of the sensor array.

If we are to keep the basic concept of how to select an element of the sensor array, we need some way to decrease the coupling between wires in the array. We can do this by using all of the unconnected wires as *shield* wires. If we connect all of them to a fixed voltage, then there will be no current flow between any two of the shield wires since they are at the same potential. Since the voltages applied to the selected row wire and, the selected column wire are not equal, there *will* be a current flow between the shield wires and either the selected row or column wire. Recall, however, that we only measure the current through the row wire. Therefore, if we keep the shield wires at the same voltage as the wire we are sensing the current through (the row wire), there will be no current flow between them and the sense wire. As a result we read only the resistance between the selected row and column wires.

When we first talked about how to measure the resistance across the tactile array, we mentioned that it could be done either by applying a known current across it and measuring the voltage, or by applying a known voltage and measuring the current, but we gave no reason for selecting the second method over the first. The reason this selection was made is that, in order to use the unselected wires as shield wires, we must know what voltage to apply to them. If we pass a current through the sense wire, and then use a voltage follower to apply the same voltage to the shield wires as appears on the sense wire, it is very easy to obtain positive feedback and get no information at all. If we measure the resistance by applying a fixed voltage to the sense wire, there is no problem applying the same voltage to the shield wires.

Since we need to use all the unselected wires as shield lines, it is no longer enough to have only one analog switch per row and column of the pad. We now need two switches, one to connect the wire as a sense line, the other to connect it as a shield line. This is done using a 2 to 1 analog multiplexer. A detailed schematic is shown in figure 3-2. The analog multiplexers are controlled by a 74138, which pulls the control line low for exactly one pair of sensor wires, causing them to be connected to +5 and the sense wire. All of the remaining wires are connected as shield wires.

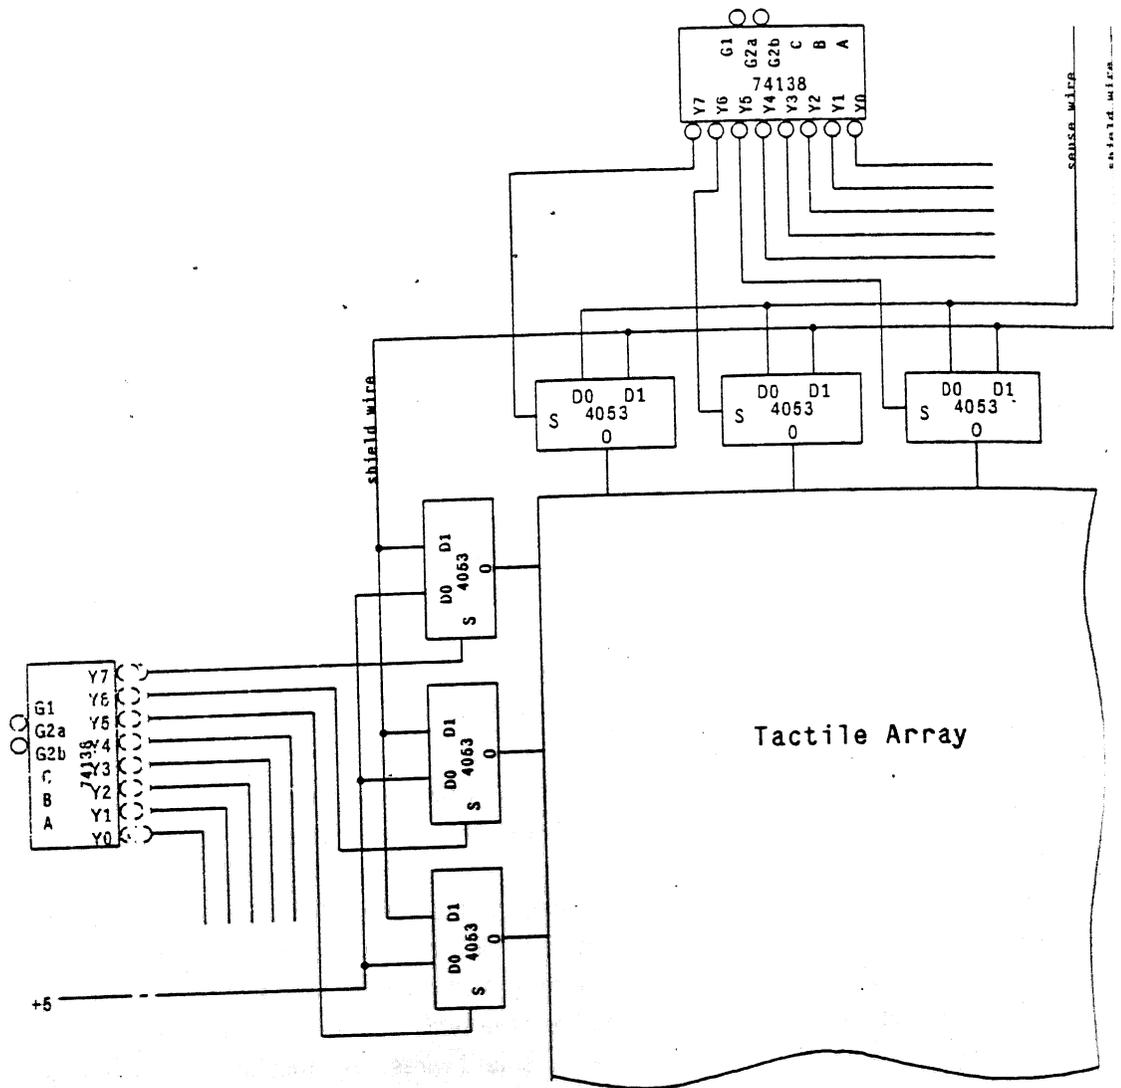


Figure 3-2: Detailed View of Connection to the Sensor Array

3.3. Input /Output Circuitry

Once the sensor value has been obtained, it needs to be communicated to the LSI-11/23. This is done through an off the shelf parallel interface board plugged into the backplane of the LSI-11/23. The amount of circuitry required to talk to this board is minimal, input and output latches plus a small amount of control logic. A circuit diagram for this portion of the circuit is contained in figure 3-8

The basic timing diagram for an input cycle for the parallel interface is shown in figure 3-4. The parallel interface pulls the line SEND DATA L low, signalling that the LSI-11/23 is ready to

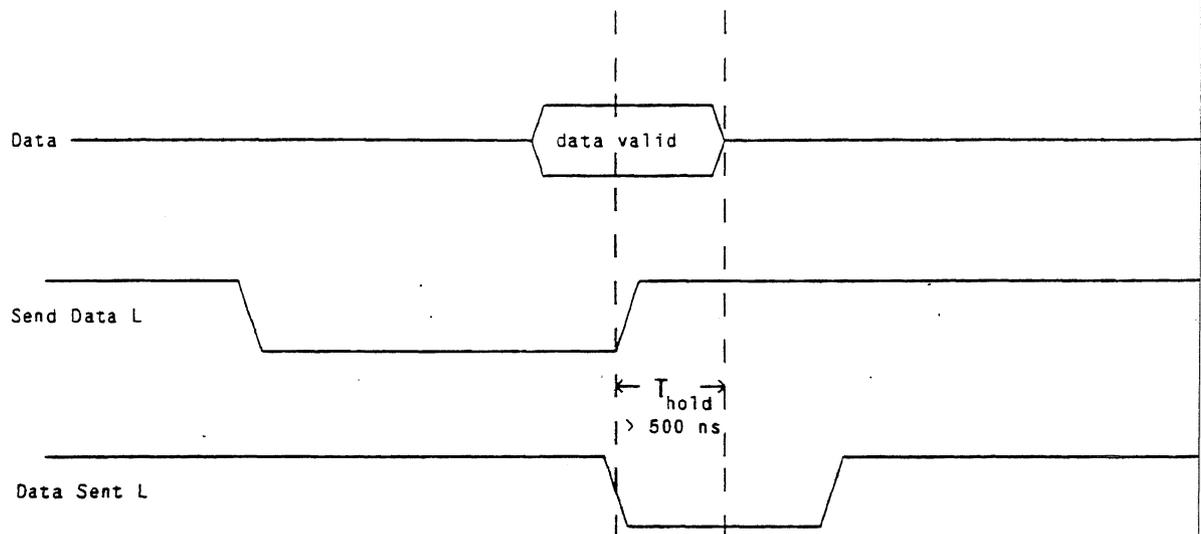


Figure 3-3: Input Timing Cycle

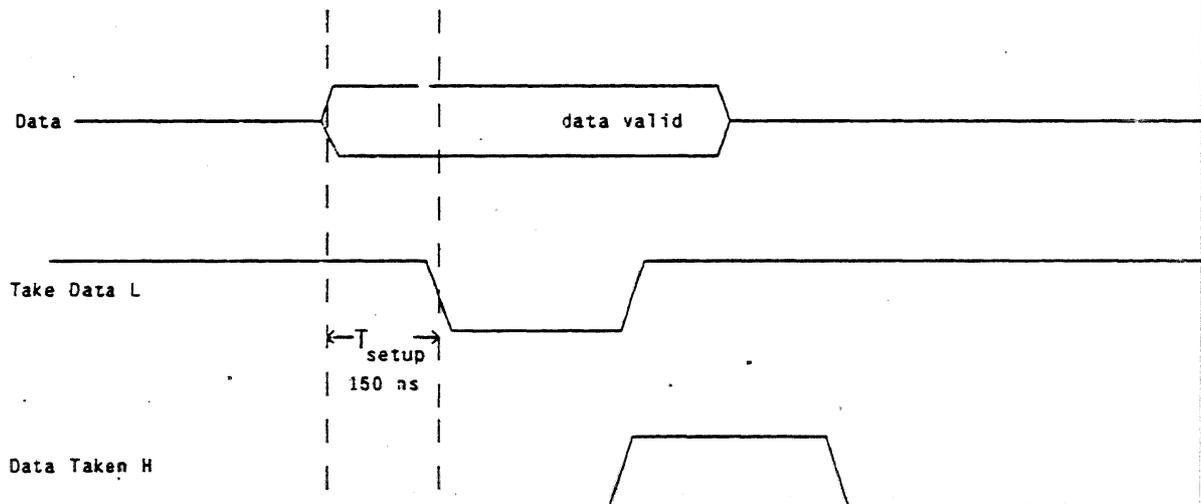


Figure 3-4: Output Timing Cycle

more input data. This signal is actually generated just after the 11/23 reads the previous data, so it does *not* mean that the 11/23 is waiting for data. When it sees this signal, the sensor interface puts new data on the ribbon cable, and then some time later pulls the line DATA SENT L low, telling the 11/23 that there is valid data on the cable. DATA SENT L is kept low until the next time the 11/23 asserts SEND DATA L. All of the timing for the input cycle is handled by the control logic which will be discussed in section 3.4.

The timing for an output cycle is almost identical. When the LSI-11/23 outputs data to the serial port, the line TAKE DATA L is pulled low. The sensor board responds by pulling DATA

TAKEN H high, signalling that the data has been received. The data on the ribbon cable is guaranteed to be valid as long as the signal TAKE DATA L is asserted.

The output cycle operates asynchronously of the remainder of the control logic on the board. One clock cycle after TAKE DATA L is asserted, the internal control line INPUT DATA READY L is asserted and the line DATA TAKEN H is asserted. Asserting INPUT DATA READY L forces the control logic into a known state, as well as transferring data from the ribbon cable to the registers controlling which element of the sensor pad is selected. The sensor board remains in this state until the 11/23 removes the signal TAKE DATA L. One clock cycle after this happens DATA TAKEN H and INPUT DATA READY L are negated, allowing normal operation to resume at the element of the sensor pad specified by the write operation.

3.4. The Control Logic

There are many ways to implement a given set of control logic. In the design used, the control circuitry was implemented as a finite state machine in order to simplify the design of the circuit and allow for easy modifications to the order in which things are done. In this design methodology, there is a single register which encodes what *state* the control logic is in. Each state has a set of combinational logic which determines the next state to enter (or possibly to remain in the same state), as well as having a set of control lines which are asserted in that state.

Before the finite state machine could be designed, the number of states necessary for controlling the sensor hardware had to be determined. There are several tasks the control logic needs to perform, most of which have already been mentioned. They are listed again here, giving a rough idea of the sequence they must occur in.

1. Ramp the digital to analog converter until the output voltage matches the voltage from the sensor array.
2. Wait for the output buffer to be empty.
3. Move the result of the analog to digital conversion and the array address to the output buffer.
4. Wait for the data on the ribbon cable to settle.
5. Assert DATA SENT L, signalling that data is valid.
6. Clear the counter in the analog to digital conversion circuit.
7. Increment the column (and possibly the row) counters driving the tactile array.

If we were to do things in exactly this way, we would require a total of seven states in the state machine. Fortunately, we can overlap some of these operations and change the order slightly. The sequence actually implemented is

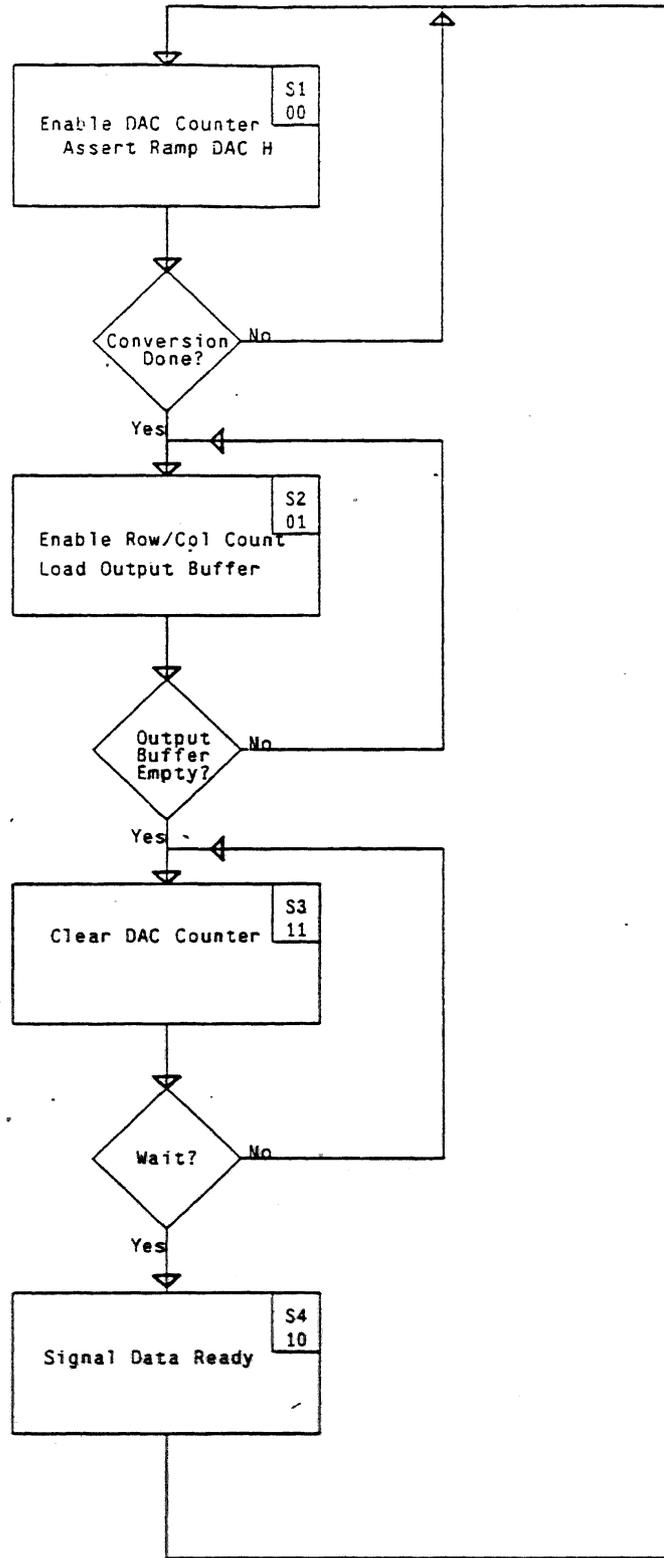


Figure 3-5: State Diagram for the Finite State Machine

1. Ramp the digital to analog converter until the output voltage matches the voltage from the sensor array.
2. Wait for the output buffer to be empty, increment the row and column counters as needed, move the result of the analog to digital conversion to the output buffer.
3. Clear the counter in the analog to digital conversion circuit, wait for the data on the ribbon cable to settle.
4. Assert DATA SENT L.

This reduces the number of states to four. A state diagram for this algorithm is shown in figure 3-5.

There are several things which need to be explained about this algorithm. In the second state, we wait for the output buffer to be empty, implying that the state machine will be in this state for an indeterminate number of clock cycles. If the state machine actually did increment the row and column counters every clock cycle it was in this state, this algorithm would not work. Instead, the row and column counters have two enable lines. One is derived from the state the control circuit is in, the other is derived from the BUFFER EMPTY H signal. The result is that the row and column counters are incremented only when the state machine *exits* this state. The same mechanism is employed to keep the output buffer from being loaded until the state is exited, i.e. not until the output buffer is empty.

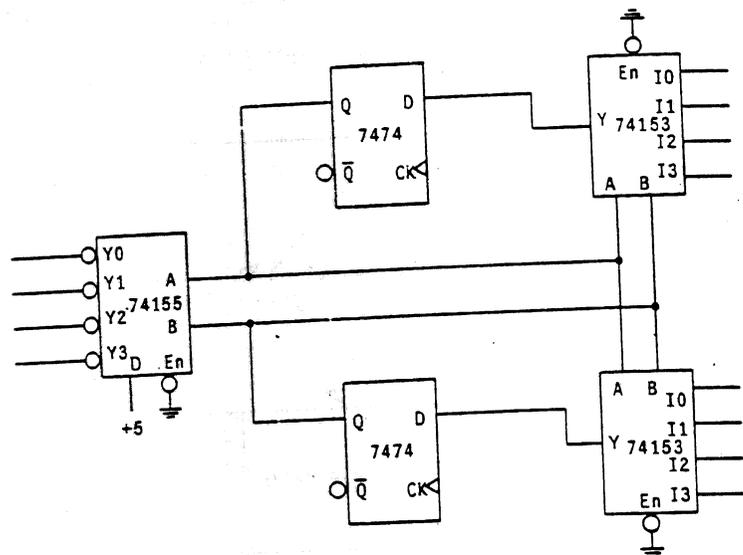


Figure 3-6: Four state finite state machine

A circuit for a four state finite state machine is shown in figure 3-6. The state is encoded by the value in the two flip-flops. It is decoded by the two-line-to-four-line decoder so that the individual state signals are available to produce control lines. The state is also used to drive

two four-line multiplexers. The output from these multiplexers is then fed into the flip-flops as the next state. All that is needed to use this circuit is an assignment of numbers to each of the four states, a list of the control lines to be asserted during each state, and the combinational logic to determine the next state for each state.

If the proper state numbers are assigned to the four states, the logic to determine the next state becomes very simple. Such a state assignment was made, and the resulting state numbers are shown in the state diagram in figure 3-5.

The only function that has not been built into the finite state machine is accepting data from the LSI-11/23. This operation may occur while the state machine is in any state, so it could not be part of the state machine. When the 11/23 writes data out to the sensor interface, the interface should immediately begin converting the value at the specified location on the pad. This means that the state machine must be forced to clear the counter in the analog to digital converter, and then start a new conversion. This can be done by pulling the preset lines on the flip-flops encoding the state low, forcing the state machine into the third state in the state diagram. The state machine will then clear the counter, and start operations at the new location.

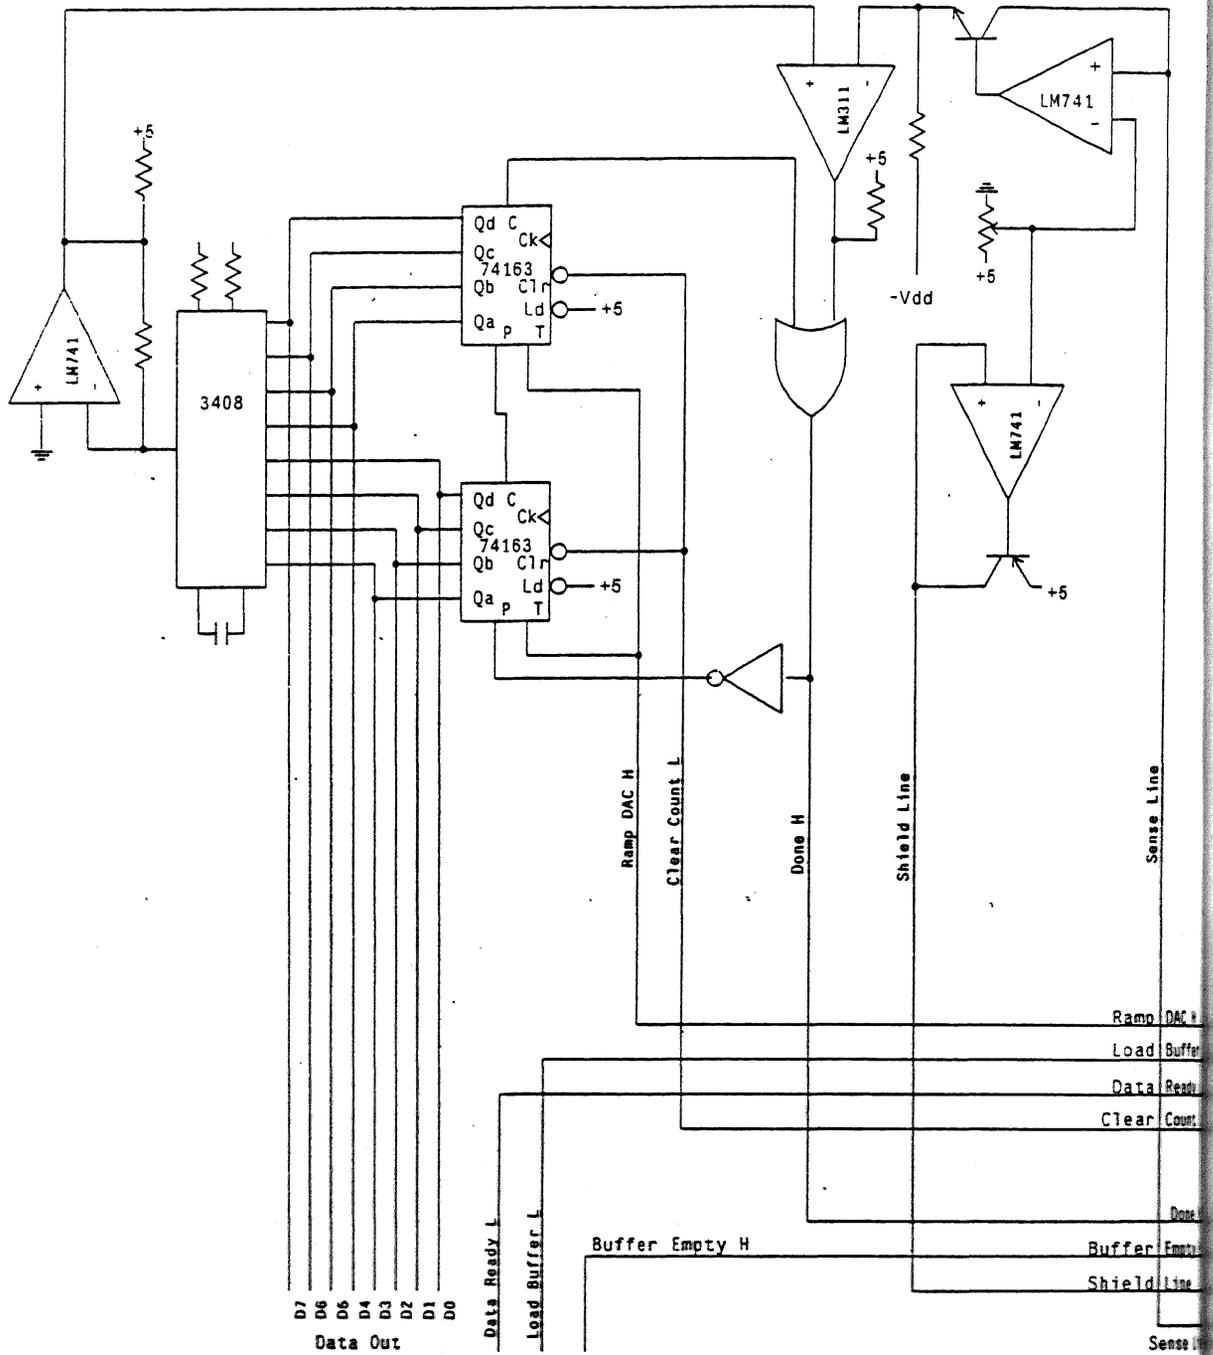


Figure 3-7:

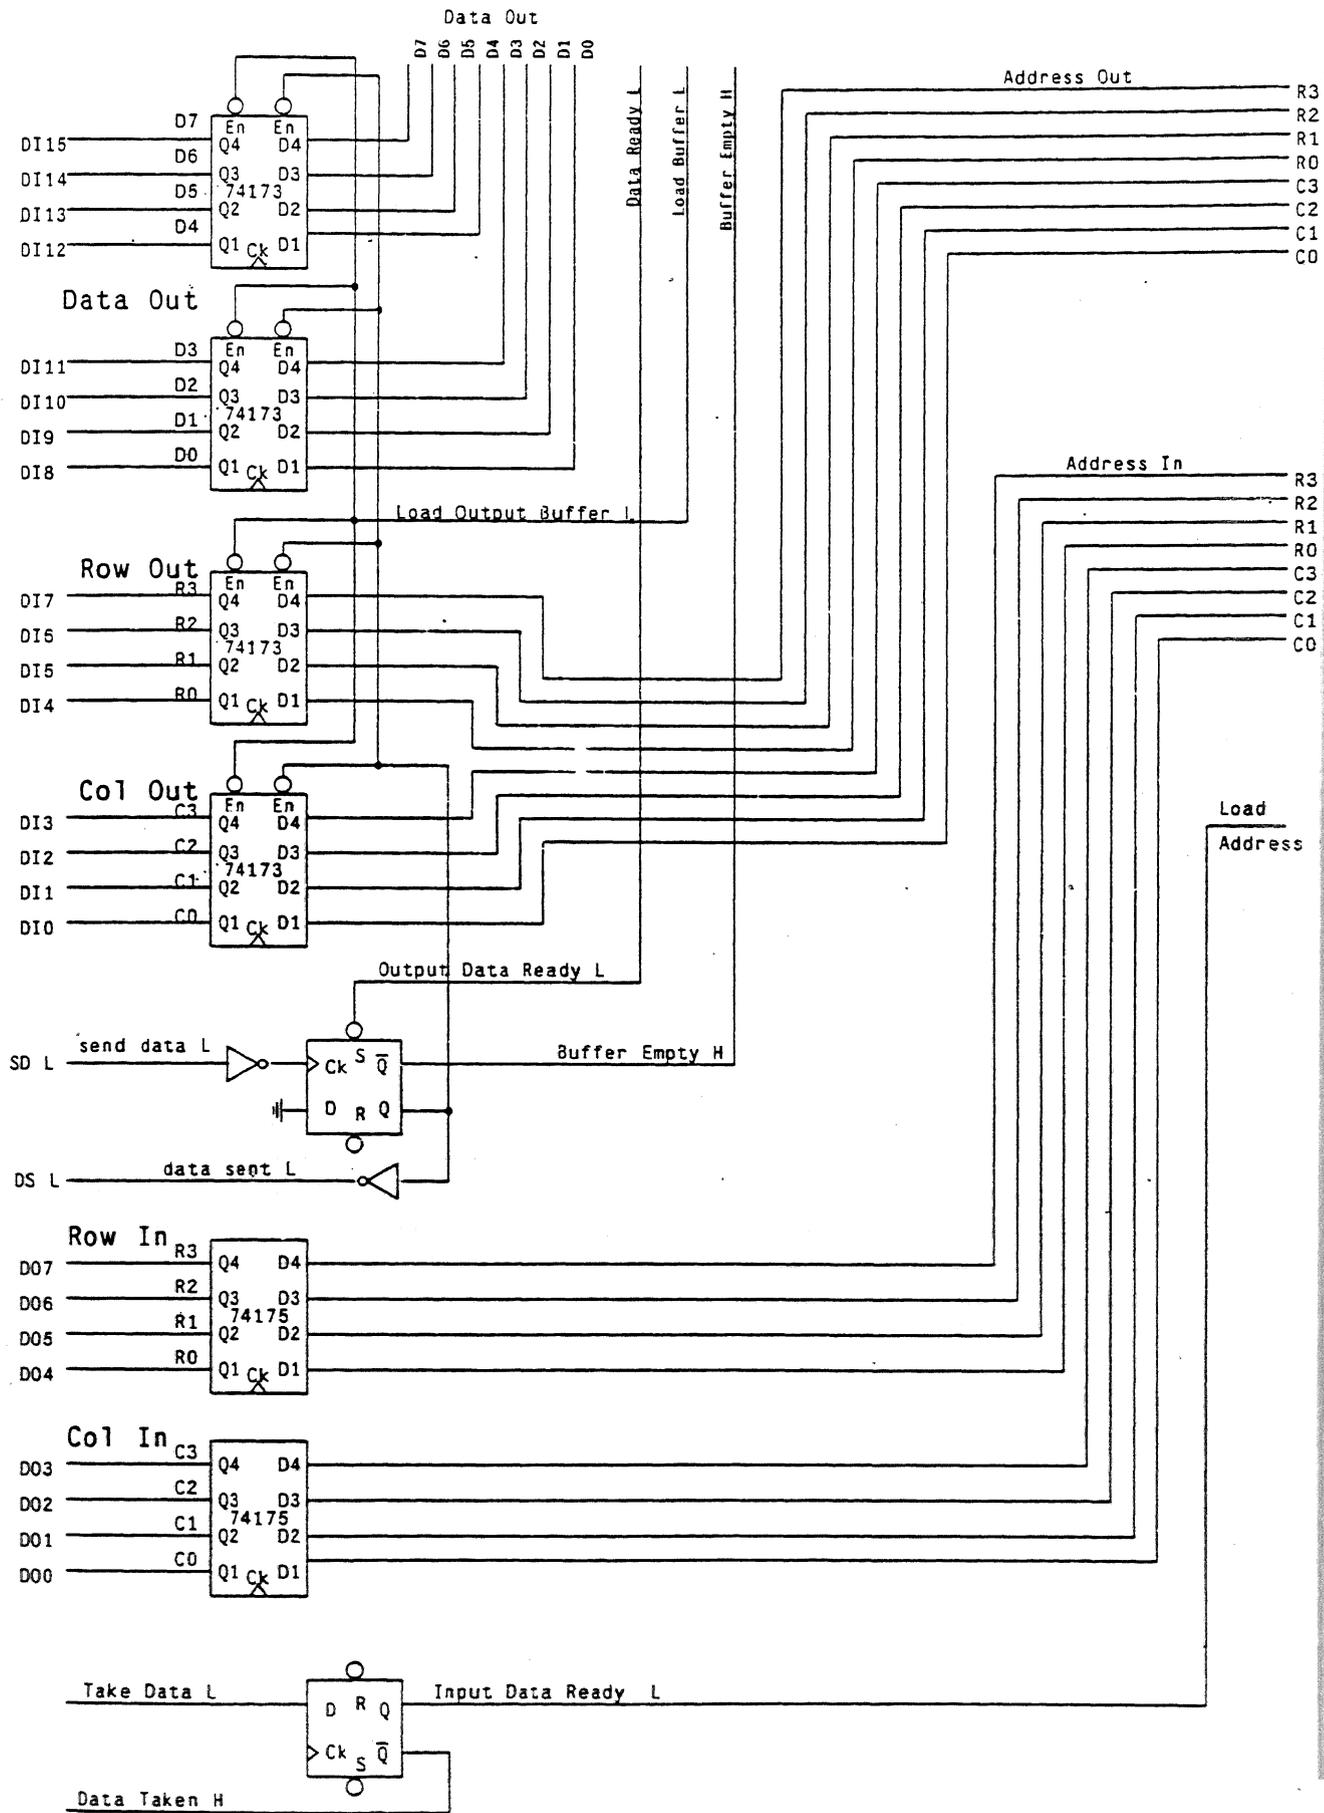


Figure 3.8.

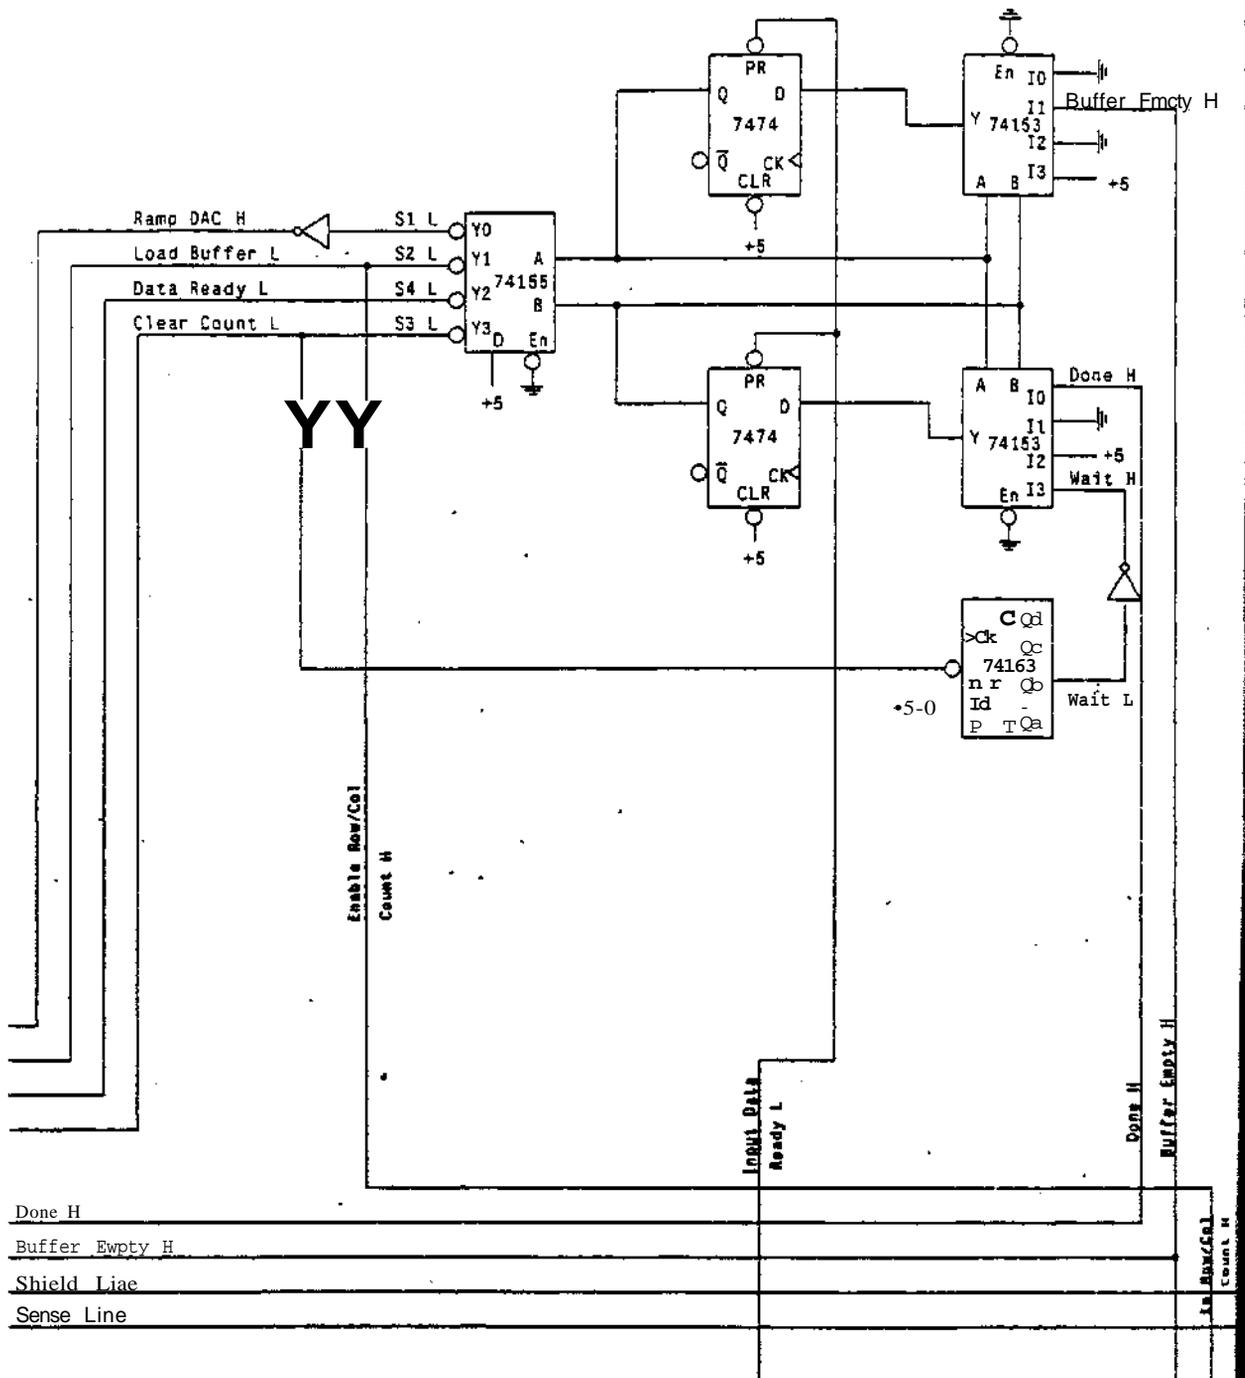


Figure 3-9:

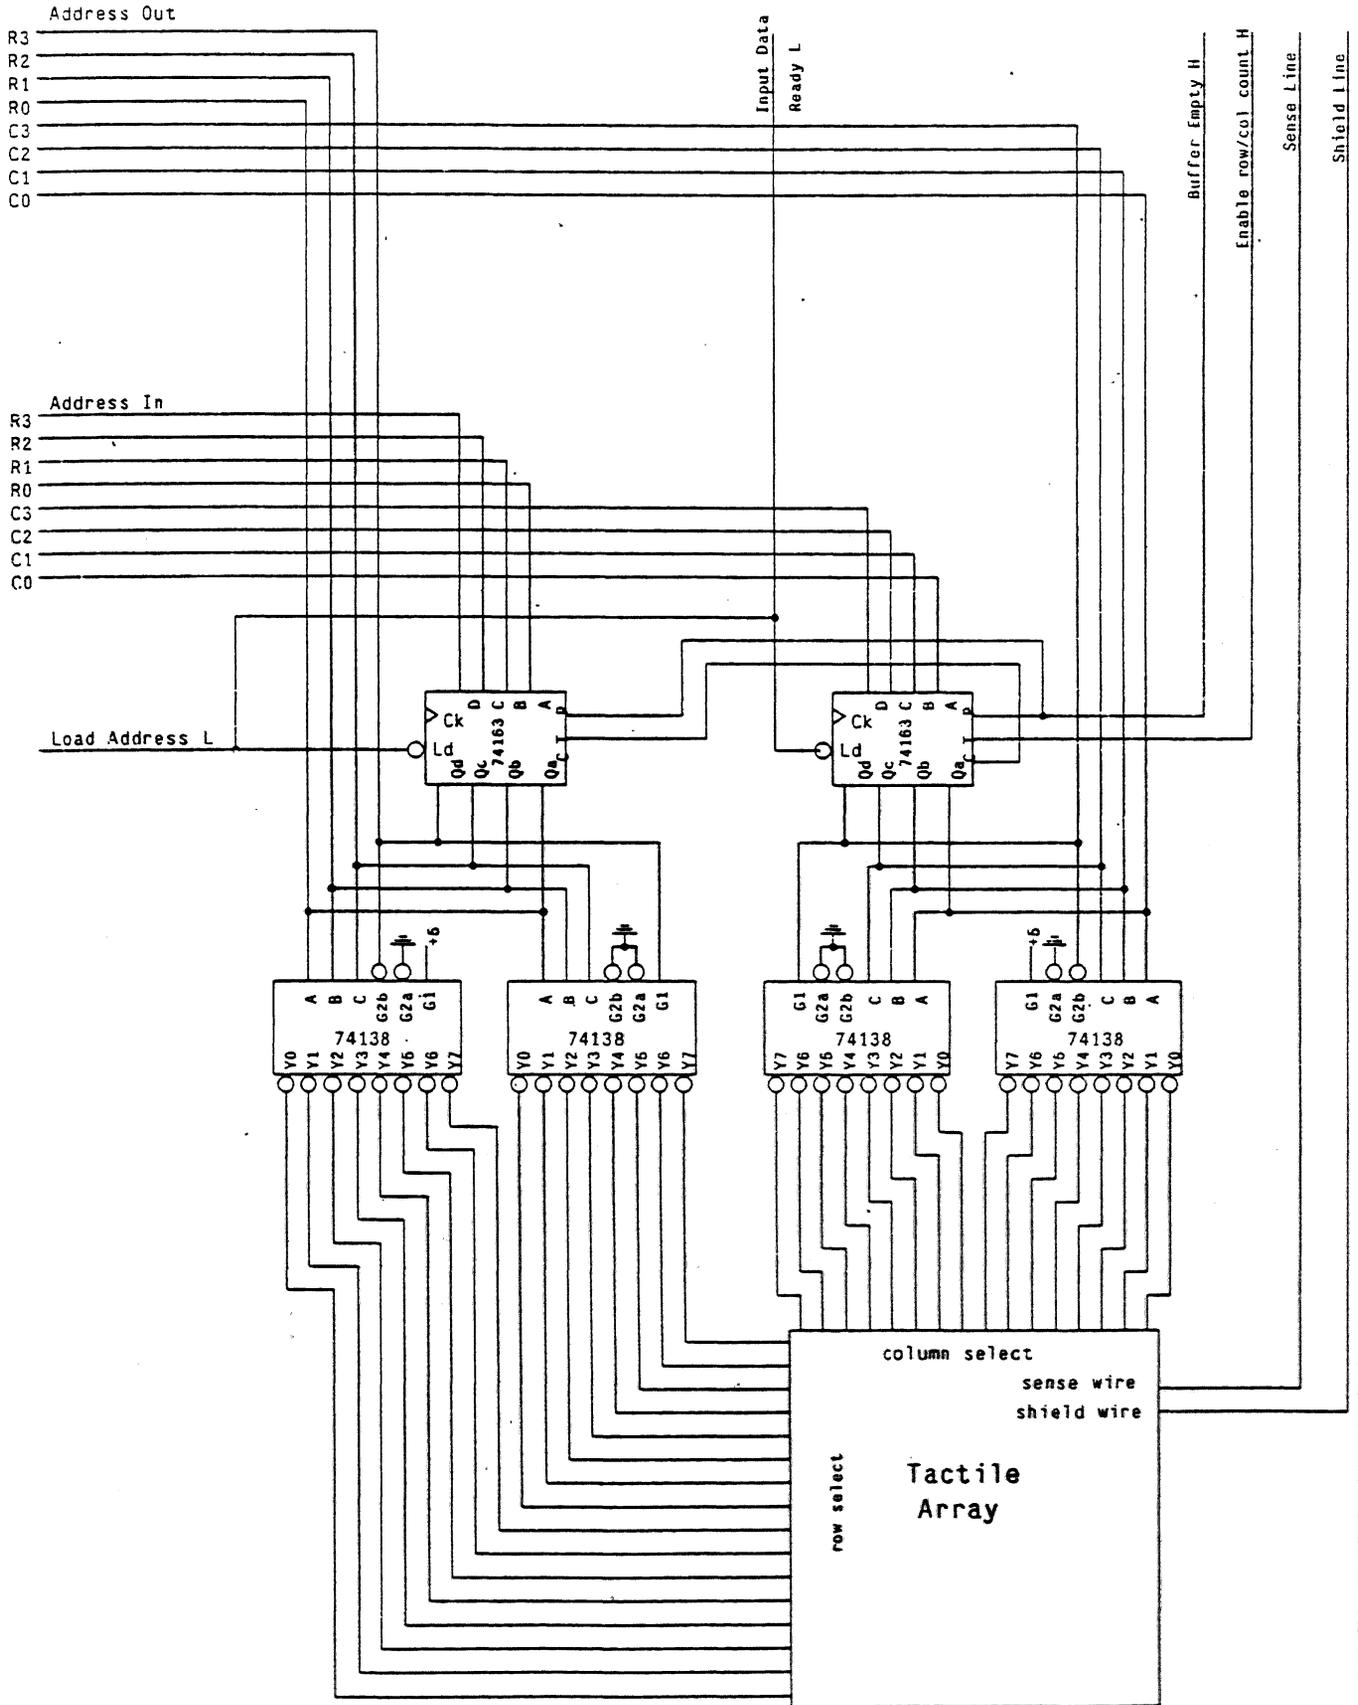


Figure 3-10:

Chapter 4

The Recognition System

The object recognition system implemented using the tactile sensor was fairly simple, yet it proved to be capable of reliably discriminating among a set of common objects. The process used to recognize the objects can be broken down into three major steps. First, the data from the tactile array must be processed to determine which elements of the array correspond to the object being sensed. Next, a set of features is extracted from the data, and then compared to the features obtained from a training set.

4.1. Segmentation of the Data

Separating objects from each other and from the background is a problem known as segmentation. In general, segmentation of an "image" such as that obtained from the tactile array can be a difficult problem. However, by insisting that only one object can be on the pad at a time, segmentation can be greatly simplified. One way to segment the image is to set some arbitrary threshold. Data points that are greater than this threshold are assumed to be part of the object, points that are less than the threshold are part of the background.

The main advantage of this method is that it is efficient and extremely easy to implement. However, it does suffer from several problems. The most obvious of these is that if the threshold is indeed arbitrary, then there is no guarantee that it will do a reasonable job of separating the object from the background. Also, this method will not yield a connected *region* for the object; it decides whether or not a point is in the object without paying any attention to the surrounding area. Hence, a single point of noise in the background may be interpreted as part of the object even though it is nowhere near the rest of the object. Alternatively, if an isolated point in the interior of the object falls below the threshold, it will be considered to be part of the background.

With a few modifications, this was the method used in the project. One of the changes is due to the fact that the resistance of the pad is not constant, so a different threshold is used at each point. This threshold was determined by watching the pad with no object on it, and is set slightly higher than the highest value seen at a given point. This yields a first estimate as to where the object is. The mean and variance of the data points greater than the threshold are then calculated, which gives information about how far into the pad the object is pressed, and how flat its surface is. The data is then thresholded two standard deviations below the

calculated mean, giving the final segmentation of the data. This procedure is an attempt to set the threshold so that any major depressions will be below the threshold, even though they are not deep enough to allow the pad to relax totally. This threshold also tends to eliminate the spreading of sharp edges caused by the mechanical properties of the pad.

4.2. The Feature Set

There are many possible sets of features which could be used to do the object classification; which features should be used is determined in part by the segmentation method used, and in part by the ease with which the features can be calculated. The segmentation algorithm yields a set of points which are classified as being part of the object. Although it would be possible to extract the edge of the object from these points, it is not a trivial process. Hence, features such as the perimeter of the object or what angles and lines there are in the boundary are not particularly suitable. In addition, since there may be some errors in the segmentation process, the features should not be extremely sensitive to leaving a few points out.

The initial feature set selected consisted of the area of the object, and its first and second moments (m_{10} , m_{01} , m_{20} , m_{02} , and m_{11}). These features have drawbacks, however. Both the first and second moments depend on the position of the object on the pad. In order to eliminate this dependence, we can calculate the moments using the object's center of area as the origin. The center of area can be found as

$$coa_x = \frac{\sum_{i=1}^{16} \sum_{j=1}^{16} iX_j}{\sum_{i=1}^{16} \sum_{j=1}^{16} X_j} \quad coa_y = \frac{\sum_{i=1}^{16} \sum_{j=1}^{16} jX_i}{\sum_{i=1}^{16} \sum_{j=1}^{16} X_i}$$

where X_{ij} is the number of pixels in the object at position (i, j) and 0 if (i, j) is part of the background. Once this is done, the first moments become identically zero (they are, in fact, used to find the center of area), but the second moments no longer depend on the object's position on the pad.

The second moments still depend on the orientation of the object on the pad. This dependence can be expressed by the equation

$$m'_{pq} = \sum_{i=0}^n \sum_{j=0}^n (i^p j^q) (\cos \theta)^{p+q} (\sin \theta)^{q-p} m_{ij}$$

where θ is the angle of rotation of the object, m'_{ij} is the rotated moment, and m_{ij} is an unrotated moment. Although it is possible to define any orientation of the object as being at the angle $\theta = 0$, there is a preferred orientation where the coordinate axes are parallel to the major and minor axis of the object. In this orientation, the moment m_{11} is, identically zero. One way to see this is to interpret m_{20} and m_{02} as the variance of the object along the x and y axis respectively, in which case m_{11} is the covariance of the object along the axis. If we

calculate the moments according to this preferred orientation, the feature set is reduced to the area of the object, m_{2Q} , and m_{Qr}

We could find the amount we need to rotate the axis to obtain $m_{1,7} = 0$ from the above equation, but there is no real need to find the actual angle since we only need the rotated values of the moments. A simple way to obtain these is, once again, to view the second moments as the variance and covariance of the object. This gives the matrix

$$M_2 = \begin{bmatrix} m_{2,0} & m_{1,7} \\ m_{1,1} & m_{0,2} \end{bmatrix}$$

We can then apply the same methods which are used to diagonalize covariance matrices of random variables to diagonalize this. The resulting matrix will have $m_{1,7} = 0$, and the rotated values of m_{2Q} and m_{Q2} on the diagonal. But the diagonal entries of the resulting matrix are just the eigenvalues of the original matrix M_r . Therefore, in order to find the rotated second moments, we only need to find the eigenvalues of the 2 by 2 matrix M_r

4.3. The Decision Rule

Although an optimal decision rule in terms of minimizing the probability of classification error is well known, a suboptimal but simpler decision rule was implemented. The decision rule used was a variant of the nearest mean classifier. For each object o_j , we need to know the expected value $\mu_{i,j}$ and variance $\sigma_{i,j}^2$ for each of the features. Given an observation with features $x_{i,j}$ we compute the distance to the mean for each object as

$$d_j^2 = \sum \left(\frac{\mu_{i,j} - x_{i,j}}{\sigma_{i,j}} \right)^2$$

and classify the observation as the object o_j with the smallest corresponding d_j^2 . Thus, the observation is classified by the nearest mean normalized to the standard deviation.

As mentioned, this is not the optimal decision rule, so it is worth investigating how they differ. The decision rule used does not take the covariance of the features into account. In fact, if we assumed independent features and then simplified the optimal classifier, it would reduce to the decision rule used. Unfortunately, there is no reason to expect the features used to be independent. If a point of the object is incorrectly labeled as part of the background during segmentation, it decreases the area of the object, as well as decreasing both of the second moments. Therefore, we may expect the features to be highly correlated. The decision rule used cannot be justified as being essentially the same as the optimal decision rule; it is used here because it is significantly simpler computationally.

4.4. Results

Before the recognition system could be tested, a set of objects needed to be selected. Rather than trying to obtain a special set of objects for the test, a variety of objects which were found in the lab were used. Five such objects were used. They included a meter (a 2 inch by 1 inch rectangle), a thin bar (about three inches long), the top of a round container (a 1.5 inch circle), a battery lying on its side (a 1.75 inch long cylinder), and a roll of tape (a donut with outer radius of 2 inches, inner radius of 1 inch). In figures 4-1 through 4-3, there are three representations of each object. First, there is the raw data received from the tactile array. The next diagram has the symbol '*' if the data is more than one standard deviation above the mean, '+' if it is less than one standard deviation above the mean, and '-' if it is less than one standard deviation below the mean; this diagram may be interpreted as an isometric plot of the data. The last diagram in each figure plots the output of the segmentation routine when it is run on the data given.

A training program was then run, and each object was presented to it ten times. This allowed the mean and variance of each feature to be calculated for the object recognition routine. Scatter diagrams for each pair of features are plotted in figures 4-4, 4-5, and 4-6. There is also an ellipse drawn for each object. The length of the axes of each ellipse are determined by the standard deviation of the corresponding feature for each object. The ellipses ignore the covariance of the features, so they represent the model the recognition system has of the scatter plot for a given object. A comparison of the orientation of the scatter plot to the ellipse therefore gives an idea of how good the model is.

After the mean and variance had been determined for each feature of each object, the object recognition system could be run. Two sets of 25 trials were run, with each object being presented 5 times in each trial. A sample of the output is included in figure 4-7. The results of the trials are presented in figure 4-8; and scatter plots are shown in figures 4-9, 4-10, and 4-11.

4.5. Sources of Error

No pattern recognition system is error-free, and the use of a sub-optimal decision rule can only increase the number of errors. However, there are several additional causes of error in this system. We have already mentioned one type of error which is introduced by the segmentation process, but there is an additional type of error it introduces. An object is represented by the points of the pad which fall in its interior; this can be viewed as approximating the object by a collection of squares centered at the points of the pad which lie within the interior of the object (see figure 4-12). Since most objects cannot be represented in this way, this introduces an error into the calculations of the area and moments of the object.

In addition to the error introduced through the discrete approximation to the objects area, there is an error introduced in the calculations of the moments from the data. The moments are calculated from the formula

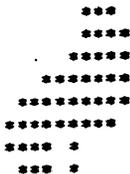
opening meter.pic.

```

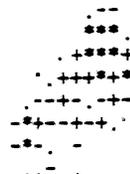
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 3a 75 7a 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 2d ff ff ff 3a 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 25 a7 ee ff ff b0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 9 53 bb c0 b4 fa d6 db 0 0 0 0 0 0 0 0 0 0
0 0 7 55 6c 7f 9f 6c 43 8f 98 67 0 0 0 0 0 0 0 0 0 0
0 0 8c fe ab 8c 78 af 8f b4 3a 0 0 0 0 0 0 0 0 0 0 0
0 0 7f f4 6e 53 1d 64 1b 1f 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 47 3f 65 1f 37 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 1f 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 22 0 0 b 1a 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

```

opening meter.pic.
Mean: 126.912277, Std: 75.720001.
thresh: 51, area: 46



Mean: 150.500000, Std: 64.790886.



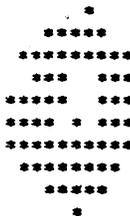
opening tape.pic.

```

0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 17 0 30 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 2c 5a 96 5e bc 88 14 4 0 0 0 0 0 0 0 0 0
0 0 0 2e 47 50 3f 33 5e cf cb 5c 0 0 0 0 0 0 0 0 0 0
0 0 0 1f 47 5e 3a 1a 1d b8 ff f6 0 0 0 0 0 0 0 0 0 0
0 0 38 7f 7c 57 35 14 0 66 ff ff 1 17 0 0 0 0 0 0 0
0 0 40 bc b0 63 1 30 19 ef ff ff 1 0 0 0 0 0 0 0 0
0 0 5a d0 8f 94 2e 9f 5e ff ff 75 0 0 0 0 0 0 0 0 0
0 0 0 35 7f ff c3 ff e3 ff 98 9 0 0 0 0 0 0 0 0 0
0 0 0 0 0 4f 73 f7 98 aa 9 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 21 8b 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

```

opening tape.pic.
Mean: 125.095894, Std: 79.965790.
thresh: 45, area: 61



Mean: 144.918030, Std: 72.423233.

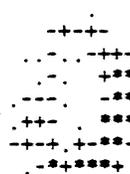


Figure 4-1: Data for the meter and tape

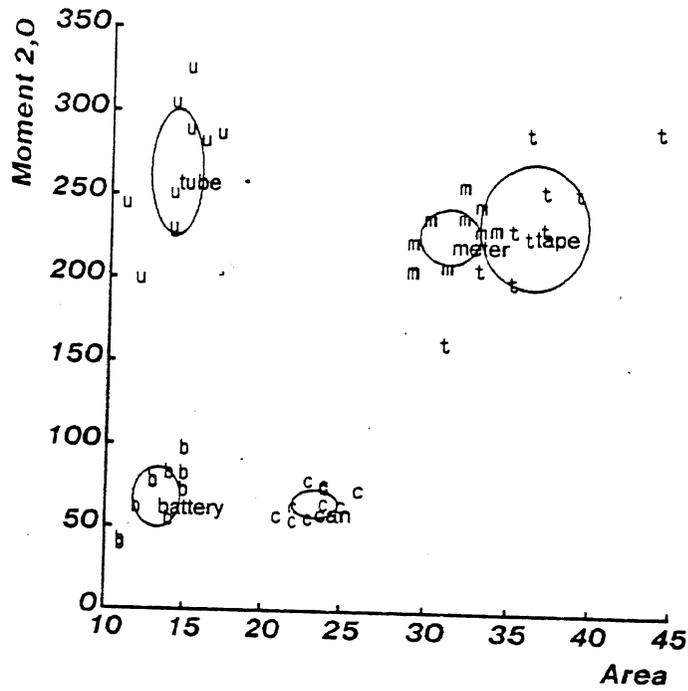


Figure 4-4: Scatter diagram of $m_{2,0}$ versus Area for the training set

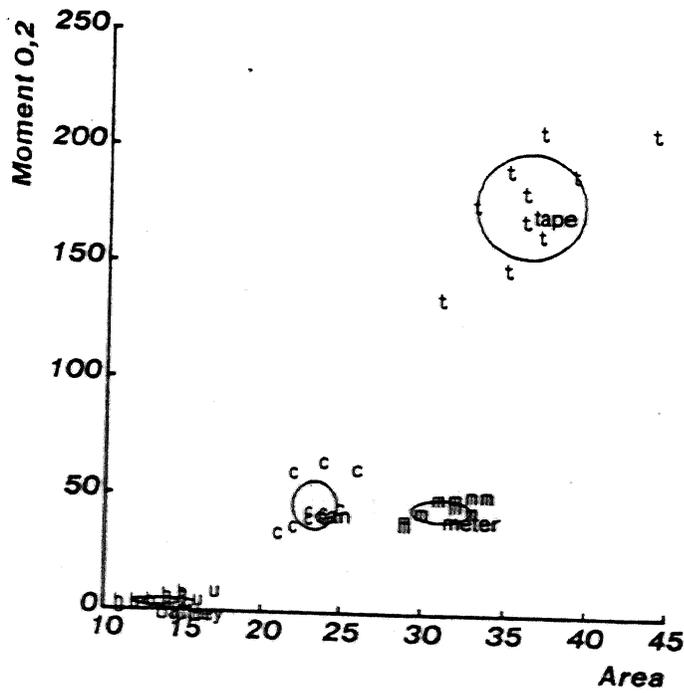


Figure 4-5: Scatter diagram of $m_{0,2}$ versus Area for the training set

t: 7E A: 31 e1: 129.447402, e2: 132.036463
 can d1: -6.660770, d2: -6.992915, d3: -1.966156, distance: 9.854906.
 tube d1: -40.665016, d2: 2.173405, d3: -8.148232, distance: 41.530242.
 tape d1: 1.655303, d2: 2.052449, d3: 2.333333, distance: 3.638135.
 battery d1: -31.354251, d2: -9.235520, d3: -9.814959, distance: 34.127956.
 meter d1: -3.754688, d2: 1.792424, d3: 1.435531, distance: 4.397211.
 I think it is a tape.
 The minimum distance is 3.638135.
 The next closest mean is meter, with a distance of 4.397211.

t: 84 A: 30 e1: 207.203033, e2: 42.963672
 can d1: 1.210397, d2: -8.159755, d3: -1.556540, distance: 8.394611.
 tube d1: -12.223998, d2: 1.620320, d3: -7.632521, distance: 14.501963.
 tape d1: 4.066128, d2: 1.747164, d3: 2.555555, distance: 5.110465.
 battery d1: -8.565190, d2: -10.747915, d3: -9.237608, distance: 16.569749.
 meter d1: 1.494365, d2: 1.372637, d3: 1.777324, distance: 2.697432.
 I think it is a meter.
 The minimum distance is 2.697432.
 The next closest mean is tape, with a distance of 5.110465.

t: 9F A: 22 e1: 51.871112, e2: 44.947071
 can d1: 1.035151, d2: 2.048139, d3: 1.720387, distance: 2.868125.
 tube d1: -12.857298, d2: 6.458884, d3: -3.506834, distance: 14.809632.
 tape d1: 4.017568, d2: 4.470385, d3: 4.333333, distance: 7.409654.
 battery d1: -9.092193, d2: 2.482997, d3: -4.618804, distance: 10.496028.
 meter d1: 1.377483, d2: 4.957582, d3: 4.511668, distance: 6.843262.
 I think it is a can.
 The minimum distance is 2.868125.
 The next closest mean is meter, with a distance of 6.843262.

t: 6F A: 12 e1: 57.401649, e2: 2.848350
 can d1: 4.754849, d2: 1.684690, d3: 5.816547, distance: 7.699265.
 tube d1: 0.584856, d2: 6.286609, d3: 1.650274, distance: 6.525865.
 tape d1: 5.048297, d2: 4.373426, d3: 6.555555, distance: 9.358818.
 battery d1: 1.669208, d2: 2.011915, d3: 1.154701, distance: 2.857865.
 meter d1: 3.858358, d2: 4.829942, d3: 7.929599, distance: 10.054641.
 I think it is a battery.
 The minimum distance is 2.857865.
 The next closest mean is tube, with a distance of 6.525865.

t: 55 A: 19 e1: 267.847125, e2: 6.247614
 can d1: 4.454502, d2: -13.448272, d3: 2.349235, distance: 14.468684.
 tube d1: -0.500531, d2: -0.885500, d3: -1.959701, distance: 2.207957.
 tape d1: 4.965071, d2: 0.336648, d3: 5.000000, distance: 7.054459.
 battery d1: 0.800278, d2: -17.800008, d3: -2.886752, distance: 17.853124.
 meter d1: 3.658039, d2: -0.483952, d3: 5.537047, distance: 6.653898.
 I think it is a tube.
 The minimum distance is 2.207957.
 The next closest mean is meter, with a distance of 6.653898.

Figure 4-7: Sample output from the recognition routine

Symbol	object	guess	Area	$M_{2,0}$	$M_{0,2}$
a	tape	meter.	29	175.71	104.77
b	tape	tape.	34	184.11	131.92
c	tape	tape.	31	227.35	143.23
d	tape	tape.	41	243.56	215.61
e	tape	tape.	32	204.13	195.81
f	meter	meter.	31	261.03	42.84
g	meter	meter.	35	246.99	47.30
h	meter	meter.	32	213.67	41.51
i	meter	meter.	37	279.79	56.05
j	meter	meter.	34	234.90	51.84
k	can	can.	24	74.42	44.87
l	can	can.	28	73.69	59.85
m	can	can.	26	68.97	56.60
n	can	can.	25	76.66	45.74
o	can	can.	28	83.75	61.03
p	battery	battery.	14	65.57	3.36
q	battery	battery.	12	104.16	1.51
r	battery	battery.	11	47.34	2.30
s	battery	battery.	12	59.20	2.38
t	battery	battery.	16	86.14	6.80
u	tube	tube.	16	260.10	2.09
v	tube	tube.	12	158.39	0.77
w	tube	tube.	16	216.38	5.12
x	tube	tube.	15	187.72	4.55
y	tube	tube.	14	248.88	1.19
A	tape	tape.	34	201.82	167.59
B	tape	tape.	37	232.55	216.80
C	tape	tape.	27	211.26	128.84
D	tape	tape.	44	287.16	244.72
E	tape	tape.	28	189.60	141.95
F	meter	meter.	31	216.45	39.93
G	meter	meter.	32	173.77	37.09
H	meter	meter.	28	276.73	64.83
I	meter	meter.	34	163.01	15.78
J	meter	meter.	24	185.58	21.83
K	meter	meter.	27	71.00	65.29
L	can	can.	28	63.48	47.65
M	can	can.	23	66.04	42.39
N	can	can.	19	54.01	23.57
O	can	can.	29	79.49	70.51
P	battery	battery.	13	99.37	4.48
Q	battery	battery.	12	45.84	3.07
R	battery	battery.	14	46.61	5.68
S	battery	battery.	14	104.80	2.49
T	tube	tube.	18	137.23	9.71
U	tube	tube.	24	357.84	7.62
V	tube	tube.	16	289.42	3.01
W	tube	tube.	16	246.41	3.09
X	tube	tube.	15	251.95	2.98
Y	tube	tube.	18	225.67	3.55

Figure 4-8: Summary of the trials of the recognition system

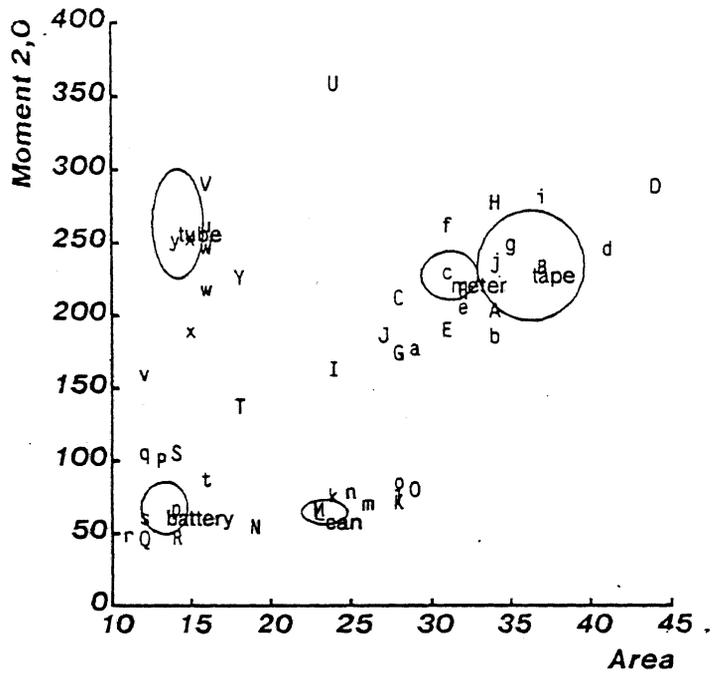


Figure 4-9: Scatter Plot of $m_{2,0}$ versus Area for test data

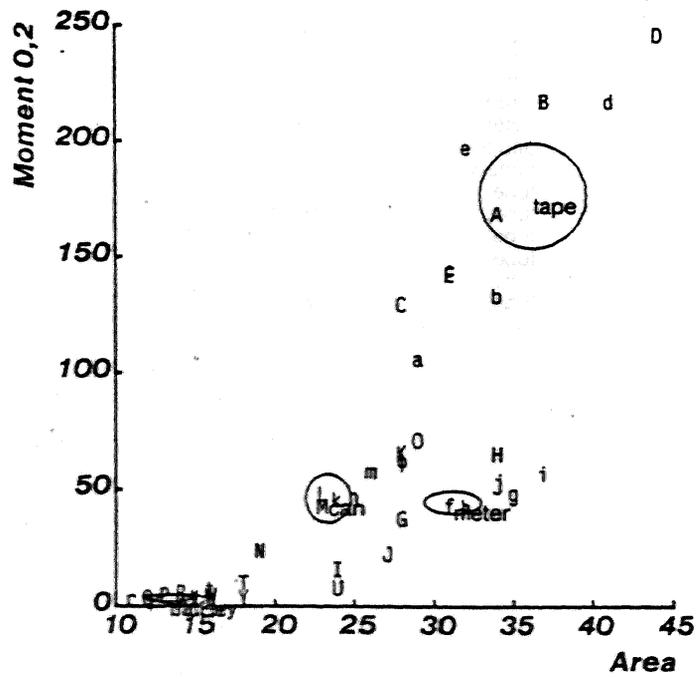


Figure 4-10: Scatter plot of $m_{0,2}$ versus Area for test data

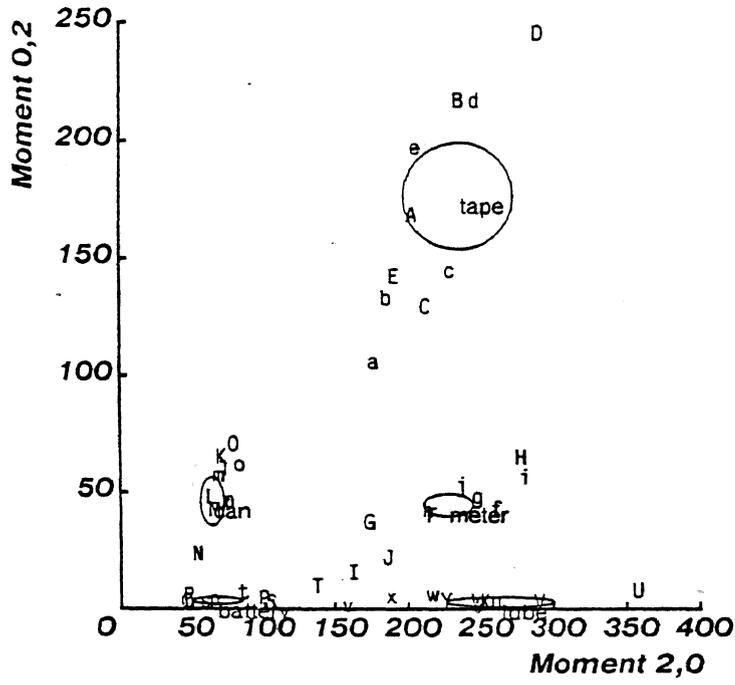


Figure 4-11: Scatter plot of $m_{2,0}$ versus $m_{0,2}$ for test data

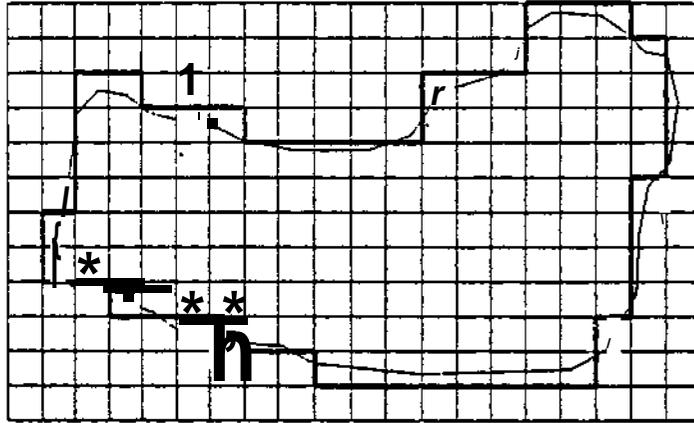


Figure 4-12: Discrete approximation to the area of an object

$$m_{p,q} = \sum_{i=1}^{16} \sum_{j=1}^{16} i \cdot j \cdot x_{ij}$$

where x_{ij} is once again the characteristic function of the object (i.e. the output of the segmentation process).

This equation, however, is only a discrete approximation to the real integral formula

$$m_{p,q} = \int_1^{16} \int_1^{16} x^p y^q \chi_{x,y} dx dy$$

This approximation could be improved, but it was felt to be accurate enough.

Another source of error was the way in which the object was presented to the pad. Since the objects were placed on the pad by hand, the pressure exerted was not at all constant from one trial to the next. The thresholding techniques used¹ in the segmentation process reduce the impact of this, but they cannot eliminate it entirely. There was another, more serious, source of error introduced by changes in pressure, however. There was no way of ensuring that the pressure exerted on the object would be uniform over its entire surface area, in fact examination of the isometric plots of the data shows that more pressure was exerted on one side of the objects than on the other, with the problem getting worse with the larger objects. With an object such as the roll of tape, this could have serious consequences. If the pressure difference is great enough, only one side of the roll of tape will be registered by the segmentation routine, in which case the tape looks very much like the meter. Similarly, if one side of the meter is depressed further than the other, it can look like the battery or the tube.

Conclusion

Although there are many potential applications for a tactile sensor array, such a sensor has not yet made it out of research labs. This is due to the many problems involved in their design. The sensor needs to be very durable, have a high resolution, have a small physical size, be relatively insensitive to noise, and have a compliant surface. In addition there is both physical and electrical coupling between elements of the array which should be eliminated or reduced as much as possible.

In order to investigate some of these problems, a prototype tactile sensor was constructed. The sensor was made of a sheet of conductive foam sandwiched between layers of conductors. When the foam is compressed at some point, the resistance through the foam decreases. By selecting the appropriate conductor on each side of the foam, the resistance at any one of 256 points could be measured. The spatial resolution of the sensor was $\frac{1}{8}$ inch.

An interface between the sensor and an LSI-11/23 was designed and built in order to evaluate the performance of the sensor. The interface allowed a program on the LSI-11/23 to examine a single element of the sensor array or to scan the entire array with a minimum of overhead in software.

As a sample application, an object recognition system was implemented using the sensor. There are many questions involved in how to build a recognition system using a tactile sensor array. These include how to separate the object from the background, what features to use, and what forms of preprocessing to perform on the tactile image. The system implemented was a first step at answering the first-two questions. When an object was presented to the sensor array, the resulting data was thresholded in order to separate the background from the object. Three features were then computed from the object, the area and the second moments along the major and minor axes. These three features were then used in a suboptimal decision rule (the nearest mean normalized by the standard deviations) to classify the object. A total of 50 trials were performed using 5 objects; the results are summarized in figure 4-8 on page 41. Two classification errors were made.

The recognition system also helped to point out some problems in the application of tactile sensor arrays. When an object was pressed into the tactile array, one side of the object tended to be depressed further than the other. Due to the thresholding of the image, this could result in only half of the object being seen. Another problem was that the contact between the conductors and the foam is very noisy when there is no object pushing the conductor into a firm contact. One reason why the image was thresholded was that the

resistance of the foam varies somewhat from point to point, so some form of calibration would be necessary before the actual values from the sensor could be used. Future research needs to investigate these problems.

In some ways, the object recognition system is a simpler problem than many of the commercial applications for tactile sensors. The objects used were only two-dimensional surfaces. When the entire three-dimensional object is involved new questions arise. A three dimensional representation of an object from tactile images of different surfaces must be developed. A search strategy must be developed to determine what part of the object to touch next when trying to recognize an object. Many problems need to be solved before tactile sensors are commonly used in commercial applications, and work on them has barely started in the laboratory.

Bibliography

- [1] Bejczy, A. K.
Effect of hand-based sensors on manipulator control performance.
Mechanism and Machine Theory 12:547-567, 1977.
- [2] Engelberger, J. F.
A robotics prognostication.
In *Proceedings of the 1977 Joint Automatic Control Conference*, pages 197-204. 1977.
- [3] Fukunaga Keinosuke.
Introduction to Statistical Pattern Recognition.
Academic Press, 1972.
- [4] Gurfinkel, V. S. et al.
Tactile sensitizing of manipulators.
Engineering Cybernetics 12(6):47-56, 1974.
- [5] Harmon, Leon D.
Touch-Sensing Technology: A Review.
Technical Report, Case Western Reserve, 1980.
- [6] Harmon, Leon D.
Automated Tactile Sensing.
Technical Report, Case Western Reserve, 1981.
- [7] Hill, J. W. and A. J. Sword.
Manipulation based on sensor-directed control: an integrated end effector and touch sensing system.
17th Annual Human Factors Society Convention, Oct., 1973.
- [8] Hillis, William D.
Active Touch Sensing.
Technical Report, MIT Artificial Intelligence Laboratory, 1981.
- [9] Larcombe, M.
Carbon Fibre Tactile Sensors.
Technical Report, University of Warwick, .

- [10] Larcombe, M.
Tactile sensors, sonar sensors and parallax sensors for robot applications.
In *3rd Conference on Industrial Robot Technology and the 6th International Symposium on Industrial Robots*, pages C3-25 - C3-32. University of Nottingham, 1976.
- [11] Page, C. J. et. al.
Novel techniques for tactile sensing in a three-dimensional environment.
In *3rd Conference on Industrial Robot Technology and the 6th International Symposium on Industrial Robots*. University of Nottingham, 1976.
- [12] Pauly, John.
The Use of Moments in Image Pattern Recognition.
Master's thesis, Carnegie-Mellon University, 1981.
- [13] Purbrick, John A.
A Force Transducer Employing Conductive Silicone Rubber.
Technical Report, MIT Artificial Intelligence Laboratory, .
- [14] Snyder, W. E. and J. St. Clair.
Conductive elastomers as sensor for industrial parts handling equipment.
IEEE Transactions on Instrumentation and Measure IM-27(1):94-99, March, 1978.
- [15] Tanner, John E., Raibert, Marc H. and Eskenazi, Raymond.
A VLSI Tactile Sensing Array Computer.
Technical Report, CalTech Jet Propulsion Laboratory, 1981.
- [16] Wang, S. M. and P. M. Will.
Sensors for computer controlled mechanical assembly.
The Industrial Robot :9-18, March, 1978.