EGM 4001

ENGINEERING DESIGN II

Implementation of Sensor and Control Designs for Bioregenerative Systems

Prepared for

National Aeronautics and Space Administration
Kennedy Space Center, Florida

and

Universities Space Research Association

May 1990

Prepared by

EGM 4001 Engineering Design
Department of Aerospace Engineering,
Mechanics and Engineering Science

University of Florida
Gainesville, Florida 32611
(904) 392-0961

Instructor

Dr. Gale E. Nevill, Jr.

Editor

Pedro R. Rodriguez

Assistant Editors

Leslie Alnwick
Richard Kern
Mark Marchan
Robert Schneider
Andrew Speicher

(NASA-CR-186552) IMPLEMENTATION OF SENSOR AND CONTROL DESIGNS FOR BIOREGENERATIVE SYSTEMS Final Report (Florida Univ.) 118 p
CSCL 06?
N90-25479
Unclas
G3/51 0289203

EXECUTIVE SUMMARY

The goal of the Spring 1990 EGM 4001 Design class was to design, fabricate, and test sensors and control systems for a closed-loop life support system (CLLSS). The designs investigated were to contribute to the development of NASA's Controlled Ecological Life Support System (CELSS) at Kennedy Space Center (KSC). Designs included a seed moisture content sensor, a porous medium wetness sensor, a plant health sensor, and a neural network control system.

The seed group focused on the design and implementation of a sensor that could detect the moisture content of a seed batch.

The porous medium wetness group concentrated on the development of a sensor to monitor the amount of nutrient solution within a porous plate incorporating either infrared reflectance or thermal conductance properties.

The plant health group examined the possibility of remotely monitoring the health of the plants within the Biomass Production Chamber (BPC) using infrared reflectance properties.

Finally, the neural network group concentrated on the ability to use parallel processing in order to control a robot arm and analyze the data from the health sensor to detect unhealthy regions of a plant.

The EGM 4001 class feels NASA will benefit from this cooperative venture. NASA received the interest and enthusiasm of engineering students. Recommendations provided will benefit future study in these areas.

ACKNOWLEDGEMENTS

The members of the EGM 4001 Design class appreciate the cooperation of the National Aeronautic and Space Administration and the Bionetics Corporation, particularly the following persons:

Dr. William Knott

Dr. John Sager

Mr. Ralph Prince

Dr. Ray Wheeler

Dr. Richard Strayer

Mr. Thomas Dreschel

In addition, the supporting grant from the Universities Space Research Association is appreciated.

The contributions of the following faculty of the Department of Aerospace Engineering, Mechanics, and Engineering Science at the University of Florida was very beneficial in research performed by the Spring Design class:

Dr. Harold Doddington

Dr. Robert Hirko

Dr. David Jenkins

Dr. C.E. Taylor

The following faculty at the University of Florida were also of great assistance:

Dr. Roy Harrell, Department of Agricultural Engineering

Dr. Jose Principe, Department of Electrical Engineering

Dr. Tom Shih, Department of Mechanical Engineering

Mr. Art Taylor, Department of Agricultural Engineering

Dr. Fedro Zazueta, Department of Agricultural Engineering

The assistance of the following graduate students associated with the University of Florida is also greatly appreciated:

Larry Carr, Physics Department

Wilhelm Schwab, Department of Aerospace Engineering, Mechanics, and Engineering Science

The guidance of the following assistants and consultants at the University of Florida is also appreciated:

Ken Anderson

Ara Manukian

Kent Tambling

We also like to extend a special thanks to the following individuals:

Mr. Don Vinton, Department of Aerospace Engineering, Mechanics, and Engineering Science

Dell Optics, Fairview, N.J.

Microcoatings, Westford, Mass.

Omega Company, Stamford, CT

Finally, special thanks to Dr. Gale E. Nevill, Jr., for his support and guidance throughout the semester.

1990 EGM 4001 DESIGN CLASS

Leslie Alnwick
Lisa Bean
Judd Bishop
Amy Clark
Gavin Clark
Alan Clayton
Patricia Debs
Barry Finger
Tom Good
Richard Kern
Don Koedam
Steve Kropf
Mark Marchan
Walt Marchand
Michael Medley
Pedro Rodriguez
Robert Schneider
Parker Severson
Andrew Speicher

INTRODUCTION

The EGM 4001 Design class has been working in conjunction with the National Aeronautics and Space Administration (NASA) and the Controlled Ecological Life Support System (CELSS) project, supported by a grant from the Universities Space Research Association (USRA). The research being done at the CELSS facility has focused on the development of a closed-loop environment capable of sustaining plants and humans for a long-term space mission. The Spring 1990 Design class has concentrated on constructing and testing sensors and controllers.

Goals

After conducting studies in sensor systems in the Fall 1989 semester, the class determined which areas of the life support system have a strong need for control and sensing. The areas of interest included - moisture sensing, porous medium wetness sensing, plant health sensing, and controlling with a neural network. The overall goal was to develop a working prototype of the CELLS environment.

Class Organization

The students of the class were divided into the four sensing and control areas. Each group then pursued their respective topics in order to accomplish the class goal.

Report Structure

The remainder of this report is divided into four sections comprised of the four final group reports. Within each report are conclusions and recommendations for future research.

TABLE OF CONTENTS

I. DETERMINATION OF SEED MOISTURE CONTENT USING INFRARED DIFFUSE REFLECTANCE

Prepared By:

Leslie Alnwick
Patricia Debs
Barry Finger
Steve Kropf
Michael Medley

TABLE OF CONTENTS

SUMMARY

Last semester, the Plant Propagation group identified several sensing needs for soybean crops in a closed-loop life support system (CLLSS). Seed processing after harvest is an important step in ensuring the growth and high yield of future crops. As a result of their investigations, the Seed Group was formed to focus on the design and implementation of a sensor that could detect the moisture content of a seed batch. The technique employed is based on the property of water to absorb certain wavelengths in the infrared spectrum; however, the measured quantity of the system is infrared reflectance. The system is constructed to be versatile, allowing the moisture content of several mediums to be determined.

The initial testing of the moisture content sensor involved the optimization of parameters such as light intensity, bias voltage, and power transmission. The next phase of testing focused on detecting the infrared reflectance of seed batches of known moisture contents, and analyzing the mathematical data. As a final test, seeds of unknown moisture content were evaluated using the sensor, and validation of its performance was made using a calibrated digital moisture meter. Although the original goal of determining seed viability has not yet been proven, the members of the Seed Group feel that measuring the moisture content of freshly harvested seeds for extremes in percent moisture can indicate the germination potential of a seed batch when used in conjunction with another method (seed density, GADA, TZ test). Nevertheless, the present moisture content design successfully differentiates seed samples according to percent moisture.

INTRODUCTION

In any long term space mission incorporating a plant growth unit as a source of food, the importance of seed processing cannot be overlooked. It is unlikely that the total mission requirements for seeds can be satisfied with a stored supply; therefore, seed harvesting, selection, drying, and storage methods must be adequately developed and controlled.

The primary objective of the Seed Group was to develop a remote sensing technique that, when fully automated, could provide a means to determine the moisture content of a representative sample of seeds. Measuring the moisture content in seeds may be used to determine harvest time of seeds, duration of seed drying, or perhaps even seed selection. The Seed Group designed a moisture content sensor that employed infrared reflectance of a seed sample based on the ability of water to absorb certain wavelengths of the infrared spectrum.

This report gives a brief background on infrared reflectance, followed by a physical description of the Seed Group's moisture content sensor. Further detail is included on the filters, detector, and lenses used in the configuration. Specifics of the procedures used to optimize the system and test seed samples are outlined. Also included are recommendations for future developments.

BACKGROUND

Sensing Moisture Content

Moisture content of soybean seeds was chosen as the property that could best be sensed for three reasons: 1) the test is non-destructive to the seeds, 2) the process of varying moisture content in the seeds for testing purposes is a simple procedure, and 3) the method is easily automated. A moisture content sensor is important in a closed-loop life support system for several applications. If the device could be constructed to sense the moisture of seeds or seed pod while still attached to the plant, the optimum harvest time may be determined. If not used to determine harvest time, a moisture content sensor could be used to determine whether a seed batch is retained for storage and replanting, or sent to food processing, based on extremes in moisture content. Seeds too high in moisture may cause problems in storage as they are conducive to growth of mold and fungi and are more susceptible to mechanical damage during harvesting and handling. Seed batches of low percent moisture may have dried on the plant too long, reducing their vigor and storage life. The measurement of seed moisture content may also be instrumental in controlling the drying times of seed batches before storage (Figure I-1).

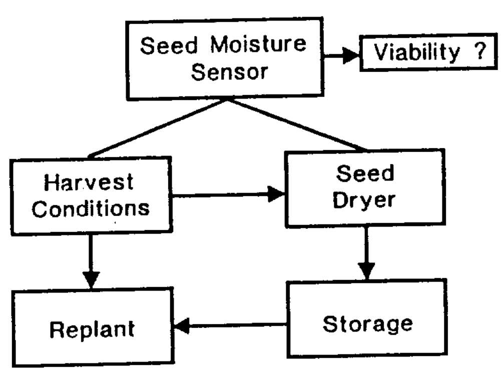

Figure I-1. Moisture content sensor integration.

Infrared Reflectance

The concept of using infrared reflectance as a means to determine the moisture content of a seed batch is a technique based on the property of water to absorb certain wavelengths in the infrared spectrum. The band most widely used in on-line moisture determination is the 1.94 micron band within the infrared spectrum (Carr-Brion, 1986). Because light rays that interact with a rough surface are scattered, the best method of sensing absorption properties is by diffuse reflectance.

In the design of this sensor system, two filters must be used--one that passes a water absorption band, and another that passes a reference wavelength. The passed wavelengths interact with the surface of the seed sample producing a reflectance inversely proportional to the amount of infrared energy absorbed. A solid with a high moisture content will absorb more infrared energy in the water absorbing band, consequently reflecting less. Collecting the reflected energy and focusing it on a detector produces a corresponding voltage output. The signal produced by the detector for each filter is recorded, and the ratio of the two reflectances is calculated. Because a reference filter was used for each measurement, the results are relatively independent of sample positioning, temperature, and other external effects.

The infrared moisture content sensor's design was based on specifications in Moisture Sensors in Process Control. The two filters chosen were a 1.8 micron reference filter and a 1.94 micron bandpass filter. The detector type is lead sulfide, which is capable of detecting both wavelengths of interest for the system. Using infrared reflectance techniques to measure moisture content in solids is a feasible method for implementation in a closed-loop automated system (Carr-Brion, 1986).

Seed Conditioning

Seed batches of 250 grams were prepared for testing by two methods. Oven drying, microwaving, and sun baking created seed samples of relatively low moisture content (approximately 6%). To obtain high moisture contents, seed batches were soaked in liquid water for different time intervals to achieve samples ranging from 20-35% moisture. The seeds were placed in sealed plastic bags and refrigerated to allow time for the moisture contents to reach equilibrium. The moisture content of each 250 gram batch was then measured by a Burrows Model 700 digital moisture meter and the values were used as the basis for test validation.

PHYSICAL COMPONENTS

The moisture content sensor (MCS) uses infrared diffuse reflectance to non-destructively measure the percent moisture in a seed sample. The measurement of the amount of light (a specific wavelength) reflected from the seed sample relates directly to the water concentration in the sample. The ratio of the reference filter voltage to the selective filter voltage provides a correlation to the percent moisture.

The moisture content sensor designed to accomplish this task is comprised of four main components (Figure I-2):

1. Infrared source
2. Optical configuration
3. Detector and circuit
4. Computer interface

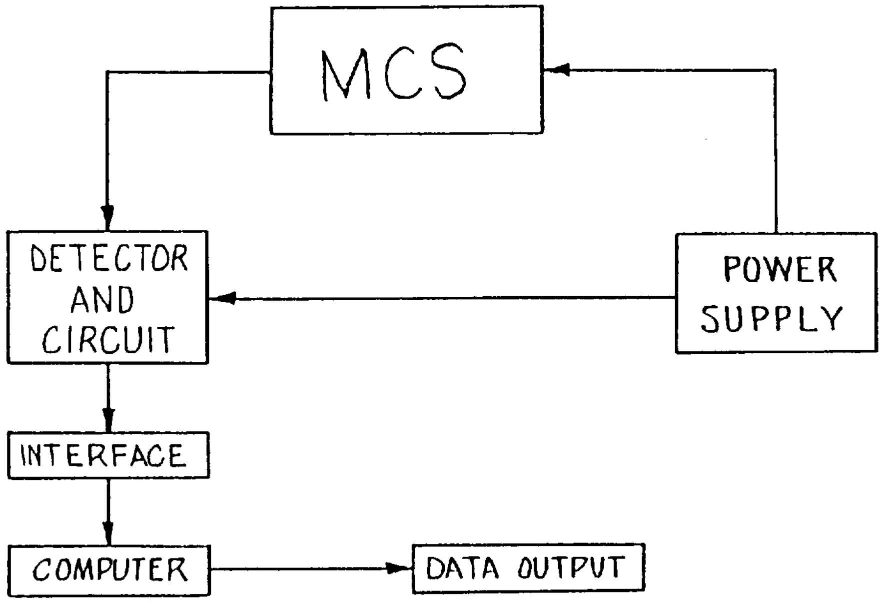

Figure I-2. MCS system block diagram.

Infrared Source

The infrared light source used in this design is a small quartz-halogen projector bulb (500 Watts). The bulb intensity is controlled by a dimmer switch and a parabolic reflector directs the light rays toward the sample. The intended design approximates a point source by allowing light to exit the housing through a small hole, producing a concentrated light beam. A cooling fan was added to dissipate excess heat, which may damage the filters in the optical system.

Optical Configuration

The optical configuration of the MCS consists of four basic units (Figure I-3)

1. Plano-convex collimating lens
2. High pass filter
3. Narrow bandpass filters with cartridge
4. Plano-convex focusing lens

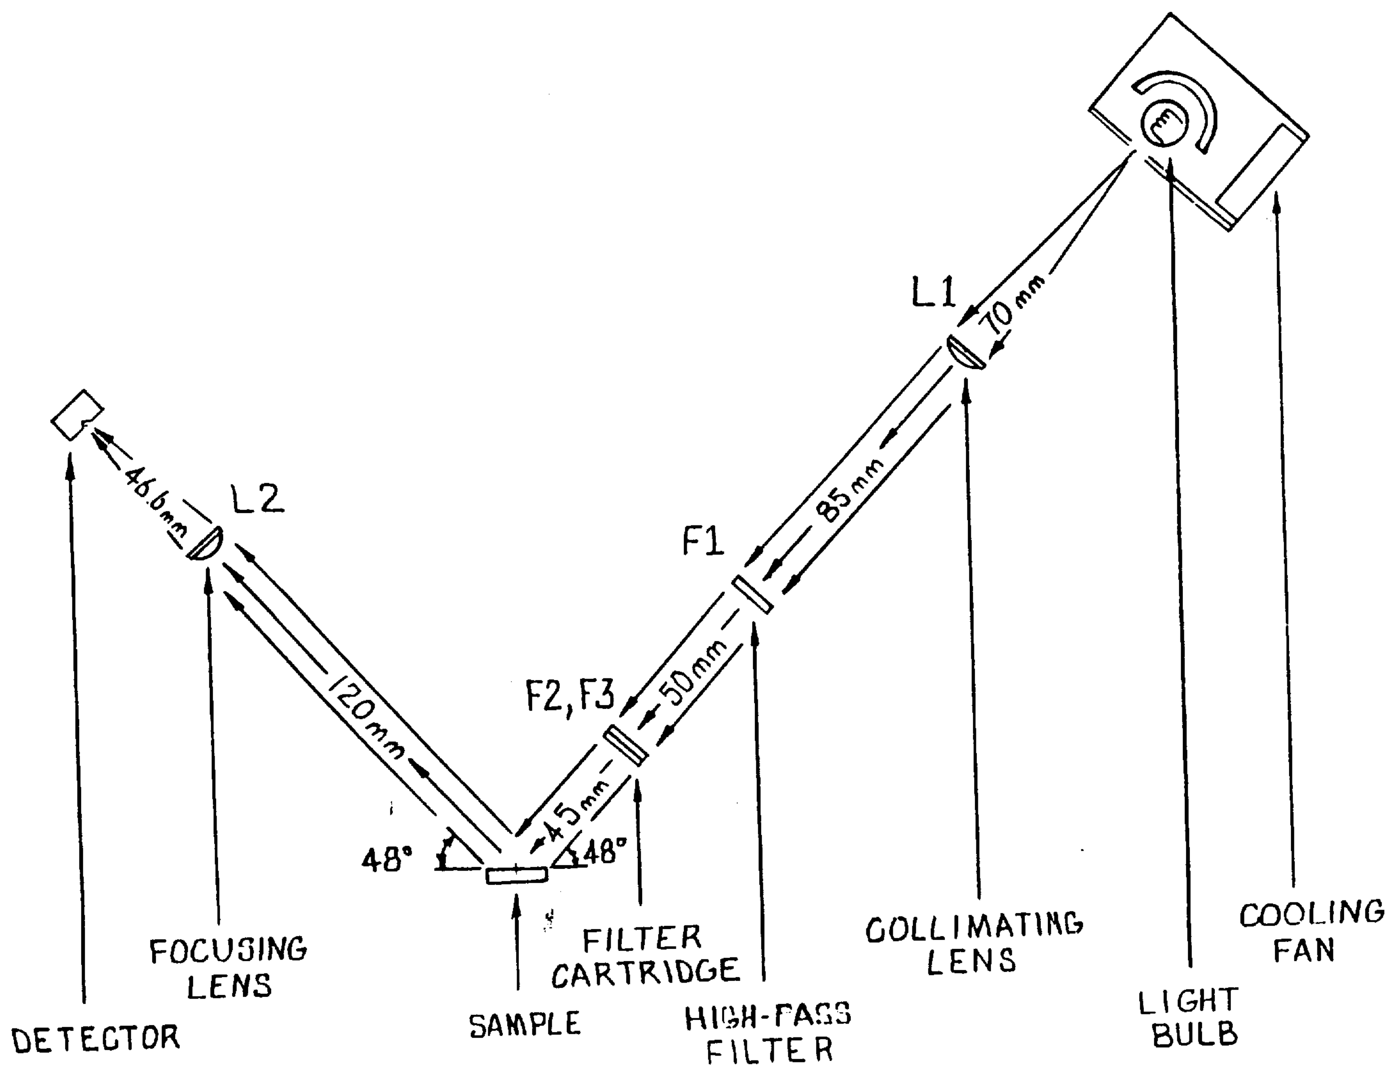

Figure I-3. MCS system configuration.

The light beam from the source diffuses as it passes through the pinhole. A plano-convex collimating lens, placed at a distance equal to its focal length from the point source, gathers the diffuse light beam. The controlled column of light is then directed through a high-pass filter passing wavelengths greater than 1.1 microns, in effect blocking out unwanted energy from ultraviolet and visible light.

The wavelengths that are able to pass through the high-pass filter encounter the filter cartridge holding the narrow bandpass filters. After passing through each filter, the infrared wavelengths are intercepted by the seed sample.

The reflected infrared energy from the sample is collected by a plano-convex focusing lens which focuses the energy on to the detector. The detector is placed at the focal length of the focusing lens and integrated into a circuit to provide an analog readout.

Detector and Circuit

The detector and associated circuitry perform the task of producing a voltage proportional to the intensity of light reflected. The detector behaves like a variable resistor when used as a circuit component connected to a bias voltage supply (Infrared Industries, 1990). Changes in circuit voltage relate to the intensity of the light reflected from the seed sample. A simple series circuit using a low impedance voltage supply, and a resistor with a value corresponding to the average value of the detector yielded acceptable results (Figure I-4).

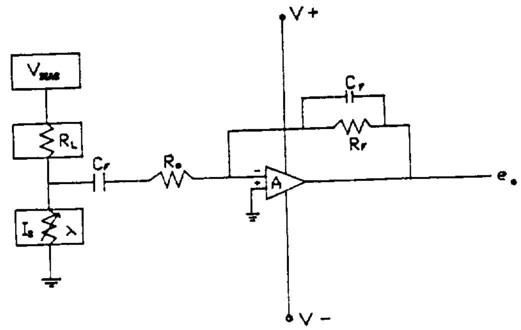

Figure I-4. Circuit diagram for detector. (Source: Infrared Industries, 1990)

Computer Interface

In a completely automated system, a computer would be interfaced with the detector circuitry, converting the analog voltage signal from the detector to a digital signal. A computer program would calculate the ratio of voltages from the reference and water band filters and determine which values corresponded to specific seed moisture contents, concluding whether the seed batch would be stored for replanting or processed as food.

MCS Configuration

The optical components of the MCS were mounted on adjustable arms so that many different configurations could be tested. Parameters that were optimized include intensity, power transmission, and bias voltage (refer to Testing Procedures). Preliminary tests were performed and data was taken in the laboratory to evaluate the success of various configurations of the MCS. Once testing was concluded, the components were mounted on a portable platform to immobilize the components, minimizing the need for reconfiguration after each test. Appendix A gives a detailed list of the specific components of the MCS.

FILTERS

Narrow Bandpass Infrared Interference Filters

Two filters are employed in the infrared reflectance technique. One filter of wavelength 1.94 microns specifically passes a water absorption band, while the other is a reference filter of 1.80 microns. These wavelengths were selected because they fall within the sensing range of the infrared detector and ordinary glass can be used for lenses. The transmission curves of the two filters are presented graphically in Figures I-5 and I-6. Because of the filters' sensitivity, special precautions were taken to ensure that they were not damaged or handled unnecessarily.

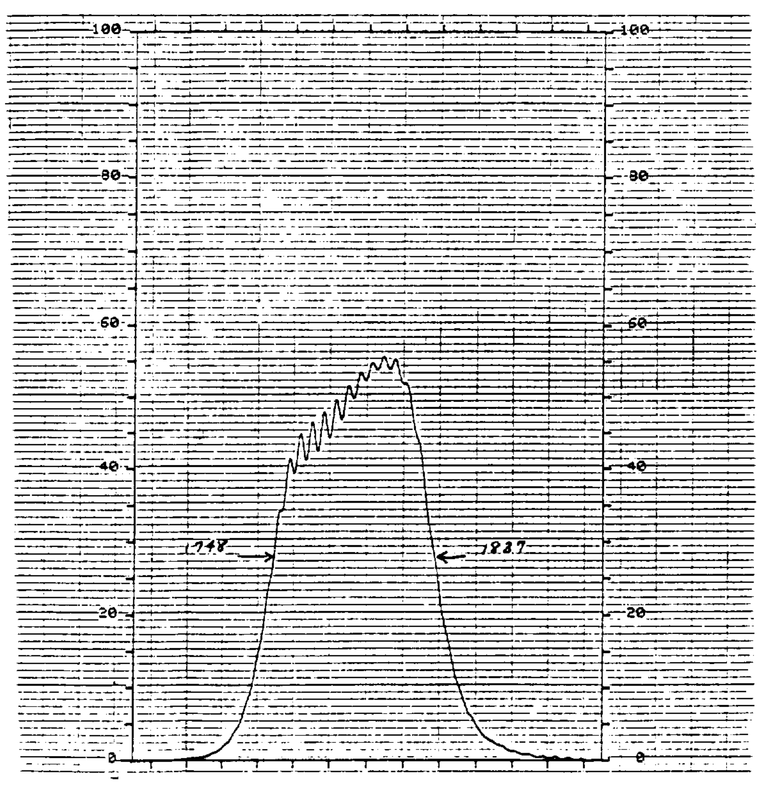

Figure I-5. Transmission curve for reference filter. (Source: Microcoatings, 1990)

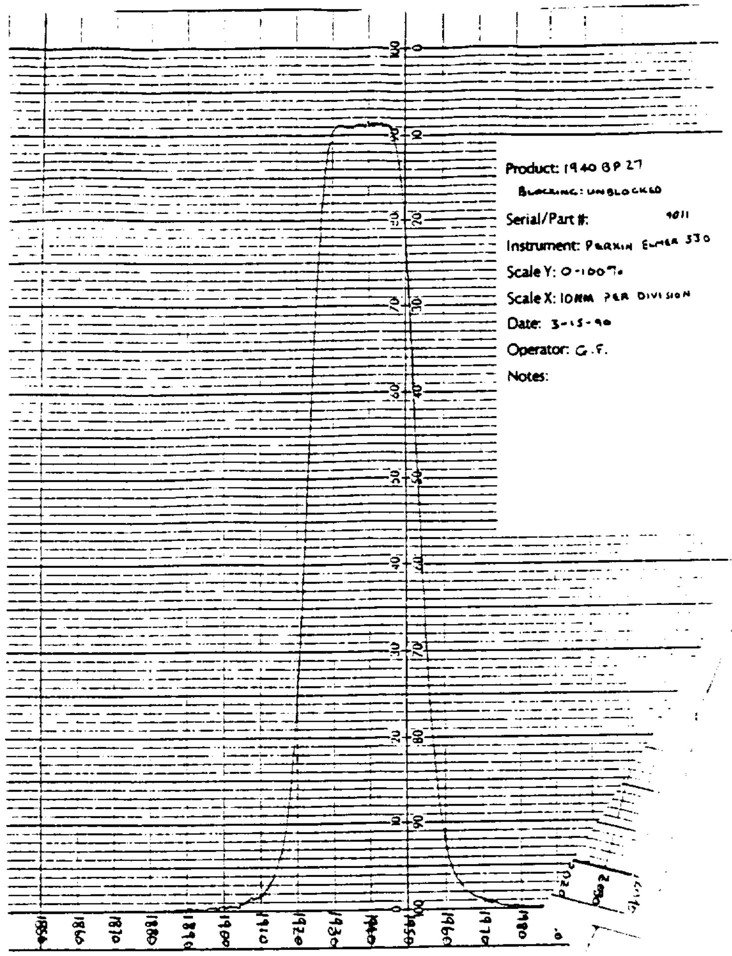

Figure I-6. Transmission curve of bandpass filter. (Source: Dell Optics, 1990)

ORIGINAL PAGE IS OF POOR QUALITY

Peak Wavelength Shift

If conditions other than 25°C and a 0° incident angle are intended for use, changes in the filter's spectral characteristics must be taken into account, especially for the narrow bandpass filters. As the angle of incidence is increased, optical system performance will be affected by a shift toward shorter wavelengths. An increase in temperature will shift the filter's peak wavelength toward longer wavelengths. The exact amount of shift depends on many factors, including the filter's design and its refractive index. The peak wavelength increases 0.005 microns per degree Celsius above room temperature (Microcoatings, 1990). Typical wavelength shifts approach 1% for a 25° angle of incidence and an operating temperature of 150°C. Although not measured directly, the group feels that these shifts do not significantly effect the accuracy of the results.

LEAD SULFIDE DETECTOR

The lead sulfide (PbS) detector is a thin film photoconductor which has been chemically deposited on a quartz substrate. The basic detecting element is the sensitized film on a substrate with two electrical leads attached to the circuit. Photons induce changes in conductivity which modulate the current flowing through the detector. The detector is passive until energized. Biasing with a matched series load resistor to a constant voltage source was used to activate the detector (Infrared Industries, 1990).

The lead sulfide detector chosen has a spectral response curve which shows good detectivity from 1.4 - 3.0 microns at room temperature. This range includes the two wavelengths of 1.80 and 1.94 microns used in the design apparatus. The detector has an active area of 1 mm x 1 mm. Lead sulfide detectors exhibit excellent stability and therefore are suitable for use on extended space missions.

OPTICAL PROPERTIES

In designing an infrared reflectance prototype, basic optical properties and formulas were used to determine the final placement of all components of the system. A basic understanding of optical terminology and properties was needed in order to design an optical system that would maximize the light energy transmitted and keep measurements accurate. Areas of concern in designing an optical system include specific lens characteristics such as focal length, convergency or divergency, and collimation of light.

General Concepts and Theory

The focal length of a lens is the measured distance from the object to the point where the image is precisely focused. There are several ways to determine the focal length of a lens, one of which is Bessel's method. Bessel's method makes use of the distance between the object and the image, L, and the distance d between the two positions of focus (Figure I-7). The formula for Bessel's method is:

$$f = \frac{L^2 - d^2}{4L}$$

The advantage of using this method is that only L and d, rather than the object distance and image distance, need to be measured (Meyer-Arendt, 1989).

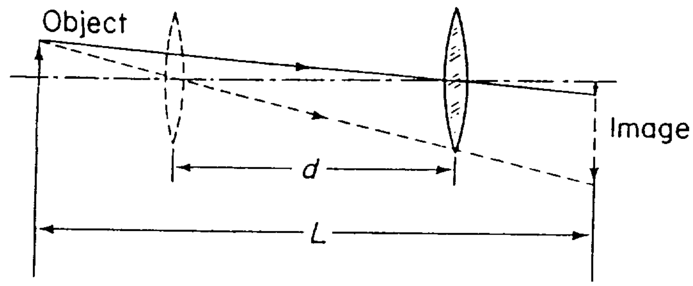

Figure I-7. Bessel's Method. (Source: Meyer-Arendt, 1989)

There are two general classes of lenses, those that cause light to diverge or those that cause light to converge. A converging (positive) lens is thicker in the center than at the periphery, while a diverging (negative) lens is thinner in the center. A converging lens can be used to collimate light. Collimation of light involves taking light from a point source and producing parallel light beams. This is accomplished by placing the light source at the focal length of the lens. Another property of converging lenses is the ability to focus parallel rays of light to a point. By placing the object (source) at infinity, the light rays are focused at the second focal point (Figure I-8).

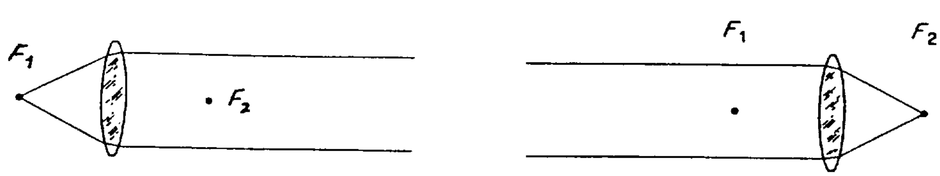

Figure I-8. Focal points of a converging lens. (Source: Meyer-Arendt, 1989)

Optical Design

In an infrared reflectance system, the radiation from a light source is directed by a collimating lens to the surface of a sample. With the use of a second lens, the diffuse reflected light is collected and focused on the detector.

Plano-convex lenses are convergent. These lenses collimate light according to the placement of the object (light source). The asymmetry of a plano-convex lens minimizes spherical aberrations in situations where the object and image are at unequal distances from the lens (Figure I-9). To obtain the sharpest focus, the curved surface of the lens should be oriented toward the distant object. To obtain parallel beams, the opposite configuration would be used (Newport Catalog, 1990).

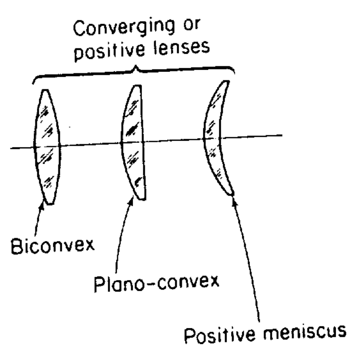

Figure I-9. Types of converging lenses. (Source: Meyer-Arendt, 1989).

Knowing these specific properties, two plano-convex lenses were incorporated into the design of the sensor prototype. The first lens, L1, is termed the collimating lens and has a focal length of 70 mm. The lens is placed 70 mm from the light source (the flat side placed toward the light source), passing a collimated beam of light through the infrared filters and onto the sample.

The second lens, L2, is also a plano-convex lens and is termed the focusing lens. The focusing lens has a focal length of 46.6 mm. This lens is used to collect the diffuse reflected radiation and focus the energy on the lead sulphide detector. This lens is mounted such that the spherical side of the lens is facing the sample and is outside of the focal length. The detector is placed 46.6 mm from the flat side of the lens (Figure I-3).

Infrared Effects on Optical Properties

In the design of the infrared reflectance moisture sensor, it is necessary to consider the optical properties of the media which lie in the path of the radiation. A concern in designing this prototype was the effect infrared has on the focal lengths and transmissivity of the lenses. However, it was found that the infrared wavelengths passed through the system, specifically 1.94 and 1.80 microns, do not have a discernible effect on the results.

TEST PROCEDURES

In order to determine whether or not the MCS design was producing meaningful results, tests initially were performed on the system using a black construction paper background. Water droplets were added one at a time, and voltages recorded for the reference and water band filters. The test showed a decrease in voltage ratios for increases in water; however, a saturation point was reached and voltages no longer decreased, but rather oscillated about a ratio. This did not present a problem as the samples to be tested would contain much smaller amounts of water.

The next set of tests for the MCS involved the optimization of parameters such as light intensity, bias voltage and power transmission. To transmit the maximum amount of energy on the detector, the focal lengths of the lenses needed to be determined, and a cylinder containing the first lens, L1, was placed at the source to ensure that light was not scattered before collimation. The focal length of the second lens, L2, determined the position of the detector in the system.

Different bias voltages were tested without any discernable difference in accuracy or sensitivity of the system, so a 40 V bias was used to ensure that the multimeter would not be overdriven. To determine the source intensity at which further tests should be run, readings were taken of seed samples exhibiting different moistures at several different intensity levels. Another parameter varied was the amount of ambient light in the laboratory. Because the reference filter allows a ratio of intensities to be calculated, the ratios stayed approximately within 1% for each seed sample. The tests run on seed samples of 6% and 25% moisture content showed that at low intensities the two samples were not easily distinguished, while at higher intensities a significant difference between voltage ratios occurred (Figure I-10). This allows a better differentiation between seed samples. Approximately 450 Watts was chosen because at even higher wattages voltage ratios began to decrease, and there was still some concern of heat damage to the filters. If a fixed, rather than variable source intensity were used, the repeatability of tests would be more evident.

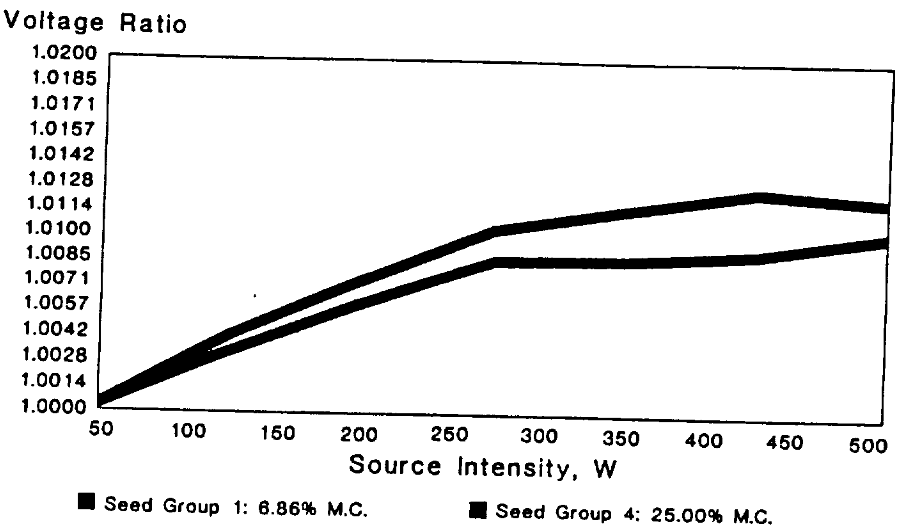

Figure I-10. Voltage Ratios vs. Intensity at two moisture contents.

The next phase of testing involved measuring the reflectance of seed samples of different moisture contents. The samples were taken from seed batches of known moisture contents determined by a Burrows model 700 digital moisture meter. The sample holder was constructed of plexiglass. The reflectance of the empty seed holder was measured, and a continual increase was observed. This occurred due to the holder retaining infrared energy or "heating up" if it remained under the light for longer than a few seconds. Therefore, each voltage measurement had to be made separately and instantaneously while the readings were stabilized.

Another concern was the effect of sample geometry on measurements recorded. Tests showed that a single seed layer did not affect the reference filter readings substantially for batches of similar moisture content. The seed samples were taken from randomly arranged seed batches, and two voltage readings for each filter were made to obtain the average values for a particular seed sample. Room temperature and differences of pressure appeared to slightly affect the readings from the detector.

CONCLUSION

The final tests performed under the optimal conditions resulted in reflectance measurements of seed samples that successfully differentiated between batches of different moisture contents. Occasionally, the readings of some samples were not ordered properly, perhaps because the actual moisture content of the seed batches had varied slightly from the time of measurement with the digital moisture meter. The following graphs summarize the results of two tests.

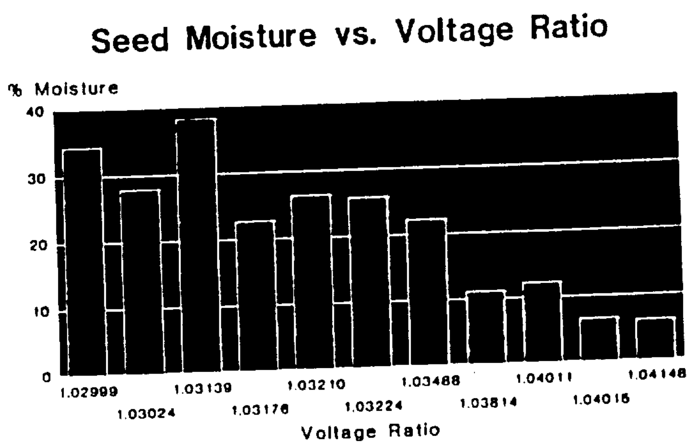

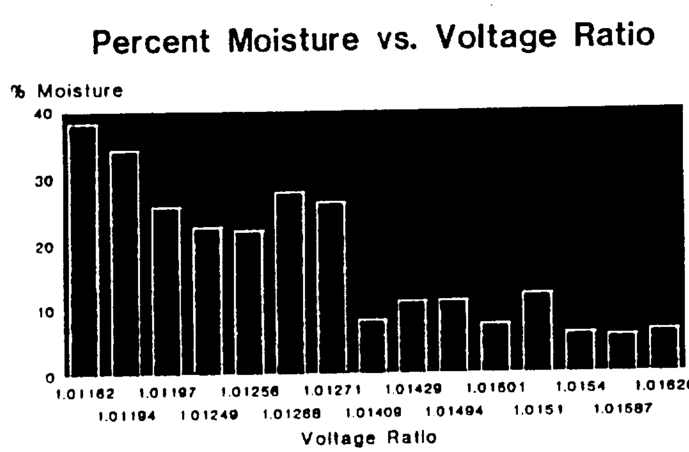

Figure I-11. Test results.

The effectiveness of the sensing prototype was verified by determining the moisture content of unknown seed batches. Several samples were correctly differentiated by the sensor under the optimal testing conditions (outlined in the previous section). The MCS was finally tested for repeatability and the voltage ratios remained within 5% of the previously measured values corresponding.

RECOMMENDATIONS FOR FUTURE DEVELOPMENTS

Further improvements in the MCS design are necessary if the system is to be implemented in a bioregenerative closed-loop system. The most important addition involves total computer control of the data acquisition process, including alternating the filters in the cartridge between the interference and reference filter. Fiber optic technology integration would minimize the sensor's physical size, as well as provide an unobstructed and direct light pathway through the system. Infrared light emitting diodes (LEDs) are designed to emit a specific wavelength without the use of filters. Improved signal processing will also increase the sensitivity of the voltage outputs, thereby permitting differentiation between batches within 0.5% moisture.

REFERENCES

Carr-Brion, K. (1986). Moisture Sensing and Process Control. England: Crown Publishing, pp. 53-4.

Dell Optics. (1990). Company literature on 1.94 micron water bandpass filter.

Infrared Industries. (1990). Company literature on lead sulfide detector.

Kruse, P.W., McGlauchlin, L.D., and R.B. McQuistan. (1962). Elements of Infrared Technology: Generation, Transmission, and Detection. New York: John Wiley & Sons, Inc.

Meyer-Arendt, Jurgen R. (1989). Introduction to Classical & Modern Optics. New Jersey: Prentice Hall.

Microcoatings. (1990). Company literature on 1.80 micron reference filter.

Newport Catalog: Precision Laser and Optics Products. (1990). U.S.A.: Newport Corporation. p. N-43.

APPENDIX A

MCS COMPONENT LIST AND SPECIFICATIONS

Lenses

L1 : Collimating Lens
Source : Optics Laboratory, Department of Aerospace Engineering, Mechanics and Engineering Science, University of Florida
Type : Plano-convex
Diameter : 65 mm
Focal Length : 70 mm

L2 : Collecting/Focusing Lens
Source : Optics Laboratory, Department of Aerospace Engineering, Mechanics and Engineering Science, University of Florida
Type : Plano-convex
Diameter : 54 mm
Focal Length : 46.6 mm

Filters

F1 : High Pass Filter
Source: Physics Department, University of Florida
Type: Single crystal layer silicon glass
Wavelengths: 1.10 microns and above
Diameter: 50 mm

F2 : Water Absorbing
Source : Dell Optics
Center Wavelength: 1.94000 microns
Half Power Bandwidth: 0.02700 microns
Half Power Points: 1.95350 microns
1.92650 microns
Maximum Transmission: 91.00 %
Diameter: 1.00"

F3 : Reference
Source : Microcoatings, Inc.
Center Wavelength: 1.78590 microns
Half Power Bandwidth: 0.08741 microns
Half Power Points: 1.82960 microns
1.74220 microns
Maximum Transmission: 56.07001 %
Diameter: 1.00"

Other Components

Cooling Fan
Source: Rotron
Type: sentinel series 747
Flow: 50/60 cps
115 VAC / 14 w

Infrared Lamp
Source: Sylvania
Type: 500 w DAY
120 V

Detector
Source: Infrared Industries
Type: Lead sulfide
Part #: 2303

Soybean seeds
Source: Mother Earth Market, Gainesville, Florida

II. POROUS MEDIUM WETNESS SENSING

Prepared by

Lisa Bean
Gavin Clark
Barry Finger
Parker Severson
Andrew Speicher

TABLE OF CONTENTS

SUMMARY

This section will describe the work done by the porous-medium wetness group. A variety of sensor designs were studied to develop a sensor that would accurately and repeatably give measured results based on the amount of moisture present on a porous surface.

The first step was to determine a difference in properties between water (the nutrient solution) and the porous ceramic plates that are used. With these results, the group proceeded to design several sensors based on these differences. A testing platform was built to give a uniform testing ground for each of the designs to minimize errors in the testing procedures.

Two of the sensors proposed gave satisfactory results in preliminary tests: an infrared reflectance sensor and a thermistor based sensor. They were developed into full working models. The results of these two models, as well as descriptions of other sensors studied, will be described.

INTRODUCTION

In order to maintain an efficient nutrient delivery system in a CLLSS, the ability to monitor the amount of nutrient solution available to a plant through the porous medium is necessary. The most important factor in plant growth and productivity is soil water, which controls the uptake of most of the nutrients required for plant growth (Phene, 1988). Nutrient uptake of a crop varies throughout the crop season, as well as daily. Monitoring the wetness of the porous plate will provide data necessary to help control the amount of nutrient solution available to the plants throughout their growth cycle. This control can help to maintain healthy plants, prevent leakage to the CLLSS atmosphere due to excess solution in the medium, and prevent air entering into the delivery system due to a lack of solution in the medium.

Plant water stress refers to situations where plant turgor is reduced enough to interfere with normal processes. Effects of plant water stress depend on plant stress timing, magnitude and duration. When a plant is under water stress, the leaf stomata close to provide protection. The transpiration rate and entry of CO_2 declines, thus decreasing the rate of photosynthesis and consequently decreasing the crop yield. By controlling the amount of water available and reducing plant water stress, crop yield can be maximized.

Soil water is usually measured as water content (either mass or volume fraction) or as soil water potential. Professor D. Hillel (1982) refers to wetness as the intensity of water in the porous medium. The degree of wetness pertains to the relative concentration of water in a porous body, independent of the body size. Two sensing methods were designed for monitoring the wetness of the porous medium: one which uses infrared reflectance, and one which is based on heat dissipation. These two methods and design results are presented.

INFRARED DIFFUSE REFLECTANCE

A technique for measuring surface moisture using the Infrared absorption properties of water is discussed in the Seed Moisture section of this report. Refer to that section for a detailed description of the sensor design. The infrared moisture sensor was successfully used to monitor the surface wetness of the porous plate test platform. This technique is especially suited to sensing porous plate wetness since it measures surface moisture. This is ideal since it is the surface moisture that should be monitored.

Test Results

A testing platform was designed and built to allow the use of several types of porous plates to see the effects of different materials on the sensor (Figure II-1). The testing platform was a chamber 2.75 in. wide, 8.25 in. long, and 0.5 in. deep. All sides of the chamber except for the top are permanently sealed. The ends have valves that can be opened and closed to provide dynamic flow or static pressure. The top surface has a pressure seal to allow the changing of plates by simply removing the four set screws that secure the porous plate.

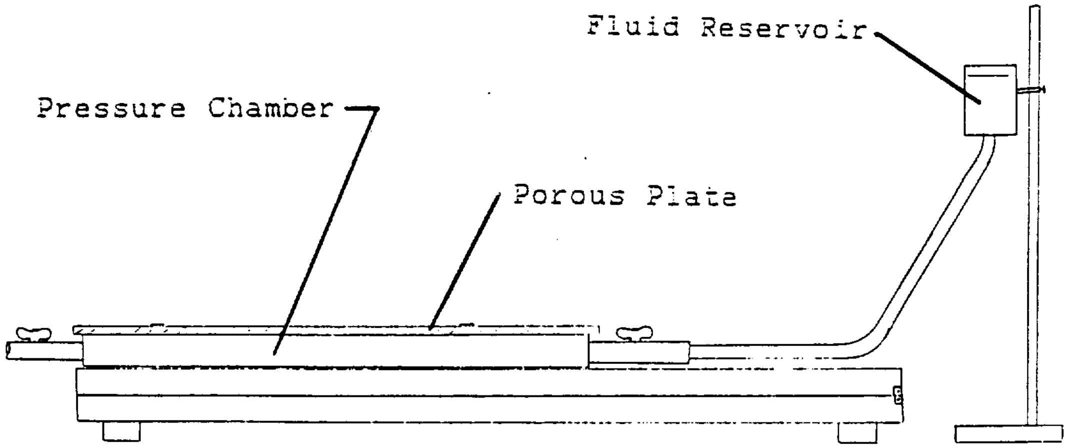

Figure II-1. Testing Apparatus.

The pressure inside the testing chamber was varied by changing the height of the solution reservoir with respect to the test platform. All of the tests performed concerning infrared reflectance were performed with a static pressure that was varied depending on the type of plate used.

The porous plate test platform was configured to provide 1.01 psi of water pressure, and the infrared source intensity was set to 450 Watts. Data was taken every thirty seconds over a time period of approximately thirty minutes. The 1.01 psi pressure forced water out through the porous plate. Over time, surface moisture built up on the top of the plate. It was this changing water concentration that the infrared sensor monitored. The results of two separate runs are shown in Figure II-2. Both curves follow the same general behavior, showing a definite relationship between surface moisture and the infrared sensor output. There is a voltage ratio offset of approximately 0.07 between the two runs despite the fact that they were thought to be conducted under identical conditions. This discrepancy may be due to a variation of source intensity between the two runs since it is known that source intensity has a significant effect on sensor output (see Seed Moisture section).

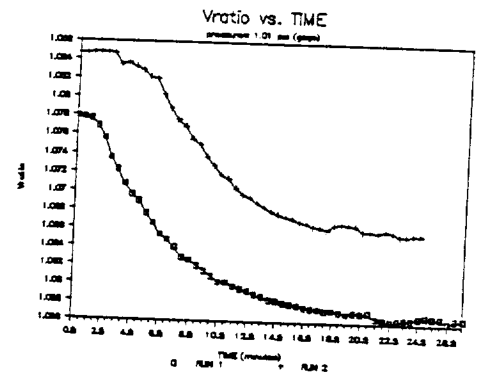

Figure II-2. Infrared sensor output for porous a plate.
Note wetness increases with time.

Future Development

The IR moisture sensing technique appears to be an excellent candidate for further development since it is a non-invasive technique that monitors surface moisture. The data taken suggest that definite, repeatable results are possible using an improved sensor design. See the Seed Moisture, Future Development section for a discussion of possible sensor improvements.

INVESTIGATIONS IN ELECTRICAL PROPERTIES

The difference in conductivity between a dry and saturated porous ceramic plate initiated an investigation into the viability of sensing moisture using this contrasting property as applied to the dielectric properties in a capacitor. As shown in Figure II-3, the sensor was envisioned to consist of two plates electrically insulated on their surfaces, an RC circuit and an oscillator.

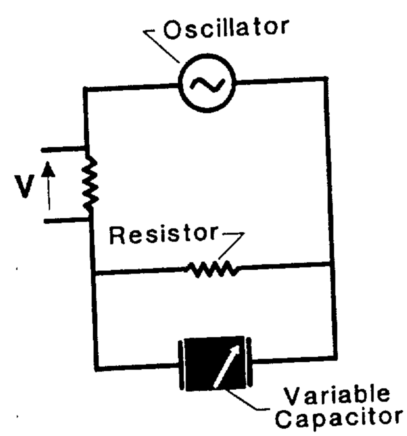

Figure II-3. Capacitance sensor.

The moisture in the porous plate would act as the dielectric medium between the two capacitor plates. An increase in concentration of moisture in the plate would raise the medium's dielectric constant, as well as increase the capacitance in the plate. The change in capacitance would cause a deviation in the natural frequency of the RC circuit to a new frequency which could be measured by a frequency-to-voltage converter. Ultimately, a change in voltage would signal a change in moisture concentration. The measurement may be correlated to a moisture concentration function varying through the thickness of the plate.

In testing, the group set out to show a repeatable direct relationship between the concentration of wetness in the porous growth plate and the capacitance of the sensory plates. Aluminum tape strips were used as the sensory plates and they were placed on the surface of ceramic tubing provided by NASA.

The tube was filled with water which was allowed to seep through the pores along the length of the tube. Capacitance was monitored by a regular capacitance meter from the instant water was added to the tube until it began to bead on the outside. Unfortunately, the capacitance varied only 2 to 3 nF and the relationship of capacitance to wetness was not clear. The capacitance at first increased and then decreased as the porous tube became saturated. Once the plate was saturated, excess water may have electrically bridged the two sensor plates, causing the decrease in capacitance.

Agricultural articles describing the use of conductance between probes in solution as a determinant of ion concentrations in solutions initiated the development of a resistance-type moisture sensor. Figure II-4 shows that the sensor was constructed with four probes: the outside probes connected to a constant current source and the inside probes connected to a voltage meter. As soon as the probes were set down on the plate, the moist plate acted to close the circuit inducing current flow. The voltage meter read the voltage drop between its leads on the plate. Given the current and the voltage drop, an expression for the resistance in that section of plate between the center probes can be determined. This derived expression represents the surface moisture of the porous plate.

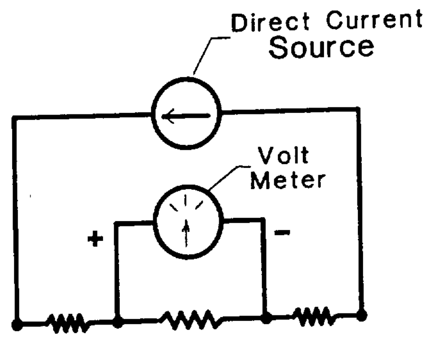

Figure II-4. Resistance sensor.

More effort was taken in locating and refining a probe. A square corner of a printed circuit board was cut to expose four metal strips that could be used as contact leads. The two center strips were glued to the circuit board backing in order to insure the voltage drops recorded were always taken over the same length of plate. This circuit was chosen to act as the probe because the metal strips were a small enough distance apart (approx 1.5mm) to force the current to choose the path across the surface. This was the least resistive path, rather than down into the plate and through the nutrient solution. Voltages across the center probes were recorded for ceramic earthenware from a damp state to saturation. It should be understood that the voltages could not be generated at drier states because of voltage limitations of a real current source. The data showed a general inverse relationship between the amount of wetness and the voltage drop between the center probes. This ranged form several megohms for a damp condition, to 20,000 Ohms for saturation. Although this behavior was expected, the sets of data were not repeatable. Pressure, area, and duration of the probe's contact with the surface seemed to shift the data. Inconsistencies in data may be explained by recognizing that variable contact between any two conductors varies the resistance. The phenomenon of voltage rising with time while the probe was in contact is evidence that resistance was possibly rising due to the moisture being heated in the path of the current.

Some alterations to the probe were considered. To remove pressure and contact area as factors, the new probe would be spring-loaded permitting a constant pressure contact. The metal strips would be replaced by four metal conical protrusions that would contact the surface. Also, a sample and hold amplifier could be timed to register an immediate voltage value so the surface moisture does not have time to heat.

INVESTIGATIONS INTO THERMAL PROPERTIES

Other possible physical characteristics to employ in a moisture sensor are the thermal properties of heat conductance and heat capacitance of water. The thermal properties of water differed quite significantly from those of air or ceramics. This contrast provided an obvious opportunity to sense moisture levels in a porous ceramic plate.

The technique of measuring heat dissipation to sense the degree of wetness in the porous medium is based on the fact that the rate of heat dissipation in a porous medium of low conductivity is sensitive to water content. As the water content decreases, a larger temperature gradient is needed to dissipate a given quantity of heat. The rate of heat dissipation is dependent on the specific heat, thermal conductivity, and density of the porous medium (Phene, Hoffman, Rawlins, 1971). In results provided by Phene, Hoffman, and Rawlins (1971), the conductivity of four porous materials was found to be higher than that of water, since their saturated conductivities were all greater than that of water. This provides a basis for developing a sensor which relies on the thermal conductivity of the medium. The wetness of the porous medium can be correlated to the temperature gradients measured by the sensor for a given heat dissipation.

One criticism of measuring moisture content based on applying heat is that moisture will move away from the heat source (Bloodworth, Page, 1957). The amount of movement depends on the temperature of the heat source and the length of time the heat is applied. However, if the temperature difference between the medium and the heat source and the time interval are small, then the movement of moisture can be considered negligible.

A thermistor-type sensor consists of thermistor elements (approx. 100,000 ohms at room temperature) in a bridge circuit (Figure II-5). One thermistor touches the plate while the other remains away from the plate exposed to the air. A high current (approx. 10mA) runs through the thermistors forcing them to self heat. As a result, the thermistors rise in temperature a few degrees above the air and plate temperature. Once the bottom thermistor contacts the moist plate, the water present around the thermistor functions as a heat sink drawing the heat away from the thermistor. The thermistor's temperature drops closer to ambient, and the thermistor's resistance is pushed above the other resistances in the bridge circuit. Finally, the presence of moisture is read as a differential voltage across the bridge. The differential voltage was seen to climb nearly 100% from sensing moisture on a saturated plate.

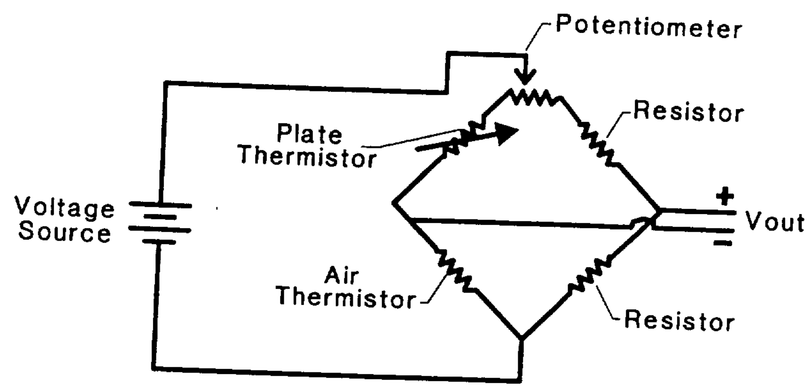

Figure II-5. Bridge circuit.

While the exact moisture levels in the test plates were not quantified, general wetness trends were visually noted. Five ceramic samples were soaked and set out to dry for different periods. The samples were arranged from driest to wettest. Near saturation, a further addition of water formed a reflective film on the surface of the piece.

Although the test results seemed favorable, the circuit required 25 volts to drive a current. This pushed the thermistors out of their working mode into self heating. Alternate temperature sensing elements were studied because of the slow reaction time of the thermistors. The high voltage requirements necessitated the group to search for a better way to heat the sensing elements. Upon suggestion to heat the thermistors externally, the team located a 820 Ohm resistor (Nevill, 1990). The geometry of the resistor made installing temperature sensors on either side of it convenient. To overcome the slow reaction times of the thermistors, cement-on thermocouples were purchased from the Omega Company. The sensing junctions of the thermocouples were only 0.0005 inches thick. This small thickness allowed millisecond reaction times. The heating resistor was sandwiched between two thermocouples, shielded from abrasion by aluminum foil and bonded together by Omegabond 101, a thermally conductive epoxy.

One thermocouple rested against the porous growth plate, while the other one, on the opposite side of the resistor, faced the air (Figure II-6). This sensor unit was based on the same heat transfer principle as the thermistor unit. While the thermocouples were self-generating and did not demand excitation in a bridge circuit, their output did require amplification. Increasing the resistor's temperature 5°C creates a voltage change of 900 microamps across the thermocouple leads (table T in Omega catalog). It was assumed that a change of 5°C would not have detrimental effects such as drying, evaporation, or precipitation. Figure II-7 displays the thermocouples in their complimentary circuit, which furnishes a gain of ten-thousand and filters out most of the higher frequency noise.

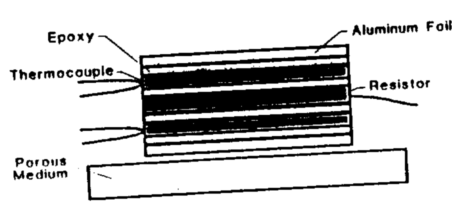

Figure II-6. Thermocouple probe.

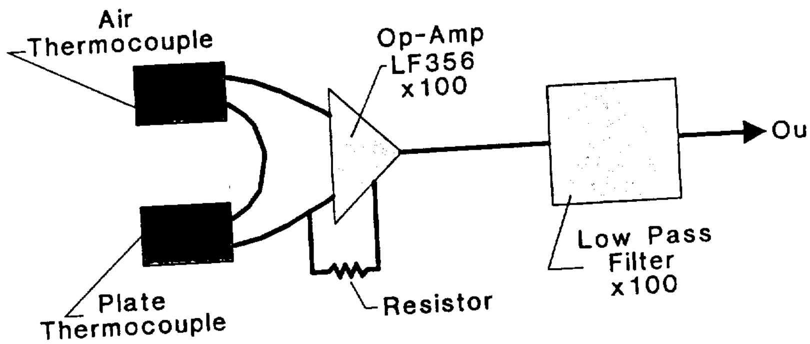

Figure II-7. Thermocouple circuit.

In theory, if the two thermocouples had been simultaneously raised 5°C, and then one of them was lowered 2°C due to its contact with the moist porous plate, the differential voltage change between the two thermocouples would have been 360 microvolts. The amplifier would have subsequently produced an output of 3.6 volts. The output was monitored by an oscilloscope. It displayed the signal fluctuating with a small noise effect superimposed. Unfortunately, the up and down shifts did not follow any patterns of heating or cooling.

After some effort to discover the cause of this fluctuation, it was concluded that the aluminum shielding may have played a part in desensitizing the sensor. Unfortunately, a stronger output was needed to plot any kind of reaction to moisture.

The bead thermistor, a much more sensitive temperature sensing element, was chosen to be employed in the previous bridge circuit similar to Figure II-7. The probe itself was designed in the same fashion as the thermocouple probe; a thermistor on either side of a resistive heating element. The thermally conductive epoxy surrounding the thermistors focuses temperature changes in the environment on the bead. The insulating epoxy was basically for support, but more importantly it did not become an extra thermal mass to heat.

The hardware for the thermistor sensor was made of plexiglass and was designed to hold the sensor in place with set screws. The sensor holder allows the sensor to be moved while still ensuring consistent pressure and contact area between the sensor and porous surface. The use of the set screws minimized the heat transfer between the sensor and its holder. A six-wire wiring harness allowed different sensors to be used in the holder without much difficulty.

When the circuit was excited, the voltage across the bridge was virtually zero. Heat was removed from the thermistor by the moisture, lowering its temperature and raising its resistance thus a differential voltage appeared across the bridge.

Tests for the new thermistor setup were conducted as mentioned previously. As seen in Figure II-8, the peak voltages varied with respect to the moisture levels in the tiles. The voltage peaks ranged from approximately 0.20 to 0.43 volts. These peaks generally occurred in 45 seconds which seems to be a reasonable reaction time for a moisture sensor. Actually the moisture characteristics of the plates could be differentiated as early as 30 seconds. Therefore this method of porous plate moisture sensing seems promising and warrants further development.

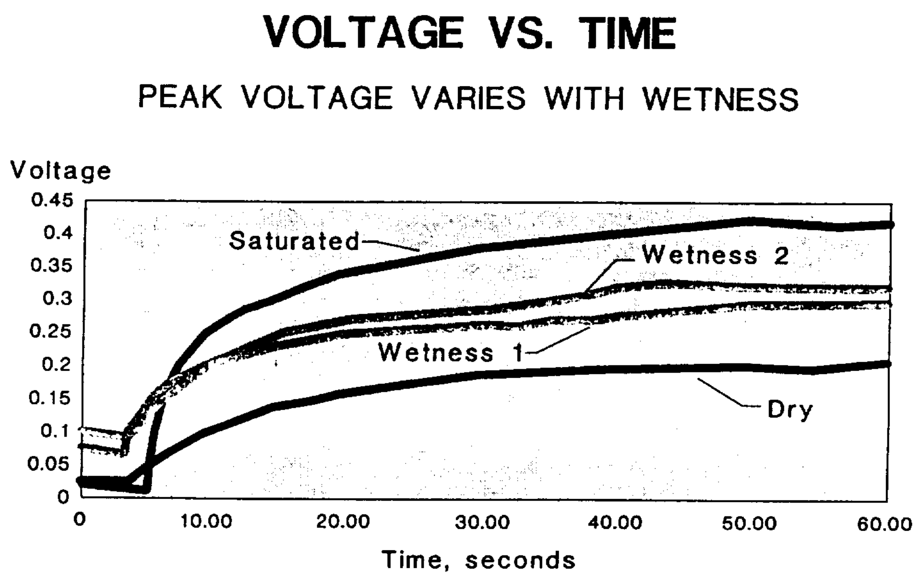

Figure II-9. Thermistor probe test data.

CONCLUSION

Various sensing methods were designed and tested for moisture sensing this semester. Each of the different sensors was based on the differences of certain properties between water and a porous ceramic. The two most successful sensors were the infra-red diffuse reflectance and thermistor-based models. Each of these sensors has strengths and weaknesses inherent in its design that may not be solvable, but further research is needed to determine if the designs may be improved to an acceptable level.

The infrared reflectance sensor was a successful approach with some problem areas that must be worked out. These problems included offset problems caused by the variations in intensity of the light source.

The thermistor-based sensor demonstrated the viability of the technique of thermal conductance method. It clearly showed a difference between a saturated porous plate and one that is dry.

REFERENCES

Bloodworth, M. E., and Page, J. B. (1957). Use of Thermistors for the Measurement of Soil Moisture and Temperature. Soil Sci. Soc. Amer. Proc., 21: 11-15

Hillel, D. (1982). Introduction to Soil Physics. Academic Press, New York, N.Y.

Jorgensen, G., Bundy, E., and Miyaski, M. (1988). Sensors and Techniques for Irrigation Management. pp. 25-33, 75-83, 127-139. California Agriculture Technology Institute, Fresno, Ca.

Phene, C.J., Hoffman, C.J., and Rawlins, S.L. (1971). Measuring Soil Matric Potential in situ by Sensing Heat Dissapation within a Porous Body: Theory and Sensor Construction. Soil Sci. Soc. Amer. Proc., 35:27-33.

Zoldoske, D.F., and Miyasaki, M.Y. (1987). Micro-Irrigation Methods and Materials Update. California Agricultural Technology Institute, Fresno, Ca.

III. PLANT HEALTH SENSING USING INFRARED DIGITAL IMAGING

Prepared by:

Amy Clark
Donald Koedam
Mark Marchan
Pedro Rodriguez
Robert Schneider

TABLE OF CONTENTS

SUMMARY

Within a closed-loop life support system (CLLSS), it is essential that the health of the plants be monitored. Therefore, the fabrication of a plant health sensor which can be automated and monitor remotely and effectively is required. In order to accomplish this, it may be necessary to sense the crop as a whole as well as individually. This project will focus on the development of a plant health sensor which can examine the Biomass Production Chamber (BPC) as a whole. The system which may be best suited would be a visual system utilizing infrared (IR) radiation. This system would incorporate the use of an infrared digitizing camera system, as well as a computer system for analysis.

By using IR absorption as a means to measure chlorophyll content, it is possible to quantify the health of a plant. Using the absorption band of chlorophyll-a at 0.671 microns a well-defined image may be obtained through a digitizing camera. This camera will produce an array of numbers representing the image which may be used in a computer analysis. This analysis may not only reveal areas which are unhealthy, but also the extent of the damage.

INTRODUCTION

In order to maintain an extended space mission, it may be necessary to grow crops to support the astronauts physically and psychologically. The plants should be grown efficiently, with minimal interaction with the crew. This requires a number of automated systems to care for the crops. One such system should be able to monitor the health of the plants. This would be accomplished on many different levels. The first of these levels would consist of a primary health sensor which could scan the entire crop in order to identify trouble areas. The secondary sensor would be able to examine a trouble area more indepth, to specify the location and extent of damage. Finally, the tertiary system would be able to analyze the trouble area for a specific cause and a possible solution for the problem. The primary and secondary sensors would be noninvasive, while the tertiary sensor may be destructive; however, the extent of the destruction could be minimized by the information gathered by the secondary system.

This project centers upon the design and fabrication of a primary sensor system. To sense the plant health, various factors may be examined. In this case, the amount of chlorophyll located in the plant tissue provided a direct correlation to plant health. It may be seen that health decreases as the chlorophyll concentration decreases. This may readily be seen in chlorotic and necrotic tissue. Chlorotic tissue shows a deficiency in chlorophyll through a yellowing of the tissue. Necrotic tissue is dead tissue, which is brown and dry in appearance.

Chlorophyll is detected through the use of IR radiation absorption. Absorption is measured through the use of a digitizing camera system. Light is reflected off the leaf and collected by the camera. The light is first passed through a bandpass filter in order to pass a wavelength of 671 nm, this wavelength corresponds to the peak absorption of chlorophyll-a. As the amount of chlorophyll decreases, the amount of 671 nm light reflected will increase. This will be observed by the camera and recorded by a computer within an array of integers. This array represents the light intensities located within an image. Each number in the array represents a grey level and as the numbers increase, the intensity of light increases.

Although ideally this system would operate on the entire crop under ambient light conditions, for this project the sample was restricted to a single leaf, laid flat, and used an additional light source. With additional modifications these restrictions may be overcome.

BACKGROUND

Multi-Component System

In a closed-loop life support system, there is a need for a sensor that can accurately monitor the health of the crops within the biomass production chamber (BPC). In order to perform this task, it may be necessary to sense on several different levels, with the sensor being able to scan the entire crop and detect if there is any deviation from the healthy status of the plants. If there is a problem, the device should be able to locate the diseased plant(s) and perform a specific analysis in order to ascertain the exact nature of the problem. Because these stages require different types of analysis, the task of finding one device capable of detecting a problem area and analyzing that area is difficult.

To overcome this dilemma, a multi-component system containing two or more sensors has been proposed as the best method for determining the health of the crop (Appendix A). The different levels of the sensor would include the following:

Primary. The primary sensor should be fully automated and able to perform a general overview of the crop. The output should be real time and easily digitized for analysis in a computer data base.

Secondary. The secondary sensor would only be utilized if the primary sensor detects a problem. This stage should be able to pinpoint one specific area, and focus on the leaf or stem surfaces of the plant(s) in question.

Tertiary. The tertiary level would use a sensor that could perform a detailed analysis of the plant. The goal is to determine the exact nature of the problem, whether it be a deficiency in a specific nutrient or an infection from a certain pathogen. This stage may need to be destructive, however, the amount of destruction could be minimized by employing the secondary system to determine specifically which portion of the plant should be removed for analysis.

Digital Camera

In designing this health sensor system, it is logical to begin with the primary sensor which will scan the crop and look for changes in the status of the BPC. A visual device, such as a digital camera, is recommended as a way to obtain remote and automated sensing because the camera may be connected directly to a computer. When a picture is taken, it can then be digitized and stored in a file as an array of numbers. The numbers represent data which can be used to reconstruct the image on a computer screen or to perform a computer analysis. In a CLLSS situation, this technology can be utilized to take pictures of the crop and then analyze the data in an attempt to find discrepancies in the normal status of the BPC.

The camera takes a picture of an object by recording the light intensities that pass through its lens. In the process of digitizing, the intensities are represented by numbers, known as grey levels, ranging from 0 to 255. A zero defines a situation in which no light passes through the camera, while a 255 represents the highest intensity possible for the sensitivity of the camera. In order to visualize the image, these numbers may be placed into 16 groups which are represented by a color (pseudocolor), which is independent of the actual color (Schwab, 1990). An example of this process is shown in Table III-1.

TABLE III-1. Computer Vision Digitizing Process.

digitized intensity	grey-level	color
0 to 16	1	black
17 to 32	2	violet
33 to 48	3	blue
.	.	.
.	.	.
.	.	.
239 to 255	16	white

In addition to visual inspection through pseudocolor representation, it is possible to analyze the actual data. Computer systems may be used to examine the grey level values to discover any regions of interest. An expert system may be used to detect unhealthy portions of the plant through gradient operations. A neural network may use pattern recognition, this will be covered in detail in the following section. It is the computer analysis method that would be most useful for a CLLSS application by allowing the digital vision system to be automated as well as remote.

Infrared Sensing

It is possible to make the digital camera system even more applicable for plant health sensing by specifying the type of light allowed to pass through the camera lens. Leaves demonstrate very selective reflectance properties when exposed to different wavelengths of light due to the chlorophyll content within their tissues.

Plant leaves contain two types of chlorophyll, chlorophyll-a and chlorophyll-b. Both types have two absorption bands; chlorophyll-a has bands centered around 450 and 670 nanometers, while chlorophyll-b has bands at approximately 520 and 700 nanometers (Schalkoff, 1989). However, chlorophyll-a tends to have better absorption effects than chlorophyll-b (Figure III-1).

In combination with the internal structure of the leaves, the amount of chlorophyll gives a plant a very distinctive reflectance pattern in the near-infrared wavelengths (Schalkoff, 1989). When healthy, a plant contains normal amounts of chlorophyll and will absorb light at 450 and 670 nm (Raven, 1986). As it develops a deficiency or disease, the chlorophyll breaks down, leaving internal air cavities within the leaf (Schalkoff, 1989). If allowed to continue, the plant will develop external symptoms such as chlorosis (yellowing or whitening of a leaf) and necrosis (death of tissue). The disintegration of chlorophyll and increase of internal spaces cause a sharp increase in light reflection in the wavelengths mentioned.

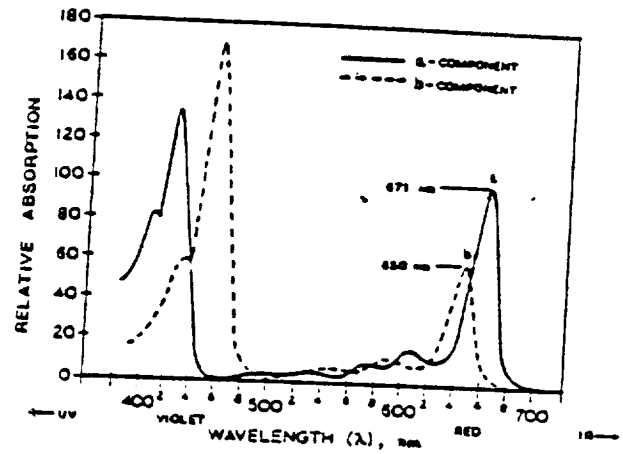

Figure III-1. Absorption peaks for chlorophyll-a and chlorophyll-b.

Thus, to increase the effectiveness of a visual plant health sensor, it is advantageous to filter out all light except that in the chlorophyll absorption bands. By using a bandpass filter to eliminate any light lower than 660 nm and higher than 680 nm, the camera can be used to record only those wavelengths of light which may be absorbed by the plants. Because of this, a change in plant health may be indicated by an increase in the light intensities of that range. In a digitized image this may be demonstrated as an increase in magnitude of the numerical data. A computer may be utilized to monitor the data in order to detect these changes and alert the sensor system to any problems that may exist. A secondary sensor may then be employed to discover the exact location of an unhealthy area.

PRELIMINARY DESIGN

In order to accomplish the objectives of a primary sensor, the system would use the infrared sensing technique previously described. The Plant Health Sensor group developed a preliminary design prior to any testing. The group envisioned a sensor that would ultimately fulfill all its requirements. However, for initial development, certain restrictions were applied:

1. The sensor would concentrate on one leaf.
2. The leaf would be removed from the plant and laid flat.
3. The sensor need not be real time.

These restrictions would simplify the project making the initial fabrication of a sensor for the BPC more attainable. Continued development would allow these restrictions to be overcome.

Based on these criteria, the theoretical sensor was developed. This sensor would consist of, ideally, an IR camera connected to a computer containing a digitizing program. The camera lens would be used in conjunction with a bandpass filter that would pass wavelengths within the chlorophyll absorption bandwidth. In essence, this would be the entire sensor. The analysis of the input from the sensor would be accomplished by a separate program.

EQUIPMENT

The following equipment was used for the sensor:

Digital Video Camera

The camera used by the plant health group was a Sony black and white video camera (MODEL # XCM-38). This camera measured light intensities with a CCD element. Since the camera was not designed specifically for IR, filters were needed to block out visible light, while allowing IR wavelengths to pass.

The signal received by the computer from the camera is converted to a 512 x 384 array of integers ranging from 0 to 255, these correspond to certain light intensities. The values in the array are grouped into 16 preset divisions which are represented by 16 pseudocolors. The image can then be displayed on a monitor or saved for later analysis.

Filters

Although an infrared digital camera would have been ideal, this was not available. Therefore, filters were required to block undesired light waves. Various natural and glass filters were investigated and tested to determine which allowed the greatest detail and accuracy in image acquisition.

Computer System

Either an expert system or a neural network may be used to analyze the data provided by the IR camera. The information can be compared with data of a known healthy plant in order to determine its state of health.

Miscellaneous

Other items used during the design and fabrication of the sensor included different light sources, black felt, and a leaf holder.

VARIOUS DESIGNS

1.1 micron Filter

The first design used is shown in Figure III-2. In this design the filter used was a silicon disc, with the intrinsic property of blocking all light below 1.1 microns (a high-pass filter). The leaf remained on the plant and the background was not controlled. This produced no useful image due to the limited sensitivity of the camera in the IR range.

In order to overcome this obstacle, a direct light source was added. This produced an image, however, these images possessed higher reflectance intensities at the center of the leaf and lower at the perimeter. This was due to the curvature of the leaf and the position of the light source. Again, the background was not controlled and the leaf was attached to the plant. Even with various backgrounds, a "hot" spot was displayed in the center of the leaf due to a concentrated amount of light at that point. One possible solution was to use diffuse light, another was to investigate the effect of different filters.

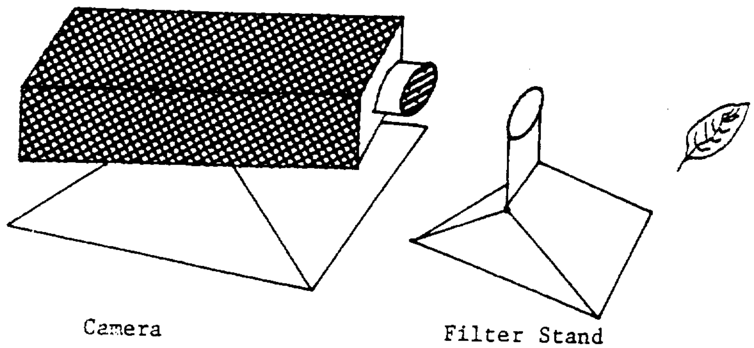

FIGURE III-2. Initial design of sensor.

0.66 micron Filter

The first modification to the original design was to replace the silicon disc with a Corning Glass Filter. This filter blocked all light below 0.66 microns. It produced images using only the ambient light which was produced by the fluorescent lights in the room. Although the image contained good details, the curvature and position of the plant continued to influence the image. Higher intensities of light were located at the top of the leaf as a result of the reflectance from the overhead lights. Although some differences could be seen between healthy and chlorotic regions, the intensity discrepancies due to curvature still dominated the image. Therefore, the leaf was removed from the plant and placed on a white background, however, this background reflected light producing additional interference.

0.66 micron Filter with Holder

To overcome both the curvature and background noise problems, a leaf holder was constructed to hold the leaf flat, provide it with an absorbing background, and shelter it from the direct light from the fluorescent source (Figure III-3). The design of the holder utilized an overhang that shielded the leaf from any direct light and produced a constant light intensity distribution. Black felt was used as an absorbing background to minimize noise interference. A plexiglass window was used to hold the leaf flat against the felt.

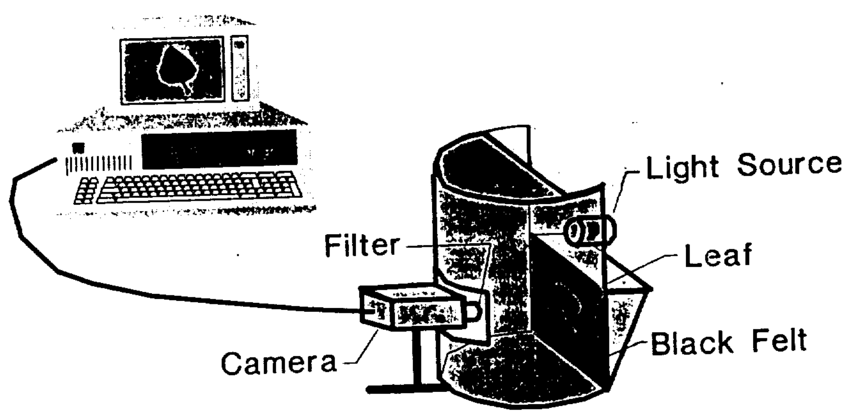

Figure III-3. Leaf holder.

With the leaf holder and the 0.66 micron filter, more consistent results were produced. The black felt absorbed the majority of the IR producing a "black" background in the image. The details noticed on the leaf were no longer a result of curvature and represented unhealthy portions in the leaf.

The data from the camera was numerically analyzed, using the computer program shown in Appendix B. This program printed the full array of integers for inspection. The arrays produced showed significant intensity gradients between the leaf and background, as well as between healthy and unhealthy regions on the leaf. However these gradients were not very large. For the background the integers ranged from 15 - 20, while the leaf ranged from 30 - 45.

Filter Combination

To produce greater detail, the bandwidth of interest was reduced to 0.7 - 1.3 microns (Figure III-4). The filter which achieved this bandwidth, however, possessed a second pass band at 0.36 microns. To exclude this second peak, the 0.66 micron high-pass filter was used in conjunction with the bandpass filter. The combination of filters required an additional light source to achieve good image definition. In order to disperse this light evenly, a white poster board was attached to the holder. A 150 Watt light source was directed onto the board which then reflected the light towards the leaf. The plexiglass was removed due to reflection. This dispersed light was reasonably uniform across the leaf. The modification of the second filter produced greater detail by making the gradients between the healthy and unhealthy parts of the leaf more noticeable.

Upon inspection of the data array, the numbers were found to be reversed from the expected values; that is, the healthier regions had higher values whereas the unhealthy regions had lower numbers. It is believed that the reversal was due to factors which may be affected by the increased light intensity, such as water deficiencies and nutrient deficiencies. Another cause may have been heating of the protoplasm in the leaf due to IR absorption. These changes may have been displayed in the image due to the large IR range being detected. In order to remedy this, 15, 25, 40, 60 ,75 and 150 Watt light sources were tested, however, none of these sources produced the intensity distribution expected.

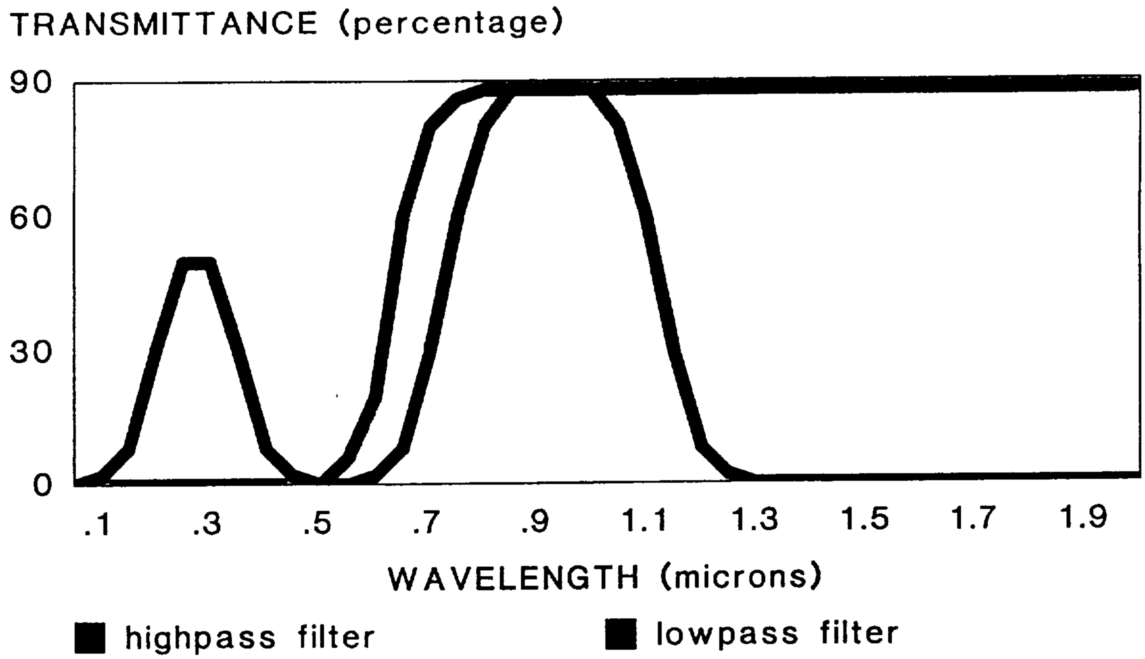

Figure III-4. Filter bandwidth.

0.671 Narrow Band Filter

To restrict the number of factors which would change the absorbance and reflectance data, a smaller bandwidth was used. A new filter of 0.671 micron bandwidth was used.

This filter produced better results, a healthy leaf was nearly invisible on the felt background, while unhealthy leaves were very visible. Furthermore, damaged areas were able to be seen in great detail allowing for a determination of the health status of the plant by a neural network. Integers were between 40 - 50 for the background, 50 - 55 for the healthy areas, 90 - 155 for chlorotic areas, and 155 - 190 for necrotic regions.

CONCLUSION

Using the 0.671 micron filter to analyze the chlorophyll content of the plant tissue, it was possible to obtain data which clearly displayed unhealthy regions on the leaf. Furthermore, the numerical values of the data would enable an evaluation of the leaf's health by a computer analysis such as a neural network, or possibly an expert system. An increase in the pixel intensity corresponds to a decrease in plant health.

As this IR imaging system may adequately detect plant health while remaining remote, it has promising applications in a CLLSS. Although, this system was limited to a single leaf with a controlled background and illumination, and it neglected curvature and positioning of the leaf, improvements could be made such that the sensor may operate under the normal conditions of the BPC.

RECOMMENDATIONS FOR FUTURE DEVELOPMENT

At this point, the sensing system has many self-imposed limitations that have been discussed. In order to further advance this type of sensor it is necessary to develop a computer system capable of analyzing a three-dimensional image, differentiating background noise and object periphery, and increasing light sensitivity. Three-dimensional analysis would account for leaf curvature and, in conjunction with background insensitivity, may correct for shadows and overlap. Light sensitivity would allow well defined images to be obtained at ambient light levels, thereby removing the need for additional light sources.

Three-Dimensional Analysis

By enhancing the computer system (neural network), it may be possible to enable the system to deal with three-dimensions. This would remove the need to hold the leaves flat. The system would be able to distinguish differences in light intensity due to surface contours from those due to plant disease. Neural networks would be ideal for this procedure because of their abilities in pattern recognition.

Background Distinctions

In examining plants within the BPC, there will be the inevitable problems with foliage overlap, shadows, and background noise due to chamber walls and light sources. Therefore, it is important that the image of interest may be distinguished adequately in order to perform an analysis with as few errors as possible. To accomplish this it will be necessary for the analyzing computer system to distinguish between the background and the image of interest. Again, this may be accomplished by a neural network.

Camera Sensitivity

In this project a regular black and white SONY camera was used in the digitizing system. However, as this camera was not specifically designed to operate in the IR range its sensitivity to the wavelengths of interest was limited. This required the need for an additional light source and IR filters. If an IR camera were to be used it is possible to eliminate the need for both filters and additional light sources. This would allow the system to operate at ambient light levels, thereby reducing the effects of glare caused by increased lighting as well as the residual heat.

System Efficiency

The ability of the system to sense reduced health must be examined, as it is essential that any problems can be discovered before they become irreversible. To accomplish this it would be necessary that a deficiency be slowly induced into a crop to discover at what point this system would reveal a problem.

REFERENCES

Corning Glass Works. (1960). Catalogue. Corning, New York.

Raven, Peter H. & Johnson, George B. (1986). Biology. Missouri: Times Mirror.

Schalkoff, R. J. (1989). Digital Image Processing and Computer Vision. Wiley Inc. New York.

Schwab, Wilhelm. University of Florida. Department of Aerospace Engineering, Mechanics and Engineering Science. Personal Interview. 27 January 1990.

Appendix A

Multi-component Sensor

Multicomponent System

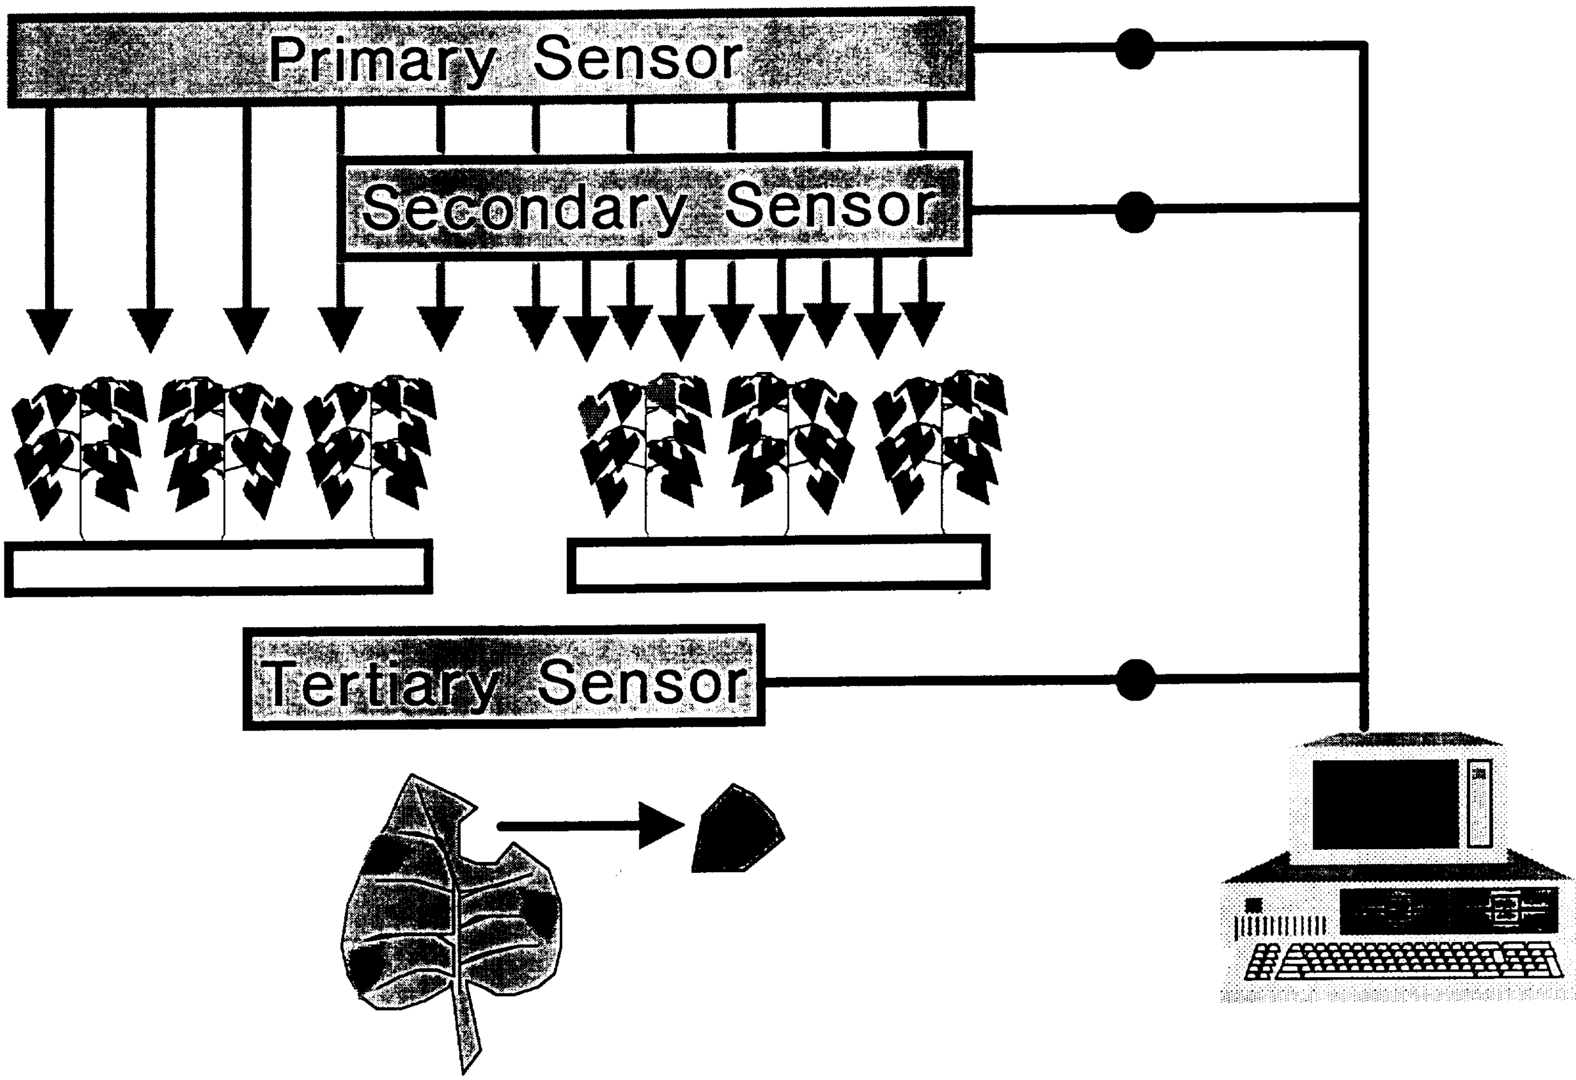

Appendix B

Data Analysis Program

```
/*****************************************************************/
/********** This program is used to view the digitized numerical   ******/
/********** data in an image file that is produced during the      ******/
/**********                infrared imaging process                ******/
/********** -------------------------------------------------------- ******/
/********** The image file is an unformatted list of unsigned      ******/
/********** characters, which represent the numerical data for     ******/
/********** which represent the numerical data for a specific pixel ******/
/**********               on the computer screen.                  ******/
/********** -------------------------------------------------------- ******/
/********** This program reads the image file, converting the      ******/
/********** characters into integer values.  It then writes        ******/
/********** the numbers into 24 files, each 16 columns by 512      ******/
/********** rows.  These files are arranged so that when placed    ******/
/********** side by side, the result is a numerical representation ******/
/**********                  of the image.                         ******/
/********** -------------------------------------------------------- ******/
/********** This becomes an array that is 384 by 512.              ******/
/*****************************************************************/

#include "stdio.h"
#include "stdlib.h"

#define ROW 512
#define COLUMN 384

void main(void)
{
/***********************Declare All The Variables ******************/

   FILE *fp, *fp1, *fp2, *fp3, *fp4, *fp5, *fp6, *fp7, *fp8, *fp9, *fp10;
   FILE *fp11, *fp12, *fp13;

/***** i & j = counter variables that maintain the current position ******/
   unsigned long int i, j;

/***************************k, m, & n = counters    ******************/
   int k, m, n;

/****vector[] = one dimensional array used to take a segment       ******/
/*********out of the array and place it in the appropriate file    ******/
   int vector[16];

/************* answer = the variable that is used to read the file ******/
   unsigned char answer;
```

```
/**************    Open The Image and Placement Files    ************/

/**************    The original image file    ***********************/
/*******************from which the program will be reading    ******/

  if ((fp=fopen("leaf_a.img", "rb"))==NULL)  {
    printf("Open Error!");
    exit(1);
  }

/******These are the files into which the numbers will be placed    ******/

  if((fp1=fopen("file_a", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
  }
  if((fp2=fopen("file_b", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp3=fopen("file_c", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp4=fopen("file_d", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp5=fopen("file_e", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp6=fopen("file_f", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp7=fopen("file_g", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp8=fopen("file_h", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp9=fopen("file_i", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
```

```
if((fp10=fopen("file_j", "w"))==NULL) {
   printf("Cannot Open File!\n");
   exit(1);
}
if((fp11=fopen("file_k", "w"))==NULL) {
   printf("Cannot Open File!\n");
   exit(1);
}
if((fp12=fopen("file_l", "w"))==NULL) {
   printf("Cannot Open File!\n");
   exit(1);
}
if((fp13=fopen("file_m", "w"))==NULL) {
   printf("Cannot Open File!\n");
   exit(1);
}

rewind(fp);

/***********    Reading The Image File     *************************/

n=0;
j=0;
k=0;
m=0;

do {
   printf("Row %i\n", n);

   for(i=j; i<j+24; i++) {

      for(k=0; k<16; k++) {
         if(fread(&answer, sizeof(unsigned char), 1, fp)!=1) {
            printf("Read Error!\n");
            exit(1);
         }
         vector[k] = (int) answer;
      }
```

```
/***********Writing The Numerical Data Into The Proper File       ******/

        if(i==j)  {
            for(m=0; m<16; m++)  {
             fprintf(fp1, "%4i ", vector[m]);
            }
            fprintf(fp1, "\n\n");
        }
        if(i==j+1)  {
            for(m=0; m<16; m++)  {
             fprintf(fp2, "%4i ", vector[m]);
            }
            fprintf(fp2, "\n\n");
        }
        if(i==j+2)  {
            for(m=0; m<16; m++)  {
             fprintf(fp3, "%4i ", vector[m]);
            }
            fprintf(fp3, "\n\n");
        }
        if(i==j+3)  {
            for(m=0; m<16; m++)  {
             fprintf(fp4, "%4i ", vector[m]);
            }
            fprintf(fp4, "\n\n");
        }
        if(i==j+4)  {
            for(m=0; m<16; m++)  {
             fprintf(fp5, "%4i ", vector[m]);
            }
            fprintf(fp5, "\n\n");
        }
        if(i==j+5)  {
            for(m=0; m<16; m++)  {
             fprintf(fp6, "%4i ", vector[m]);
            }
            fprintf(fp6, "\n\n");
        }
        if(i==j+6)  {
            for(m=0; m<16; m++)  {
             fprintf(fp7, "%4i ", vector[m]);
            }
            fprintf(fp7, "\n\n");
        }
```

```
        if(i==j+7)  {
            for(m=0; m<16; m++)  {
             fprintf(fp8, "%4i ", vector[m]);
            }
            fprintf(fp8, "\n\n");
        }
        if(i==j+8)  {
            for(m=0; m<16; m++)  {
             fprintf(fp9, "%4i ", vector[m]);
            }
            fprintf(fp9, "\n\n");
        }
        if(i==j+9)  {
            for(m=0; m<16; m++)  {
             fprintf(fp10, "%4i ", vector[m]);
            }
            fprintf(fp10, "\n\n");
        }
        if(i==j+10)  {
            for(m=0; m<16; m++)  {
             fprintf(fp11, "%4i ", vector[m]);
            }
            fprintf(fp11, "\n\n");
        }
        if(i==j+11)  {
            for(m=0; m<16; m++)  {
             fprintf(fp12, "%4i ", vector[m]);
            }
            fprintf(fp12, "\n\n");
        }
        if(i==j+12)  {
            for(m=0; m<16; m++)  {
             fprintf(fp13, "%4i ", vector[m]);
            }
            fprintf(fp13, "\n\n");
        }
    }
    j=j+1;
    n=n+1;
  }  while(n!=ROW);

  fclose(fp);
}
```

```
/*****************************************************************/
/*********This is the second part of the program that writes the ******/
/**********          digitized image file into useable arrays.   ******/
/*****************************************************************/

#include "stdio.h"
#include "stdlib.h"

#define ROW 512
#define COLUMN 384

void main(void)
{

  FILE *fp, *fp14, *fp15, *fp16, *fp17, *fp18, *fp19, *fp20, *fp21, *fp22;
  FILE *fp23, *fp24;
  unsigned long int i, j;
  int k, m, n;
  int vector[16];
  unsigned char answer;

  /* open the image file */

  if ((fp=fopen("leaf_a.img", "rb"))==NULL)  {
    printf("Open Error!");
    exit(1);
  }
     if((fp14=fopen("file_n", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp15=fopen("file_o", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp16=fopen("file_p", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp17=fopen("file_q", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
   if((fp18=fopen("file_r", "w"))==NULL)  {
     printf("Cannot Open File!\n");
     exit(1);
   }
```

```
if((fp19=fopen("file_s", "w"))==NULL) {
  printf("Cannot Open File!\n");
  exit(1);
}
if((fp20=fopen("file_t", "w"))==NULL) {
  printf("Cannot Open File!\n");
  exit(1);
}
if((fp21=fopen("file_u", "w"))==NULL) {
  printf("Cannot Open File!\n");
  exit(1);
}
if((fp22=fopen("file_v", "w"))==NULL) {
  printf("Cannot Open File!\n");
  exit(1);
}
if((fp23=fopen("file_w", "w"))==NULL) {
  printf("Cannot Open File!\n");
  exit(1);
}
if((fp24=fopen("file_x", "w"))==NULL) {
  printf("Cannot Open File!\n");
  exit(1);
}
  /* reading the image file  */
n=0;
j=0;
k=0;
m=0;

do {
  printf("Row %i\n", n);

  for(i=j; i<j+24; i++) {

  for(k=0; k<16; k++) {
      if(fread(&answer, sizeof(unsigned char), 1, fp)!=1) {
        printf("Read Error!\n");
        exit(1);
      }
      vector[k] = (int) answer;
  }

     if(i==j+13) {
        for(m=0; m<16; m++) {
           fprintf(fp14, "%4i ", vector[m]);
        }
        fprintf(fp14, "\n\n");
     }
```

```
if(i==j+14) {
    for(m=0; m<16; m++) {
        fprintf(fp15, "%4i ", vector[m]);
    }
    fprintf(fp15, "\n\n");
}
if(i==j+15) {
    for(m=0; m<16; m++) {
        fprintf(fp16, "%4i ", vector[m]);
    }
    fprintf(fp16, "\n\n");
}
if(i==j+16) {
    for(m=0; m<16; m++) {
        fprintf(fp17, "%4i ", vector[m]);
    }
    fprintf(fp17, "\n\n");
}
if(i==j+17) {
    for(m=0; m<16; m++) {
        fprintf(fp18, "%4i ", vector[m]);
    }
    fprintf(fp18, "\n\n");
}
if(i==j+18) {
    for(m=0; m<16; m++) {
        fprintf(fp19, "%4i ", vector[m]);
    }
    fprintf(fp19, "\n\n");
}
if(i==j+19) {
    for(m=0; m<16; m++) {
        fprintf(fp20, "%4i ", vector[m]);
    }
    fprintf(fp20, "\n\n");
}
if(i==j+20) {
    for(m=0; m<16; m++) {
        fprintf(fp21, "%4i ", vector[m]);
    }
    fprintf(fp21, "\n\n");
}
if(i==j+21) {
    for(m=0; m<16; m++) {
        fprintf(fp22, "%4i ", vector[m]);
    }
    fprintf(fp22, "\n\n");
}
```

```
        if(i==j+22) {
            for(m=0; m<16; m++) {
                fprintf(fp23, "%4i ", vector[m]);
            }
            fprintf(fp23, "\n\n");
        }
        if(i==j+23) {
            for(m=0; m<16; m++) {
                fprintf(fp24, "%4i ", vector[m]);
            }
            fprintf(fp24, "\n\n");
        }
    }
    j=j+1;
    n=n+1;
  } while(n!=ROW);

  fclose(fp);
  exit(1);
}
```

IV. ROBOT ARM CONTROL AND PLANT HEALTH RECOGNITION USING NEURAL NETWORKS

Prepared by:

Judd Bishop
Alan Clayton
Tom Good
Richard Kern
Walt Marchand

TABLE OF CONTENTS

SUMMARY

In a closed-loop life support system (CLLSS) it is desirable to have an automated system monitoring and maintaining food crops growing in the Plant Growth Unit (PGU). One method of accomplishing this is by applying the pattern recognition properties of neural networks to determine plant health. The first phase of the project consisted of using the Back-propagation algorithms to train and control a robot arm, and in the process gain basic knowledge and understanding of neural networks in order to tackle the second phase of the project. In phase two, Competitive Learning and Back-propagation algorithms were used to try to determine the health of a plant leaf using data supplied by the Plant Health Group.

INTRODUCTION

Neural networks are today where conventional computers were in the 1950's. A whole new class of devices can be built out of semiconductors which are just as powerful as computers but follow entirely different rules (Roberts, 1988). The kinds of problems best solved by neural networks are those that people solve well: association, evaluation, pattern recognition, and working with obscure or incomplete data. Neural networks also handle problems that are difficult to compute and do not require perfect answers, but do require fast, good answers, such as industrial controllers or real-time robotics applications.

All neural networks have 3 things in common: distributed processing elements (neurodes), the connections between them (network topology), and the learning rule. Neural networks ideally operate as parallel distributed processors, all operating simultaneously to converge on a solution. They are not specifically programmed, but rather trained, a process which involves modifying the connection weights until the system converges, or finds a minimum-energy state, based on gradient descent methods.

The main purpose of this project is to try to adapt neural network theory to CLLSS applications, or more specifically, to plant health. In order to do this, sequential programs are used to simulate the parallel distributed processing of a neural network. This project utilizes a Parallel Distributed Processing software package (McClelland and Rumelhart) in conjunction with 'C' programs written by the authors.

THEORY

Artificial neural networks are loosely based upon a simplified model of the human brain. This simplified model is organized into networks of simple neuron-like processing elements that pass signals to each other. Since they operate similarly to the human brain, they demonstrate behavioral characteristics analogous to human intelligence. However, they are not good at solving the kind of problems at which conventional computers traditionally excel, such as precise numerical computations.

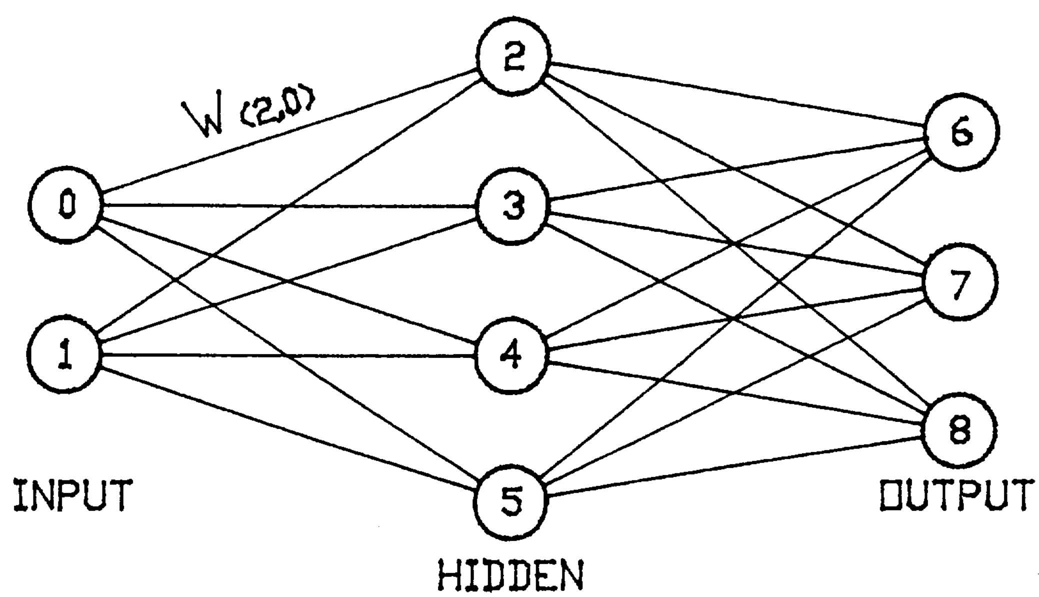

Figure IV-1. Back-propagation network structure.

Figure IV-1 shows a graphic representation of a simple three-layer BP network, identified as a 2:4:3. Each circle represents a processing element, or neurode. Each arrow represents a connection and its associated weight with each column of neurodes being called a layer. The first layer is the input layer, the last is the output layer, and optionally, those in between are called hidden layers. Each neurode, which may be both a sending and a receiving neurode, is connected to every neurode in the next layer, but only to those neurodes. The value exiting a neurode along an arrow, the activation level of the neurode, is typically a value between -1 and +1, and is sent to the receiving neurode. The values received by a neurode are processed (more on this below), and a new activation value is calculated and sent simultaneously along all outgoing connections. There are two special cases of activation values. First, the input neurodes' activation values are simply the input values. Second, the output neurodes' activation values represent the completed transformation (answers).

In this example, neurodes are numbered from 0 to 8. The variable w(r,s) identifies the weight of the connection between neurode r, the receiving neurode, and neurode s, the sending neurode. For example, the weight of the connection between neurodes 0 and 4 is written w(4,0). There is only one input and output layer, but there can be n optional hidden layers. Each layer can have an arbitrary number of neurodes. There appears to be some relationship between the number of layers, the number of neurodes in the layers, and the quality of the solution, but the relationship is presently unknown, except for some special cases.

The feed-forward operation of the network is straightforward. Feed-forward means that all data flow is from left to right, and there are no feedback loops. The input layer neurodes can accept analog or discrete values. The output of an input layer neurode is exactly equal to the input. The activation value of a sending neurode is sent along its connections to receiving neurodes, where these multiply each activation by its associated weight and sums them, thus giving the total input. The input is operated on by a transfer function, yielding the new activation value for that neurode, which is then sent along its connections to the next layer. The transfer function must be continuous and possess a derivative at all points. The function used here, and in most implementations, is a sigmoid function, shown in Figure IV-2.

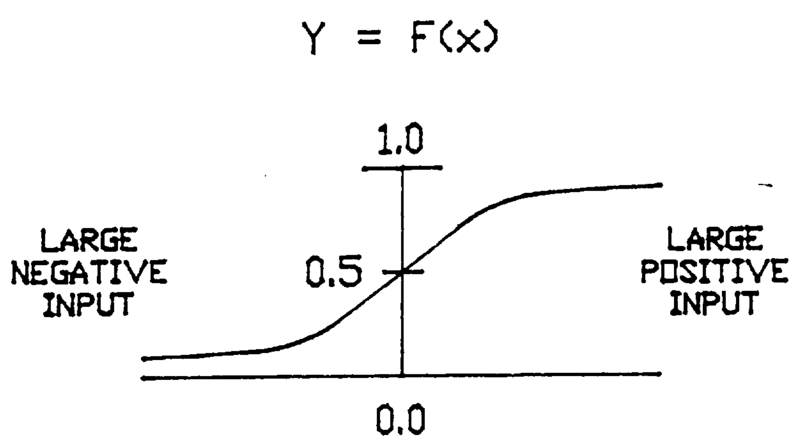

Figure IV-2. Sigmoid Transfer Function.

The sigmoidal function is stated as:

$$f(x) = 1/(1+(e^{-net}))$$

This assures an activation value between 0 and 1.

The feed-forward operation in effect transforms an input state to an output state. For a meaningful transformation, there must be some method of determining and modifying the values of the weights of the connections. This method is called the learning rule, and is used to train the network. The learning rule used for this project is the generalized Delta rule given by Rumelhart(1986).

Back-propagation assumes that there exist some arbitrary number of training patterns, each made up of a set of input values and their associated output values, and that the output values, called targets, are correct for the inputs. Since the correct outputs are known, this is called supervised learning. Training is a process in which the weights of the connections are changed in response to a given input and output pair according to the learning rule.

To summarize the operation of the network during training, the inputs are presented to the network. The resulting activation flows from the input layer to the output layer. The output neurodes' values are then compared to the target values, and an error term is calculated. Starting with the output layer, the weights are adjusted, using the Delta rule, layer by layer, propagating each layer's error back to the previous layer, computing weight changes along the way. The weights are not actually changed until after the error has been propagated back to the previous layer. Once the activation has flowed forward and the error has propagated backward, one iteration is complete, and the network is ready for the next pattern. When all the patterns have been processed once, one epoch of training has been accomplished.

The Delta rule is expressed by the equation:

$$W(new) = W(old) + Beta * E * X / |X|^2$$

where W is the weight vector before and after the weights are adjusted, Beta is the learning rate, E is the error term, and X is the input vector. After the feed-forward of the input signals, all the information needed for application of the Delta rule is available. The error term E is easily calculated by the equation:

$$E = (target - output).$$

After the Delta rule is applied to the output layer, the matter becomes more complicated. The problem is that the correct activation value for a hidden layer neurode is not known. Since the error in each output neurode can be spread over its input vector, the exact proportion of error attributable to each connection weight can only be approximated. The error term for a hidden layer neurode is:

$$e(i) = f'(I) * [summation\ w(ij) * E(j)]$$

where e(i) is the error in the i^{th} hidden layer neurode, and the sum is taken over j, where j is the j^{th} output layer neurode, and f'(I) is the derivative of the sigmoid function for the hidden layer neurode for the net input it received. It can be shown that

$$f'(I) = f(I) * (1-f(I)).$$

By applying the derivative of the transfer function, only relatively small changes are made in the weights coming into a hidden layer neurode, by keeping these error terms smaller. Now that the error term is available, the Delta rule can be applied and the weights adjusted.

This process is repeated for every hidden layer. When the weights connecting the input layer to a hidden layer are adjusted, the back-propagation of the error is completed for this iteration. The process continues until some desired value of total error is reached, whereupon the network is said to be trained.

Once the network is trained, the weights are frozen and the network can be used for its intended application by applying input signals and allowing them to flow through to the output state. This operation is typically very fast, even in sequential simulation. Transformations which in a conventional program are too computationally complex for real-time applications can be accomplished by using a module containing a trained neural network.

ROBOT ARM CONTROL

The Robot

The robot arm used is a small Pro-Arm with 5 degrees of freedom. It's range of motion is 240 degrees about the base, 144 degrees about the shoulder, 120 degrees about the elbow, 180 degrees for wrist pitch, 360 degrees for wrist roll. The gripper can open about 4 inches. For the purposes of this project, the wrist pitch and wrist roll are held constant. The object to be grasped is assumed to be the stem of a soybean plant about 2 inches above the growing surface.

The robot has its own microprocessor that receives instructions from a control program through a parallel port. An instruction is typically a character followed by 6 integers i[1] to i[6] where $-2000 <= i[n] <= 2000$, representing the number of steps available to a particular stepper motor. There are six stepper motors, one for each joint. All motors operate simultaneously, but can accept a zero motion command.

Program

The authors developed a program in Turbo C under MS-DOS 3.x that allows real-time interactive control of the robot arm. This was necessary since the neural network requires accurate input and output pairs for supervised training. A base of plywood was built, with the robot centered on it, accommodating the two-dimensional range of motion of the robot arm. Grid overlays were then mapped over the surface, one in polar and one in rectangular coordinates. Figure IV-3 shows the base overlaid with the polar and the Cartesian coordinate grid. Using the control program, sets of values were generated for a polar coordinate (r,0) associated with the number of steps required (base, shoulder, elbow) to move the gripper to the proper position. The range for the base axis was limited to +- 90 degrees.

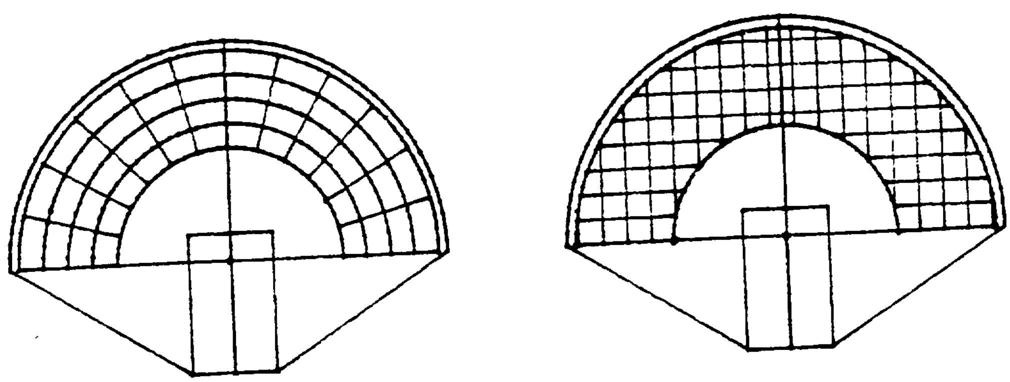

Figure IV-3. Robot base with overlays.

Back-propagation requires all analog input to be normalized between 0 and +1. Equations to normalize input and output were coded in a module of the control program to generate the training sets for points on the polar coordinate grid. A program to demonstrate the operation of the network was also written. It allows interactive entry of a desired rectangular coordinate, calculates the corresponding polar coordinate, normalizes all four inputs, and passes the values to a module simulating the feed-forward operation of an artificial neural network. The output values are then rescaled, and sent to the robot arm, along with the commands required for the gripper.

BP Program

The training is done using the back-propagation program supplied with Rumelhart and McClelland's Parallel Distributed Processing. Network architecture is nearly unrestricted in this program, and it allows the user to specify and display any or all variables. It reads in 5 files to set up a training session. The .net file defines the network architecture; the .pat file contains the training sets; the .str file sets variables; the .wts file contains any weights saved from a previous session; and the .tem file specifies the screen layout. To test the network's ability to generalize solutions to inputs it has not been trained on, a pattern is run through a single feed-forward operation, and the error is calculated. An epoch of training consists of one iteration of back error propagation for each pattern in the training set. Training is computationally intensive and normally takes hours or days, depending on the size of the network, the number of epochs, and the speed of the computer.

Networks Trained

The first preliminary step was to design and train a network to transform two dimensional cartesian coordinates to polar coordinates. Networks for polar to cartesian coordinate transformations were also developed. The first network designed was a (2:4:2) network. The notation gives the number of neurodes in each layer, and has the form (i:j:k) where i=input, j=hidden, and k=output neurodes. This network was given the task of transforming cartesian coordinates to polar coordinates. This network was trained for 1,000 epochs giving a tss (total sum of squares) of 0.0227, with no bias included. A (2:3:2) and a (2:3:3:2) network were developed to transform polar coordinates to cartesian coordinates. The (2:3:2) network was trained for 4000 epochs and achieved a tss of 0.0165. The (2:3:3:2) network was trained for 11,245 epochs, yielding a tss of 0.0001. This level of error produced output values that differed from the targets by no more than 0.006. Both of these networks were presented with 11 target patterns. No test patterns were evaluated for generalization at this point because evaluation methods were not known.

The acquired knowledge was then applied towards the construction and training of neural networks that mapped two dimensional polar and/or cartesian coordinates into robot units corresponding to base, shoulder, and elbow positions. The first network mapped polar coordinates to robot units. This was done by constructing a (2:5:5:3) network. This network was trained for 6000 epochs with 10 target patterns and yielded a tss of 0.0003. Next, a (2:4:4:3), (3:4:4:3), and a (4:8:8:3) network were constructed and trained. The networks had difficulty with inputs and outputs near zero, so the origin of the coordinate system was changed such that all points were within the first quadrant and no point was described as a zero. The (4:8:8:3) had inputs of both polar and rectangular coordinates (x,y,r,0). The (3:4:4:3) had inputs of (x,y,r) and the (2:4:4:3) had inputs of (x,y) only. This was done to compare how the transformations are affected by different inputs and varying numbers of hidden elements. The (2:4:4:3) ran 20,000 epochs with a tss of 0.0026, the (3:4:4:3) ran 20,000 epochs with a tss of 0.0001, and the (4:8:8:3) ran 30,000 epochs with a tss of 0.00001. All of these networks had 11 training patterns. These networks were also presented with test patterns to determine the accuracy of learning. The best transformations and generalizations were produced by the (4:8:8:3) network. Individual output values varied from the target with a maximum deviation of 0.001 on the trained patterns. The maximum deviation on a generalized output was 0.091, but almost all other generalized outputs were within 0.02.

Other networks that were developed were (3:5:3), (3:8:3), and (3:32:3). These networks accepted inputs in terms of (x,y,r) and were trained under 30,000 epochs. This was done to compare the effect of the hidden layer size on the learning rate and the final convergence state. There seems to be a minimum number of neurodes, which differs with the inputs and outputs, that will produce an acceptable error level. Two hidden layers generally perform better than one.

To measure the rate of convergence of the tss error, four networks were selected and reset to randomize the weights on the interconnections. Each network was reset four times to ascertain whether the original weights affected how well the network trained. The networks were then run for 3000 epochs. The tss was recorded at initialization and again when the networks were trained. All of the networks converged to a tss of about 0.04 at 500 epochs, the rate of convergence slowed dramatically. Since the network requires on the order of 10^4 epochs to reach an error level corresponding to a few stepper motor units, the rate of convergence was determined to be a relatively minor factor. It was found that the networks that transformed x and y coordinates to robot units did not achieve the level of accuracy shown when polar coordinates were included in the input. The (x,y) transformations, regardless of the network architecture, showed slightly worse performance on patterns they were trained on, and a much worse ability to generalize. The following networks were trained and tested: (2:5:3), (2:32:3), (2:3:3:3), (2:3:5:3), (2:4:4:3), (2:8:8:3), (2:16:16:3), and (2:32:32:3).

Results

The following results are for the (4:8:8:3) network at 30,000 epochs. Figure IV-4 shows some of the 20 training patterns.

Pat	x	y	r	θ	Base	Shldr	Elbow
06	.1	.08333	.75898	.18598	.97790	.275	.205
16	.1	.25000	.68959	.41319	.80010	.275	.205
22	.3	.33333	.66304	.58865	.71120	.400	.100
42	.1	.66667	.33696	.58865	.33560	.400	.100
56	.1	.91667	.24102	.18598	.08890	.275	.205

Figure IV-4. Training patterns.

Figure IV-5 shows target and actual values for the patterns shown above, along with the values for some test patterns.

	Target			Actual		
Pat	Base	Shldr	Elbow	Base	Shldr	Elbow
06	978	275	205	976	274	204
16	800	275	205	800	273	205
22	711	400	100	709	401	96
42	336	400	100	335	399	99
56	89	275	205	89	274	204
07	978	400	100	973	386	118
10	978	600	800	978	584	781
20	800	600	800	794	602	798
24	711	525	570	720	520	572
28	622	420	300	628	411	316
58	89	420	300	96	390	391

Figure IV-5. Target and actual values.

The team found that by giving the network more information by increasing the number of input elements, the network's ability to generalize was enhanced. The (4:8:8:3) network with 11 training patterns satisfactorily generalized test patterns, so to see if this generalization could be improved, the number of training patterns presented to the network was increased to 20. The network then ran 30,000 epochs and yielded a tss of 0.0001. This level of accuracy is considered excellent, and more than adequate for its intended purposes. The individual output error was 0.001 for trained patterns, showing that higher numbers of training patterns improved the performance of the network in this case. In addition, this network generalized much better.

PLANT HEALTH RECOGNITION

Problem Description

Phase two of the project involved using neural network techniques to determine the state of health of a plant leaf in an automated fashion. The goal was to develop a program that would read in a file containing data on the infrared reflectance of a leaf and prepare the data for use by a neural network which would be trained to distinguish between a healthy leaf and a sick leaf. Although the neural network was trained on a very limited number of unhealthy conditions, the purpose was to demonstrate the feasibility of the process . If the process can distinguish between healthy and unhealthy leaves, then phase 2 is successful.

Raw Data

The raw data preparation was done by the Plant Health Group. The basic idea makes use of the fact that chlorophyll-a absorbs light in the bandwidth 0.671 microns plus or minus 0.01 microns. The leaf is mounted on a black felt background, which is highly absorptive in this bandwidth, so that there is a definite edge. A diffuse 130 watt light source illuminates the leaf. An infra-red bandpass filter of 0.671 microns is mounted on a black and white digital camera, and the camera digitizes an image of 512 by 384 pixels, feeding it directly into a computer program that stores it in a file as unsigned characters.

Program

The first step was to extract usable characteristics from the image files. Neural networks have the ability to deal with incomplete and incorrect data. This ability is enhanced by increasing the number of inputs. The image files have a certain amount of noise, especially near the edges of the leaves. It was decided to find at least 10 significant characteristics to prepare as inputs. The following is a list of the characteristics selected as input data:

1. The number of positive slopes.
2. The number of negative slopes.
3. The average pixel intensity for pixels in the leaf.
4. The gradient of steepest ascent.
5. The gradient of steepest descent.
6. The average value for the 50 highest intensities.
7. The number of pixels in the VERY HEALTHY range.
8. The number of pixels in the FAIRLY HEALTHY range.
9. The number of pixels in the SICK range.
10. The number of pixels in the DEAD range.
11. The percent of leaf pixels > MAXPIXEL.
12. The percent of leaf pixels < MINPIXEL.

A number of other characteristics were considered and rejected because the difference between healthy and unhealthy values was either too small or inconsistent. All of the inputs selected showed significant and consistent differences between healthy and sick leaves. The program reads the file and extracts these raw values.

The next step was to normalize the raw input values between 0 and +1. Each input can be normalized individually, as long as the resolution of the data does not get lost in the process. The program normalizes the input data and stores it as a list of floating point numbers, along with an associated output value, in a pattern file format compatible with the BP program.

Training Set. The Plant Health Group provided 28 usable images. All of the images were taken in one sitting, so the lighting intensity was constant. Twelve of the leaves appeared healthy, and sixteen appeared unhealthy. The Plant Health Group looked at a color graphic computer display of each leaf which was consistent with the leaf's visual appearance. One of the leaves, however, showed a spot of some kind. That leaf was graded as suspicious. Arbitrary values were then assigned to each category of leaf. Healthy leaves were assigned 0.9, and unhealthy leaves were assigned 0.1. The suspicious leaf was not assigned a value. A training set was created, consisting of eight healthy leaves and twelve unhealthy leaves. The other eight leaves, including the suspicious one, were saved for testing.

<u>Networks Trained</u>. Figure IV-6 shows the ten networks designed and trained.

NETWORK	TSS(1000)	TSS(2000)
12:5:1	0.0972	0.0801
12:8:1	0.0430	0.0402
12:16:1	0.0423	0.0403
12:3:2:1	0.1707	0.1631
12:4:4:1	0.0031	0.0008
12:6:3:1	0.4166	0.2109
12:7:5:1	0.0225	0.0017
12:8:4:1	0.0055	0.0014
12:8:8:1	0.0797	0.0754
12:4:4:4:1	0.0020	0.0007

Figure IV-6. Networks trained.

Each network was trained for 2000 epochs. The tss error was recorded at 1000 and 2000 epochs. The 12:4:4:1 was chosen for further training and testing.

The 12:4:4:1 network was trained for a total of 20,000 epochs and achieved a tss of 0.00001, which is very small in this case. The target output of each leaf was compared to the actual output, and 19 of the patterns had an error of less than 0.001. One pattern had an error of 0.003. This means that if the target value was 0.1, the network typically graded it between 0.099 and 0.101.

Results. The test set consisted of the 8 leaves on which the network had not been trained. The test set was presented to the network, and output was produced for each pattern. The output is shown in Figure IV-7.

12:4:4:1 20,000 EPOCHS

PAT	SIGHT GRADE	NETWORK GRADE	SIGHT DECISION	NETWORK DECISION
0	.1	.761	SICK	SUSPICIOUS
4	.1	.099	SICK	SICK
15	.1	.083	SICK	SICK
21	.1	.099	SICK	SICK
23	.5	.898	SUSPICIOUS	HEALTHY
10	.9	.908	HEALTHY	HEALTHY
16	.9	.902	HEALTHY	HEALTHY
27	.9	.898	HEALTHY	HEALTHY

Figure IV-7. Results.

The network performed very well. For every leaf except leaf 0 and leaf 23, the network graded the leaf to within 0.017 of the grade assigned.

Leaf 0 was graded much closer to healthy than was expected with a value of 0.761. Visually, the leaf was healthy looking except for a brown edge. The value of 0.761 was considered far enough from 0.9 to assign a network grade of suspicious. Leaf 23, the one that had been labeled suspicious, was assigned a value of 0.898 by the network. Since the leaf looked completely healthy, this grade was accepted reasonable.

Competitive Learning

Early detection of plant health is a serious consideration when plants are used as a vital source in an, enclosed environment. The Plant Health Group has addressed the issue of how to obtain data relevant to plant health, while the concerns of the of the Neural Network group is how to process this data into useful information.

The Network group has approached this matter of data processing with two distinct types of learning mechanisms. The first model utilized back propagation in the pattern associator learning paradigm. The main attribute of pattern association is the presentation of pattern pairs (input and target patterns), which supports supervised learning. The second model was developed in accordance with the regularity detector learning paradigm. The outstanding trait of this model is the ability to recognize important features of a set of input patterns without being given target patterns, i.e., no forced learning. The competitive learning (CL) model is the mechanism that was chosen to fit this environment. Both of these learning paragons are especially useful for feature detection, but the following discussion will concentrate on the competitive learning mechanism and its role in plant health detection.

CL Theory. The primary goal of a competitive learning network is to discover prominent, general internal features as a result of a set of input stimulus patterns. There are three fundamental concepts involved with the competitive learning scheme.

1- A set of units are such that all are the same except for some randomly distributed parameter which makes each of them respond differently to a set of input patterns.

2- The strength of each unit is limited.

3- Provisions are made so that the units may compete for the privilege to respond to specific groups of the input patterns.

The basic architecture of the competitive learning mechanism is similar to a multi-layered pattern associator (See Figure IV-8.): the output of a unit in one layer is connected to all of the units in the layer immediately above it. The marked difference in the network's physical structure is that within each layer, clusters are formed by interconnecting units in groups. There are no interconnections between clusters in a mutually shared layer.

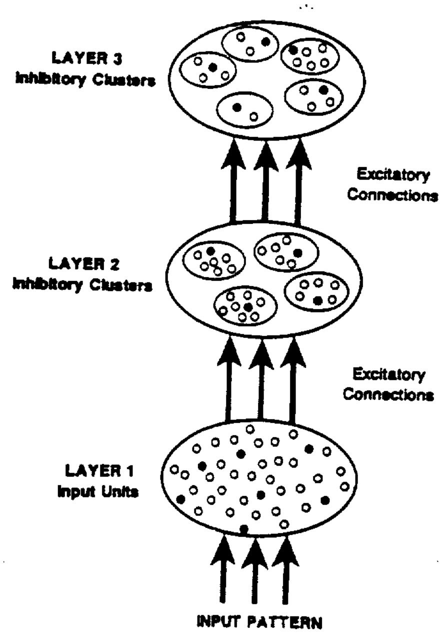

Figure IV-8. Pattern associator.

The connections between layers are excitory and connections within clusters are inhibitory. All elements in a cluster send inhibitory signals to all other elements in the cluster. This is what inspires the competitive nature of the network. Each unit competes with all other units in the cluster until one unit dominates. The stronger the response of the dominating unit, the further it inhibits all of its competing units. This continues until only one unit in the cluster is active in reaction to a particular input pattern. The implication here is that learning takes place only if a unit wins the competition with other units in its cluster.

A geometric interpretation of the competitive learning model gives a clearer insight to the overall process. For simplicity, a simple two layered network is illustrated in Figure IV-9. There are three units in the input layer and three units in the output layer. The three output units comprise a single cluster. The values of the output units are either 1 or 0, 1 being active, 0 being inactive.

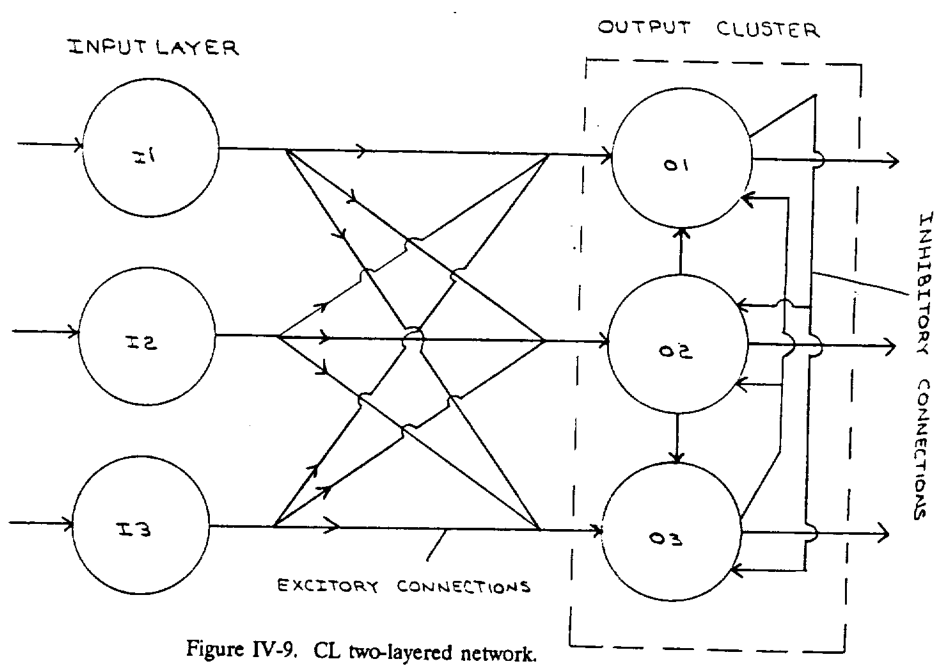

Figure IV-9. CL two-layered network.

There are eight input patterns presented to the network. These input patterns can be considered as vectors with three elements in the vector. This corresponds to 3 dimensional space. If the input patterns are normalized, then these vectors will have the same lengths where the ends will fall somewhere on the surface of the sphere. (See Figure IV-10.)

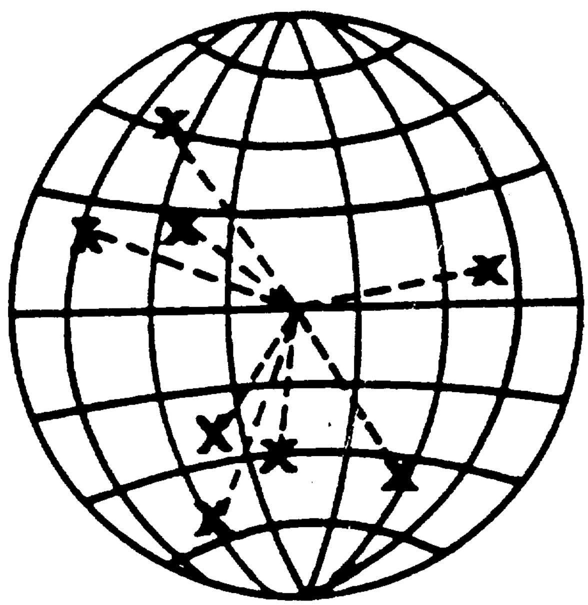

Figure IV-10. Sphere w/vectors.

The set of eight input patterns is presented to the system. For each specific input, the output units compete for the right to become active while disabling the other two. This is accomplished by adjusting the weights on the lines connecting the input and output units. It is useful to also represent these weights corresponding to each output unit as a unit vector. This allows the weight vectors to fall on the same sphere surface as do the input vectors.

When the input pattern is presented to the network, one unit in the output cluster will be slightly more active than the other units. This unit will then create an inhibitory signal that is slightly greater than the other two. The weights are readjusted to accommodate the dominant unit, thus giving that unit even stronger input values. This in turn generates a stronger inhibitory signal. This procedure continues until the system reaches stability where the weights are no longer readjusted and fully compliments the dominating unit. At this point the other two units are totally disabled.

This process occurs over a number of training runs for the entire set of input patterns. What is effectively occurring is that the weight vectors are repositioning themselves toward the center of a group of input patterns. (See Figure IV-11.) This allows an output unit to respond to an input pattern that is closest to the corresponding weight vector.

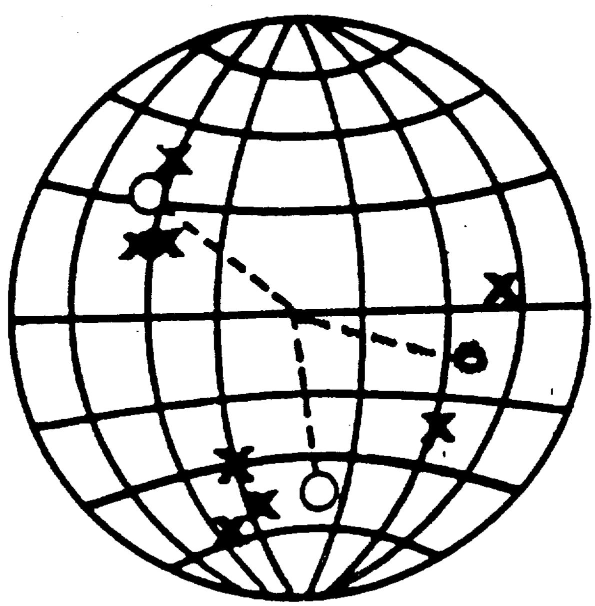

Figure IV-11. Sphere w/vector repositioning.

The groupings of the input patterns is an important characteristic of the competitive learning mechanism. Each cluster in a network creates N groups of patterns, one for each unit in the cluster. Each unit in the cluster attempts to capture approximately the same number of stimulus patterns for a group. Hence, for the example given, the network acts as a tertiary feature detector. One feature for each of the output units.

It should be noted that since this type of learning paradigm is completely spontaneous, correlations between the output of the network and the real-world scenario have to be made. This can be accomplished by intimate knowledge of the structured network and the physical system to which the network is modeled.

In an attempt to provide the Plant Health group with useful information in regards to image data, the network group designed an alternative network model using the competitive learning model. Due to limited time, this model sought to extract only the most basic features from the input data. The following is an overview of the actual design and the association between the output of the network and the image of the plant's leaf.

CL network and results. The goal of the designed competitive learning network was to obtain important basic features of the transition zones of the leaf. The transition zones were identified as those areas where the intensity either increased, decreased, or both. Locating a zone where the intensity either increased, or decreased, implied that the data set may be a representation of the beginning or ending edge of the leaf, respectively. If a zone was located where the intensity both increased and decreased within a pattern, then this was interpreted as a definite indication that the leaf was experiencing unhealthy conditions. The restriction here was that the data come from the interior edges of the leaf.

P0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
P1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
P2	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
P3	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1
P4	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
P5	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1
P6	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1
P7	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1
P8	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
P9	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1
P10	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
P11	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1
P12	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1
P13	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1
P14	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
P15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
P16	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
P17	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
P18	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
P19	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0
P20	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0
P21	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0
P22	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0
P23	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0
P24	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0
P25	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0
P26	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0
P27	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0
P28	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0
P29	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0
P30	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0
P31	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0
P32	0	0	1	1	1	1	1	1	1	1	1	1	1	1	0	0
P33	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0
P34	0	0	0	0	1	1	1	1	1	1	1	1	0	0	0	0
P35	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0
P36	0	0	0	0	0	0	1	1	1	1	0	0	0	0	0	0
P37	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0

INPUT PATTERNS

weights

unit_1			unit_2			unit_3		
7	9	0	10	4	12	12	6	4
0	8	5	11	1	1	9	9	0
2	0	0	1	10	5	5	2	8
9	9	10	4	4	5	4	2	10
11	9	5	10	11	1	7	4	10
9			2			1		

BEFORE TRAINING

weights

unit_1			unit_2			unit_3		
0	0	1	0	1	2	21	16	12
1	2	2	3	5	7	9	8	6
3	4	4	10	17	17	5	4	3
5	6	8	10	7	5	3	2	1
9	11	14	3	2	1	1	0	0
22			0			0		

AFTER TRAINING

Figure IV-12. Input patterns + weight matrix.

The network's architecture was rudimentary since it was to perform a simple task. The structure consisted of one input layer with 16 input units and one output cluster with 3 units. The network was presented with 37 input patterns. These input patterns were in a binary format. The system was trained on 40 epochs with a learning rate parameter of 0.05. The system quickly achieved stability and did not further modify the weight matrix. This was an indication that the system learned all that it could within 40 training runs. The patterns presented to the system and the weight matrix used for the connections are shown in Figure IV-12. Note the formation of the weight distribution on each output unit, and how the input patterns relate to the output of the units in the cluster. There are 3 possible outputs: 1-0-0 depicts an increase in intensity, 0-1-0 depicts an increase and decrease in intensity, 0-0-1 depicts a decrease in intensity.

For example, pattern P12 yields the following results for each of the output units: unit 1 total weight is 56, unit 2 total weight is 6, unit 3 total weight is 1. Since unit 1 has the largest weight factor, then unit 1 will have the strongest activation level and will inhibit both units 2 & 3. The output for the system will therefore be 1-0-0. This output relates to an increase in intensity across the data field.

The binary format was used to prevent the system from running without bounds. Raw data inputs caused the weight matrix to exceed limitations. This was in violation of the concept that the strength of each unit be limited. Hence, for each pattern set the first pattern would cause one output unit to overwhelmingly dominate the others. The dominating unit would remain in control regardless of the patterns that followed. This resulted in no learning. The binary format allowed fair competition among the output units.

The binary numbers were generated from the digitized image data in the following way. Raw data was taken from a row or column of the image matrix. The row or column of data was broken up into blocks 16 data points long. These 16 data points were then averaged. To smooth the data set, 10% of the average value was subtracted from each data point. Data points were then compared to the average value. If a data point exceeded the average value, then a 1 was used for input to the network. If the data point was less then the average value, then a 0 was used. These 16 data point conversions comprised one input pattern. The number of input patterns presented to the network depended on the length of the row, or column, divided by 16. Thus, a column that is 320 data points long would have 20 input patterns presented to the network.

CL future potential. The spontaneous response of competitive learning models adds to the realm of information processing. Allowing a system to capture its own features from its internal representation has great potential to solve complex problems that otherwise might not be possible. Although the competitive learning mechanism that was designed for processing plant health data is basic in nature, this does not mean that there are limitations to the information that can be obtained using this data. The limitations extend as far as the designer's creativity. Adding an unlimited number of units, clusters, and layers may reveal such detail about the current condition of the plant that early detection and prevention could easily be accomplished. For enclosed environments, such as CLLSS, this is a critical factor.

CONCLUSION

In a closed-loop life support system (CLLSS) it is desirable to have an automated system monitoring and maintaining food crops growing in the Plant Growth Unit (PGU). This could be accomplished with an expert system (ie. a traditional serial computer with an enormous data base stored in memory). This type of system, however, takes up so much memory in the computer that it gets very expensive to have a data base so large. In addition to the expense, the larger the data base is, the longer it takes to access the information desired. Therefore, it is believed that the best method of accomplishing this is by applying neural network theory to determine the health of a crop growing in the PGU. From a commercially available software package, Parallel Distributed Processing by McClelland and Rumelhart, Competitive Learning and Back-propagation algorithms were used to determine the health of a plant leaf. The results of creating the networks and processing the leaf data seem promising. The "answers" given by the neural network agreed with visual inspections of the leaves. With further investigation and network building, we feel a neural network system could be developed for use as a plant health detector.

GLOSSARY

NEURAL NETWORK - A computing system made up of a number of simple, highly interconnected processing elements which processes information by its dynamic state response to external inputs.

NEURODE - An artificial neuron in a neural network, consisting of a small amount of local memory and processing power. The output from this simple processing element is fanned out and becomes the input to many other neurodes.

CONNECTION - A signal transmission pathway between processing elements, loosely corresponding to the axons and synapses of neurons in a human brain, that connects the processing elements into a network. The strength of the connection is determined by its weight.

WEIGHT - Within a processing element, an adaptive coefficient associated with a single input connection. The weight determines the intensity of the connection, depending on the network's design and the information it has learned.

LEARNING LAW - An equation that modifies all or some of the adaptive coefficients (weights) in a processing element's local memory in response to input signals and the value supplied by the transfer function. The equation enables the network to adapt itself to examples of what it should be doing and to organize information within itself.

TRANSFER FUNCTION - A mathematical formula that, among other things, determines a neurode's output signal as a function of the most recent input signals and the weights in local memory.

SUPERVISED TRAINING - Trial and error process where the neural network is supplied with both input data and desired output data.

BACK-PROPAGATION NETWORK - A network comprised of neurodes in layers. This network is always hierarchial with a minimum of 3 layers (input, output, hidden) Each layer is fully connected to the next layer and there are no interconnections within a layer. The error from the output layer is propagated back through to the previous layer and the weights are adjusted. The previous layer compares this error with its output, back-propagates the error and adjusts its weight. This process continues for any number of layers.

REFERENCES

Rumelhart, David E. and McClelland, James L. (1986). Parallel Distributed Processing, Explorations in the Microstructure of Cognition, Vol. 1: Foundations. Cambridge, Massachusetts: MIT Press.

Roberts, Markus (1988, August). Twelve Neural Network Cliches. AI Expert, pp. 103-109.

Josin, Gary (1988, August). Integrating Neural Networks with Robots. AI Expert, pp. 110-118.