ECHOLOCATION AND SCENE INTERPRETATION

BY

MICHAEL JOHN HANEY

B.S., Massachusetts Institute of Technology, 1977
B.S., Massachusetts Institute of Technology, 1977

THESIS

Submitted in partial fullfillment of the requirements
for the degree of Master of Science in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 1981

Urbana, Illinois

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

THE GRADUATE COLLEGE

May 1981

WE HEREBY RECOMMEND THAT THE THESIS BY

MICHAEL JOHN HANEY

ENTITLED ECHOLOCATION AND SCENE INTERPRETATION

BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR

THE DEGREE OF MASTER OF SCIENCE

William D. Ellis, Jr.

Director of Thesis Research

G. W. Swenson, Jr.

Head of Department

Committee on Final Examination†

Chairman

† Required for doctor's degree but not for master's.

O-517

Acknowledgement

I would like to thank my two advisors, Professor William D. O'Brien and Professor David L. Waltz, for their continuous willing assistance, and their infinite patience. I would also like to thank my many friends for helping me feel guilty enough to finish.

Preface

Techniques of simple (stereophonic and triphonic) echolocation are examined. An experiment designed to capitalize on these techniques is proposed, with the intent of developing a scene interpretation system that provides position and size information rapidly (with respect to picture processing systems), and provides this information in a form that facilitates further processing by other systems. The proposed system will attempt to yield 2-dimensional position, size, shape, and planar orientation information, with possible auxiliary information about material and texture. Only gross surface detail will be directly available. Fine surface detail may be inferred from the auxiliary information.

Table of Contents

Acknowledgement.....iii

Preface.....iv

Table of Contents.....v

Introduction.....1

Mathematics.....6

Information Processing.....26

The Experiment.....35

Results.....40

Conclusions.....51

Bibliography.....54

Appendix A: Hardware.....55

Appendix B: Software.....68

Chapter 1: Introduction

Vision is perhaps the most complex perception process in man. Although it has been established that only a small area of the retina (the fovea) is highly sensitive to detail, images are perceived as whole scenes. Vision, and scene interpretation, must rely on the integration of many small detailed regions seen by the fovea. Thus, scene interpretation is based on a large array of data points. Computer-driven vision systems also start from a large array of data points, whether taken from a television camera or collected by a single light detector and a scanning laser beam. These data points can be manipulated directly to enhance contrast or improve focussing, using signal processing methods (such as Fourier transformations), but no interpretation of the scene is made. More complex systems perform a limited amount of actual scene interpretation, but the amount of computation required to manipulate the data points is very large. This is principally because two dimensions of data must be processed. Once the data have been reduced into some standardized form (such as D. Marr's (of M.I.T.) primal sketch), programs exist that can identify surfaces and, to a limited extent, objects.

But not all creatures rely on visual images to understand their world. Bats, for example, seem to employ

highly sophisticted forms of sonar to aid their navigation and hunting. As far back as 1794, naturalist Lazaro Spallanzani postulated that bats relied on sound to perceive obstacles. (This theory was all but lost until the late 1930's, when the development of ultrasonic detectors proved that bats employed inaudible sound.) Extensive experiments with bats (Griffin) have shown that they emit bursts of ultrasonic energy, and interpret the returning echoes. The specific acoustic signals vary between species of bat. However, most bats emit a short (1 to 5 millisecond (ms)) frequency modulated pulse of sound, with initial frequencies between 40 and 80 kilohertz (kHz), dropping off 10 to 30 kHz by the end of the pulse. The rate at which pulses are repeated appears to depend on how close the bat is to an obstacle, or food insect. Experiments with blinded bats show that they can detect obstacles as small as 0.1 mm (wire). Moreover, bats show a remarkable ability to distinguish their own pulses from other ultrasonic noises present (such as pulses from other bats.) Not only can bats fly in the presence of ultrasonic backround noise, but they are also reluctantly capable of flying in the presence of a tape recorded playback of their own sounds. This indicates a highly sophisticated form of signal recognition and discrimination.

Recent studies with dolphins suggest that they too rely

heavily on sonar. Dolphins produce two characteristic noises: whistles and clicks (Kellogg). Whistles are frequency modulated pulses of nearly 0.5 second duration. The frequency of the whistle varies smoothly from 7 to 15 kHz during its duration. The clicks, on the other hand, are very short, and are repeated in trains of 5 to several hundred clicks per second. The clicks tend to contain a large spectrum of frequencies, ranging from 20 to 120 kHz. Since the velocity of sound in water is roughly 4.5 times faster than that in air, the wavelength of the whistle is comparable to the wavelength used by bats. Both animals are able to distinguish between food objects and non-food objects (fakes and obstacles) based only on returned echoes. Hence it is conjectured that dolphins use the whistle to get shape and type information, and that the clicks are used for distance and direction information. Dolphins, like bats, also exhibit a resistance to ultrasonic jamming.

There are several other animals that take limited advantage of echolocation, man included. Older ideas of "facial vision", or the sense of pressure from nearby objects, have fallen away through experiments with blind and blindfolded subjects (Griffin). Experiments with noises generated by the subject, and noises generated by nearby, uncontrolled sources, both show that people can detect the presence of reasonably small targets (6 inch discs (Dufton))

based on their echoes when there are few other distracting noises. Several blind guidance devices have been developed to capitalize on these abilities, but most require either too much attention to operate (thus making it difficult to keep track of where one is) or are unable to distinguish adequately between types of obstacles.

The production of a pulse of sound, and the detection of the returning echo is simple compared to the detailed interpretation of the component frequencies of the echo signal. Yet even the timing between the transmission and reception provide a simple measure of direction and range. More subtle details could be obtained from the echo signal using correlation methods, such as those employed in many models of pattern recognition. If a machine could imitate these limited abilities in man, it could become a useful aid to the blind. Moreover, if such a machine could present its information to a computer system designed to interact with its surroundings, then a simple form of "vision" would have been achieved. Alternatively, the information could augment a more sophisticated vision system, and increase the efficiency of image data processing by pointing out areas of interest. Also, certain visual ambiguities could be readily resolved. (Many optical illusions rely on distortions of depth and persective. With an echolocation system, depth can be determined quantitatively, eliminating the possibility of

deception.) However, a scene interpretation system based on echolocation would be useful only if it provided sufficiently detailed information to describe the scene, and if it was faster (and preferably less expensive) than current image processing vision systems.

Chapter 2: Mathematics

The following analysis of the underlying mathematics in sonar systems assumes a single point source of sound energy at a known position, and an infinitely attenuating background (i.e. free field conditions).

2.1: Sound

Sound (to the extent of its application here) is a longitudinal pressure wave that satisfies the wave equation

$$\frac{\partial^2}{\partial t^2} p = c^2 \nabla^2 p \qquad (2\text{-}1)$$

where c is the velocity of propagation in the medium, which depends on the elastic properties of the medium, and its density. (For simplicity, the medium will be modeled as an ideal, non-viscous fluid. With this assumption, c is real and constant (for a fixed temperature and pressure). The case of viscous fluids (and complex velocities) will be treated later.) The solution of particular interest is the case of a single point source radiating in an isotropic medium. Since the solutions will be spherically symmetric, the wave equation reduces to

$$\frac{\partial^2}{\partial t^2} p = c^2 \left(\frac{\partial^2}{\partial r^2} + \frac{2}{r} \frac{\partial}{\partial r} \right) p \qquad (2\text{-}2)$$

which has solutions of the form

$$p = \mathrm{Re}\frac{\underset{\sim}{P}}{r}\exp(j(\omega t - kr)) \qquad (2\text{-}3)$$

$\underset{\sim}{P}$ is the complex pressure amplitude of the wave, and describes both the magnitude and the phase of the pressure. (Only the real part of pressure is physically significant.) In this solution, $\omega = 2\pi f$, where ω and f are both referred to as the frequency of the wave (measured in radians per second and hertz, respectively). Also, $k = 2\pi/\lambda$, where λ is the wavelength, and k is the wavenumber. In terms of real quantities, this equation can be expressed as

$$p = \frac{A}{r}\sin(\omega t - kr - \gamma) \qquad (2\text{-}4)$$

For this solution, A is the peak acoustic pressure, and γ is a phase offset. The sound intensity, I, is proportional to the square of the acoustic pressure, and varies as

$$I = \frac{W}{4\pi r^2} = \frac{p^2}{2\rho_o c} \qquad (2\text{-}5)$$

where W is the radiated power.

When the sound wave reaches a large object (large with respect to λ), part of the wave is transmitted into the object, and part is reflected. For large distances (r >> the the size of the object), a spherical wave can be approximated by a plane wave. For a plane wave (normally) incident on a surface of density ρ, the intensity of the reflected wave is given by

$$I_r = I_i\left(\frac{\rho c - \rho_o c_o}{\rho c + \rho_o c_o}\right)^2 \qquad (2\text{-}6)$$

where ρ_o and c_o are the density and speed of propagation in the medium, respectively. For any material (including the medium), the value ρc (material density times speed of propagation in that material) is the characteristic acoustic impedance. If the acoustic impedance of an object is very nearly the same as that of the medium, only a very small amount of the wave energy will be reflected (the material can be thought of as "transparent"). Similarly, for materials with significantly large (or small) acoustic impedances (with respect to the medium), a large amount of the wave energy will be reflected (equivalently, such materials would be called "opaque". The difference between reflection from an object with a large acoustic impedance and reflection from an object with a small acoustic impedance is the change in phase of the returning echo. In measurements, this introduces an error of 1/2 wavelength.) Both in air and in water, most common building materials (metal, ceramic, hard plastics, bone) are fairly opaque. However, they each have a reflectivity coefficient that results in some attenuation of the reflected wave.

For large, flat, smooth surfaces, the reflected wave leaves the object at an angle equal to that of the incident wave. Corners, edges, and surface texture, on the other

hand, tend to scatter energy in all directions. Small objects also exhibit scattering. The amount of energy that is scattered depends both on the size and shape of the object. Since most objects can be modeled as a collection of appropriately connected cylinders and spheres, it is worthwhile to study their scattering properties. For a cylinder of radius b, the power (per unit length) scattered follows

$$W_s \sim 7.5(kb)^3 bI_o \qquad kb<<1 \qquad (2\text{-}7)$$

and

$$W_s \sim 4bI_o \qquad kb>>1 \qquad (2\text{-}8)$$

where I_o is the incident intensity. Since $k=2\pi/\lambda$, the conditions placed on kb are equivalent to conditions placed on the ratio of b to λ. Similarly, for a sphere of radius a, the scattered power follows

$$W_s \sim 5.6(ka)^4 I_o \qquad ka<<1 \qquad (2\text{-}9)$$

and

$$W_s \sim 2\,a^2 I_o \qquad ka>>1 \qquad (2\text{-}10)$$

(Hueter and Bolt). When the size of the object is comparable to the wavelength, the scattering equations are more complex, and exhibit a more complicated behavior (figure 2-1 (Meyer and Mayer)).

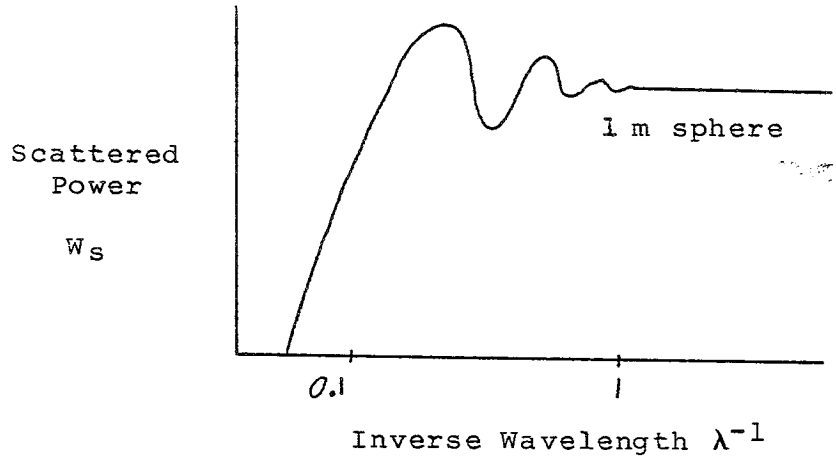

Figure 2-1: Backscattered Power from a Sphere

However, a significant qualitative observation can be made from these approximations. When the wavelength is small with respect to the size of the object, the amount of scattered power is proportional to the cross sectional area of the object, independent of small variations in wavelength. When the wavelength is large with respect to the size of the object, the scattered power decreases rapidly (λ^{-3} for cylinders, λ^{-4} for spheres) with increasing wavelength. (As will be seen later, the ability to precisely identify the mathematical details of this effect are not as valuable to sonar as the ability to simply recognize the effect.)

If the object is moving when the sound wave is scattered from it, the wavelength can undergo a Doppler shift according to

$$\frac{1}{\lambda_d} \sim \frac{1}{\lambda}\left(1+\frac{2}{c}|\vec{v}_t|\cos(\theta)\right) \qquad (2\text{-}11)$$

(Hueter and Bolt). Here, $\vec{v}_t$ is the velocity of the object,

with respect to a fixed source/receiver, as shown below (figure 2-2).

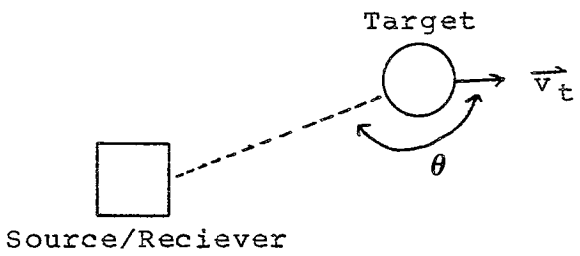

Figure 2-2: Doppler Effect - Source/Receiver and Target

The cosine term selects that component of the velocity vector that is in the direction of the source-target-receiver path. This is because the Doppler effect depends only on the relative radial velocity of the two objects.

Finally, the medium itself affects the wave. The preceding analysis was based on an ideal (non-viscous, adiabatic) fluid, which results in a real, and frequency independent value for the velocity of propagation. However, for viscous media, the velocity of propagation is more accurately described by

$$c' = c\left(1 + \omega \frac{jR}{\rho_o c^2}\right)^{1/2} \qquad (2\text{-}12)$$

(Kinsler and Frey). The physical significance of this modification is two-fold. First, the velocity of propagation is dependent on frequency (dispersion). Fortunately, in water, the effect of dispersion at frequencies below 100 MHz

is negligible, and can be ignored for most applications. The second, and non-negligible effect of a complex velocity is the absorption of wave energy. There are several mechanisms whereby energy is lost in a traveling wave (viscous loss, thermal loss, molecular energy transfer), but the net effect for a homogeneous medium can usually be modeled by

$$I = I_o * \exp(-2\alpha d) \qquad (2\text{-}13)$$

where the transmitted intensity is attenuated exponentially with respect to the total distance traveled times an "attenuation coefficient" α. Although this attenuataion effect is not negligible, the total distance traveled by the wave can be determined (below), and a corresponding gain correction can be made. (In water, it can be shown that is proportional to the square of the frequency of the wave. For pulses of varying frequency, a similar form of gain correction can be applied to remove the frequency dependent effect of absorption.)

The preceding discusion has assumed a simple point source of spherical waves. However, most common sound producing devices (speakers, whistles, and ultrasonic transducers,) are more accurately modeled as vibrating pistons. Consider the case of a circular piston (disk) vibrating along the y axis in a hole of an infinitely large baffle (rigid reflecting plane (figure 2-3)).

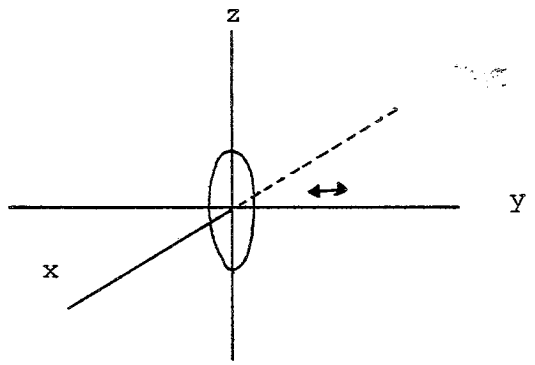

Figure 2-3: Piston Model of Transducer

If the piston has a radius of a, and is vibrating according to

$$y = Y_o * \sin(\omega t) \tag{2-14}$$

then it can be shown that

$$\underset{\sim}{P} = \left(\frac{j \rho_o c k a^2 U_o}{r} \exp(j(\omega t - kr)) \frac{J_1(ka * \sin(\theta))}{ka * \sin(\theta)} \right) \tag{2-15}$$

for r>>a (Kinsler and Frey). Here, $U_o = \omega Y_o$ is the velocity amplitude of the piston. This equation is essentially the same as eq. 2-3, the solution to the wave equation. There is, however, the additional of a directivity function

$$\frac{J_1(ka * \sin(\theta))}{ka * \sin(\theta)}$$

The function J_1 is the Bessel function of the first order. When ka is very small, $ka * \sin(\theta)$ is also small for all θ, and $J_1(ka * \sin(\theta))$ remains close to 1. However, for large ka (large compared to 1), $ka * \sin(\theta)$ varies by a large amount to produce a complicated pressure distribution. Thus if a

spherical wave source is desired, either the size of the transducer must be chosen to be small compared to the wavelength, or only the space immediately in front of the transducer (small θ) must be used for experimentation.

2.2: SONAR

SONAR is a method of using reflected sound waves to determine the size and position of target objects. A pulse of sound is transmitted from some known source, and the returning echoes are measured and timed. The number of receivers used to collect these echoes, and their relative positions, determine the capabilities of the system. For the sake of brevity, the following description of sonar systems will describe the properties characteristic of N receiver systems, for N of 1, 2, and 3. Clearly, all properties of an N-1 receiver system will also be applicable to an N receiver system.

For a single transmitter/receiver system, the radial distance from a target to the receiver can be determined directly from the amount of time required for an echo to return from the target (figure 2-4). First, assume that the distance between the transmitter and receiver is small compared to the distance between the target and the receiver ($d>>1$).

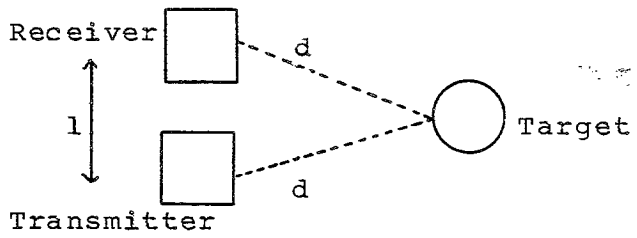

Figure 2-4: SONAR - Transmitter, Target, and Receiver

Then, if it takes t seconds for a pulse from the transmitter to bounce off the target, and be detected as an echo by the receiver,

$$d = \frac{ct}{2} \tag{2-16}$$

If the transmitter sends out a pulse which is too long, the echo will return while the transmitter is still broadcasting. It may not be possible to distinguish the echo from the transmitted pulse, either because the echo is masked by the stronger transmitted signal, or because the two signals are ambiguously similar. Hence it is important that the pulse be shorter than the round trip distance to the nearest possible target:

$$T_{pulse} < 2\,\frac{D_{min}}{c} \tag{2-17}$$

The pulse duration should not be arbitrarily short, however, since the amount of energy transmitted is proportional to the length of the pulse. If insufficient energy is transmitted, the wave will be attenuated (either by reflection, scattering, or absorption) below the level where it can be

detected accurately.

The pulse repetition period is limited by the time required for an echo to return from the most distant target:

$$T_{rep} > 2\frac{D_{max}}{c} \tag{2-18}$$

If the transmitter sends out pulses more frequently than this, it will be necessary to find a way to distinguish between old echoes and current echoes. It may be possible to ignore old echoes based solely on the strength of the signal. The strength of the returned echo (from a single target) embodies several pieces of information. First, the size of the target and its distance from the transmitter determines how much energy can be reflected:

$$I \propto I_o \frac{a}{2\pi d^2} \cdot \frac{a'}{2\pi d^2} \tag{2-19}$$

where it has been assumed that the transmitter broadcasts uniformly in a hemispherical pattern. Here, a is the area of the target, and a' is the area of the detector. This equation is only a proportionality, and should include the reflection coefficient (which depends on the material), the scattering effect (which depends on the size), and the absorption of energy by the medium (which depends on the distance traveled by the wave). Although it is important to anticipate these effects, it is not possible to calculate values for an unknown object before it has been detected.

Hence the system would have to fit parameters of a general model to several returned echoes from the same object to determine these effects if numerical values are desired.

The problem becomes more complicated when there are several objects in the field. If two objects are located at the same radial distance, their echoes can combine to form a single signal with double the expected strength. It is very difficult to separate this information in a single receiver system. Also, if two objects are located at radial distances that differ by $\lambda/2$ (or $3\lambda/2$, $5\lambda/2$, etc.), the two returning echoes will be out of phase and can cancel exactly. This can be a very serious problem if the pulse is many wavelengths long, and of constant frequency. Since the two returning echoes do not arrive at exactly the same time, a small portion of the beginning of the first echo (1/2, 3/2, 5/2... cycles) and an equal amount of the end of the second echo will not be cancelled. If the system is designed to work with very short pulses, these partial signals may be detectable. Also, if the pulse is not a wave packet of constant frequency, but rather a set of frequencies (either ordered or random), the cancellation will be minimized.

Much can be learned from investigating the frequencies used in the pulse. If the target is moving radially toward (or away from) the receiver, there will be a Doppler shift in

the frequency distribution of the echo (see eq. 2-11). If the transmitter/receiver is held fixed, the radial velocity of the target can be found from the frequency shift. (Unfortunately, this technique is more sensitive to large velocities. For studying slow moving objects, several radial positions can be measured from successive pulses and echoes, and velocity can be determined from the observed motion.) Also, if a frequency modulated pulse is used, the attenuation as a function of frequency can be studied (to a limited extent). The scattering equations (eqs. 2-7 through 2-10) suggest that the size of a target can be determined by finding the frequency dependence of the returning echoes. The amount of scattered power is independent of wavelength, for small wavelengths (small compared to the object), and falls off rapidly for large wavelengths. Hence, determining the wavelength at which the scattered power begins to show a λ^{-3} or λ^{-4} dependence will indicate the size of the object (only if enough data points can be fitted to a model (e.g. figure 2-1) with reasonable confidence). Surface textures and materials should also have frequency dependent scattering properties which can be used to identify them. The ideal frequency modulated pulse would be a chirp with an initial frequency corresponding (in wavelength) to some reasonably large size, and linearly increasing in frequency with time (see Chapter 1 and its reference to bats and dolphins). The easiest way to measure frequency dependent attenuation is to

measure points with respect to the onset of the pulse, and compare them to a model of the unattenuated pulse. In a dispersion-free medium, the frequency of the returning echo, as a function of time, should be identical to the frequency of the outgoing pulse. Unfortunately, this suggests a large amount of signal processing to correlate the pulses.

For two (stereophonic) receivers, the pulse delays give two radial distances, which constrain the target to lie on the intersection of two spheres. (Equivalently, one can measure the delay between detection in one receiver and detection in the other. This inter-receiver delay defines a hyperboloid of rotation. The average delay then defines the radial distance from the transmitter). Hence the target is known to lie on a circle of known radius and orientation, which can be seen geometrically in figure 2-5.

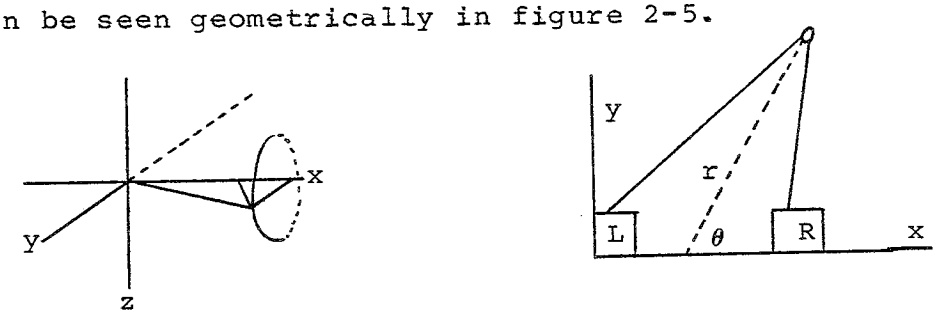

Figure 2-5: Target Geometry for Two Transducers

The circle intersects the x-y plane at

$$x = \frac{l^2 - t_R(t_R - t_L)c^2}{2l} \qquad (2\text{-}20)$$

$$y = \left(\frac{t_L^2 c^2}{4} - x^2\right)^{1/2} \qquad (2\text{-}21)$$

Or, in polar coordinates,

$$r = \frac{t_L c}{2} \qquad (2\text{-}22)$$

$$\theta = \cos^{-1}\left(\frac{l^2 - t_R(t_R - t_L)c^2}{l t_L c}\right) \qquad (2\text{-}23)$$

If the receivers can be assumed to receive signals from only one side (i.e. the forward hemisphere), the position on the circle is further limited. The problem of finding the z-coordinate still remains. One approach is to move the receivers and take multiple samples, but this is equivalent to having more than two receivers. A more elegant, but difficult approach is to draw on internal knowledge of the angular distribution of the transmitted pulse. In the earlier discussion of piston sources, it was found that the amount of radiated power varied as a function of the angle from the central axis. If the 3-dimensional position of the target is known, then it is possible to correct for the angular distribution of power, and find the normalized (with respect to a spherical source) reflected signal. If, on the other hand, the normalized signal is already known (as from previous on-axis observations of the object), a comparison of the actual signal and the normalized model will give the effect of the angular distribution. This effect defines a locus of points that satisfy eq. 2-15, which can be

intersected with the "known" circle of position. This method is fairly involved, and is only applicable to "well known" objects.

But even if this can not be done, a great deal can be learned by viewing the world as a horizontal plane, where the x-y positions of the targets are known. For relatively large flat surfaces, the orientation of the surface can be determined by adding the signal from the two receivers directly, and allowing the phase differences to produce an angular interference pattern (figure 2-6).

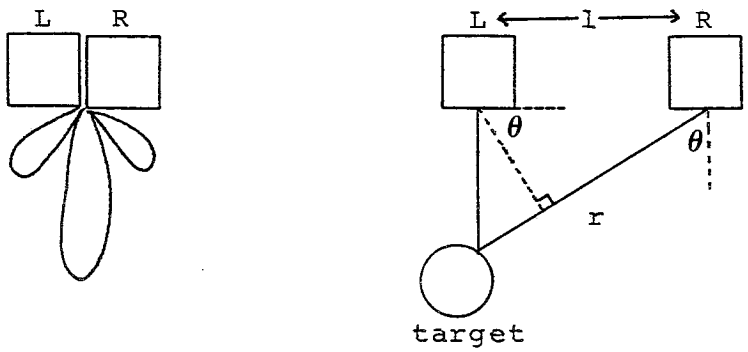

Figure 2-6: Phase Cancellation with Two Transducers

By varying the phase delay in one channel, the beam of primary node can be swept back and forth. The interference patterns satisfy

$$I \alpha \cos^2\left(\frac{l*\sin(\theta)}{\lambda}\pi\right) \qquad r>>l \qquad (2\text{-}24)$$

The delays from the secondary nodes determine the plane orientation (figure 2-7)

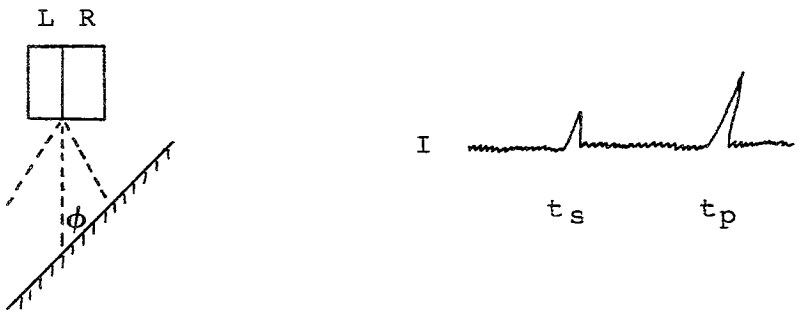

Figure 2-7: Orientation Determination from Phase Cancellation

according to

$$\sin(\phi) = \frac{\lambda t_s}{l}\left(t_s^2+t_p^2-2t_s t_p\sqrt{1-\left(\frac{\lambda}{l}\right)^2}\right)^{-1/2} \qquad (2\text{-}25)$$

It is important for this technique that the primary beam be centered on the target, and that the target be large enough to reflect at least one secondary beam. Moreover, the higher order beams must be distinguished and ignored.

Unfortunately, the problem of matching pulses together is nontrivial. Both the problems of constructive and destructive interference, as described earlier, can cause in one or both of a pair of echoes arriving simultaneously at one receiver to disappear. (As before, this problem is minimized by employing a pulse of varying frequency.) In terms of processing, the range of time which must be checked

between channels is limited by the amount of time required for sound to travel from one receiver to the other (figure 2-8).

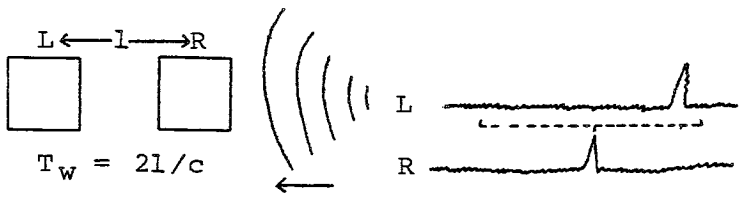

Figure 2-8: Pulse Matching Window

In order for a signal to be the counterpart to a pulse in channel A, it must arrive in the window shown for channel B. However, if two or more pulses are found in the window, and can not be uniquely identified, the system can detect "ghosts" (figure 2-9).

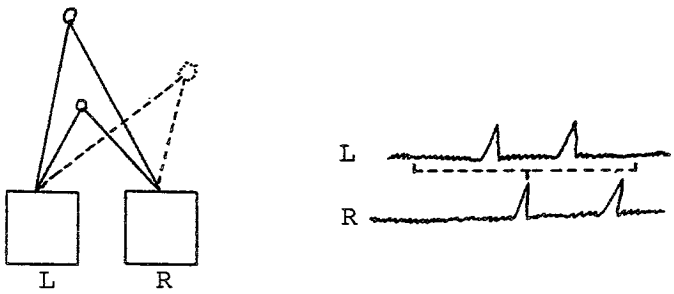

Figure 2-9: Pulse Matching Error

It may be possible to correctly match pulses according to special characteristics (e.g. pulse height, pulse width, or frequency effects) by cross correlation methods. However,

cross correlation only indicates a degree of similarity. If two identical objects are involved, the four apparent objects will be indistinguishable. Another alternative is to move the receivers and take a second measurement. Since the real objects will keep their positions, while the "ghosts" will move, the two can be separated. However, comparing measurements taken from different positions is equivalent to employing additional receivers.

For three (triphonic) receivers, the complete 3-dimensional position of a target can be found immediately by using information from the three sets of position data. Three circles can intersect (all at the same points) each other in at most 2 points, and one is eliminated immediately since it is behind the receiver. Moreover, the problem of ghosts is effectively removed, since each pair of receivers will detect a different set of ghosts. Only those points detected by all three possible pairs will correspond to real targets. However, the amount of data to be processed increases, and minor differences in the reception quality of the channels can prove to be a problem.

For more than three (polyphonic) receivers, the amount of processing required to interpret the data (in pairwise combinations) and correlate the results probably outweighs the advantages gained by the simple data system. The extreme

example of this approach is acoustic holography, which samples a large array of reception points, and then employs picture processing techniques to produce an image. An important exception, however, is the case of several receivers arranged in a carefully spaced array. If the ouputs from the receivers are passed through appropriate phase delays, the sum of the outputs appears the same as one large receiver. The use of phased arrays increases the system gain, enhances the directional characteristics of the receivers (if desired), and reduces the effects of destructive interference (phase cancellation). However, in terms of signal processing, such a system would appear to consist of a small number of physically large receivers, and could be considered as a practical implementation of one of the previously described systems.

Chapter 3: Information Processing

Scene analysis consists of four processes: transduction, feature extraction, classification, and representation. Transduction is the conversion of energy from one form to another. In this case, the energy contained in returning echoes is converted to electricity (by the piezoelectric effect), measured, and transformed into digital data. The details of this process and the selection of parameters for controlling the conversion are engineering considerations, and are not appropriate to the discussion here. (A particular system, and the consequences of certain design decisions will be studied in the next chapter.) It will be assumed, however, that the digital data contain sufficient information to permit the measurements indicated in the following paragraphs.

Once the raw data has been collected, the important characteristics of the echoes must be measured. These characteristics, or features, fall into two groups. The first group, level-0 features, consists of measurements made directly from the individual echoes. This group includes arrival time, rise time, peak pulse height, pulse width, pulse energy (area), fall time, and for frequency modulated pulses, a table of pulse height with respect to frequency. There are no hard-and-fast rules for which level-0 features

should be collected, and which can be ignored. A complete set of level-0 features would permit a simulation routine to construct the raw data purely from the measured features. (Although this is true in theory, there is little to be gained by the exercise of recreating the data.) In practice, the set of level-0 features need not be this elaborate.

The second group of features (level-1), are values that can be measured directly from sets of echoes. The most useful level-1 features for pairs of pulses from opposite channels are average arrival time, and the difference between arrival times. These times will determine the positions of the objects in the scene. (As was shown in Chapter 2, when an echo is detected in one channel, it is necessary to check the other channel only over a limited temporal window to find the counterpart. Hence a list of pairwise level-1 features will not include values for every possible interchannel echo pair, but only for those pairs that could physically refer to real objects.) Another important level-1 feature is the relative magnitude of echoes. When comparing echoes in opposite channels, this will help identify cases of constructive and destructive interference. For echo pairs in the same channel, the relative magnitudes will help identify absorption and scattering effects. The relative timings and magnitudes for planar orientation measurements also fall in this group.

The next step is to group the features into sets that correspond to the objects that generated the features. For a sonar system, it would be appropriate to form a set for each level-1 echo pair, and associate with each set a copy of the level-0 features that describe the constituent echoes. For convenience, let the vector $\vec{f}_i$ denote the i-th set of features, and let the scalar f_{ij} denote the value of the j-th feature in that set. It is necessary to classify the I objects in the scene in terms of the J features associated with each object. This classification process should produce a list of K properties for each object. (Let $\vec{p}_i$ and p_{ij} be analogously defined for these properties.)

The distinction between features and properties is not sharply defined. An important property of an object is its position. But this property can be calculated directly from the feature vector by means of the equations in Chapter 2. Moreover, the position property itself may be an important feature. Defining a feature as a quantity that can be measured directly from the data, and a property as a quantity that is of interest permits overlap between the two sets of quantities, and avoids conflicts in interpretation.

Although some of the properties are directly calculable from the features, other important properties can only be guessed at intelligently. For example, as explained earlier,

the problem of correctly matching echoes does not lend itself to an exact solution. By cross-correlating the data points of several suspect echoes, a measure of similarity can be obtained. However the final decision on echo matching still remains. For this problem, a fundamental concept of parametric classification can be applied. Consider the more general problem of a system that has S possible states. Let the a priori probability of state s be denoted by P(s). Then the conditional probability of the system being in state s, given that a specific feature vector $\vec{f}_i$ has been observed is determined by Bayes Theorem:

$$P(s|\vec{f}_i) = \frac{P(\vec{f}_i|s)*P(s)}{P(\vec{f}_i)} \qquad (3\text{-}1)$$

where

$$P(\vec{f}_i) = \sum_{s=1}^{S} P(\vec{f}_i|s)*P(s) \qquad (3\text{-}2)$$

(Duda and Hart). For matching echoes, the states are the various possible combinations of echo pairs, and the a priori probabilities P(s) are all equal. Initially, the conditional probabilities $P(\vec{f}_i|s)$ can be generated heuristically (a pair of echoes with identical features should have a much higher probability of matching than two dissimilar echoes). It is important to recognize that these conditional probabilites need not be strictly calculated and fixed, for the following reason. Let G(d|s) represent the relative gain (or loss) of making a decision d given that the system is in state s.

Then the risk function R(d) is defined by:

$$R(d) = \sum_{s=1}^{S} G(d|s) * P(s|\vec{f_i}) \qquad (3\text{-}3)$$

(Duda and Hart). According to the Bayes Decision Theorem, the best choice for the decision d is the choice that minimizes R(d). This predicted choice can be trimmed to match a desired choice by modifying the conditional probabilities $P(\vec{f_i}|s)$. These adjustments are a limited form of learning, and this method can help remove the burden of precisely determining $P(\vec{f_i}|s)$.

Another important property is the degree of similarity between the target and the elements of a set of "ideal" objects. (An important corollary to Bayes Decision Theorem is that any information about the conditional probabilities is better than none.) If the echoes could be described as "sphere-like," "cube-like," or similar to some other archetype, then unmeasurable quantities can be guessed at using Bayes Theorem. Non-parametric classification is applicable to making such distinctions. Consider the J-dimensional vector space described by the feature vectors $\vec{f_i}$. Each of these I vectors represents one point in this space. Graphically, non-parametric classification is the process of partioning this space into a group of subvolumes that are associated with each of the archetypes. Thus when a new vector $\vec{f_i}'$ is measured, the subvolume that contains it

will determine the archetype that the object matches. As an example, let the feature vector consist of only two quantities: a measure of frequency dependence in the reflection, and a normalized intensity measurement (normalized for distance to target, and for this example, normalized for size.) Since cubes have sharp edges and corners, they should produce large frequency dependent effects. Spheres, on the other hand, should have a characteristic frequency response (eqs. 2-9 and 2-10) for wavelengths comparable to the diameter of the sphere. Due to reflection, the measured intensity from a cube will vary, depending on the orientation. A sphere has no orientation. One would expect that a series of measurements taken for test cubes and spheres would produce a graph similar to figure 3-1. For this simple example, it is clear that the space defined by the feature vector can be partioned into 3 regions, such that if a new measurement was made of an unknown object ($\vec{f_i}$), the region that it fell into would define its archetype.

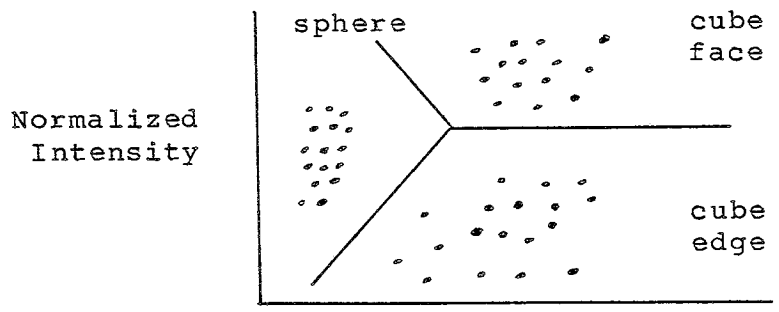

Figure 3-1: Classification Example

In the more general non-parametric classification problem, not all of the feature coordinates need to be used. An initial selection process can be used to isolate the relevant coordinates. Also, the space may not conveniently subdivide into exactly as many regions as there are archetypes. Often, the space becomes fragmented into many regions, each with an associated archetype. Occasionally, two regions with the same archetype, and undefined space between, can be combined to form one region. The intervening space then takes on the archetype of the region. This corresponds to a form of learning called generalization.

There are several approaches for defining the boundaries of the regions, and for determining the region that contains a given measured feature vector. One approach is to find the center of mass of each region by

$$F^{s}_{Ij} = \sum_{i=1}^{I} \frac{f_{ij}}{I}\Bigg|_{s} \qquad (3\text{-}4)$$

and to determine the best fit according to the minimum distance:

$$\min_{s}(\|\vec{f}_i - \vec{F}^{s}_{I}\|) \Rightarrow \text{type } s \qquad (3\text{-}5)$$

(Fu). This method tends to divide the feature space into regions separated by straight line segments, or decision boundaries, as in figure 3-1. Another appoach is to assert that a new object is most like its nearest (already

classified) neighbor:

$$\min_i(\|\vec{f}_i' - \vec{f}_i\|) \Rightarrow \text{same type as } \vec{f}_i \qquad (3\text{-}6)$$

(Fu). More complex methods, such as nonlinear regression can be applied to this problem, but the key lies in finding an appropriate set of archetypes that are both easy to distinguish and descriptive to use.

After all of the properties have been determined from the feature vectors, the information must be combined into some usable representation. There are several alternatives. For a system whose primary purpose is to produce maps of the area, a list of key points, ordered by x-y coordinates, is the most appropriate. This form would also assist a television-based system that employed sonar as a source of auxilliary information. However, if the data are being used for navigation, an x-y list is not nearly as useful as a list ordered by radial distance and direction. For a scene interpretation system, the ordering of the data with respect to position may not be as important as the order of size. Alternately, it may not be necessary to order the data at all, but merely form groups of objects with similar archetypes.

However, all of this processing remains an academic exercise unless it can be done fast enough. Current

television based scene interpretation systems are incapable of real-time processing due to the large number of points and computations that must be handled. Sonar data, on the other hand, have two distinct advantages. First, the incoming data are temporally limited. Only a relatively small and well-defined set of echo pairs need to be processed. Although the incoming data are serial in nature, the size of the window for parallel data processing is small. Television data are also serial, but there is no natural window size for parallel processing. The second important advantage to sonar data is the relative availability of information in the data. A television picture contains a large amount of potential information (512 * 512 dots per frame * 8 bits per dot * 30 frames per second = 63 million bits per second), but also requires a large amount of processing to determining the position of an object, without identifying the object. Sonar data have a much smaller amount of potential information (2 channels * 8 bits per sample * 256 samples per scan * (up to) 500 scans per second = 2 million bits per second (for the following experiment)) but position information can be extracted quite readily. The question of whether real-time processing is possible can be determined only by experiment.

Chapter 4: The Experiment

The following proposed and experimental systems are designed to demonstrate the feasibility of employing sonar for scene interpretation. Only one receiver was used in the experiment. (A higher order system could be simulated by moving the receiver, and combining successive measurements.) The experiment was performed underwater, but there is no reason why the system could not be modified to work in air.

4.1: Proposal

An ideal sonar system for microprocessor based scene interpretation would consist of two receiving transducers with a transmitting transducer mounted midway between them. As noted previously, however, an acceptable compromise is to use one of the two receivers for both transmission and reception. The incoming data are converted from analog to digital, and passed to a microprocessor for feature extraction and classification. Range data can be found immediately from the time-of-flight for the echoes. Direction data can be found by matching pulses. Size and orientation require further processing.

The range resolution of the system depends on the rate at which the incoming echoes can be sampled and converted to

digital form. If the velocity of propagation is c, and the rate of conversion is r, then the minimum resolvable difference in range is c/r. The wavelength determines the minimum size object that is readily detectable.

The angular resolution depends on the separation between the receivers. With a base line of b, the data window becomes b*r/c data points wide (see figure 2-9). If the angle calculation (eq. 2-23) were linear with respect to time difference, the 180° of arc in front of the receivers would be divided into b*r/c segments. This angle, at the maximum range of the system, gives a rough measure of the worst case transverse resolution. However, the angle is not a linear function, and is more dense around the 0° line. Hence, the spatial resolution at maximum range is better than this estimate, and of course better at all points closer. Another important angular measurement is the planar orientation of the target. If the data from the two channels are added in phase, the first antinode determines the minumum angular size required (eq. 2-24).

Keeping track of the phase in the digital representation would involve a large memory overhead. In order to compare phase differences between the channels, it could be beneficial to have three channels: right, left, and sum. Planar orientation measurements through interference patterns

can be investigated, but beam steering (by means of introducing a phase delay in one channel) would not be practical. Only objects in the center of the field would be applicable to orientation measurement.

4.2: Experiment

In order to investigate some of the properties of the above proposal, a simplified experimental system has been built. Only one transducer was used as both transmitter and receiver. The experiment was performed in a water-filled tank for the sake of noise immunity.

At 20°C, the velocity of sound in water is 1483 m/s. (Although this value is temperature dependent, the changes can be corrected for in the information processing program.) With a pulse of 1 MHz ultrasound, the wavelength is 1.483 mm. In terms of back scattered energy (eqs. 2-7 through 2-10), this implies "minimum" detectable object size of roughly 1.5 mm.

The test tank (and its dimensions) are shown below (figure 4-1). For convenience, the minimum distance to the first object was chosen as 5 cm. This implied a maximum pulse width (eq. 2-17) of 67.4 μs (or 67 cycles per pulse). A maximum distance of 25 cm implied a minimum repetition

period (eq. 2-18) of 337 μs. Within these boundaries, the spatial resolution was deterimined by the rate at which data points could be collected. This rate depended on the speed with which the analog data could be converted to a digital form (1.0 μs per 8-bit sample, using a Datel ADC-G8B converter), and the speed with which the data could be stored in memory (1.2 μs, using readily available 2102 random access memory (RAM)). These times suggested a sample rate of 1.32 μs (256 samples per scan), which resulted in a radial resolution of 0.977 mm.

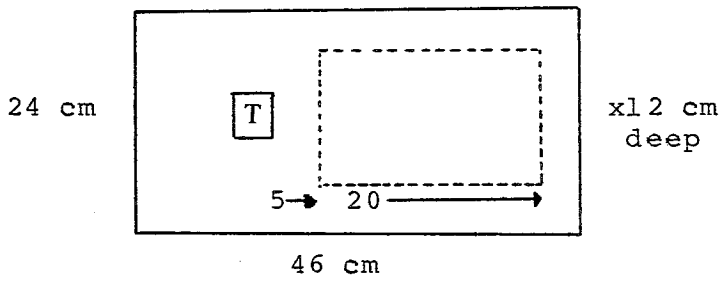

Figure 4-1: Test Tank

A critical factor in the equipment was the angular power distribution of the transducer. At 1 MHz, k=4.24 mm^{-1}. The radius of the transducer was a=11 mm. For these values, $J_1(ka*\sin(\theta))/(ka*\sin(\theta))$ has its first zero at $\theta=4.7^{\circ}$. This resulted in an effective beam width of 9.4°.

Pulses were generated by a Panametric 500 pulser receiver, and broadcast by the transducer. The returned echoes were then collected and fed to a pulse envelope

follower (a preamplifier with a high frequency cut-off below 1 MHz), and then to an analog-to-digital converter.

In principle, the digitized envelopes could be processed as soon as they are generated by the converter. However, due to the high incoming data rate, the data were buffered in a read-write memory, and later transmitted to a large computer for detailed analysis. (In particular, the data were sent to a CDC Cyber 175, by way of a PDP-11 which read the read-write memory.) Details of the hardware involved can be found in Appendix A.

Detailed frequency analysis would be prohibitively time and space consuming at this point. The envelope follower removed the explicit carrier frequency from the echo, but limited frequency effects were studied in the rising and trailing edges of the echoes. Although the rise and fall times appear as features in the processing programs (see Chapter 3) rather than as directly measured frequency dependent attenuations, the information is still present in a limited extent.

Chapter 5: Results

Data were taken for 3 different metal spheres (diameters 2.5, 2.7, and 3.2 mm) placed around the center of the beam. Several other small objects were also scanned, for comparison. After transmision to the CYBER computer, the data was processed by a feature extractor to determine, for each echo,

1) the time at which the pulse was distinguished from the background;
2) an estimated time of arrival (fitted from the data preceeding (1));
3) the time that the leading edge of the pulse began to level;
4) the exponential rise time (logarithmic fit, plus variance);
5) the time of the signal peak;
6) the value at the signal peak;
7) the time at which the signal dropped below 1/2 of the peak value;
8) an estimated 3dB point (from (7));
9) the exponential fall time (logarithmic fit, plus variance);
10) the area of the pulse from detection to the 1/2 peak point;
11) an estimate of the total area under the pulse

(energy);

12) the background noise level.

The output of the feature extractor is suitable for further processing, and would serve as input to a pulse matching algorithm if multiple channels were being processed.

The feature extraction program, EXTRCT, can be found in Appendix B. Pulse discrimination was performed using a parameter sensitive ad-hoc algorithm. A running calculation of the background noise level and standard deviation were maintained for a data window of 5 points. When the signal level rose above twice the standard deviation, a pulse was declared. A better approach would have been to apply local smoothing to the data, and then threshold the gradient to find the leading edges of the pulses. This, however, would dramatically increase the amount of required processing, and further reduce the possibility of real-time data reduction. If time is not a serious constraint, preprocessing to reduce noise and enhance the pulse could be very helpful.

Range data for the pulses can be readily calculated from either (1) or (2) above. Figure 5-1 compares the ruler measured distance and data echo time of flight (1) for the spheres. The data includes individual spheres and pairs of different sized spheres separated by distances from 5 to 0.5 cm. As can be seen, the data is reasonably linear, as

predicted by eq. 2-16. (All of the following plots are generated by the the program FPLOT, which can be found in Appendix B.)

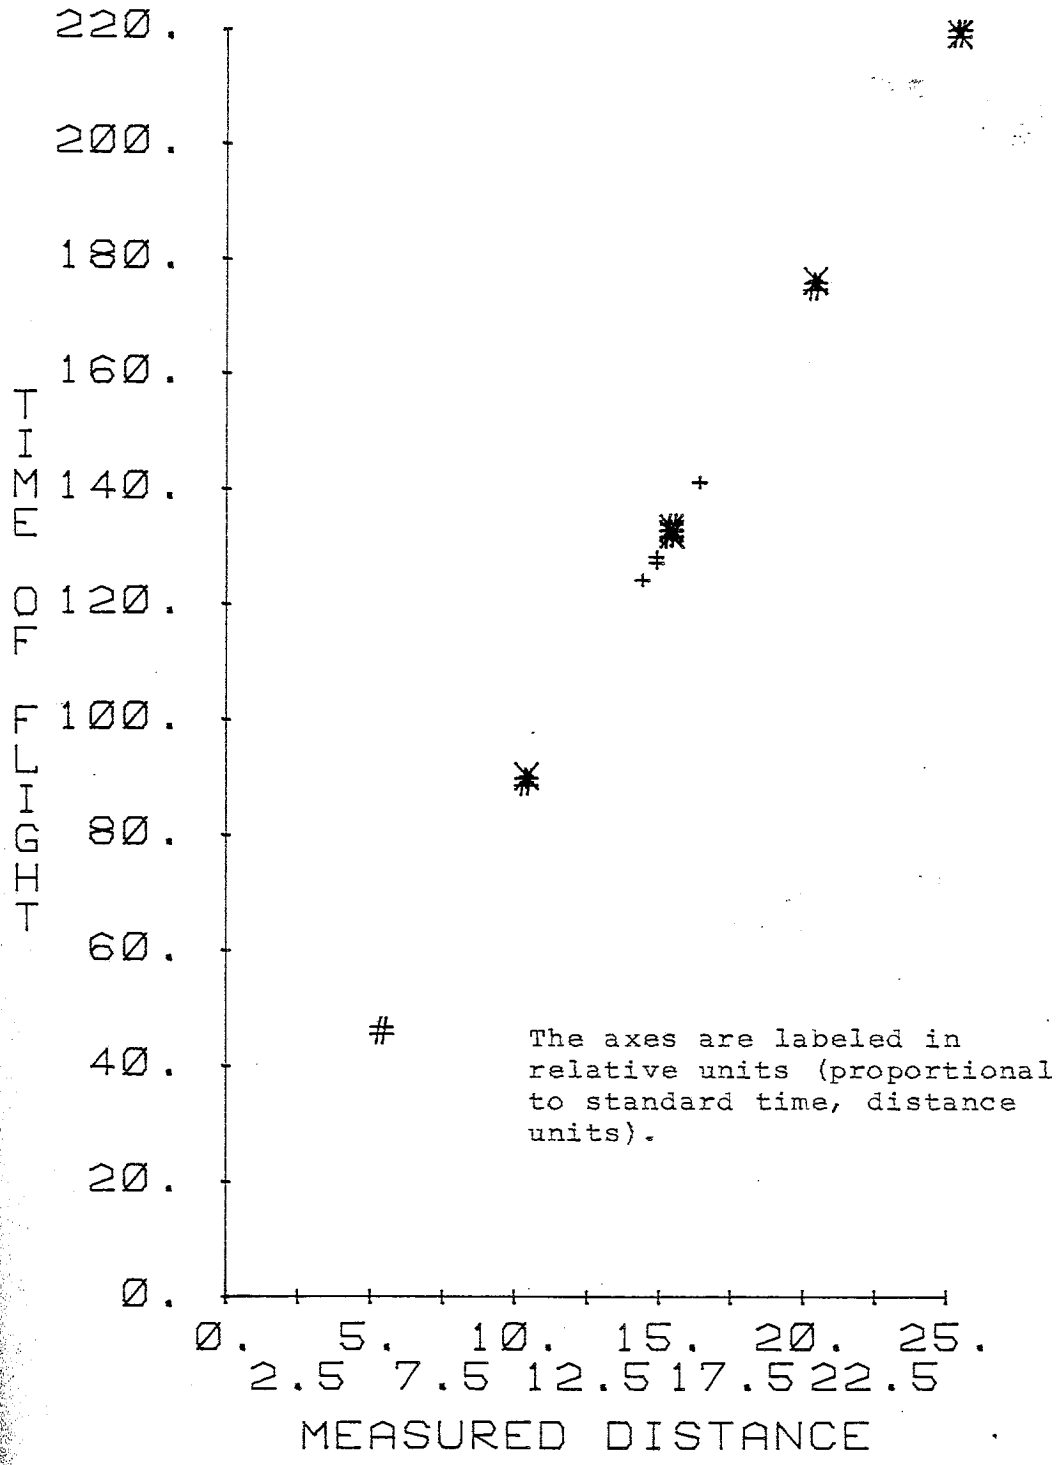

Figure 5-1: Time of Flight Measurements

Since the velocity of propagation is temperature sensitive, and the rate of data sampling may not be exactly known, the constant of proportionality relating "time" (measured by the number of samples taken up to that point) and distance (expressed in conventional units) must be determined empirically. When samples at known distances are available, this determination is staightforward. (The method of least squares can be used to fit a line to a graph similar to figure 5-1. Periodic calibration would require only a minimum amount of processing in exchange for a uniform degree of accuracy.) If known references are not available, then analytic approximation is the only alternative. For many applications it may be sufficient to work with a distance scale determined by the hardware, without regard to conventional distance units.

Target size classification appears possible from the data. Figure 5-2 shows the signal area (energy (10)) compared to the signal peak (6) for both sphere and non-sphere data. (The non-sphere objects included pieces of ceramic, metal, and metal screen, all with different densities and surface texture.) With the exception of one (possible) noise point, the sphere data cluster is a readily discernible distance from the non-sphere data. The difference in signal peak might be attributable to the fact that the non-sphere test objects were somewhat larger than

the spheres. However, the shape of the clusters for sphere and non-sphere data are dramatically different, indicating that the non-spheres indeed have (detectable) non-sphere properties.

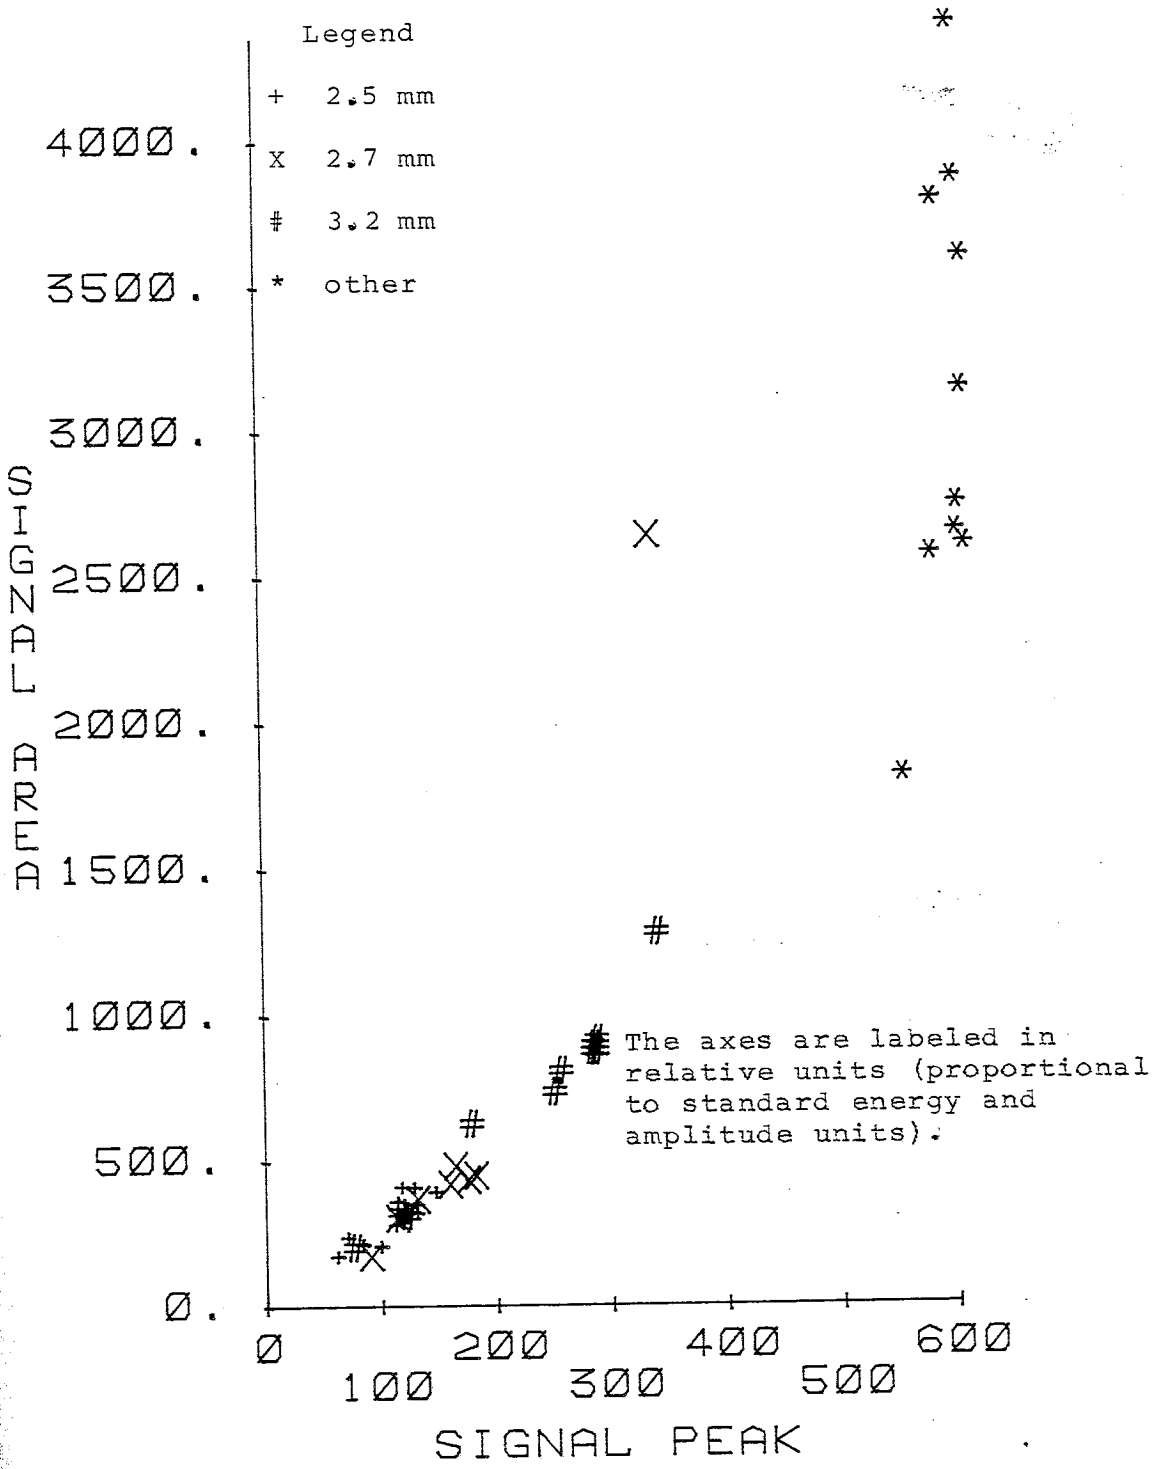

Figure 5-2: Feature Plot (all data)

Figure 5-3 shows the same features for the sphere data only. Three data point symbols are used to distinguish the different sized spheres. Again, this plot includes both individual sphere and sphere pair data. Although the clusters show more overlap, it is apparent that the plot does show some tendancy for data from each sphere to lie together. For this situation, a center of mass classification scheme (eq. 3-4) might be most appropriate.

As with distance determination, size determination would require known references in order to provide conventionally scaled numerical results. However, the general noisiness of figure 5-3 suggests that size determination would be less reproducible than size discrimination. Thus it would be easier to classify a sphere as "bigger than" rather than "x times bigger than" another sphere.

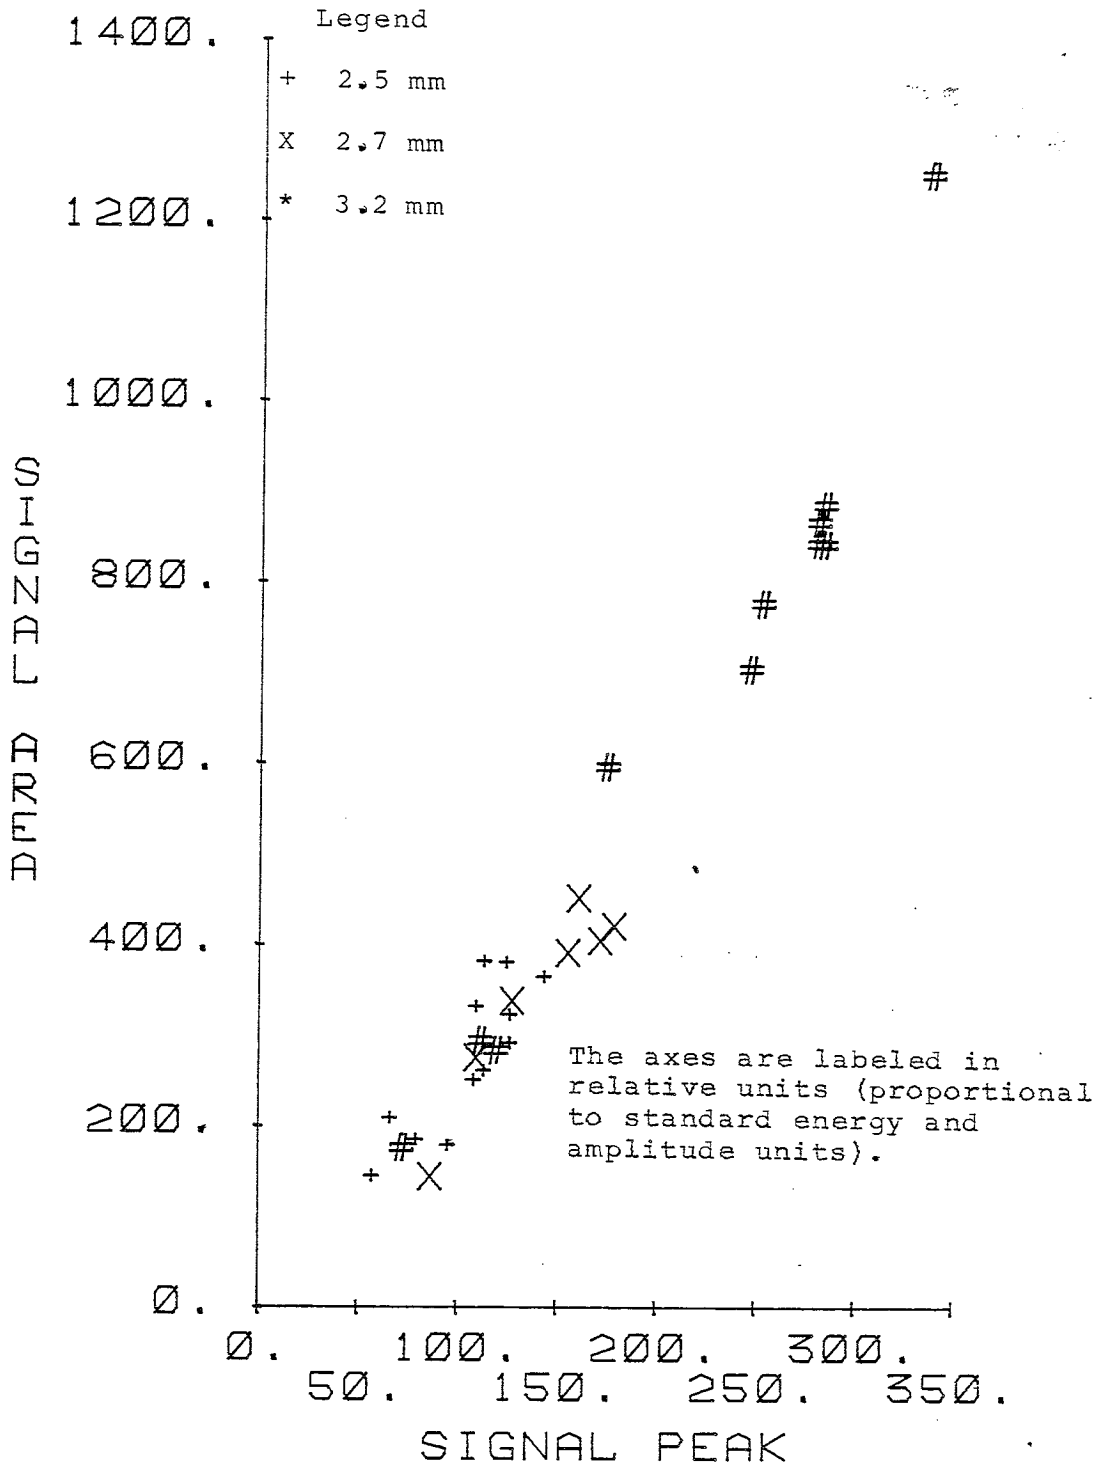

Figure 5-3: Feature Plot (spheres only)

Due to the large velocity of propagation of the signals, it was not possible to sample the pulse in sufficient detail to permit a careful study of the pulse rise time and fall time effects. The experiment produced very short rise times, which resulted in only a small number of data samples. This made accurate rise time measurement difficult or impossible. Fall times were equally troublesome. Although a larger number of data samples were available, the end of the falling edge was difficult to determine, resulting in a low accuracy for calculated fall times.

In principle, the rise time and fall time features should have provided a more sensitive measure of size and shape-complexity (smooth vs. irregular). To study these features, two major improvements must be made to the experiment. First, the sampling rate must be increased. This could be accomplished either by obtaining faster equipment (very costly), or by working in a medium with a smaller velocity of propagation (e.g. air). Second, real-time processing is required to identify and extract the pulses as they are detected in order to minimize the amount of data memory space required. (Only the pulse data need to be saved; the background can be thrown away.)

One other important feature apparent from oscilloscope measurements was the character of the top of the pulses.

Although the outgoing signal was reasonably flat-topped, the returning echoes showed varying numbers of peaks. These were virtually lost in the sampling process. Only the larger peaks were evident in the data. Pulse top information could prove valuable in multi-channel pulse matching, and in tracking an object from one data scan to the next. However, the only way to obtain this information would involve the same changes required for more detailed rise/fall time study.

Chapter 6: Conclusions

Several difficulties hampered the experiment. Although a water tank was chosen for noise immunity, the naturally smaller wavelengthes (which are characteristic of immersion tranducers) result in effectively large transducer areas, which in turn lead to a small beam width. This introduces problems associated with a non-uniform angular distribution of energy.

In order to normalize the incoming signals (to correct for attenuation), a time varying gain is required for the input preamplifier. This can be simulated in the information processing, but at the cost of data level resolution and the introduction of quantization noise.

Due to the high velocity of propagation in water, the incoming data rate is too large for real-time processing. For most conventional microprocessors, the data rate is too large for conventional program-controlled input, so that an external device must handle the buffering of data. This "direct memory access" complicates both the hardware and the programs that use it.

The difficulty of extracting a pulse envelope from the noise indicates a major problem in using sonar data. Due to

the nature of the data, a burst of noise may be indistinguishable from a real echo. This large noise sensitivity can be reduced only by filtering carefully, and combining the results of several "scans." (It should be noted that conventional imaging schemes do not suffer as badly in this regard, since a burst of noise becomes a thin horizontal line which does not match the image above or below it.) A superior approach would be to take sufficient data points to be able to detect the carrier frequency in the data, and digitally filter for that frequency. Unfortunately, this would require a prohibitively large number of data points, and an unreasonably fast data rate.

Using an air-based sonar system would overcome both the angular distribution and data rate problems, at the expense of introducing ambient room noise. The Polaroid Corporation has recently begun to market an "Ultrasonic Ranging Unit" for experimental purposes. The unit consists of an electrosatic transducer, and the appropriate driving circuitry, which includes a time varying gain amplifier. This device may open the door for "robot hobbyists" to study echolocation.

Despite these problems, the feature extractor is capable of producing data suitable for use by some other system. Range data are readily available, and limited size classification appears possible. Therefore, a sonar system

could serve as an adjunct to a conventional vision system. Although a stand-alone scene interpretation system is not practical, due to the large amount of processing required and inherently large noise sensitivity, there is no reason why obstacle avoidance and limited target recognition systems could not be developed.

Bibliography

Bruinsma, A. H. Practical Robot Circuits. Translated by E. Harker. New York: Macmillan, 1960.

Carne, E. B. Artificial Intelligence Techniques. Washington: Spartan Division, Books, Inc., 1965.

Clemenson, Gregory Duane. Simple Aspects of Echolocation. Bachelors thesis, Department of Electrical Engineering, M.I.T., 1974.

Duda, Richard O., and Hart, Peter E. Pattern Classification and Scene Analysis. New York: John Wiley and Sons, 1973.

Dufton, Richard, ed. International Conference on Devices for the Blind. London: St. Dunstans, 1967.

Fu, K. S., ed. Digital Pattern Recognition. Berlin: Springer-Verlag, 1976.

Goldman, Richard. Ultrasonic Technology. New York: Reinhold Publishing Co., 1967.

Griffin, Donald Redfield. Listening in the Dark. New Haven: Yale University Press, 1958.

Hueter, Theodor F., and Bolt, Richard H. Sonics Techniques for the Use of Sound and Ultrasound in Engineering and Science. New York: John Wiley and Sons, Inc., 1955.

Kellog, Winthrop Niles. Porpoises and Sonar. Chicago: University of Chicago Press, 1961.

Kinsler, Lawrence E., and Frey, Austin R. Fundamentals of Acoustics. New York: John Wiley and Sons, Inc., 1962.

Lin, C. C., and Segal, L. A. Mathematics Applied to Deterministic Problems in the Natural Sciences. New York: Macmillan Publishing Co., 1974.

Meyer, Daniel P., and Mayer, Herbert A. Radar Target Detection. New York: Academic Press, 1973.

Appendix A Hardware

The following drawings describe the equipment that was built to perform the experiment outlined in Chapter 4. In most cases, both the circuit diagram and the physical component layout have been included.

The block diagram (figure A-1) shows the relationships of the individual modules.

The preamplifier (figure A-2) conditions the signals from the pulser-receiver, and produces a signal envelope.

The A-D converter (figures A-3, A-4) encodes the analog data in digital form.

The controller (figures A-5, A-6) routes the data from the A-D to the memory, and handles the communication with the LSI-11 Processor.

The memory (figures A-7, A-8) is 2048 bytes by 8 bits wide, and stores the data until it can be transfered.

The bus terminator (figures A-9, A-10) provides current for the bus, and maintains idle bus lines in an inactive state.

A definition of the bus (table A-11) is included for reference.

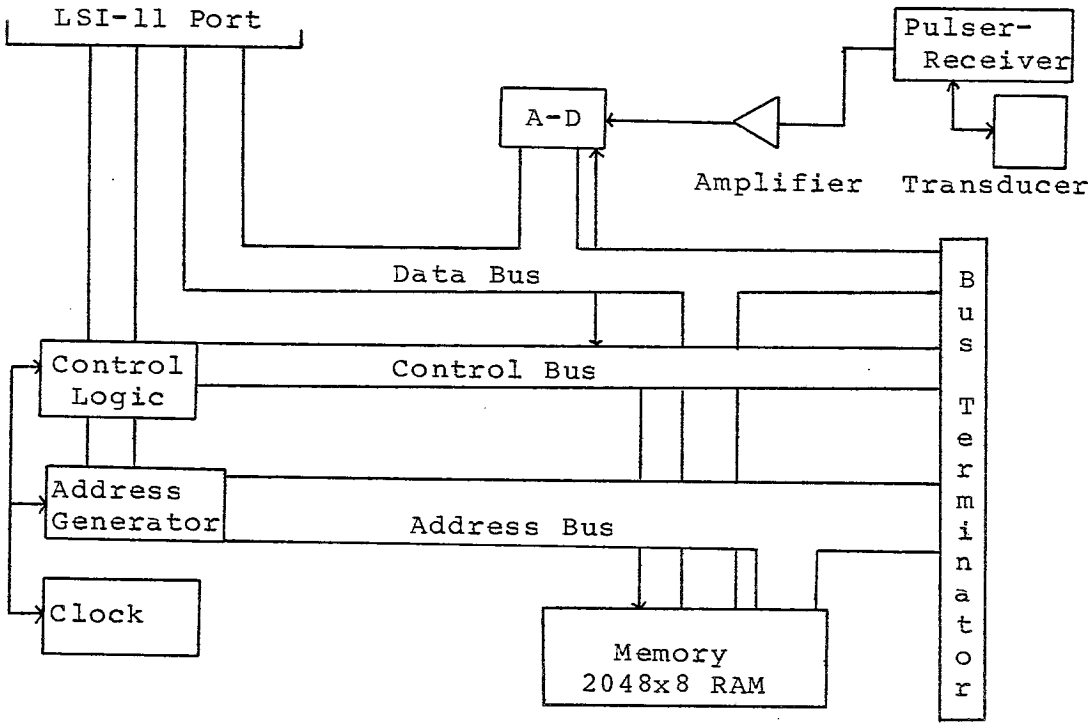

Figure A-1: Block Diagram

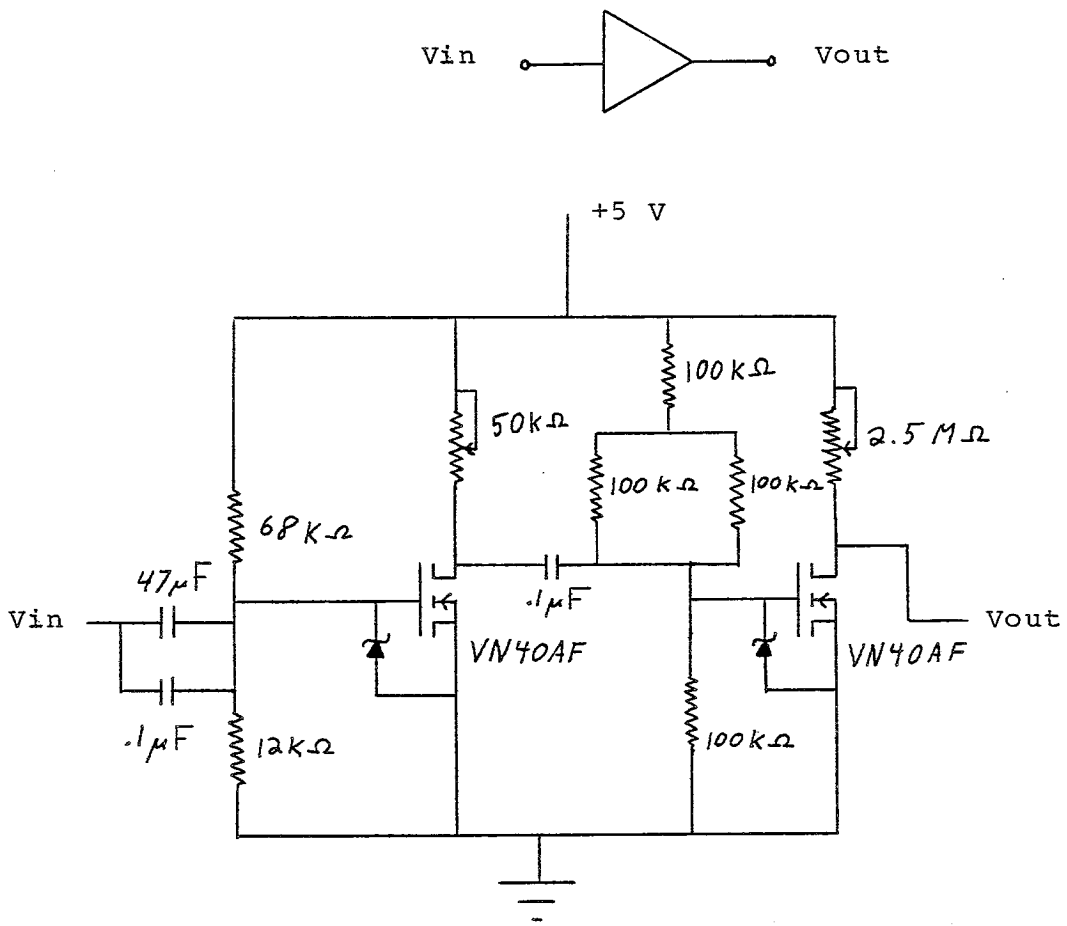

Figure A-2: Preamplifier Circuit

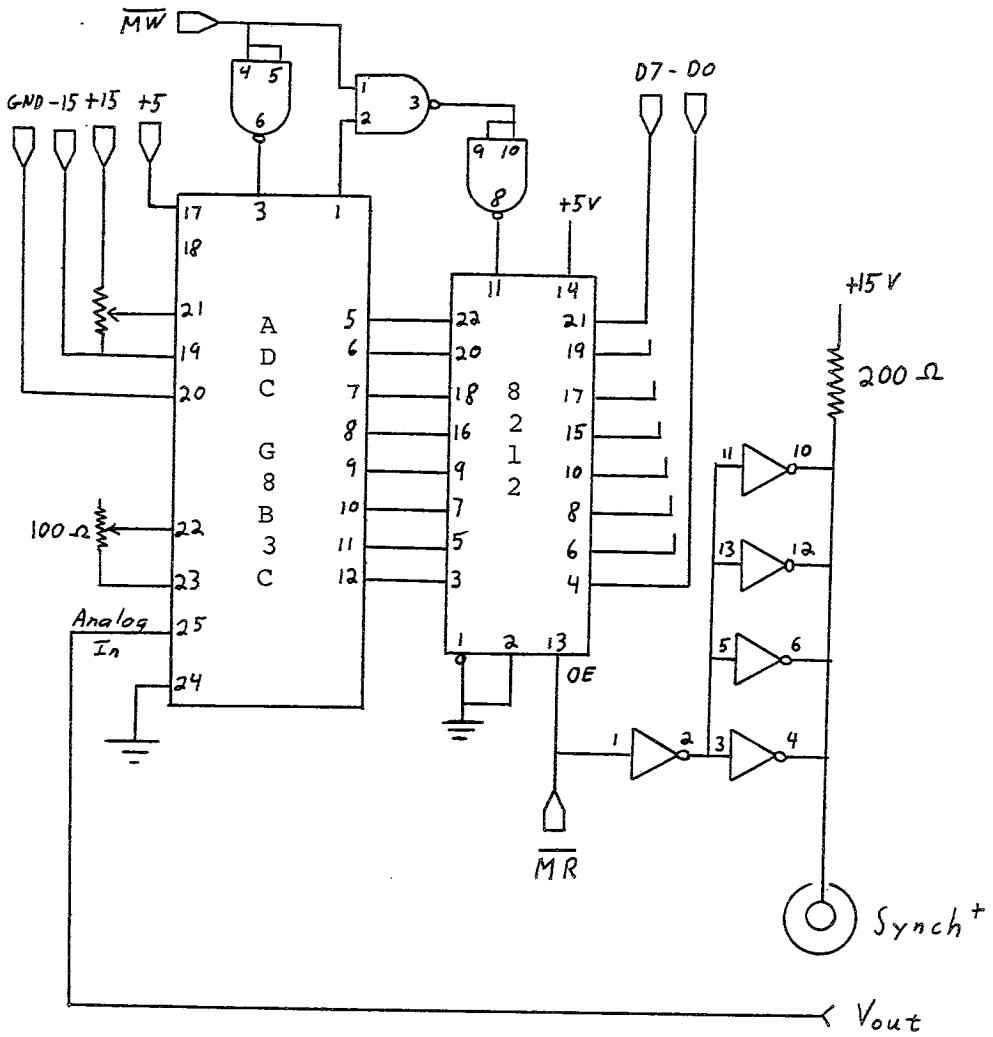

Figure A-3: A-D Logic

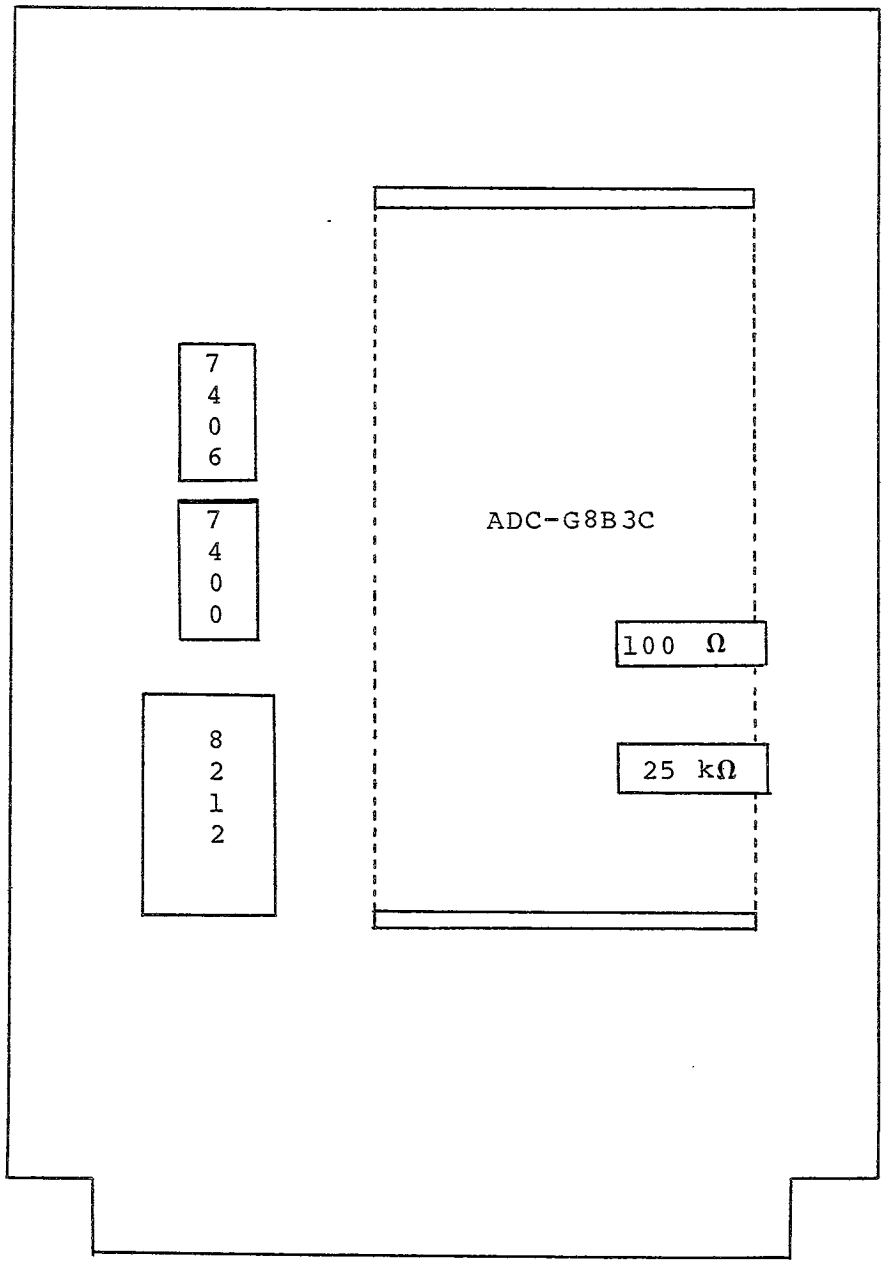

Figure A-4: A-D Card

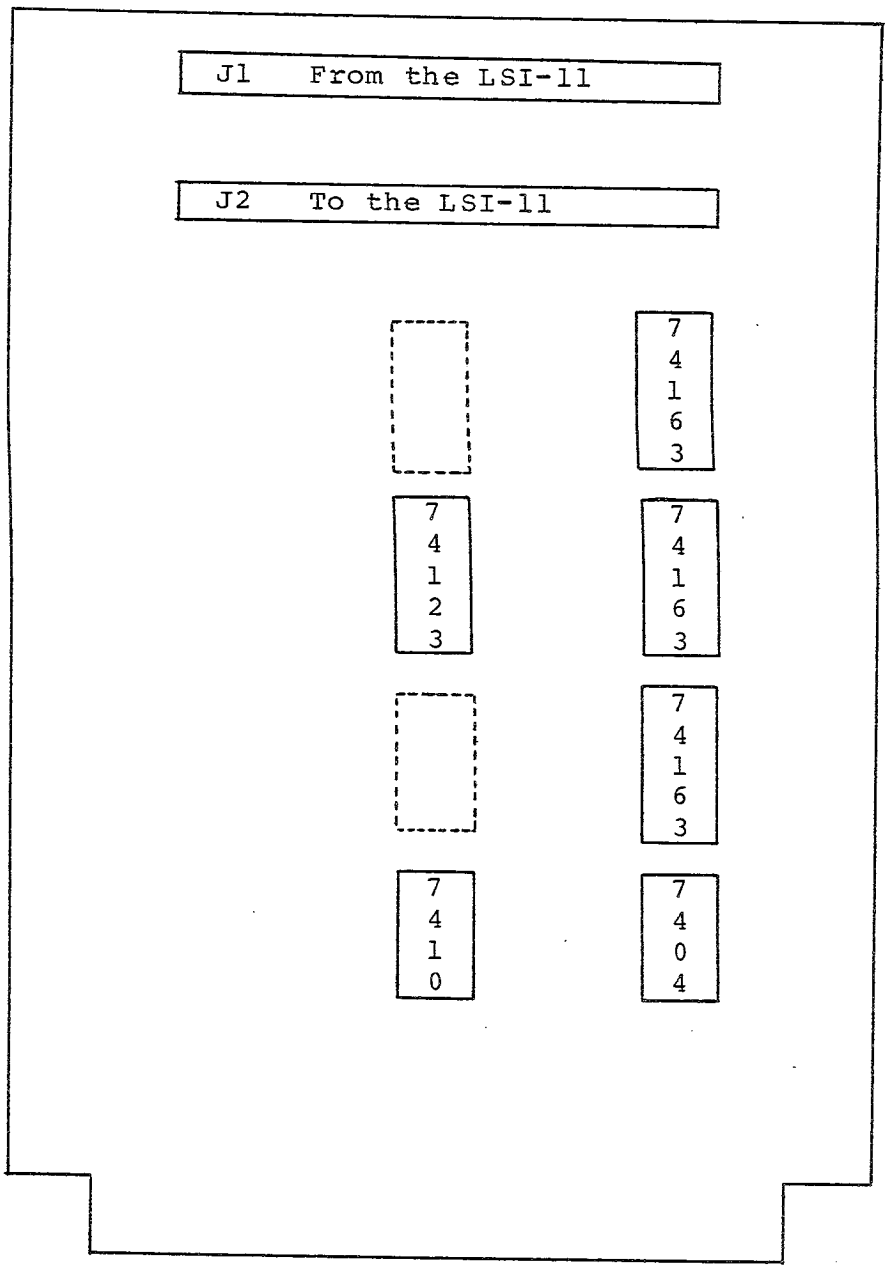

Figure A-6: Controller Card

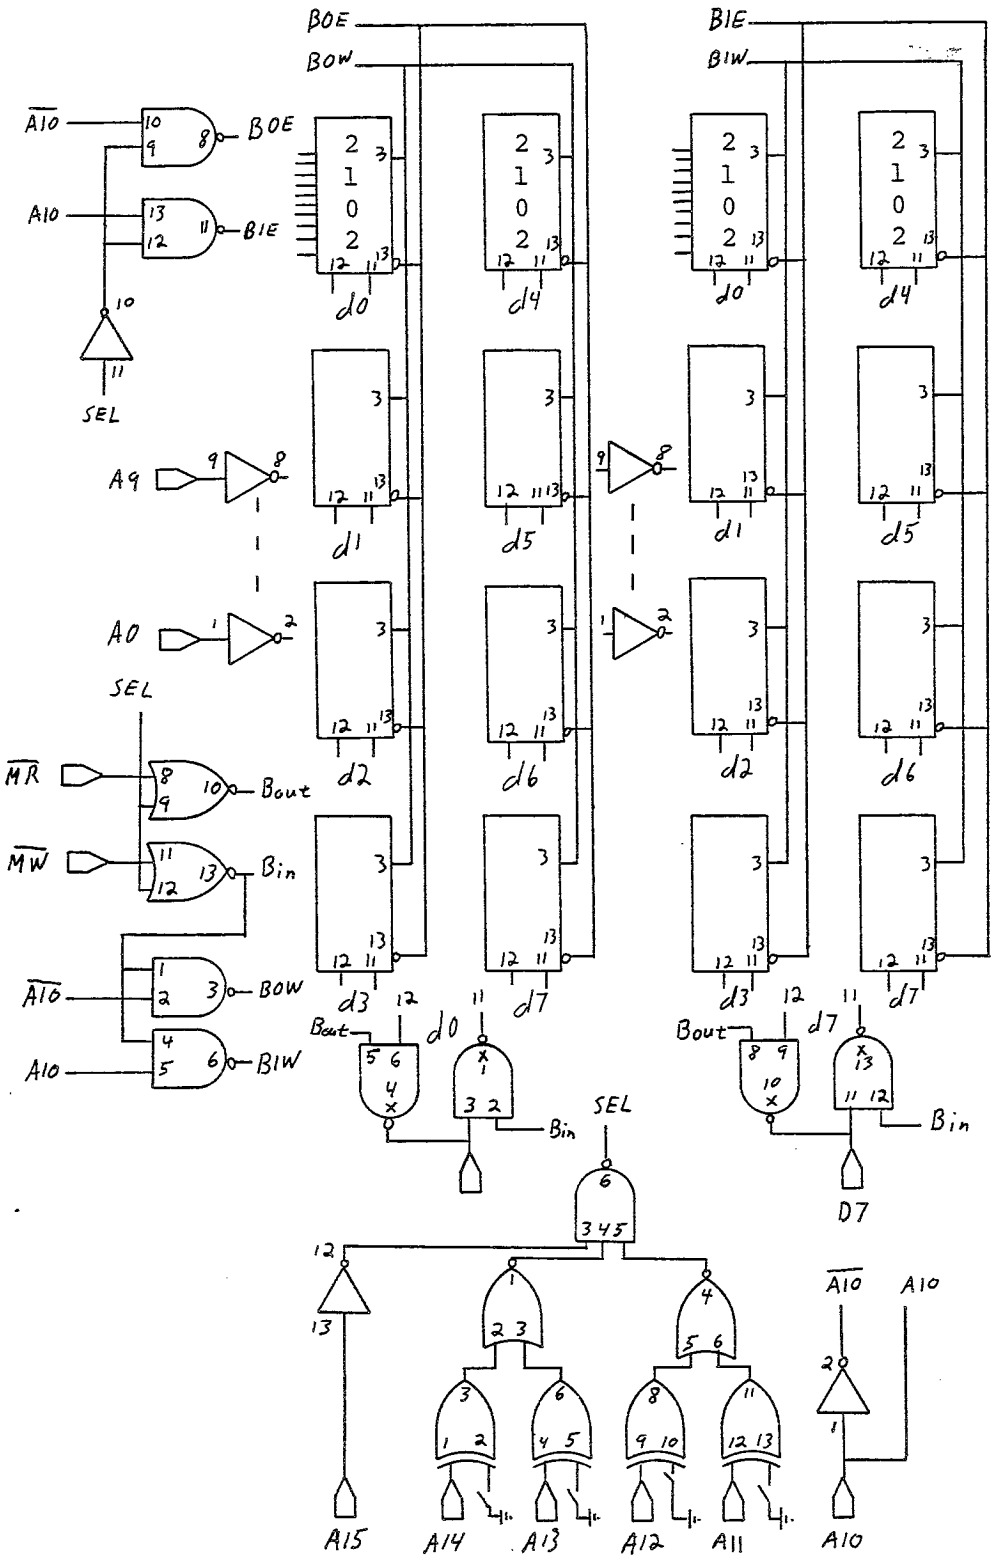

Figure A-7: Memory Logic

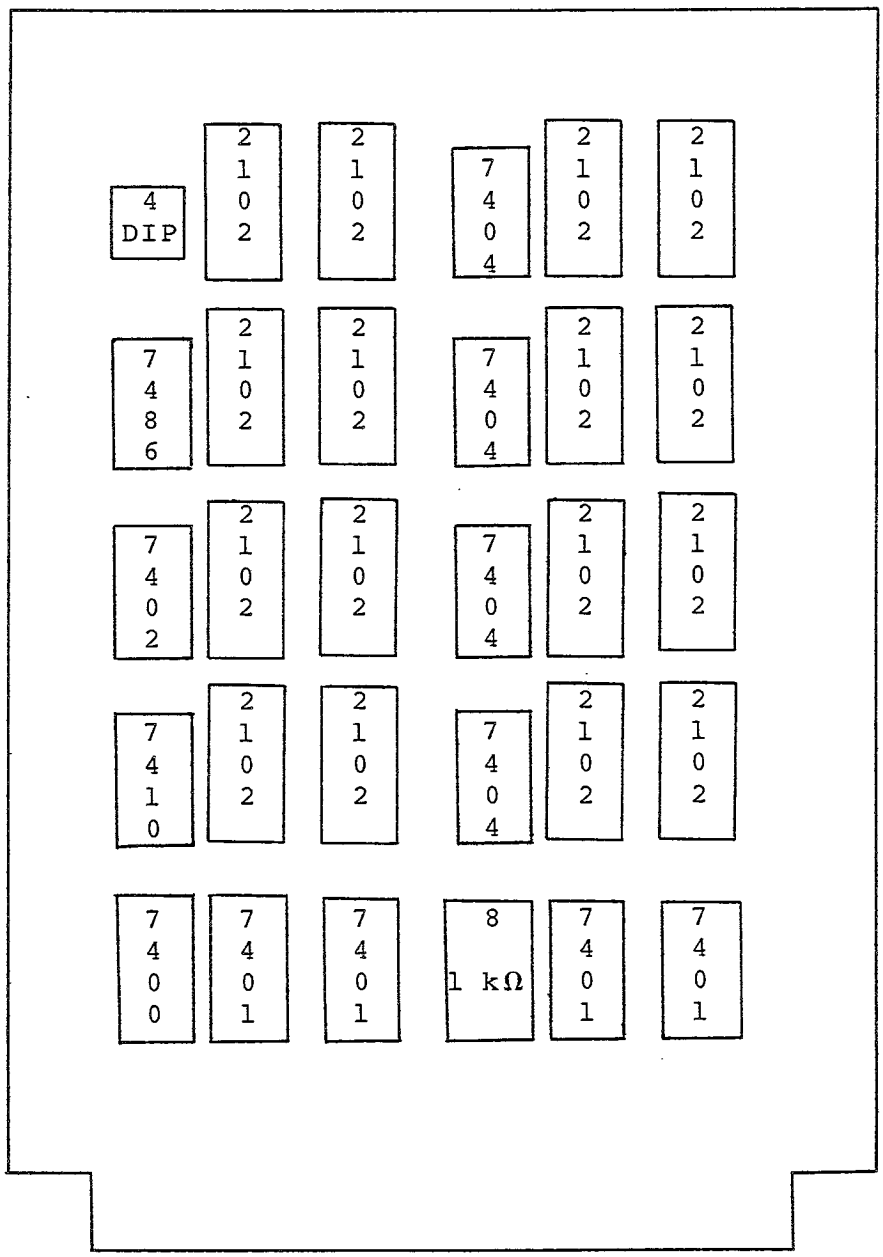

Figure A-8: Memory Card

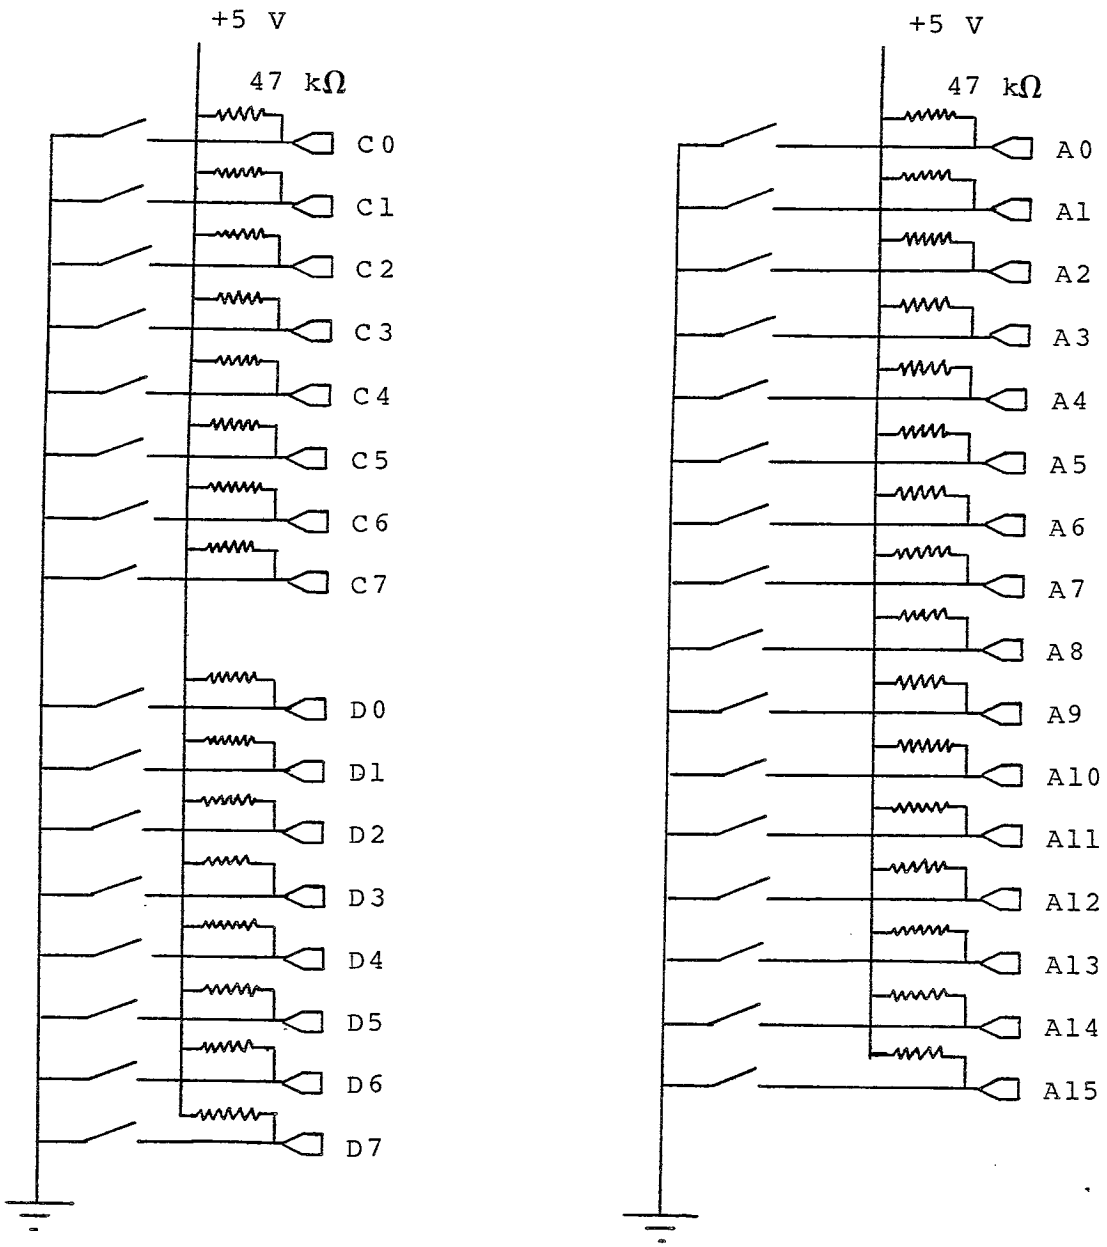

Figure A-9: Bus Terminator Circuit

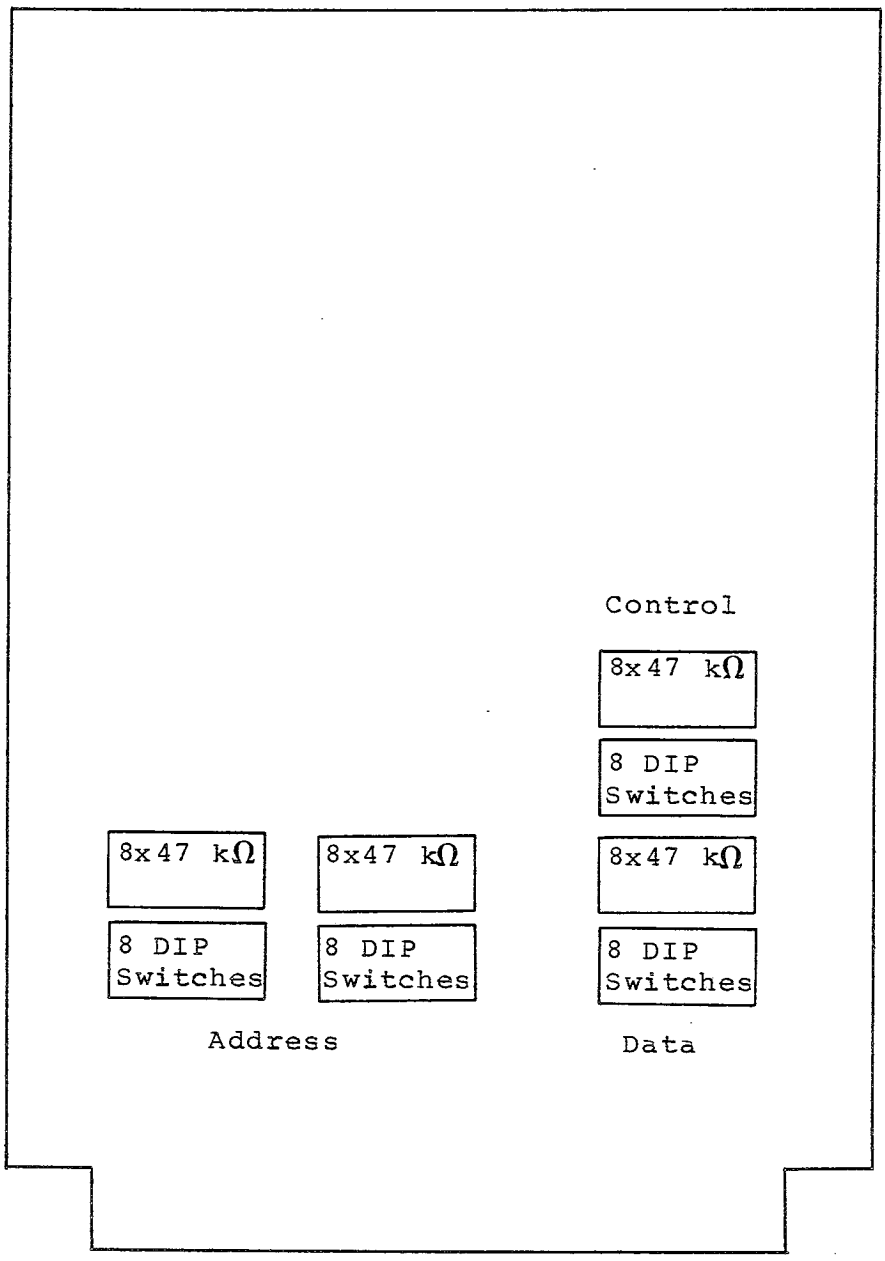

Figure A-10: Bus Terminator Card

Table A-11: Bus Definition

LINE	SIGNAL	LINE	SIGNAL
1)	IB0 - CLOCK	A)	IB3
2)	+12	B)	IB4 - RESET (L)
3)	+5	C)	A0
4)	GND	D)	A1
5)	C0 - CLOCK	E)	A2
6)	C1	F)	A3
7)	C2 - RESET (L)	H)	A4
8)	C3	J)	A5
9)	D0	K)	A6
10)	D1	L)	A7
11)	D2	M)	A8
12)	D3	N)	A9
13)	D4	P)	A10
14)	D5	R)	A11
15)	D6	S)	A12
16)	D7	T)	A13
17)	C4	U)	A14
18)	C5 - WRITE (L)	V)	A15 - I/O
19)	C6	W)	-12
20)	C7 - READ (L)	X)	-5
21)	IB1 - WRITES (L)	Y)	GND
22)	IB2 - READ (L)	Z)	IB5

Appendix B: Software

The following are the two key programs used to process and display the data from the experiment. Although several other interface programs were involved in communicating the data from the data acquisition equipment (Appendix A) to the CYBER computer, their role is of limited interest to the problem of data interpretation, and hence have not been included.

EXTRCT (figs. B-1.1 through B-1.9) is the feature extraction program, as described in Chapter 5. Some of the output from this program has been written into the program FPLOT (figs. B-2.1 through B-2.5), which produced the plots in Chapter 5. Although this information could have been passed from EXTRCT to FPLOT by way of a file, it was more expedient to include the information directly.

```
      PROGRAM EXTRCT(DATA,OUTPUT,FEATUR,TAPE5=DATA,
     -               TAPE6=OUTPUT,TAPE7=FEATUR)
C
C
      INTEGER ISEQ,OSEQ
      INTEGER SIGNAL(256),T
      INTEGER I,J,K
      INTEGER IVECT(8)
      INTEGER SBUF(5),SPTR,ST5
C
C
      REAL RVECT(7)
      REAL SVECT(3)
      REAL NOIS,NOIS2,SIG,SIG2
      REAL OSLOPE,SLOPE,TIME
      REAL TEMP
C
C
C     IVECT(
C           1) ISEQ
C           2) T OF NOTICE (SIGNAL> 2*SIGMA)
C           3) T OF SHOULDER (SLOPE(T)<SLOPE(T-1))
C           4) T OF PEAK
C           5) SIGNAL AT PEAK
```

Figure B-1.1: EXTRCT

```
C           6) T OF SIGNAL DOWN 3DB (1/2 PEAK)
C           7) SIGNAL DOWN 3DB
C           8) AREA OF PULSE FROM NOTICE TO 3DB
C
C
C     RVECT(
C           1) T OF ARRIVAL (BACK FIT OF T-NOTICE TO NOISE)
C           2) EXPONENTIAL RISE TIME CONSTANT (LOG FIT)
C           3) SIGMA**2 OF RISE TIME FIT
C           4) T OF 3DB (BACK FITTED)
C           5) EXPONENTIAL FALL TIME CONSTANT (LOG FIT)
C           6) SIGMA**2 OF FALL TIME FIT
C           7) ADJUSTED AREA OF PULSE
C
C
C     SVECT(
C           1) NOISE LEVEL
C           2) NOISE**2 LEVEL
C           3) SIGMA OF PULSE DISCRIMINATOR
C
C
C     INITIALIZATION
C
      OSEQ=0
```

Figure B-1.2: EXTRCT

```
C
C       GET NEXT SEQUENCE OF DATA
C
 100    READ (5) ISEQ
        IF (ISEQ.EQ.999) GOTO 999
        READ (5) SIGNAL
C
C       NOW PROCESS THE SEQUENCE
C
        OSEQ=ISEQ
        NOIS=0.0
        NOIS2=0.0
        SPTR=0
        DO 110 T=5,9
        SPTR=SPTR+1
        SBUF(SPTR)=SIGNAL(T)
        NOIS=NOIS+SIGNAL(T)
        NOIS2=NOIS2+SIGNAL(T)**2
 110    CONTINUE
        NOIS=NOIS/5.0
        NOIS2=NOIS2/4.0
        T=10
C
C
```

Figure B-1.3: EXTRCT

```
 200  IF (T.GT.240) GOTO 100
      SPTR=SPTR+1
      IF (SPTR.EQ.6) SPTR=1
      ST5=SBUF(SPTR)
      SBUF(SPTR)=SIGNAL(T)
      NOIS=NOIS+(SIGNAL(T)-ST5)/5.0
      NOIS2=NOIS2+(SIGNAL(T)**2-ST5**2)/4.0
      SIG2=NOIS2-(5.0/4.0)*NOIS**2
      IF (NOIS.GT.SIG2) SIG2=NOIS
      SIG=SQRT(SIG2)
      T=T+1
 205  IF (FLOAT(SIGNAL(T)).LE.(NOIS+2.0*SIG)) GOTO 200
      IF (FLOAT(SIGNAL(T)).LE.(SIGNAL(T-1)+SIG))
     -   GOTO 200
C
C
      SVECT(1)=NOIS
      SVECT(2)=NOIS2
      SVECT(3)=SIG
C
      IVECT(1)=ISEQ
      IVECT(2)=T
      IVECT(8)=SIGNAL(T)
      SLOPE=SIGNAL(T)-SIGNAL(T-1)
```

Figure B-1.4: EXTRCT

```
      TIME=T
      RVECT(1)=TIME-(SIGNAL(T)-NOIS)/SLOPE
      RVECT(2)=0.0
      RVECT(3)=0.0
      RVECT(7)=SIGNAL(T)-NOIS+ABS(RVECT(1)-T+1)
     -          *(SIGNAL(T-1)-NOIS)
      I=0
      GOTO 211
C
  210 IVECT(8)=IVECT(8)+SIGNAL(T)
      RVECT(7)=RVECT(7)+SIGNAL(T)-NOIS
  211 TEMP=ALOG(SIGNAL(T)/NOIS)/(TIME-RVECT(1))
      RVECT(2)=RVECT(2)+TEMP
      RVECT(3)=RVECT(3)+TEMP**2
      I=I+1
      OSLOPE=SLOPE
      T=T+1
      IF (T.GT.250) GOTO 100
      TIME=T
      SLOPE=SIGNAL(T)-SIGNAL(T-1)
      IF (SLOPE.LT.SIG) GOTO 212
      IF (SLOPE.GE.(OSLOPE-SIG)) GOTO 210
  212 IVECT(3)=T-1
      RVECT(2)=RVECT(2)/I
```

Figure B-1.5: EXTRCT

```
      IF (I.EQ.1) GOTO 215
      RVECT(3)=RVECT(3)/(I-1)
      GOTO 220
C
C     SPECIAL CASE OF ONLY ONE RISING EDGE DATA POINT
C
 215  RVECT(3)=-1.0
C
 220  IVECT(5)=SIGNAL(T-1)
      IVECT(4)=T-1
 225  IF (SIGNAL(T).LE.IVECT(5)) GOTO 230
      IVECT(5)=SIGNAL(T)
      IVECT(4)=T
 230  IF (FLOAT(SIGNAL(T)).LT.((IVECT(5)+NOIS)/2.0))
     -   GOTO 250
      IF (T.GT.(IVECT(3)+10)) GOTO 400
      IVECT(8)=IVECT(8)+SIGNAL(T)
      RVECT(7)=RVECT(7)+SIGNAL(T)-NOIS
      T=T+1
      IF (T.GT.250) GOTO 100
      GOTO 225
C
C     NOW WE ALLOW FOR THE POSSIBILITY IF A NEW PULSE
C
```

Figure B-1.6: EXTRCT

```
250   IVECT(6)=T
      IVECT(7)=SIGNAL(T)-NOIS
      IVECT(8)=IVECT(8)-NOIS*(T-IVECT(2))
      TIME=T
      RVECT(4)=TIME-(SIGNAL(T)-(IVECT(5)+NOIS)/2.0)
     -         /FLOAT(SIGNAL(T)-SIGNAL(T-1))
      RVECT(5)=ALOG(2.0*SIGNAL(T)/(IVECT(5)+NOIS))
     -         /(TIME-RVECT(4))
      RVECT(6)=RVECT(5)**2
      I=1
255   IF (FLOAT(SIGNAL(T)).GT.(SIGNAL(T-1)+2.0*SIG))
     -   GOTO 800
      IF (FLOAT(SIGNAL(T)).LT.(NOIS+2.0*SIG)) GOTO 800
      TIME=T
      TEMP=ALOG(2.0*SIGNAL(T)/(IVECT(5)+NOIS))
     -     /(TIME-RVECT(4))
      RVECT(5)=RVECT(5)+TEMP
      RVECT(6)=RVECT(6)+TEMP**2
      I=I+1
      RVECT(7)=RVECT(7)+SIGNAL(T)-NOIS
      T=T+1
      IF (T.GT.250) GOTO 800
      GOTO 255
C
```

Figure B-1.7: EXTRCT

```
C
C       THIS IS SOMETHING OF A PATCH. IT IS ASSUMED THAT PULSES
C       ARE NOT WIDER THAN 15 FROM SHOULDER TO 3DB. THIS
C       ASSUMPTION CLEANS UP THE PROBLEMS OF A SLOWLY GROWING
C       BASE NOISE.
C
 400    NOIS=0.0
        NOIS2=0.0
        SPTR=0
        DO 410 I=1,5
        J=6-I
        SPTR=SPTR+1
        SBUF(SPTR)=SIGNAL(T-J)
        NOIS=NOIS+SIGNAL(T-J)
        NOIS2=NOIS2+SIGNAL(T-J)**2
 410    CONTINUE
        NOIS=NOIS/5.0
        NOIS2=NOIS2/4.0
        GOTO 200
C
C       ADJUST RVECT(7) (AREA) FOR MISSING FALLING EDGE
C
 800    RVECT(5)=RVECT(5)/I
        IVECT(5)=IVECT(5)-NOIS
```

Figure B-1.8: EXTRCT

```
      RVECT(7)=RVECT(7)+IVECT(5)/ABS(2.0*RVECT(5))
      IF (I.EQ.1) GOTO 810
      RVECT(6)=RVECT(6)/(I-1)
      GOTO 900
C
 810   RVECT(6)=-1.0
C
C     NOW WRITE OUT FEATURE VECTOR
C
 900  WRITE (6,910) IVECT(1),(SVECT(I),I=1,3),
     -              (IVECT(J),J=2,8),(RVECT(K),K=1,7)
 910  FORMAT(1X,I5,3F15.5/6X,7(5X,I10)/6X,7F15.5/)
      WRITE (7) IVECT,RVECT,SVECT
      GOTO 205
C
C     END OF DATA
C
 999  IF (OSEQ.NE.51) STOP
      IVECT(1)=999
      WRITE (7) IVECT,RVECT
      STOP
      END
```

Figure B-1.9: EXTRCT

```
      PROGRAM FPLOT(INPUT,OUTPUT)
C
      INTEGER I,J,K
C
      REAL TM(12),TT(12),TP(12),TA(12)
      REAL SM(13),ST(13),SP(13),SA(13)
      REAL LM(9),LT(9),LP(9),LA(9)
      REAL NP(11),NA(11)
      REAL X(60),Y(60)
      REAL PTS(4),OPTS(4)
C
      DATA TM/ 15., 15., 15., 5., 10., 20., 25., 15.,
     -         15., 15., 15., 15./
      DATA TT/131.,131.,131., 44., 87.,173.,217.,130.,
     -        130.,130.,130.,130./
      DATA TP/260.,275., 67.,331.,169.,278.,114.,247.,
     -        241.,106.,278.,275./
      DATA TA/854.,828.,162.,1236.,582.,871.,270.,762.,
     -        690.,281.,827.,852./
C
      DATA SM/ 10., 15., 20., 25., 10., 14., 14.5, 16.,
     -         20., 10., 14.5, 16., 20./
      DATA ST/ 87.,130.,174.,218., 88.,122.,125.,139.,
     -        174., 87.,126.,139.,174./
```

Figure B-2.1: FPLOT

```
      DATA SP/ 76.,138.,119., 52., 74.,121.,108., 90.,
     -         61.,103.,121.,108.,104./
      DATA SA/233.,352.,368.,132.,172.,279.,369.,166.,
     -        196.,238.,310.,248.,319./
C
      DATA LM/  5., 10., 15., 20., 25., 15., 15., 15.,
     -         15./
      DATA LT/ 40., 88.,131.,174.,217.,131.,130.,130.,
     -        130./
      DATA LP/168.,150.,155.,122.,104., 81.,173.,166.,
     -        327./
      DATA LA/426.,378.,438.,325.,262.,131.,407.,391.,
     -        2598./
C
      DATA NP/ 576., 575., 593., 593., 589., 599., 600.,
     -         571., 598., 594., 545./
      DATA NA/2398.,3746.,3824.,2610.,4355.,3554.,2564.,
     -        2530.,3098.,2706.,1774./
C
      DATA OPTS/"N#","N+","NX","N*"/
      DATA PTS/  12., 13.,  9., 11./
C
      J=1
      DO 10 I=1,12
```

Figure B-2.2: FPLOT

```
      X(J)=TM(I)
      Y(J)=TT(I)
      J=J+1
 10   CONTINUE
      DO 20 I=1,13
      X(J)=SM(I)
      Y(J)=ST(I)
      J=J+1
 20   CONTINUE
      DO 30 I=1,9
      X(J)=LM(I)
      Y(J)=LT(I)
      J=J+1
 30   CONTINUE
C
      CALL USTART
      CALL USET("XBOTH")
      CALL UPSET("XLABEL","MEASURED DISTANCE;")
      CALL USET("YBOTH")
      CALL UPSET("YLABEL","TIME OF FLIGHT;")
      CALL UPLOT(X,Y,3.0,PTS,OPTS)
      CALL UPAUSE
C
      J=1
```

Figure B-2.3: FPLOT

```
      DO 40 I=1,12
      X(J)=TP(I)
      Y(J)=TA(I)
      J=J+1
 40   CONTINUE
      DO 50 I=1,13
      X(J)=SP(I)
      Y(J)=SA(I)
      J=J+1
 50   CONTINUE
      DO 60 I=1,9
      X(J)=LP(I)
      Y(J)=LA(I)
      J=J+1
 60   CONTINUE
      DO 70 I=1,11
      X(J)=NP(I)
      Y(J)=NA(I)
      J=J+1
 70   CONTINUE
C
      CALL UERASE
      CALL URESET
      CALL USET("XBOTH")
```

Figure B-2.4: FPLOT

```
      CALL UPSET("XLABEL","SIGNAL PEAK;")
      CALL USET("YBOTH")
      CALL UPSET("YLABEL","SIGNAL AREA;")
      CALL UPLOT(X,Y,4.0,PTS,OPTS)
      CALL UPAUSE
C
      PTS(3)=PTS(3)-1.0
      CALL UERASE
      CALL URESET
      CALL USET("XBOTH")
      CALL UPSET("XLABEL","SIGNAL PEAK;")
      CALL USET("YBOTH")
      CALL UPSET("YLABEL","SIGNAL AREA;")
      CALL UPLOT(X,Y,3.0,PTS,OPTS)
      CALL UPAUSE
C
      CALL UEND
      STOP
      END
```

Figure B-2.5: FPLOT