MACHINE VISION SORTING AND INSPECTION

IN COMMERCIAL AUTOMATIC

DISHWASHING

By

ANTHONY K. JOHNSON

Bachelor of Science

Oklahoma State University
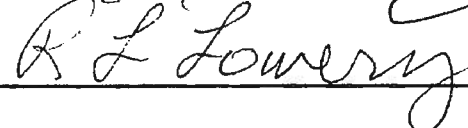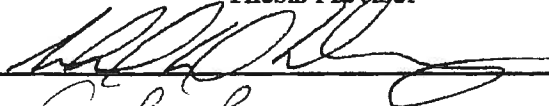
Stillwater, Oklahoma

1991

MACHINE VISION SORTING AND INSPECTION

IN COMMERCIAL AUTOMATIC

DISHWASHING

Thesis Approved:

_____
Thesis Adviser

_____

_____

_____
Dean of the Graduate College

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Background

Commercial dish washing operations, widespread throughout the United States in large food service operations, such as in hospitals, educational institutions, airlines and hotels, could benefit greatly by automation. In most current practice, dishes are loaded manually into continuous feed dish conveyors, passed through a continuously operating, serially staged dishwasher, and manually removed and inspected for cleanliness as they exit the machine. The work is tedious, repetitious and monotonous and, coupled with undesirable hot and wet working conditions, leads to a high turn-over rate among low-paid employees. High turn-over causes added expense of recruiting and training new employees and degrades the efficiency of the system.

As an example of generic commercial dish washing operations, we consider the dish-washing operation of a private 700 bed hospital in the midwestern U.S. This hospital, typifying current practice, operates 3 two-hour shifts, each processing approximately 600 trays of dishes. Each tray may consist of a cup, three silverware pieces, and up to six dishes of five different types. As dishes are brought to the dishroom, they are passed through a prescrub sluice, stacked on trays, and presented for loading into a continuous-feed dishwasher. Both dishes and trays are loaded manually into conveyor racks, while cups and silverware are placed in removable racks that can be placed on top the moving conveyor rack. They are then passed through a continuously operating, serially staged dishwasher; manually removed and inspected for cleanliness as they exit the machine; and stacked into baskets to await the next food serving shift.

1

Robotics and Machine Vision Automation

In this work, we focus on the repetitious work of loading, unloading, and inspecting dishes, because these operations seem particularly appropriate for automation: robot arms can replace humans both loading and unloading dishes, while machine vision can perform sorting and inspection tasks. Wash-down robots now available for food-grade service are ideal for wet working conditions, as well as the repetitious and monotonous tasks, conditions which lead to the high turnover rate of their human counterparts. In addition, robots require only a one-time cost with minimal service charges, need only be "trained" once, and can out-perform humans due to higher endurance and speed of operation.

Robots have long been applied to those operations where either working conditions are undesirable or even hazardous to humans, or where speed, efficiency, production costs, or labor costs prohibit manual labor. The Tony's Pizza plant in Salina, Kansas uses robots to palletize packing cases of frozen pizzas and other products in temperatures undesirable to humans("Palletized Production", 1992). Industrial robots have become common place in the automotive manufacturing industry, particularly in hazardous welding applications, and in assembly lines as press feeders and unloaders (Vaccari, 1992; Stephen, 1993). Innovative tooling of end effectors allows specialized functions in manufacturing such as deburring bulkheads and other F-16 fighter parts as done by General Dynamics to increase production quality and speed (Koelsch, 1991). Precision robotic chamfering of titanium and Inconel aircraft-engine turbine blades and hubs by Pratt & Witney reduce scrap and increase productivity ("Tech Update," 1992). Material handling applications include palletizing packaged products, feeding blanks to vacuum molds, and placing printed wrappers on boxes(Schwind, 1992; Langenfeld, 1992). New niche areas for robotic handling and automation are constantly being identified. Commercial dish washing

operations, currently untouched by any form of automation other that the continuous feed type dishwashers, are a good candidate for robotic automation.

Likewise, machine vision systems are making significant improvements in product quality and reducing manufacturing costs as they perform sorting, recognition, gauging and automated inspection tasks, and process control in real-time (Zuech, 1987). Some recent specific applications include a seven-second evaluation of 6 dimensional tolerances of 130 key guides in a keyboard manufacturing company, and recognition of scratches, defects, and occlusions in the 4 mm diameter lenses used in compact disc player scanning heads (Phillip, 1992). In addition to manufacturing, machine vision has been successfully applied to food handling applications such as sorting lemons according to grades of ripeness and size and identifying culls from a stream of freshly shelled peanuts. Both Hershey Foods Corp. and Mars use machine vision to check for defects in their candy bars just prior to wrapping them (Griswold, 1993). Machine vision has proven effective in many instances with conditions no more demanding than those in the our dishwashing setting, and should improve the quality of inspection for cleanliness, since the rejection of unclean dishes can be monitored and quantified, rather than analyzed qualitatively by a human inspector.

Scope of Automation

To automate a dish handling operation such as described above, two robots are needed: one to load dishes before they enter the dishwasher and one to unload them after they exit the final rinse section, as shown in Figure 1.1. As dishes are brought to the dishroom, they may be typically passed through a prescrub sluice, stacked in trays, and presented to the loading robot. This robot loads the dishes into the dishrack that conveys them through the dishwasher. Though the tool handling the dishes can accommodate each type of dish, optimum efficiency requires that only one type of dish be handled during each loading pass. After being washed, the dishes are sorted and inspected for cleanliness by

ADEPTONE ROBOTS

CLEAN DISH BASKETS

BI-DIRECTIONAL POWER CONVEYOR

VISION HARDWARE

STACKED DIRTY DISHES

DISHWASHER

SOAK TANK

→ CONVEYOR MOTION

LOADING AREA

UNLOADING AREA
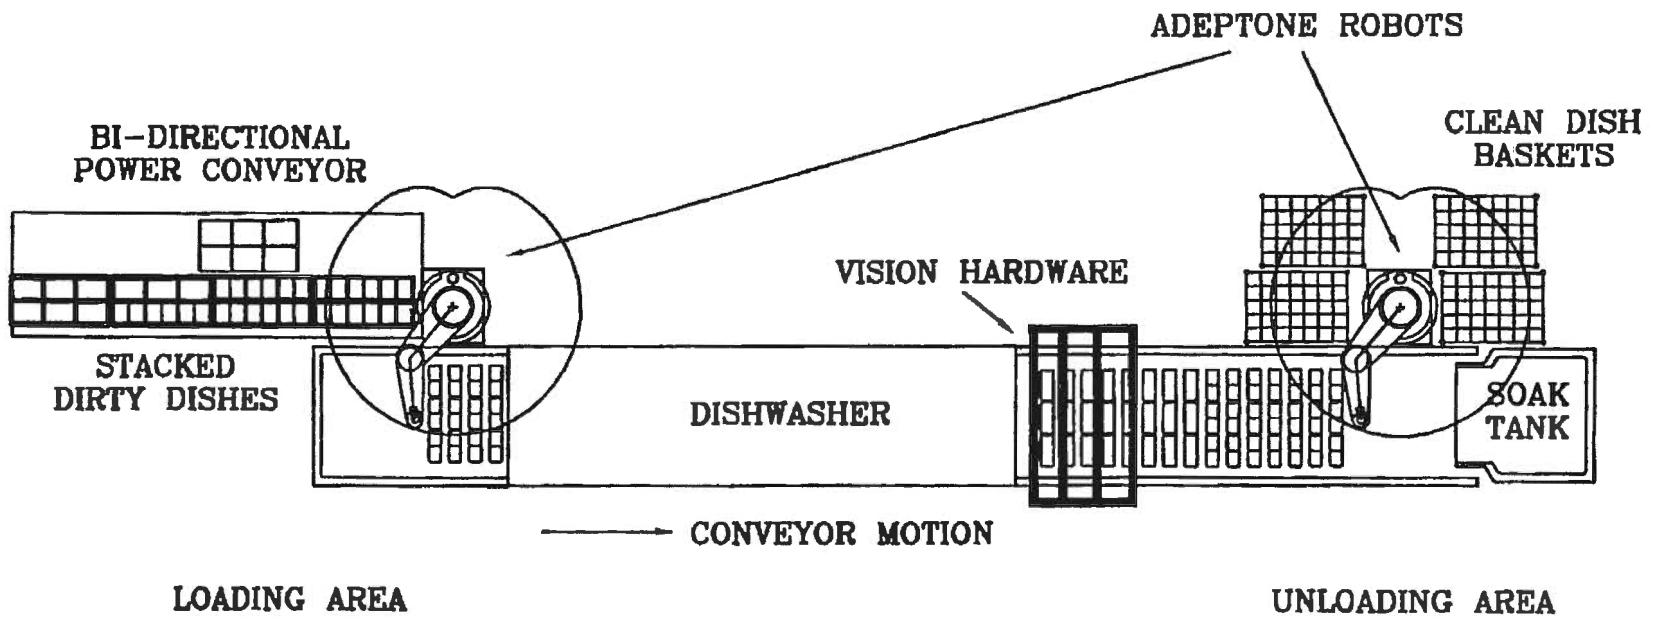
Figure 1.1 Schematic of Dishwashing System Layout Integrated with Robotic and Machine Vision Automation

machine vision, removed from the dishrack by the second robot and placed into their respective baskets for storage until the next meal.

Robotics implementation to load and unload dishes involves material handling, presentation of dishes, gripper design, motion control programming, synchronization with the conveyor, and simple error recovery. This thesis, however, is concerned with a particularly challenging portion of this problem, namely identifying and inspecting the dishes as they come into reach of the unloading robot. The robot needs to know what type dish it is going to handle so it can place it in the proper location. It is also desired to avoid handling a dish that is still unclean, allowing it to fall into a soak tank at the end of the conveyor.

Automated inspection of dishes is probably best accomplished through machine vision because it has proven effective in many inspection applications in a broad spectrum of industries.

<div align="center">Experimental Set-up</div>

To simulate the dishroom setting described above with integrated automation, only a few elements need be replicated: the dishrack conveyor, dishes, a robot (controller, arm, and software), and a vision system including proper lighting for illumination. Shown in Figure 1.2 is a schematic of the experimental setup, and photographs of the set-up are given in Appendix A. A short section of a modified dishrack mounted to a ten foot conveyor operating at the same speed as the dishwasher conveyor aids simulation of the process of loading and unloading dishes from the dishwasher conveyor system. To accommodate any one type of dish in a row of the dishrack, each dishrack row is divided into three equal sections, each capable of carrying one large dish or two small dishes. Figure 1.3 shows how the dishrack accommodates both sizes. The dishes used in this work are real dishes borrowed from the hospital mentioned earlier.
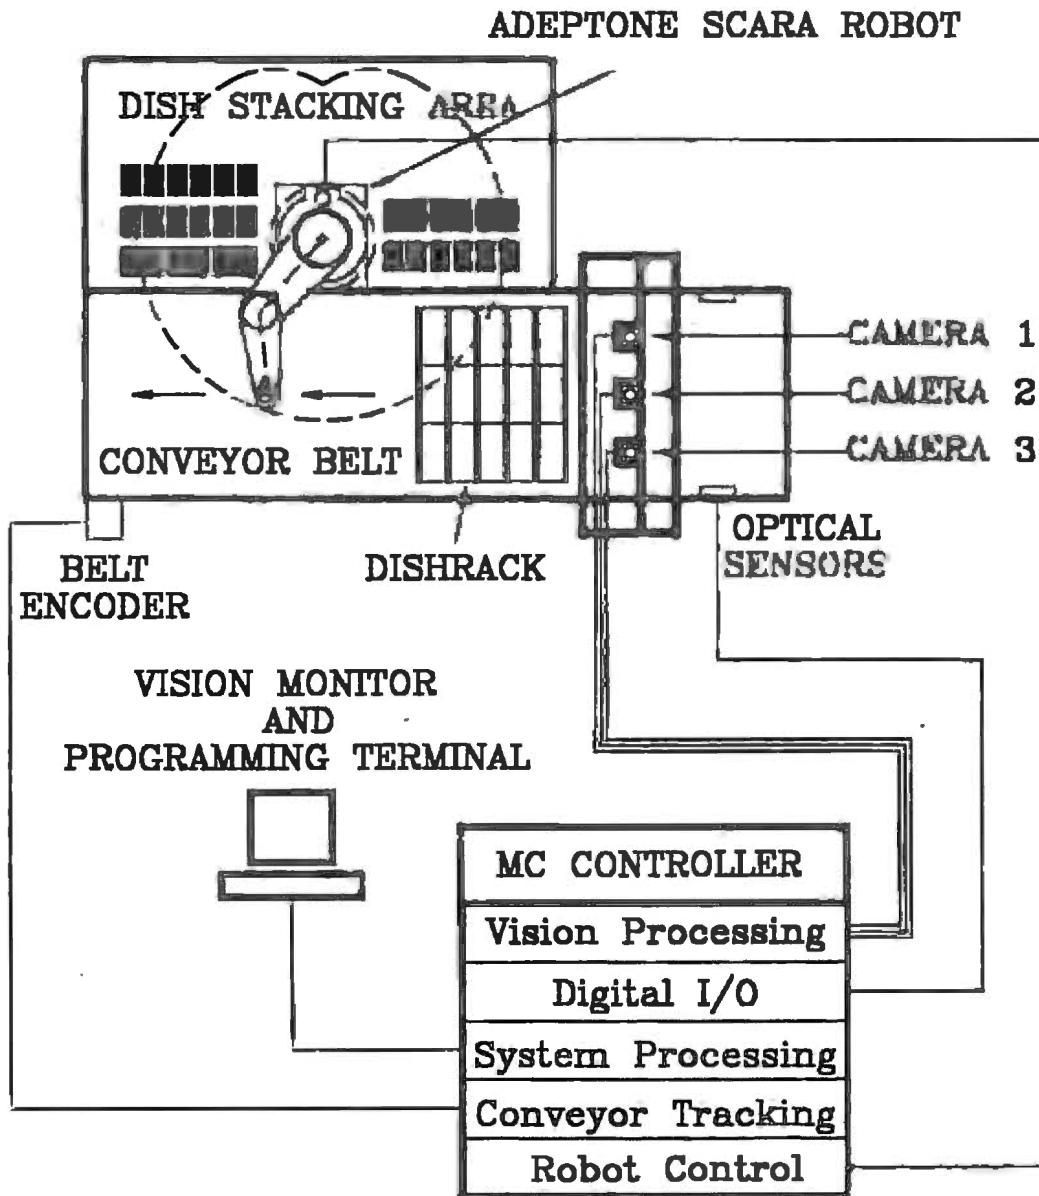
Figure 1.2 Experimental Setup
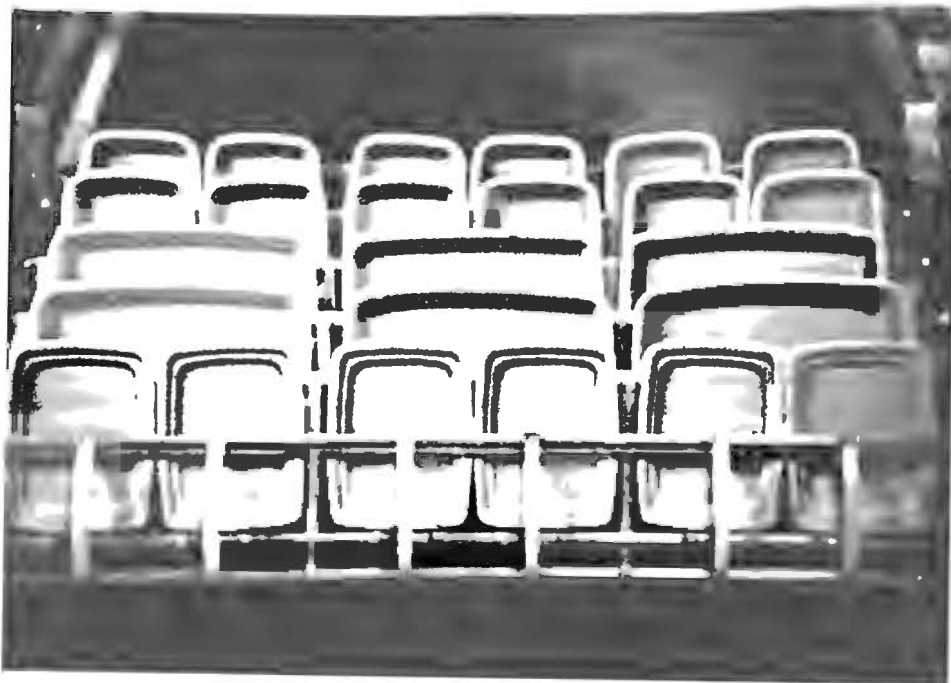
Figure 1.3 Dishrack Loaded with Both Large and Small Dishes.

Located adjacent to the conveyor is an AdeptOne SCARA robot as shown in Figure 1.2. The AdeptOne is a five-joint robot arm with a cylindrical workspace large enough to traverse the 30-inch width of the conveyor. Mounted to the robot is an end effector, or hand tool, designed to handle up to six dishes at a time using six flexible suction cups, each individually controlled. The conveyor is equipped with a belt encoder integrated with the robot controller to provide it with continuous information about the location of any point on the conveyor.

The AdeptOne robot controller is also integrated with its own vision system, allowing ease of interaction between the two systems. An array of three cameras is mounted above the conveyor, upstream and outside the robot workspace. Each camera has a field of view (FOV) large enough to scan one section of the dish rack row (Figure 1.3). That is, each camera views either one large dish or two smaller dishes. A light box, which is an enclosure around the cameras and light sources and open at the bottom, provides proper illumination to the dishrack containing dishes. It was recognized that a two-task program, controlling simultaneously the robot and the vision systems would be needed to implement sorting and inspection of the dishes exiting the dishwasher.

The details of the vision task program will be addressed in the following chapters. Chapter II introduces the basics of machine vision and the AdeptVision AGS system, and Chapter III discusses the implemented sorting and inspection routines. The design of the illumination system, a particulary interesting aspect of machine vision, will be discussed in Chapter IV. Experimental results are presented in Chapter V. Chapter VI gives conclusions and recommendations.

CHAPTER II

MACHINE VISION BACKGROUND

Machine Vision can be defined as automatic acquisition and analysis of images to obtain data for interpreting a scene or controlling an activity (Schaffer, 1986). Machine vision systems consist of four main components: Charge Coupled Device (CCD) cameras, a frame grabber, a host computer, and an image analysis software package. The cameras usually have a standard RS170 video format, 525 scan lines per frame, 30 frames per second, two fields per frame, and an image aspect ration of 4:3 (Wong, 1992). A frame grabber is a circuit board that acquires images on command, converts the analog signal from the camera to binary data, and sends the data to the computer for either immediate processing or for storage to be processed later. In our research, we used three Pulnix cameras, model TM-540, serial numbers 015207, 022177, and 014969 corresponding to camera 1, 2, and 3 respectively, the AdeptVision AGS system which is fully integrated with the robot controller, a frame grabber, (or vision processor) and the V/V+ software language compatible with the Adept MC controller.

To aid in the understanding of the methods used for sorting and inspection (Chapter III) using the AdeptVision AGS system, some concepts such as grayscale and binary imaging, frame store, "virtual" camera, "ping-pong" frame-grabbing, "blob", and vision "tool" are defined and discussed below in relation to the object recognition methods supported by the AGS system.

Imaging

Grayscale and binary refer to the two modes the frame grabber uses to convert the analog signal of the CCD to the digital value assigned to each picture element or pixel of the image. Grayscale imaging displays colors as different shades of gray pixels, while binary imaging shows only black or white pixels. CCD output is calibrated by three vision processor parameters: gain, offset, and binary threshold. Gain and offset affect grayscale imaging by respectively amplifying and biasing pixel intensity, aiding in the manipulation of contrast. Adjustment of these two parameters may either increase image contrast or nearly eliminate it, depending on the type of image needed for vision processing.

Binary threshold indicates the intensity level of grayscale above which pixels will be white; pixels with intensity below the threshold will be black. Varying the threshold parameter can isolate a specific intensity range needed for inspection. For instance, a light-colored object can be isolated from a group of darker objects by setting the threshold above the intensity of the dark objects and below that of the light-colored object.

Machine vision images are produced from single complete camera scans and are usually stored into vision buffers, which are blocks of memory designated specifically for vision images. The AdeptVision AGS system has two vision buffers, called frames stores. One important difference between the two types of imaging is that a binary frame store takes about 40 percent less time to acquire than a grayscale frame store (*AdeptVision Programming*, 1990). Therefore binary imaging is more desirable in cases where the highest speed is needed.

Using binary imaging, an object on a dark background will appear on the CRT monitor as a white "blob," or a group of white pixels, on a black background. Until the object is recognized, it is referred to as a blob. Finally, vision "tools" refer to system software routines that perform various measurement tasks on the frame store image. Examples include rulers and line and arc finders, which return measured distances and fit lines and arcs to edges, respectively.

Camera Designation

The Vision System supports up to eight physical cameras and up to 32 "virtual" cameras. A virtual camera is actually a data file containing pre-set parameters, such as grayscale offset and gain, binary threshold values, and the identity of the physical camera in the system to which they apply (*AdeptVision AGS*, 1990). Default identities for the first eight virtual cameras correspond to the eight physical cameras supported by the hardware. Each physical camera may be assigned to several virtual cameras, allowing for flexible programming. Every time a frame store is acquired, a virtual camera must be specified to let the system know which camera to acquire from and the values of the vision parameter settings to be used. The system allows parameters of any virtual camera to be changed during runtime. However, less processing time is required to call several different virtual cameras, each with their pre-set parameters to make the same inspection.

Frame-Grabbing

Several methods of frame-grabbing are supported by the AdeptVision system, each affecting the subsequent image processing in different ways. The default method, or mode, of frame grabbing automatically initiates region or boundary analysis on the blobs in the image as soon as the image is acquired. A second mode, called a future frame grab, initiates the image acquisition, and holds it in the frame store for future processing. No operation may be performed on this frame store except initiating the processing done in the default mode. This allows system processing to take place before the frame store is analyzed. The third mode, called a quick frame grab, because it is the fastest method, grabs an image into a frame store, but performs no automatic region or boundary analysis. This mode takes about 1/30 to 1/20 second to complete, compared to 1/4 second required by the first mode. Any of the software vision tools may be applied to the frame store acquired in this manner.

A frame-grabbing procedure called "ping-ponging" optimizes the use of the two frame store buffers. While one frame store is acquiring a quick frame grab, the other can be processed and analyzed by the appropriate vision tools, allowing images to be continuously acquired at a rate of 30 frames per second.

Identification Methods

Identification methods supported by the AdeptVision AGS system that can be applied to the dish sorting problem fall into two basic categories: prototype recognition and characteristic isolation (*AdeptVision Programming*, 1990). A discussion of how they function, as well as their respective benefits and disadvantages, is presented below.

Prototype Recognition.

Prototype recognition is based on whole-object pattern identification and requires that appropriate visual representation of each prototype be previously "taught" or stored in the system. This information consists of arcs and lines mathematically fitted to the boundary of the prototype. When a particular vision tool called the VLOCATE command is issued after a frame store has been acquired, the boundary pattern of a blob in the image is fitted with arcs and lines and then is compared mathematically and statistically to each prototype known to the system. When a match is made between the blob and a stored prototype within the statistical allowances, the blob is identified as an object, returning the name of the matched prototype. The advantage of this method is that the programming is easy: only one vision command, VLOCATE, is required to identify a "blob". The disadvantages of this method are that it is slow, requiring high overhead vision processing; the method is very sensitive to prototype training technique; and the recognition times may vary significantly with similar blobs, typically falling in the range of 100 to 400 ms for objects with only simple geometry (*AdeptVision Programming*, 1990). Recognition time

increases with increased geometric complexity and the number of prototypes trained in the system.

Characteristic Isolation.

Characteristic isolation differs from prototype recognition in that it investigates only the areas and characteristics needed for identification. The software developer has more flexibility with the analysis of the frame store, since only the information in these selected areas is processed, rather than the complete object pattern. Moreover, identification may not necessarily be affected by objects touching or overlapping each other. Vision tools provide information quickly and allow the software to sort the information, in contrast to the trial-and-error comparison of prototype recognition. In our case, only three characteristics need be identified to sort the five dishes, namely the plan area, the number of dishes present, and the radius of any one of the corners of the dish. This constitutes considerably less information to process than the complete dish boundary fitted by arcs and lines, required by prototype recognition. This, combined with the greater flexibility and accommodation of touching parts, favors sorting by characteristic isolation for dish sorting, where the dishes are distinguished by few characteristics and are more often touching than not.

Now, with a basic understanding of vision nomenclature, we proceed in the next chapter with a discussion of the sorting and inspection routines in which characteristic isolation is employed to sort dishes, and ping-pong frame grabbing and binary imaging are used to minimize processing time.

CHAPTER III

MACHINE VISION SORTING AND INSPECTION

The sorting and inspection routines presented here are written in the V+ language

introduced in Chapter II. They run sequentially, with sorting first, followed by inspection,

in the Task 1 Program, and will be discussed in that order. The sort and inspection

information required by the robot control program (Task 0) is passed through an array that

is global to both program tasks, which will be addressed at the end of this chapter.

Machine Vision Sorting

Dish Descriptions

The first concern in machine vision sorting is specific knowledge of the objects of

the sort. For the application in our dish washing operation, the system must be able to sort

among five different dishes, varying in size, shape, color and material as illustrated in

Figure 3.1. From top to bottom left to right, we have a small plastic dish, a small ceramic

dish, a small plastic spacer, a large ceramic dish, and a large plastic spacer. All five dishes

are basically rectangular in shape, with sides that form an angle of approximately 100 to

115 degrees from the open face of the dish. Viewed from above, the plan views of the two

larger dishes are approximately the same size and roughly twice the size of the plan views

of each of the other three. Specifically, the widths of the large dish and spacer are

approximately equal to the lengths of the small dishes and spacer. This characteristic

allows either 6 small or 3 large dishes in a row of the conveyor rack to be presented to the

vision system, as well as to the robot end-effector handling the dishes. If the system is
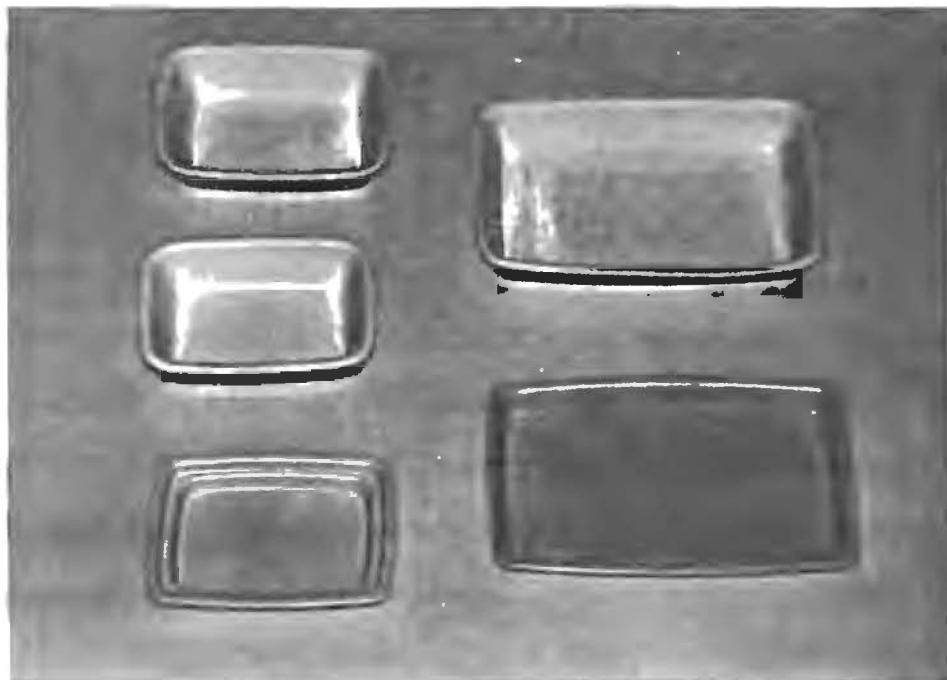
Figure 3.1 Dishes Used in This Application

constrained to keep all dishes identical in each row, only one dish from each row need be analyzed to determine the type present in that row. Any sorting performed on other dishes in the same row would be not only redundant, but would also require considerable additional processing time.

The major characteristic for distinguishing between large and small dishes is the plan projected area, easily detected by the two-dimensional camera image. After some initial study, it was determined that the large ceramic dish could be distinguished from the large plastic spacer by a larger arc radius at the corners of the ceramic dish. The three smaller dishes consist of a plastic dish, a plastic spacer, and a small version of the large ceramic dish described previously. Further study revealed that the small plastic spacer could be separated from the other two small dishes by its characteristic short-radius corner. The small plastic and ceramic dishes have essentially the same arc radius at their corners, but can be sorted on the basis of their plan projected area, the plastic dish being the smaller of the two.

Implemented Sorting Algorithm

The basic scheme of the sorting routine, developed in this work, follows the sequence of a) acquiring a frame store, b) accessing blob data processed during frame store acquisition, c) applying the vision tools, and then d) sorting the objects based on the information gained. The language used in this software is V+, provided by Adept Technology (1990). The VDISPLAY and VPICTURE commands work together to set up the CRT monitor display and acquire a frame store from a selected "virtual" camera. For our application, high speed is necessary, such that unnecessary graphics display to the monitor is eliminated, which reduces processing time. Threshold values for the dish characteristics described above are set as standards. For simplicity, when comparing area or arc radii with a standard, the system need know only whether the measured value is less than the standard or equal to or greater than the standard.

Once a frame store is acquired, the blobs, or dish images, are also placed in a queue for the VLOCATE command. In the characteristic isolation mode used here, this command finds only the centroidal location of the queued blob, its area, and the number of holes, or dark spots, in the white blob (*AdeptVision Reference*, 1990). The number of blobs, or dishes, and the area of the first blob in the queue are stored as variables to be evaluated later. Only the first blob in the image is needed to perform the sort, since all dishes in a row of the dishrack are identical. If no dishes are found in the image, a second frame, from Camera 2, is grabbed and analyzed for the presence of dishes before further processing is begun. If need requires, a third image, from Camera 3, is taken and processed. Obviously, if no dishes are present in the row, the sorting and inspection process is terminated and the system waits for the next row to come into view.

Blob area is used for several purposes. First, the V.MIN.AREA parameter is set such that any blob with an area less than the area of the smallest dish is not even processed. This is necessary since the camera will image the dishrack as small blobs when no dishes are present. Blob area also distinguishes large dishes and small touching dishes from small non-touching dishes. If the area is "small", small dishes are assumed present, and the area can be used to differentiate between the small plastic and small ceramic dishes, the ceramic being the larger of the two. If either large dishes or small touching dishes are present, they will form a blob with a "large" area. The dish-rack prevents any dish outside the FOV from overlapping one inside the FOV, thus insuring a maximum of two completely visible dishes to touch and appear as one. If this is the case, further processing is necessary to distinguish large from small dishes.

The next operation employs arc finders, using a vision tool called VFIND.ARC. This routine attempts to mathematically fit, or "find," an arc to a boundary within a sector specified and placed by the programmer. If an arc is found, a data array containing the arc radius is returned as the argument (*AdeptVision Reference*, 1990). Two arc finders are placed: one at the center of the lower edge of a blob representing a large dish or two

touching small dishes, called Arc1, and a second, called Arc2, at the expected location of the lower right hand (LRH) corner of the dish blob. The expected locations of the lower edge and the LRH corner are offset the appropriate amounts from the centroid found by the VLOCATE command. The arc finder Arc1 identifies the dish as large or small, because the lower edge of the large dish or spacer will have a very large radius at this location while two small dishes, whether touching or not, will have corners at this location with much smaller radii. If a small radius is returned, the blob area is divided by two to allow distinction between the small plastic and small ceramic dishes. The arc finder Arc2 distinguishes between dishes and spacers, regardless of size, based on the radius returned. Figure 3.2 shows a sorting diagram that illustrates this algorithm.

Figures 3.3 - 3.6 show a series of how the arc finders are placed on the blobs in this algorithm. The "located" arc in these figures does not show up well in these black and white photographs because they were taken from a color monitor which displayed the located arc in red. In Figure 3.3, the image has two blobs with small areas, indicating small dishes. Only one arc radius at the corner, Arc2, is needed to differentiate between small dishes and small spacers. Placed on the LRH corner of the right dish, Arc2 indicates that either a small plastic or small ceramic dish is present. Dividing the blob area by two and checking the results separates the two possibilities, and the system has sorted small ceramic dishes. Arc1 is not needed.

Figure 3.4 shows an image of two small, touching plastic spacers producing one blob. In this case Arc1 is needed to distinguish the small dishes from a large one. The arc finder returns the arc radius of the LRH corner of the left dish, and identifies that small dishes are present. Then Arc2, placed on the LRH corner of the right dish, indicates that a small plastic spacer is present. Using the first arc finder in this way accommodates touching small dishes that appear as one object, allowing them to be recognized as two small dishes.
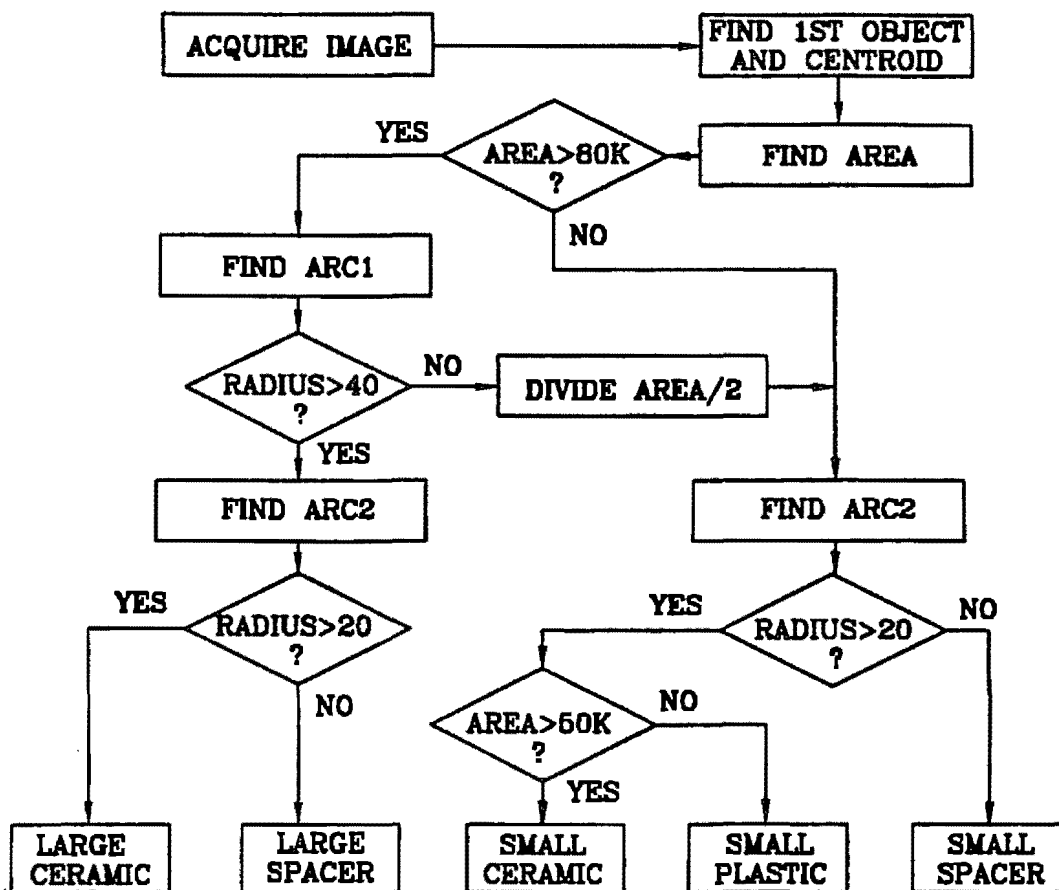
Figure 3.2 Flow Chart of Dish Sorting Logic
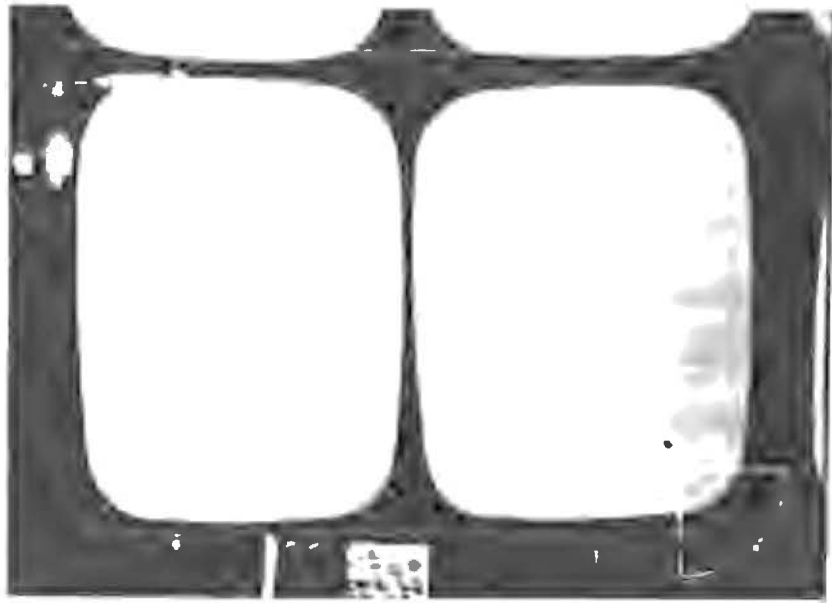
Figure 3.3 Two-Blob Image Sorted as a Two Small Ceramic Dishes
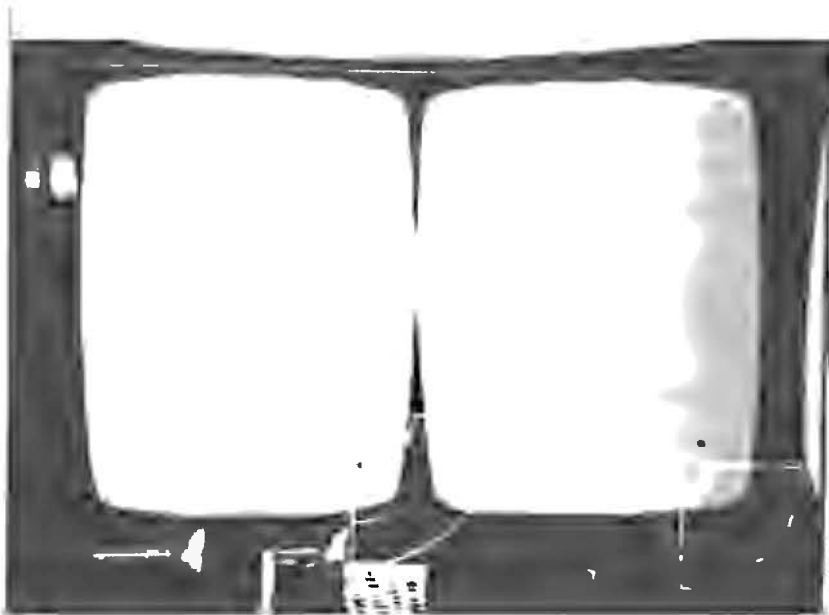
Figure 3.4 One-Blob Image Sorted as Two Small Plastic Spacers

The next two figures show a large ceramic dish and a large spacer respectively. In Figure 3.5, the Arc1 returns a large radius indicating a large dish or spacer, while Arc2, placed at the LRH corner, returns a large radius, sorting the object as a large ceramic dish. In Figure 3.6, Arc1 returns a large radius, indicating a large dish or spacer, and the radius found by Arc2 is small, indicating a spacer.

Use of this approach requires that dishes be presented each time to the cameras in the same way and in the same location for proper sorting and inspection. By immobilizing the dishes relative to the dishrack and triggering the imaging and inspection off the belt location tracker, this requirement can be satisfied. Once the dish is identified, the vision parameters for inspection are set, and inspection is initiated.
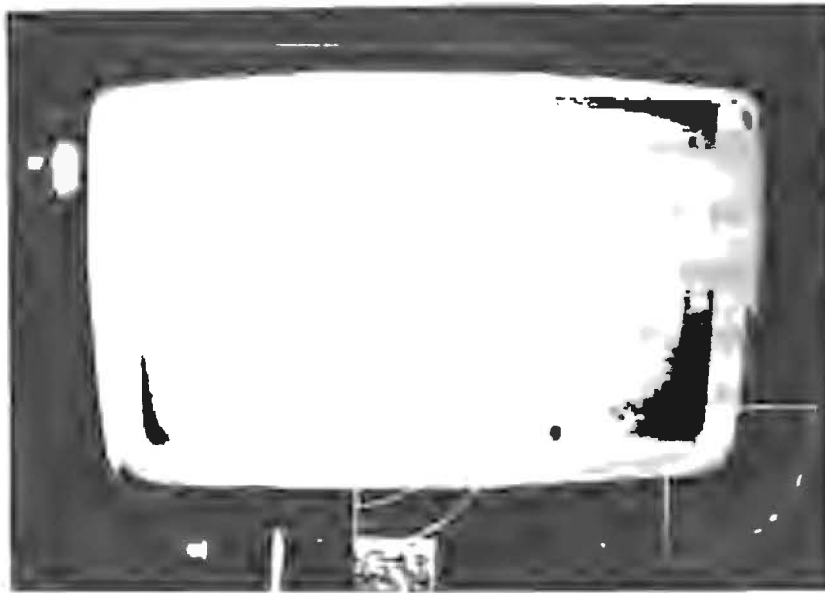
Figure 3.5 One-Blob Image Sorted as a Large Ceramic Dish

Figure 3.6 One-Blob Image Sorted as a Large Plastic Spacer.

Machine Vision Inspection

Food Particle Detection

A quick but thorough inspection completes the machine vision segment of the dish handling automation . Frequently, the dishwasher fails to completely clean every dish. To prevent the unloading robot from placing unclean dishes with those that are clean, machine vision inspection (MVI) is necessary. Therefore, the purpose of MVI is to quickly verify that dishes exiting the dishwasher at the rate of 30 rows per minute have in fact been thoroughly cleansed and to pass this information to the robot controller. Contaminated dishes identified by MVI should remain untouched by the end effector and left to ride in the moving dishrack to the soak tank.

Stray food particles present on a dish constitute contamination and can be detected by their discoloration, contrasting with the light background color of the dish. Using binary imaging, the discolorations appear as "holes" in the dish image "blob". A simple method of scanning for "holes," i.e., food particles, in a dish image is essentially all that is involved for inspection. The challenge lies in insuring the detection of even the slightest amount of stray dirt on a dish, minimizing processing time, and converting the clean/unclean information to signals controlling independently the six individual pick-up points of the end effector.

Insuring food particle detection requires proper setting of parameters for the cameras and vision system, together with proper illumination to avoid shading interference. Shading interference refers to the phenomena in which shadows or shading appearing on an object are detected in binary imaging as holes in the blob. The most obvious remedy to shading interference is to optimize the gain, offset and binary threshold parameters used in obtaining the frame stores to eliminate the shadows. With the use of virtual cameras, discussed in Chapter II, these parameters can be tuned and assigned to each dish type. Any food particle darker than the shading interference would appear as a hole and would be

detected. This also implies, however, that food particles lighter than the shading will not be detected. It is therefore necessary to reduce the amount of shading interference caused by illumination effects. Chapter IV has been devoted to this topic.

Inspection Algorithm

The first step in inspection is to know what type dish is present. Since the dishes vary in color, food particles may show up better on some than others. Proper vision parameter settings, including the binary threshold, must be set to allow for maximum sensitivity to discoloration. In the implemented algorithm, sets of three virtual cameras, corresponding to the three physical cameras, have been tuned for each type of dish. The dish type is conveniently known from the sorting routine executed before inspection, and identifies the set of virtual cameras to be used with each physical camera for inspection. When a frame store is acquired from the specified virtual camera, the vision parameters associated with that camera are used, insuring maximum sensitivity to food particles regardless of the dish type.

Two scanning routines where developed, one for the large dishes and one for the smaller dishes. Though these routines are different, the basic algorithm is the same. During image processing, "windows" specifying the region of the frame store to be processed are placed over the expected location of each dish, and only this area of the frame store is processed. Reduced area to process yields reduced processing time. In each window, the dish "blob" is located using the VLOCATE command; dirt-causing holes in the blob are automatically counted. If the number of holes counted in the dish blob is zero, the dish is clean; any number greater than zero indicates at least one dirty spot has been identified and that the dish has not been thoroughly cleansed. The hole count can be checked through the VFEATURE command. The scanning routine for small dishes has the same basic algorithm, differing only in the number and size of the windows used to

process each frame store image. Two windows, each covering the expected location and size of the small dishes are employed instead of the one window for the large dishes.

Setting the Code for the Robot End Effector

Once a dish has been inspected and labeled as clean, it is ready to be picked up by the robot end effector, or hand tool. The hand tool has an array of six suction cup grippers allowing it to hold six small dishes or three large dishes, with two suction cups per dish. The vacuum drivers are individually controlled to allow the robot to handle only those dishes it is instructed to pick-up. A binary code of six digits tells which vacuum drivers to activate when the robot moves in to remove dishes from the rack. When inspection begins, the code is set to 000000. As each dish is inspected and labeled as clean, the corresponding bit is turned on to activate that vacuum. For instance a gripper code of 111011 would indicate than a small dish in the fourth slot did not pass inspection. All vacuum drivers except the fourth would be turned on, removing all from the dishrack except the unclean dish.

This code along with dishtype information is passed from the vision task to the robot task in the form of a queuing list, actually a 10x2 array, which provides for 10 rows of dishes to be queued. The first column element contains the dishtype number, while the second contains the gripper code. The robot removes the dishes as they are listed in the 10x2 array, top to bottom. Then as the array fills up at the 10th row and the robot removes rows of dishes from the conveyor, the array index is rolled over to the top of the array and continues as before, queuing rows of dishes for the robot to remove.

With the software design for food particle detection, minimum processing time, and integrated vision-based robot control as a foundation, we must optimize food particle detection. As indicated by VanDommelen (1990), the most difficult aspect of machine vision is proper illumination, for objects appear differently depending both on how they are illuminated and how they are presented to a camera. In our application, the amount of dirt

on a dish needed for detection depends on these two factors. Chapter IV has been dedicated to the topic of illumination and discusses how the final design was selected. The speed and accuracy of the sorting algorithm and the speed of the inspection routine are detailed in Chapter V.

CHAPTER IV

ILLUMINATION DESIGN

Need for Proper Illumination

As in robotic motion control, the machine vision sorting and inspection routines can be accomplished with the assistance of Adept Technology's available software. However, a special challenge exists in properly presenting dishes to the cameras and in illuminating them. Dish presentation issues involve where the dishes appear in the camera's field of view, how the orientation angles of both dish and camera affect the image, and how to inspect completely a dish sitting in a dishrack. Dish presentation is influenced by the dishrack design. Given a dishrack design, camera positioning can be tuned such that they address these issues.

Illumination design deals with the effect of lighting on the dish and how it affects the image seen by the camera. Ambient lighting usually is inconsistent, changes in intensity, and casts shadows by moving objects, such as personnel, in the surrounding environment. In food particle detection, particles are detected because they either scatter or absorb light. However, the geometry of the dishes produces interfering shading by scattering light, such that distinction must be made between the actual dirt-causing light intensity loss and shading interference. Consistent lighting is obviously necessary to provide reliability in the inspection process. Ambient lighting, therefore, is unacceptable. Instead, an enclosure referred to as a light box, containing light sources, cameras and the objects of inspection (dishes in the dishrack) was found necessary to provide both consistent lighting and proper illumination to reduce shading interference.

Light Box Design

Design Methodology

Designing the light box from a limited base of experience and available literature requires trial and error, and some creativity. Different light sources, orientations of light sources, non-mirror-like reflectors, diffusers and several optical filters were tried in various combinations to arrive at an acceptable solution. Initially, these components were arranged and varied to get a qualitative feel for what seemed acceptable and what was not. It was soon discovered that overhead lighting produced harsh shadows along the dish sides, even when diffusers were introduced between the light source and the dishes. Another discovery was that the more diffuse the lighting, the softer the shadows became, and shading interference declined. To produce diffuse illumination, one must use a very diffuse light source (such as fiber optics), fluorescent lights shielded by a diffuser, or an indirect light source banking light off a diffusing reflector. Fluorescent lights produce the most diffuse lighting for large areas, whereas fiber optics are typically used for flood lighting from the direction of the camera for close inspections (Harding, 1993). Diffusing reflectors that also eliminate variant, ambient lighting effects can be used to stabilize the illumination as well as increase diffusivity of the light. Reflectors with a matte finish diffuse light much better than shiny reflectors. In our application the most effective location for fluorescent light tubes was found to be at the side of and parallel with the dishrack conveyor, just below the level of the dishes in the rack.

Fluorescent illumination, particularly useful in the inspection of specular surfaces (Chen and Tretiak, 1992) such as the dishes used in this application, was addressed by experimenting with ultra-violet band pass filters.

Using this basic understanding, a concept for a light box was formed and experimental set-ups were constructed such that geometric shapes could be adjusted to find the optimum combination. The goal was to eliminate, as much as possible, the shading

interference described in Chapter III. To compare results from each experimental setup, a test dish was used with varying degrees of dirt . Each experimental setup was evaluated by first adjusting the grayscale gain, while viewing a binary image, to eliminate the black pixels due to shading or shadows on a standardized dish containing food particles. A count of the black pixels remaining inside the dish boundary was made and recorded. A higher number of black pixels counted, indicated a greater ability to detect the lightest or least amount of dirt particles present on a dish. The same standardized dish was used in each evaluation of the experimental setups.

Fluorescent lamps were tested at a variety of locations relative to the FOV to find their most effective orientation. Placed overhead, they produced the largest amount of shading interference. Placing them above and to the side of the FOV, (either parallel with or perpendicular to the rows of dishes provided the same results), and shielding direct light from the dishes was an improvement, but could not reduce the shading interference as well as when the lamps were placed along side of the dishrack, just below the level of the dishes. Lowering the lamps further also increased shading interference.

Varying the top panel of the light box also had some interesting results. For instance, a flat top panel performed better than a corrugated panel, but not as well as a V-shaped trough. Corrugations aligned with the rows (traversing the dishrack) performed better than when they were aligned parallel with the dishrack.

Final Design

Figure 4.1 shows the schematic of the final light box design. Illumination is provided by two banks of fluorescent lights, placed at either side of and parallel with the dishrack as mentioned earlier in this chapter. Each bank consists of two 4 foot lamps, either F40 or F32T8. The level of the light tubes is set just below the level of the dishes in the rack to prevent any direct lighting this being found undesirable. High-frequency ballasts operating at 20 kHz are used to drive the fluorescent lamps in an effort to eliminate

Camera 1   Camera 2   Camera 3   Plastic Diffuser

16

Filters

16 — 8 — 8

48

16

Side Walls

Leading Wall

Pixel Targets

30

38

Trailing Wall

4

Motion

Dishrack Clearance

6

Fluorescent Lamps

24

Fluorescent Lamps

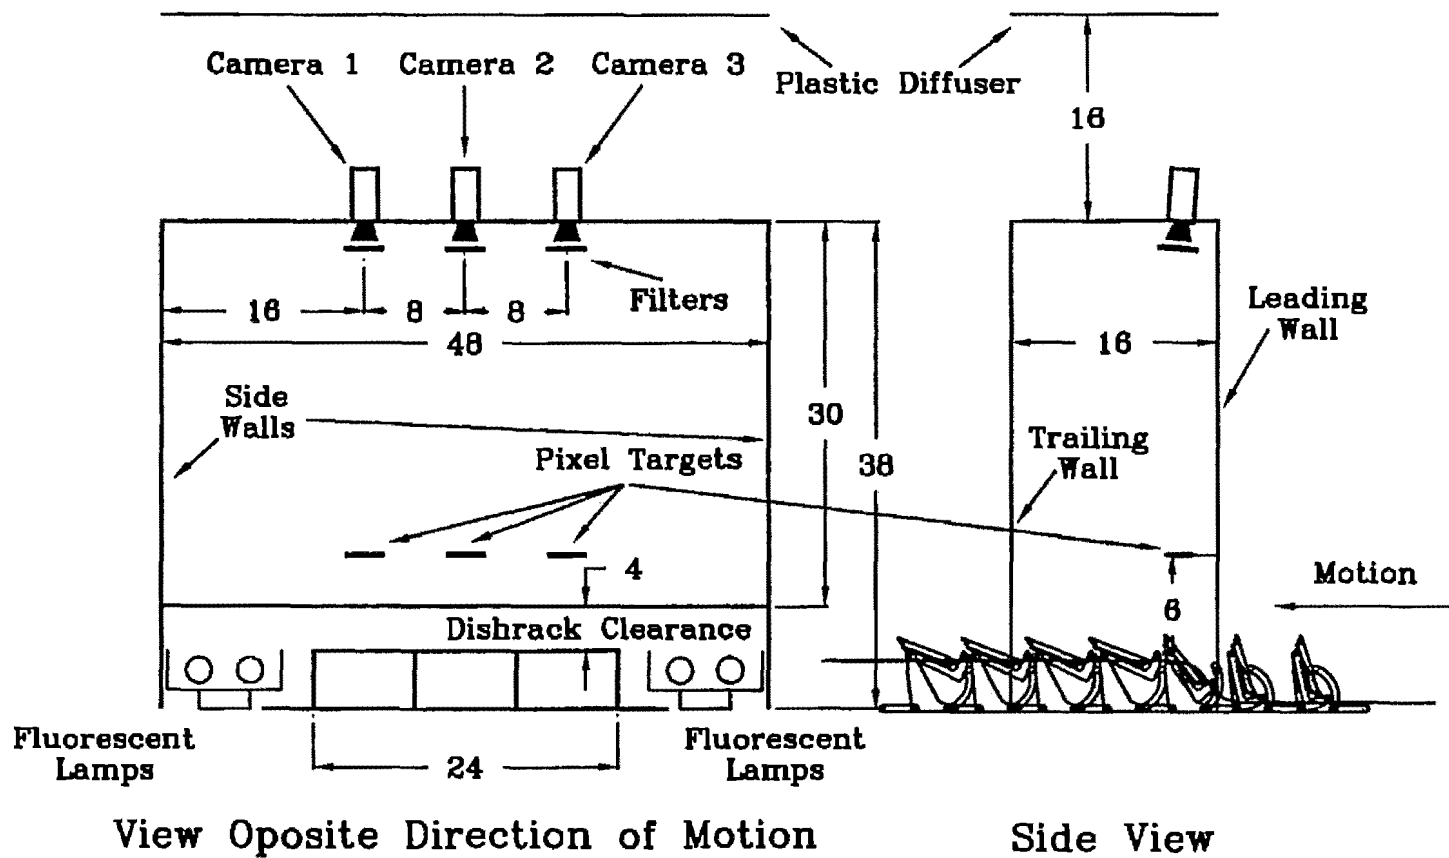View Oposite Direction of Motion   Side View

Figure 4.1 Schematic of Final Light Box Design

frequency beating between the lamp cycle period and the frame grabbing interval. Beating is particularly noticeable as a flickering image in constant output mode.

Gray scale coupons are mounted near the inside of the leading wall of the light box under each camera. These are used for automatic threshold control discussed later in this chapter. The four vertical sides of the light box are composed of white matte (non-glare) poster board, 1/16 inch thick. The top of the box consists of a white translucent plastic light diffuser, common in office lighting fixtures. Overall dimensions of the box are 36 inches high, 48 inches wide, and 16 inches deep. The leading and trailing walls are hinged near the top to allow large pots and pans to pass through the light box, if typical operation demands, and the box is open at the bottom so as not to disturb any dishes as they proceed into the FOV of the cameras. The dishrack clearance, optimized in the illumination experimentation presented in Chapter V, is measured from the top of the dishrack to the bottom of the leading and trailing walls of the light box. The other two walls, enclosing the light box by extending from the camera height to below the level of the light tubes together with the leading and trailing walls, act as diffusing reflectors, scattering light diffusely into the center.

<center>Cameras and Optics</center>

Filters and Optics Hardware

To avoid collision with large pots and pans that may pass through the dishwasher, the cameras are placed 2.5 feet above the top of the dishrack. Each camera is equipped with a 25 mm focal length lens, producing a square FOV with a side length approximately equal to 1.2 times the length of the large dishes. Ultra-Violet (UV) band pass filters placed over the camera lenses where used to accentuate the presence of food particles. Experiments with two UV filters, model numbers 1-64 and 5-59 manufactured by Kopp Glass Inc., were conducted to determine which performs best in our application. Transmittance curves for these filters are given in the appendix. The results of the

experimental testing are presented in Chapter V and show that the UV filters increased sensitivity to food particle detection over a naked lens, with the 5-59 filter being the superior of the two.

CCD Output Intensity Drift

During evaluation of various experimental setups the phenomena of intensity drift was encountered. For any one setup, data, i.e., pixel counts, could be taken for one setting of parameters at one time and within a matter of only a few minutes, data taken with the same settings yielded completely different results.

To investigate this problem, the pixel count of the standardized dish was monitored over a period of 3.5 hours with no changes made to the parameter settings for offset, gain, and threshold. Each recorded pixel count was actually the average of 10 consecutive readings, spannig a period of approximately 80 seconds. As shown in Figure 4.2, there was significant drift in the direction of increasing intensity until all but one very dark spot had completely disappeared from the images. This prompted a series of controlled experiments to isolate the cause of the problem from the following suspects: a) unstable cameras, b) fluctuating light intensity, and c) fluctuating power voltage.

The first experiment consisted of monitoring the intensity of the fluorescent light sources, the CCD camera output, and the electrical outlet voltage feeding power to the lights. Light source intensity was monitored by employing a photoresistor bridge circuit, shown in Figure 4.3. The photoresistor (PR) was mounted inside the light box and remote from the circuit, such that manual readings of circuit voltages with a hand-held voltmeter would not affect the measured lighting intensity inside the box. Since an increase in light intensity reduces the resistance in the photoresistor, it can be shown that an increase in light intensity produces a decrease in voltage (mV) across the bridge. For the purpose of this experiment, the change in the value of the voltage is of particular interest, rather than the actual value, since we sought only relative changes in intensity.
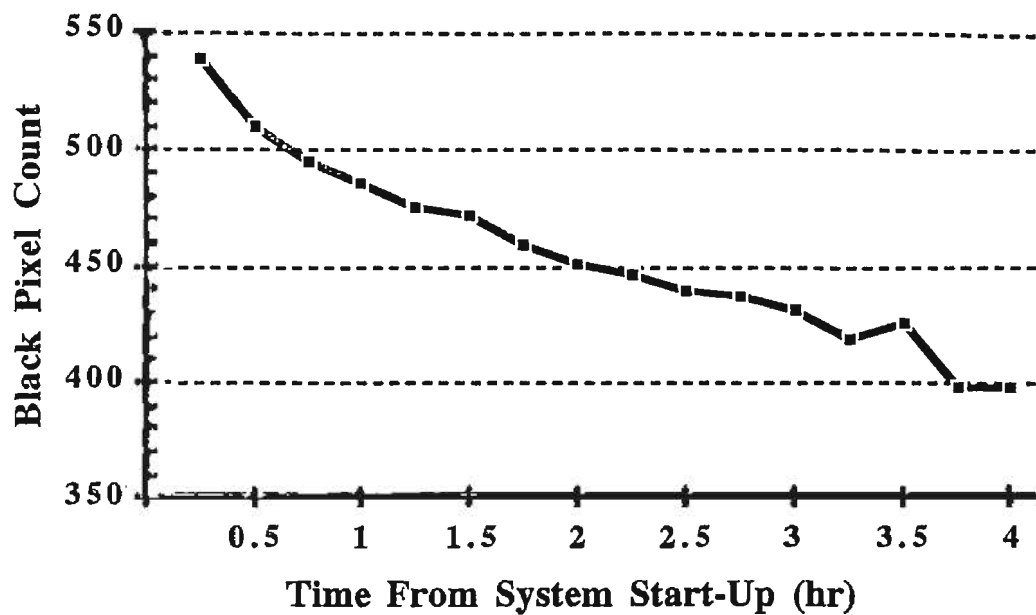
Figure 4.2 Shift of Camera 1 Output Evidenced by Pixel Counting
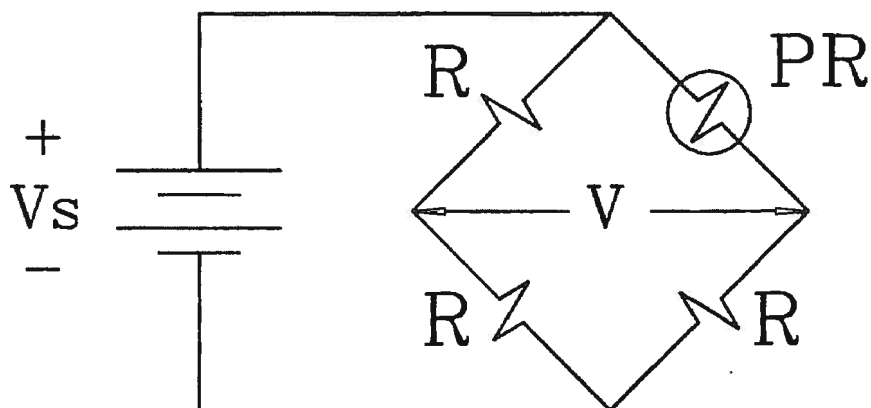Pulnix Camera, Model No. TM-540, Serial No. 015207

Figure 4.3 Schematic of Photoresistor Bridge Circuit

To monitor the CCD output from the cameras, the method of pixel counting was again employed, but with some modifications. Instead of using the standardized dish, a target of varying grayscale shown in Figure 4.4 was used to ensure a complete range, with no cutoff level. The vision imaging parameters were set such that the median grayscale and lighter shades appeared as white pixels. Also, the white pixels were counted instead of the black as in the preliminary experiment. The 120 V wall outlet voltage was measured using an RMS voltmeter and recorded along with the bridge voltage and the white pixel counts for each camera. In this experiment, the wall outlet voltage powered both the lights and the bridge circuit, while the cameras were powered from a different 240 V circuit through a step-down transformer in the AdeptOne MC controller cabinet.

Figures 4.6 - 4.7 show the results of this experiment. From this data it is hard to conclude whether the outlet voltage feeding both the lights and the 4 volt power supply to the bridge circuit is the cause of the changing light intensity and ultimately the CCD output. There is some correlation among the results from the three graphs, since the major peaks occur at the same data collection times. However, no conclusive evidence isolating one component as the problem source is evident.

A difficulty in this experiment was that the pixel count measurements where recorded and averaged manually, taking approximately 7 minutes to complete for each data point. This prompted a second experiment, in which the basic form of the first was repeated, except that the power source for the lights and the bridge circuit was drawn from the transformer inside the robot controller, and the pixel count recording was automated, decreasing the data acquisition time to 2 seconds. This second experiment was conducted twice, first as just described, and then later with the camera cables to Cameras 1 and 3 switched in an effort to isolate a defective camera from a defective I/O port in the MC controller, and with high-frequency (20 kHz) ballasts replacing 60 Hz ballasts for the fluorescent lamps. The results from both of these experiments gave essentially identical
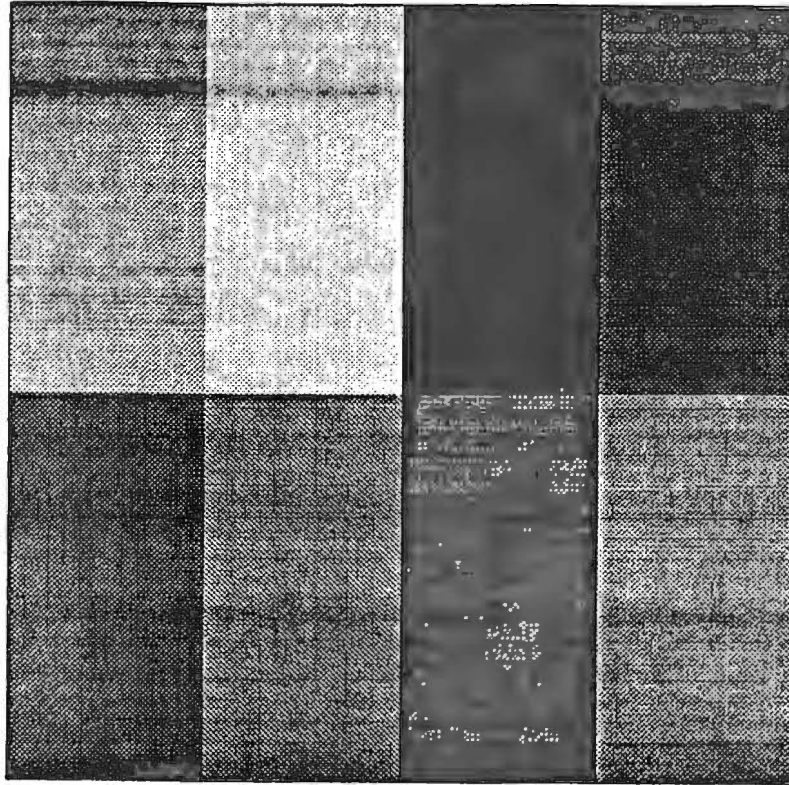
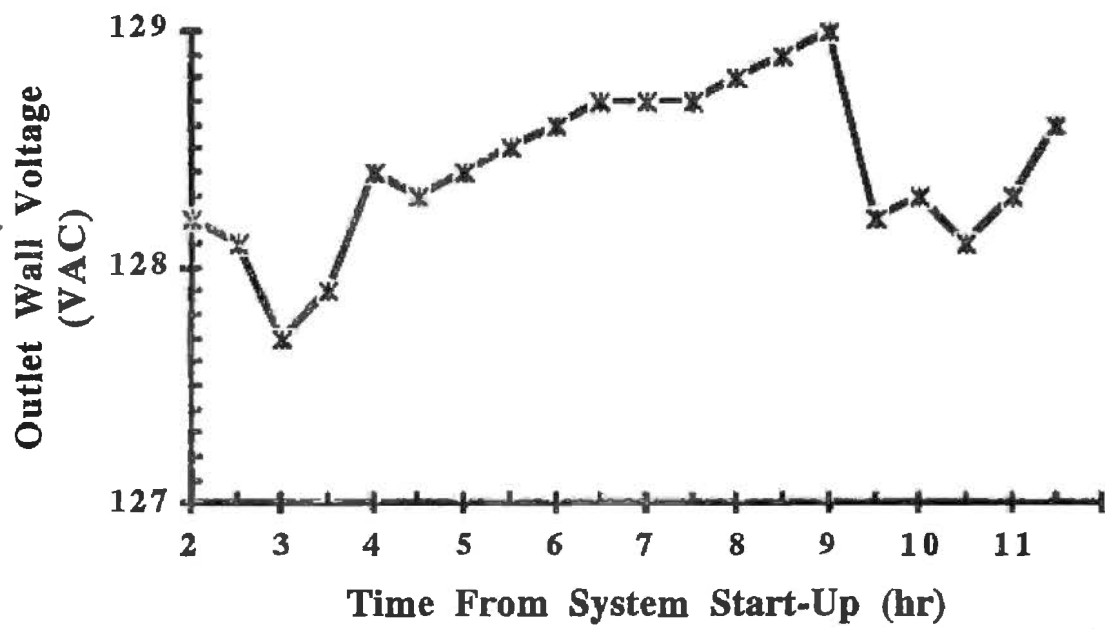Figure 4.4 Target of Varying Grayscale

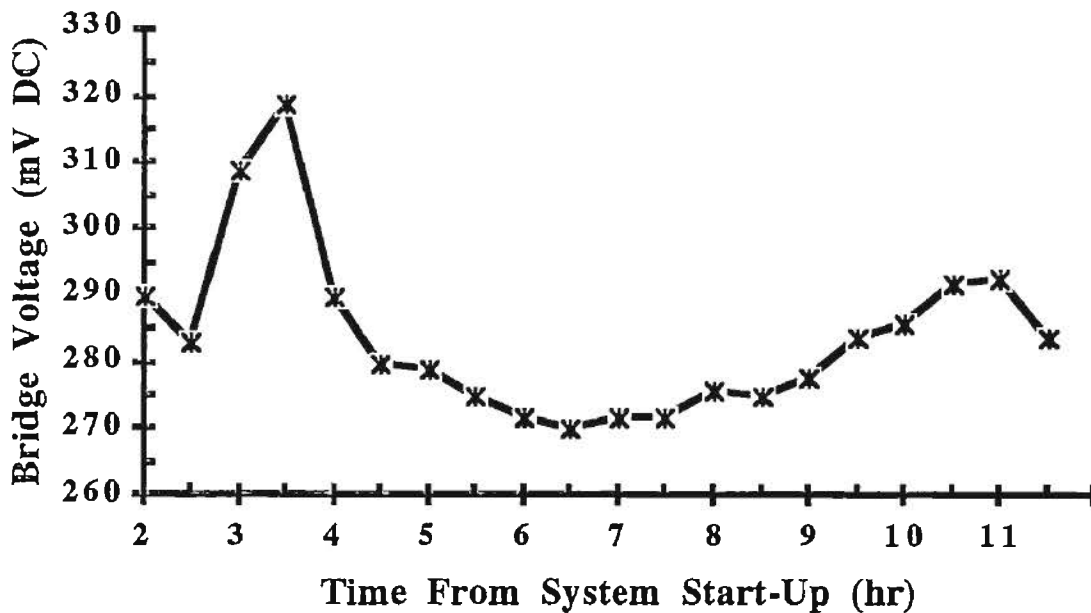Figure 4.5 Fluctuation of 120 Volt Wall Outlet Voltage

Figure 4.6 Voltage Variation Across Bridge Circuit Using F40 Fluorescent
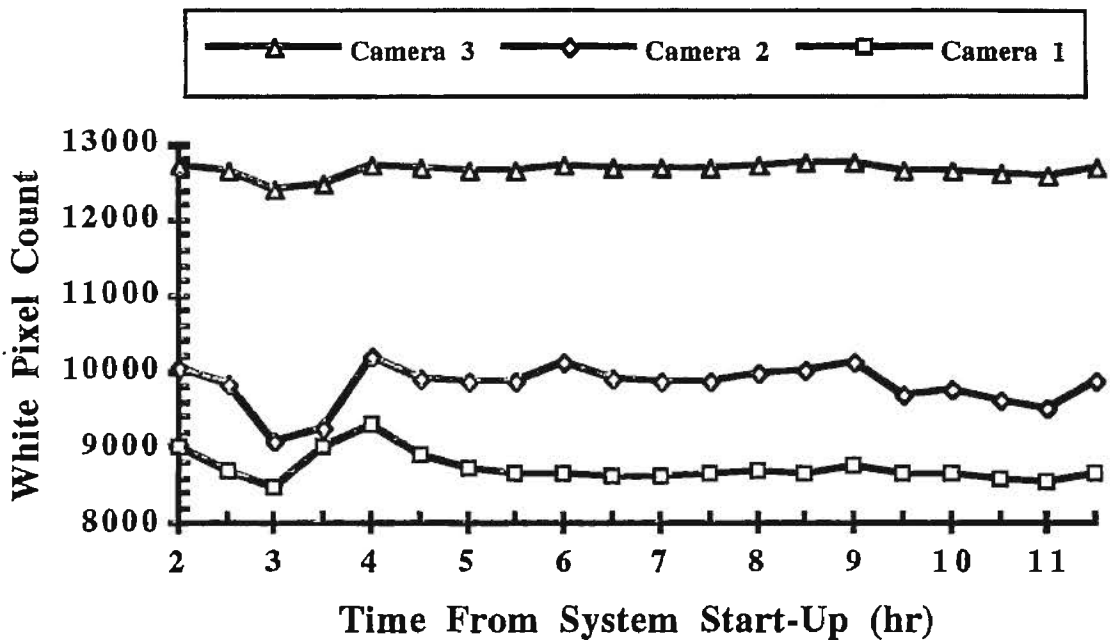Lamps and 60 Hz Ballast Powered from 120 Volt Wall Outlet

Figure 4.7 Monitored CCD Output of the Three Cameras as White Pixel Count
Camera 1: Pulnix Model No. TM-540 Serial No. 015207
Camera 2: Pulnix Model No. TM-540 Serial No. 022177
Camera 3: Pulnix Model No. TM-540 Serial No. 014969

results confirming that both power fluctuations and CCD output drift are significant problems that must be solved.

Figure 4.8 shows the power fluctuation of the outlet from the step-down transformer inside the MC controller unit over a 12 hour period beginning approximately at 9:00 am, with data points gathered every fifteen minutes. Notice that the power supply voltage fluctuates within a range of 2 volts and spans that range within a one hour interval as seen from 6.5 to 7.5 hours after system start-up. Note also that the voltage tended to increase throughout the day.

Figure 4.9 shows the data collected from the photoresistor bridge circuit. The power supply apparently filtered out any voltage fluctuation from the MC controller transformer that provided it with power, as seen by its nearly constant output. The scales in Figure 4.9 are in units of volts for the power supply and millivolts for the bridge voltage. However, they both cover the same range of 250 mV. From this, we can conclude that a change in the bridge voltage is not a function of the fluctuating outlet power from the MC controller transformer, but rather a change in the light intensity. The bridge voltage, measured in the same manner as in the previous experiment, tended to fall off during the day, indicating somewhat of a rise in light intensity. This correlates with an upward trend in the power source voltage in Figure 4.8.
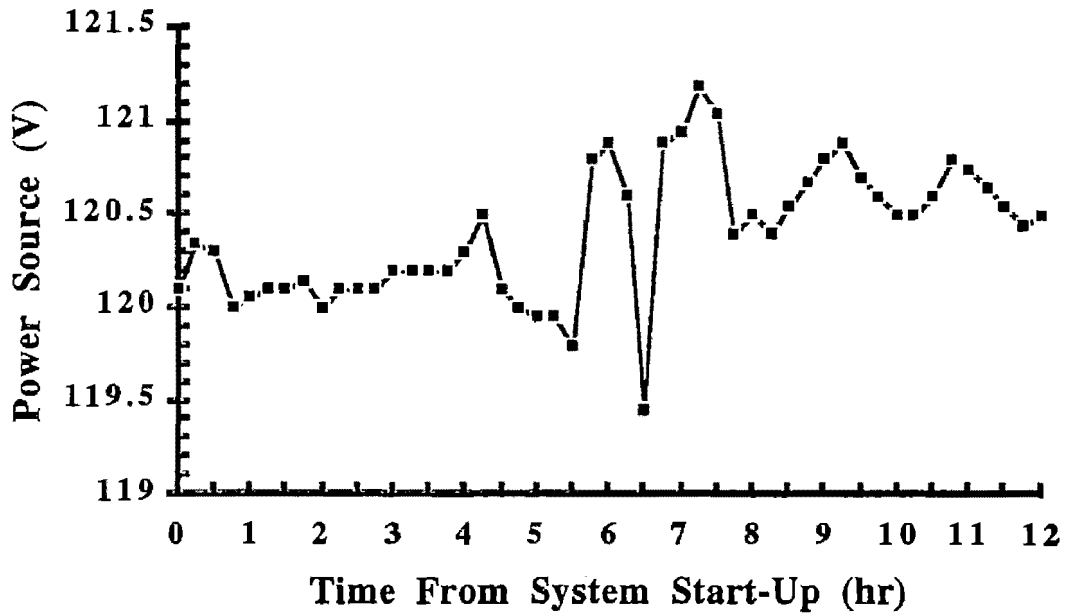
Figure 4.8 Fluctuation of Outlet Power from the Step-Down MC Controller
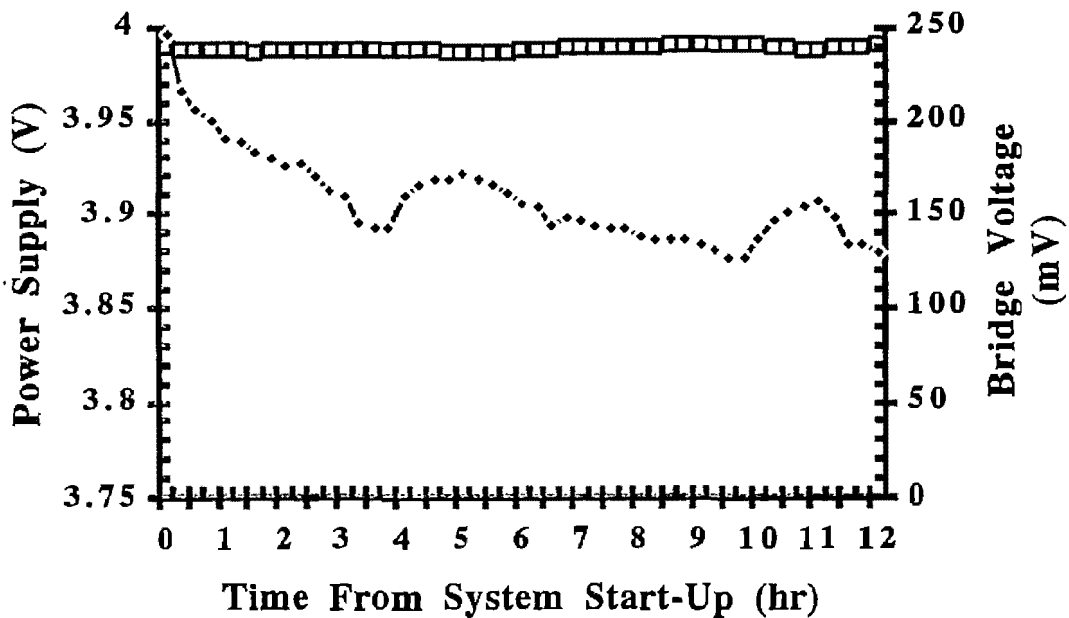Transformer

Figure 4.9 Measured Voltages of Bridge Circuit and 4-Volt Power Supply
to Monitor Light Intensity of F40 Type Fluorescent Lamps
Driven with 20 kHz Ballast

Threshold Control

In order to compensate for power and lighting fluctuations and camera intensity

drift described above, a simple automatic real-time threshold controller based on

proportional-integral (PI) control theory was developed for each camera. The controller

takes a pixel count of a grayscale grid (similar to the one used for the intensity experiments)

mounted to the coupons placed six inches above the dish rack such that they appear in the

lower central part of the FOV of each camera. The pixel count is compared to a desired set

point, yielding an error and an error sum necessary for the PI control algorithm, given by

$$TH_i = KP[E_i + KI\ ES_i] \tag{1}$$

where

TH_i	= Threshold Setting at time t_i
KP	= Proportional Gain Constant
KI	= Integral Gain Constant
E_i	= pixel count - setpoint, at time t_i
ES_i	= $\sum E_i$, from time 0 to time t_i

An increase in threshold yields a lower pixel count, such that the error is calculated

as the measured value (pixel count) - setpoint, rather than the usual setpoint - measured

value. As the intensity experiments indicate, each camera responds differently to what

appears in its FOV. For this reason, the proportional and integral gains for each camera

where calibrated by trial-and-error. The main criteria for selecting the gain constants was a

one second settling time which, as the above experiment shows, is much faster than the

dynamics of the drifting CCD output. The percent overshoot was not considered since

vision processing will occur at steady state. The results for each camera are shown in

Table I. below.

While the Task 1 program is waiting for the next row of dishes, the system rolls

through at least one increment of the PI compensator which provides for real-time control.

For every change in threshold from the PI controller, the thresholds of virtual cameras 1-3

are used as baselines from which the thresholds of virtual cameras, 4 - 12, are adjusted by

empirically determined offsets to provide maximum food particle detection.

TABLE I

PI CONTROL CONSTANTS FOR AUTOMATIC THRESHOLD
CONTROL FOR CAMERAS 1-3

Gain Constant	Camera 1	Camera 2	Camera 3
KP	0.00575	0.0023475	0.5
KI	1.35	3.057	1.45

The compensation scheme designed and implemented above was evaluated to see how well it performed. During the same time the data represented in Figures 4.8 and 4.9 above was taken, data from each camera was recorded to show the response of the each camera with and without the threshold controller. The results are presented in Figures 4.10 through 4.12.

In Figure 4.10 the white pixel count of the grayscale target shown in Figure 4.4 is obtained with constant settings of binary threshold, and grayscale gain, and offset, set at 145, 20 and 60 respectively. Note that Camera 3 has a five-hour "warm-up" time. Each camera shows a hump at hours 5.5 and 6.5, which correspond to blips in the power source (Figure 4.8). The humps at hours 10 and 11 correspond to light intensity changes (Figure 4.9). Thus, as expected, camera output is sensitive to both power source and light intensity changes. However, these are coupled effects, since both cameras and lights are powered by the same source.

The threshold was designed to maintain a constant pixel count for a constant image source. The values of the threshold setting from the controller to maintain constant pixel counts are shown in Figure 4.11 while the pixel count at the controlled threshold settings are given in Figure 4.12. Grayscale gain and offset were left unchanged from the settings above. Notice the curves in Figure 4.11 track those in Figure 4.10, as the controlled
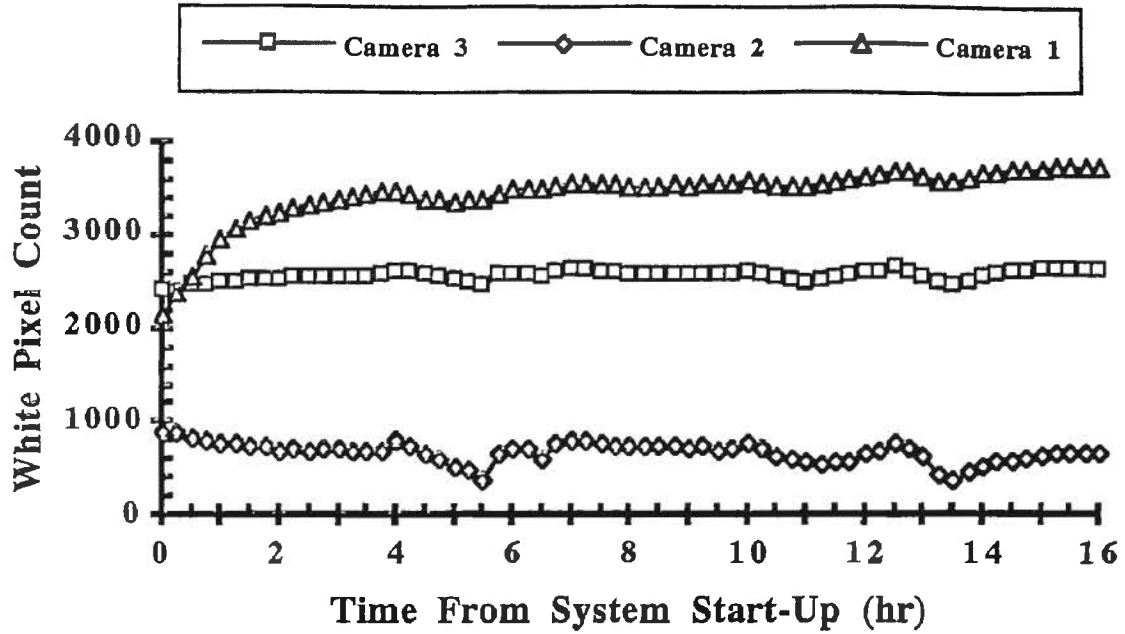
Figure 4.10 Pixel Count Obtained From Constant Imaging Parameters
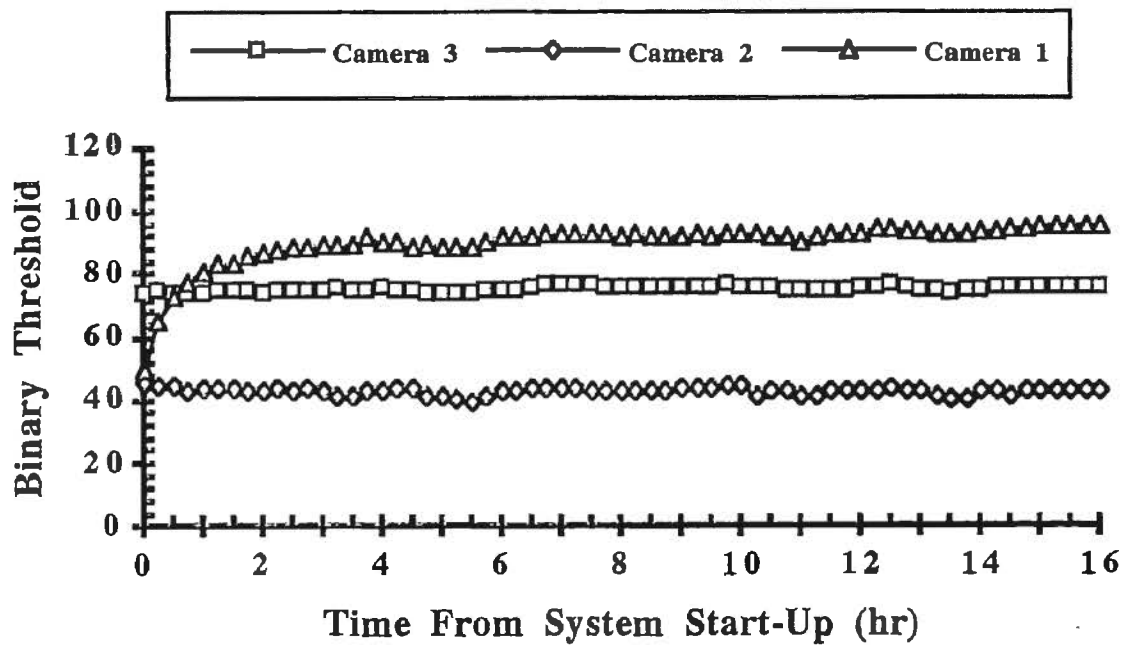

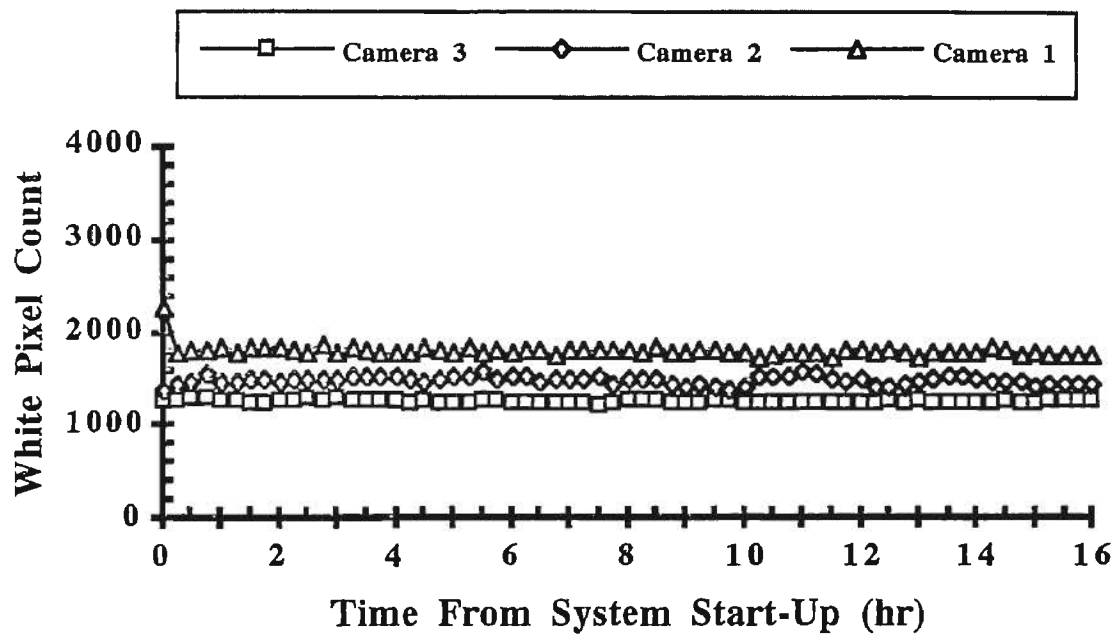Figure 4.11 Binary Threshold Adjusted by PI Compensator

Figure 4.12 White Pixel Counts Obtained from Adjusted Binary Threshold

threshold follows the uncontrolled pixel counts for each camera. The controlled pixel counts given in Figure 4.12 are much more stable than their counterparts in Figure 4.10.

A discussion of the experimental results is given in Chapter V as well as the presentation and discussion of the experimental results of the performance of the machine vision task, and the optimization of illumination design.

CHAPTER V

EXPERIMENTAL RESULTS

Sorting Performance

Accuracy

The accuracy of the sorting algorithm shown in Table I. The tabulated results in the
column labeled "overall" represent at least 100 events for each dish, an event being an
attempt by the sorting algorithm to correctly identify a dish presented in the FOV of the
camera while the complete system, robot control and machine vision, operated together.
For each event, a record was made of whether the sort was accurate and, in the case of the
smaller dishes, whether they were imaged as one or two blobs. The results in Table I
show that the sorting algorithm is very accurate with small errors occurring only in the sort
of the small plastic dish when touching, and the small spacer. The errors are due to
incorrect arc radii returned from the arc finder. We conclude that sorting accuracy would
be acceptable for commercial dishwasing operations.

Processing Time

Because processing time required to sort a row of dishes is an important
performance measure, experiments were conducted to assess this quantity. Table II shows
the results tabulated according to which dishrack slot was the first to have dishes present.
the first slot to have dishes present. "Slot One" indicates that dishes were present in the
first section of the dishbelt, requiring only an image from Camera 1; "Slot Two" indicates
that the first section was empty, requiring an image from Camera #2; "Slot Three" indicates

44

TABLE II

SORTING ERROR PERCENTAGE BY DISHTYPE

Dishtype	Imaged as One Blob			Imaged as Two Blobs			Overall		
	Events	Errors	% Error	Events	Errors	% Error	Events	Errors	% Error
Small Plastic Dish	70	1	1.43	50	0	0.00	120	1	0.83
Small Spacer	68	3	4.41	57	2	3.51	125	5	4.00
Small Ceramic Dish	50	0	0.00	67	0	0.00	116	0	0.00
Large Spacer	110	0	0.00	---	---	---	110	0	0.00
Large Ceramic Dish	125	0	0.00	---	---	---	125	0	0.00

TABLE III

PROCESSING TIME REQUIRED TO SORT DISHES

Dishtype	Slot One (ms)	Slot Two (ms)	Slot Three (ms)
Small Plastic Dish	690	591	813
Small Plastic Spacer	545	529	573
Small Ceramic Dish	652	687	682
Large Plastic Spacer	563	545	666
Large Ceramic Dish	562	593	705

that the first two sections were empty requiring an image from Camera #3 to sort. Each value presented in Table II is an average of 10 events.

We expect a longer processing interval if more than one picture must be grabbed to find a dish to sort. The results show that this is not necessarily the case. Notice that the processing time for sorting a dish first sighted in either slot one or two is essentially the same, while a dish in slot three takes considerable longer. This is because vision processing time is more greatly influenced by the complexity of an image than by the additional lines of code in the program needed to locate the first blob. Cameras 1 and 3 view the outsides of the composite three-camera FOV, which are exposed to more light than the inside. This additional light reflects off the dishbelt enough to provide some noise in the image, making it more complex to analyze the images taken from the outside cameras. Notice also, that the smaller dishes generally produce a more complex image by yielding two blobs instead of one, such that they require longer times to process. On the other hand, the results presented in Table II, indicate that this complexity has little effect on sorting reliability.

At normal dishrack speeds, the processing time for sorting and inspection is one row every 2 seconds. Since the largest sorting time is Table III is approximately 0.8 seconds, there is ample time to inspect the dishes following the sort before the next row comes into view. Since the algorithm knows which slot has the first set of dishes, it will not have to re-scan an empty slot for clean dishes, which further reduces processing time.

Inspection Performance

Food Particle Detection

The ability of the inspection process to detect the presence of food particles is represented in Figures 5.1 and 5.2. In Figure 5.1, a photograph of the test dish discussed in the previous chapter shows clearly the presence of food particles present on the dish. In the next figure, we show a binary image of the dish with shading interference eliminated by proper threshold setting and the optimum illumination design (including use of the 5-59 optical filter) discussed later in this chapter (see Figure 5.6). The black spots inside the dish indicate the presence of food particles. Notice that all but the very lightest shades of the food particles can be detected. Another note is to recall that this particular dishtype has the highest degree of shading interference to be filtered by the binary threshold. In each of the other dishes, shading interference is less, allowing more of the lighter shades of food particles to be detected.

The dish above the test dish in Figure 5.2 is a large ceramic dish that shows shading interference on the inside corners when viewed with the vision parameters for proper inspection of the small spacer dish. This illustrates the need for differing parameter settings for each dishtype, but does not pose a problem in inspecting the dish of interest, i.e., the test dish, since a window is placed around only the dish of interest. Everything outside the window, including surrounding dishes, is excluded from the inspection for holes in that particular dish.
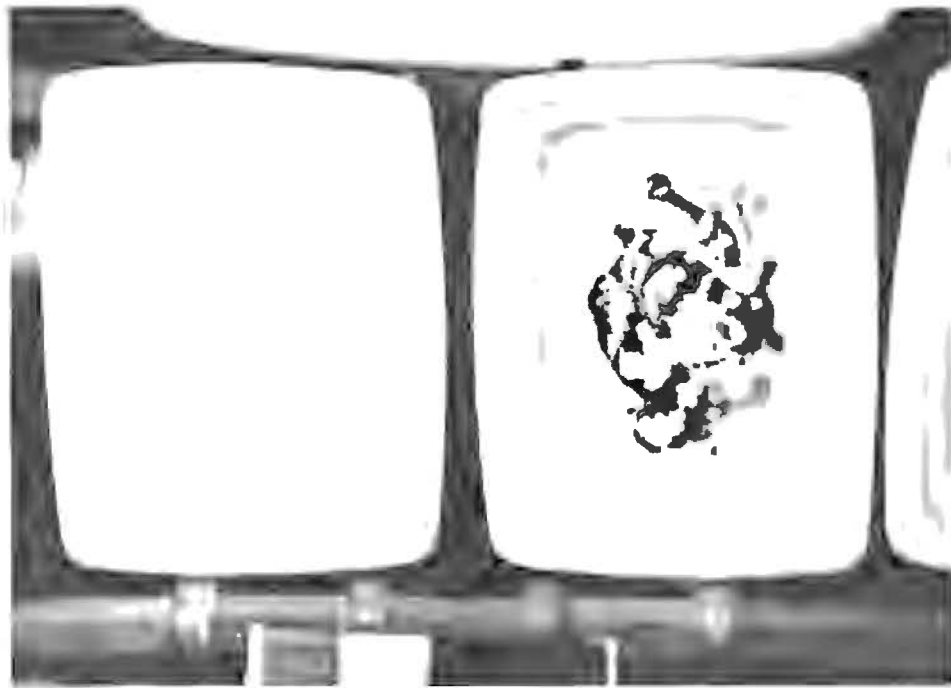
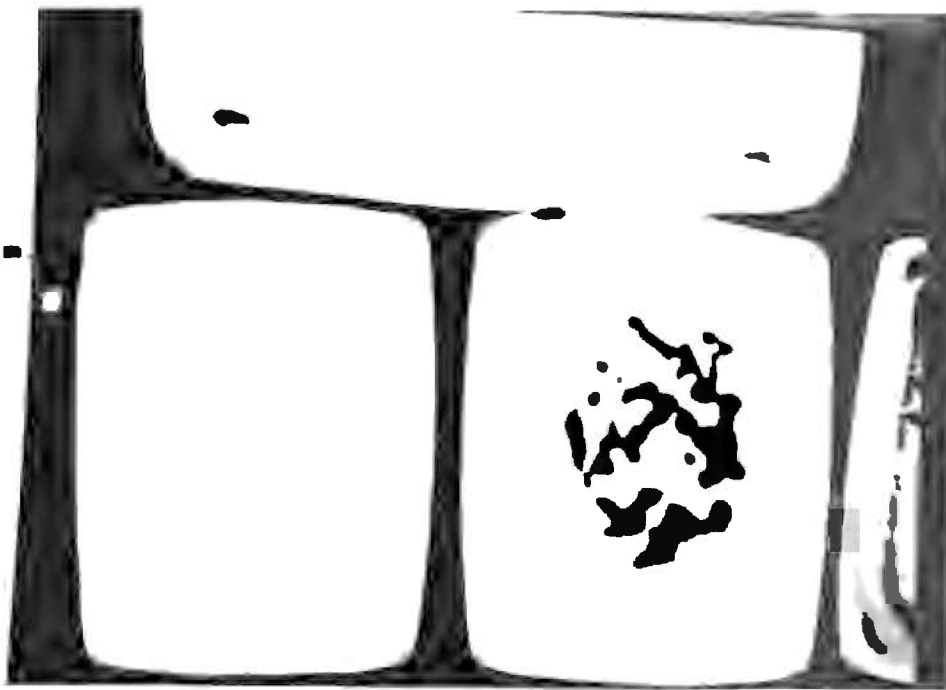Figure 5.1 Photograph of Test Dish with Food Particles

Figure 5.2 Binary Image of Test Dish with Food Particles

For comparison, Figure 5.3 shows a binary image of the dish with the same illumination as Figure 5.2, except that the filter has been removed. As in the previous figure, the binary threshold was adjusted to eliminate the shading interference. Notice how much food particle detection increases with the use of UV band pass filters in Figure 5.2.

Figure 5.3 Unfiltered Binary Image of Test Dish with Food Particles

Processing Time

Table IV presents the processing time required to inspect dishes that fill one, two, and three slots respectively. Each value is an average of 10 events. In this experiment, note that as intuition would indicate, fewer frames to analyze result in smaller processing times. Notice also that the smaller dishes, producing slightly more complex images, require more processing time to perform the vision analysis. However, even for the slowest case, the processing time is very reasonable. In fact, the combination of the longest sorting time from Table III and the longest inspection time from Table IV is approximately 1.5 seconds, well within the two second interval between rows of dishes.

TABLE IV

PROCESSING TIME REQUIRED TO INSPECT DISHES

Dishtype	Slots 1, 2, & 3 (ms)	Slots 2 & 3 (ms)	Slot 3 (ms)
Small Plastic Dish	805	343	183
Small Plastic Spacer	662	385	170
Small Ceramic Dish	878	365	190
Large Plastic Spacer	395	253	138
Large Ceramic Dish	527	222	122

Illumination Hardware

Evaluation of the illumination hardware was based on the resulting ability of the vision processor to identify food particles on a test dish as mentioned in Chapter IV. The hardware variables considered were dishrack clearance, an overhead ambient light diffuser versus no diffuser, two types of fluorescent lamps, F40 and F32T8, powered by 20 kHz ballasts, and two Ultra-Violet (UV) band pass lens filters, numbers 1-64 and 5-59, together with a naked lens. As shown in Figure 5.4, the dishrack clearance is the distance between the top of the dishrack and the bottom of the light box leading and trailing panels. A diffuser above the light box was used to diffuse the overhead ambient room lighting, which consisted of common 8 foot Fluorescent lamps found in commercial industrial room lighting applications. The UV filters were placed directly over the camera lens. The transmittance curves for these filters are given in Appendix B.

Data was collected from alternate set-ups, using the overhead diffuser with both types of lamp and then removing the diffuser and repeating the experiment with both lamp types. For each of these variations, the dishrack clearance was adjusted in increments of one inch from 2 inches to 5 inches. A minimum clearance of 2 inches is required for dishes to pass unobstructed under the light box. Again for each increment of clearance, data was taken with and without the two filters. For each filter, the lens aperture was adjusted to allow enough light to the CCD to form an image. Filter 5-59 required an aperture of 1.4; Filter 1-64 required an aperture of 2; the naked lens aperture was set at 5.6.

The data collected from these set-ups is given in Figures 5.5- 5.8. For each data point, the binary threshold was held constant while the gain was adjusted to just eliminate the shading interference. Then pixel counts were taken from a series of 20 image frames and averaged to give the values presented in the figures. The test dish, a small spacer, chosen because it is the most sensitive to shading interference, remained stationary throughout data collection to insure consistency. Also, the outside camera (Camera 1) was used to perform the imaging because shading interference is greatest at the outside.

OVERHEAD
DIFFUSER

CAMERAS

UV FILTERS

LIGHT
BOX

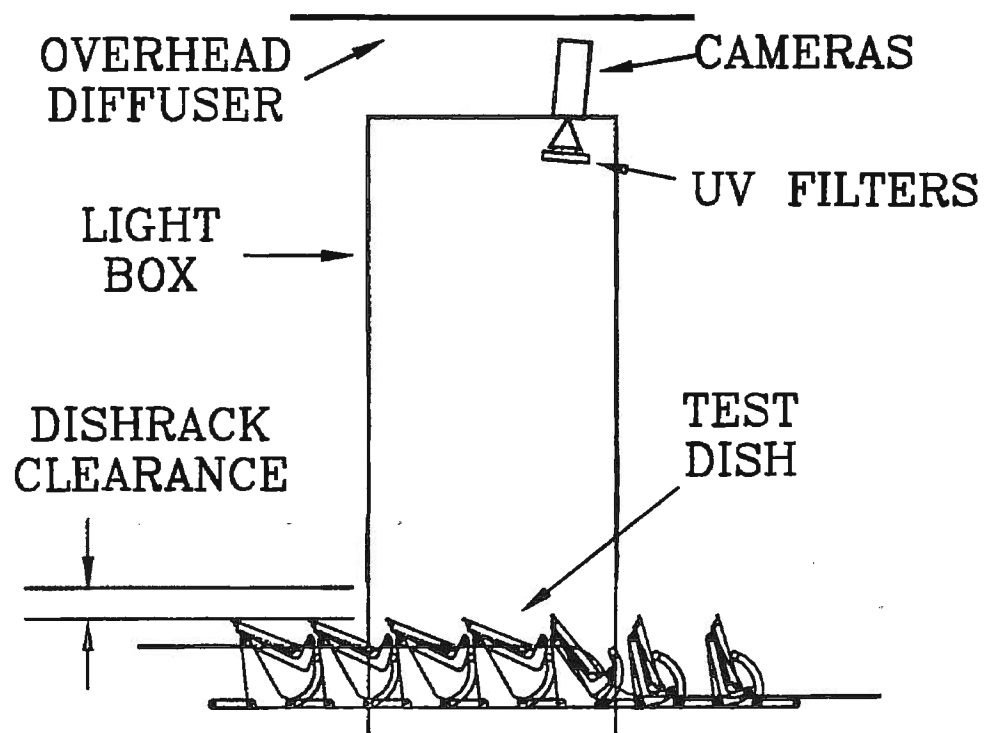DISHRACK
CLEARANCE

TEST
DISH

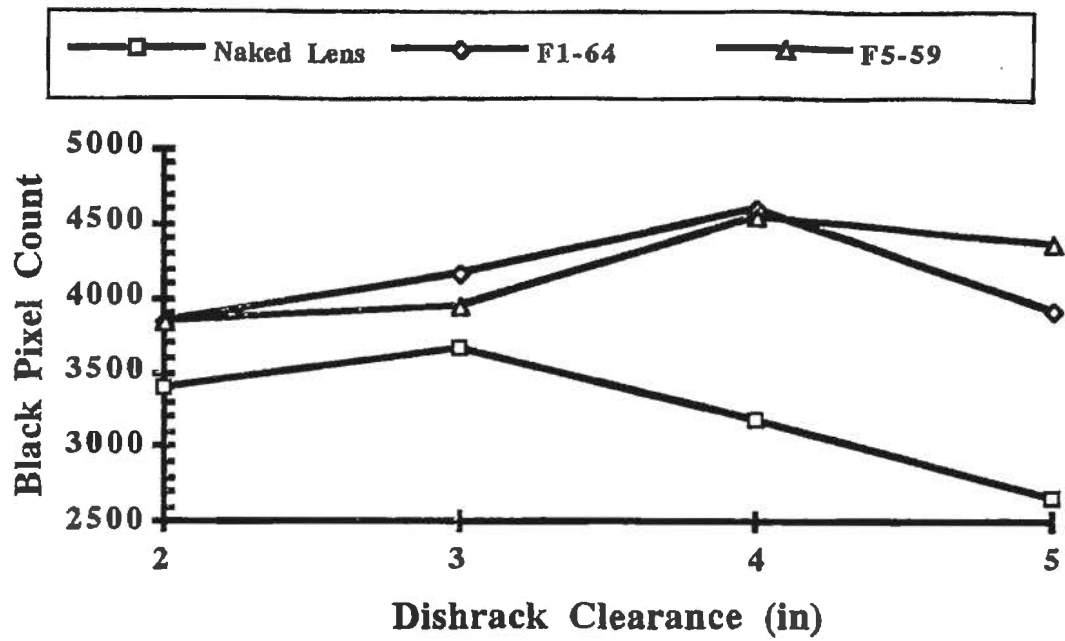Figure 5.3 Illumination Hardware

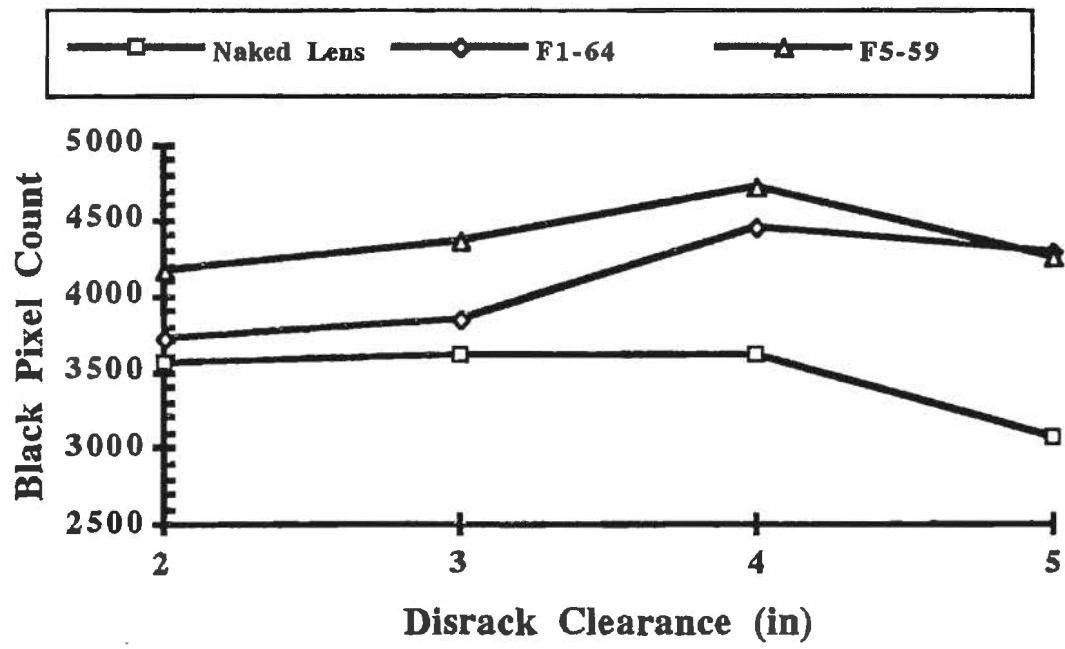Figure 5.5 Black Pixel Count vs. Clearance with F32T8 Lamps with
Overhead Diffuser

Figure 5.6 Black Pixel Count vs. Clearance with F40 Lamps with
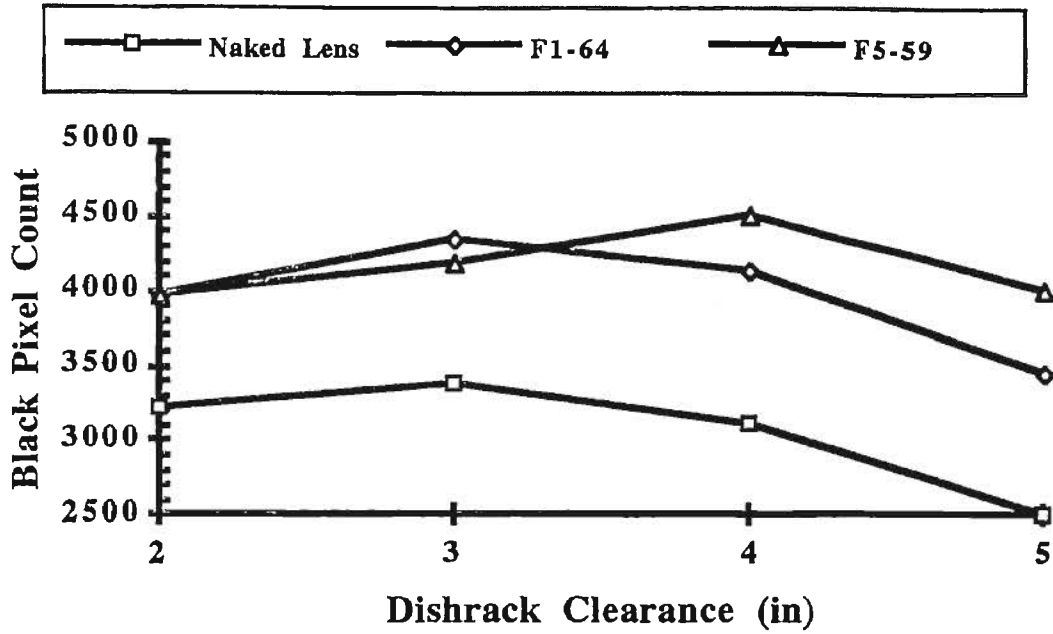Overhead Diffuser

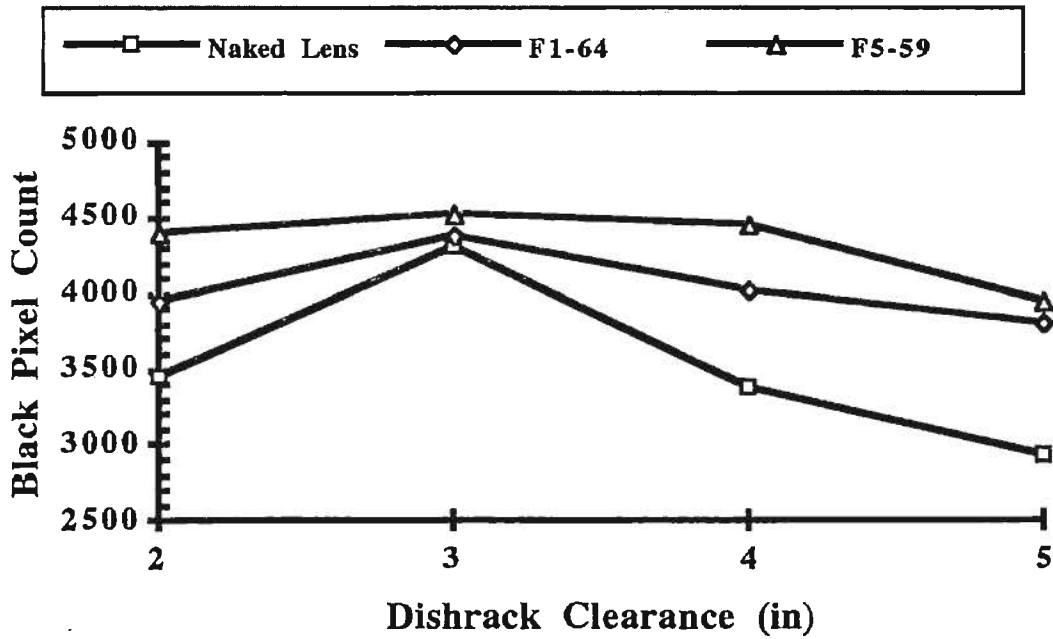Figure 5.7 Black Pixel Count vs. Clearance with F32T8 Lamps without
Overhead Diffuser

Figure 5.8 Black Pixel Count vs. Clearance with F40 Lamps without
Overhead Diffuser

To interpret the results in these figures, note that higher values of black pixel counts indicate greater ability to determine food particles, which appear black. The results show that the UV filters do indeed increase the detectability of food particles on the specular-surfaced dishes in every case. The performance of the individual filters are approximately the same, although the 5-59 filters out the shading interference slightly better. This may occur because it is a darker filter and eliminates more of the longer wavelengths of light that tend to produce shadows on specular surfaces.

The common F40 lamps provide superior results than the newer F32T8 lamps. This is most likely due to the fact that the F32T8 lamp provides a light with more longer wavelengths, similar to an incandescent lamp. Since these wavelengths are filtered by the UV glass, the illumination provides less contrast between the food particles and the dish surface, rendering it harder to detect.

The diffuser above the light box improves the contrast between food particle and dish. This is because overhead room lighting commonly causes shadowing if not diffused. The illuminating lamps where placed above the dishes during the initial stages of this work and revealed that this location was unacceptable because of the sharp shadows they produced inside the dishes.

The optimum configuration comes from Figure 5.6, indicating the use of the 5-59 UV filter, the F40 fluorescent lamps, an overhead diffuser, and a dishrack clearance of 4 inches.

CHAPTER VI

CONCLUSIONS AND RECOMMENDATIONS

Basic Results

We have focused, in this research, on the implementation of Machine Vision
Sorting and Inspection for commercial dish washing applications. The major contributions
of this research may be summarized as follows:

1. A quick reliable method of sorting dishes as they exit the dishwasher has been
   identified and implemented.

2. A fast and efficient inspection process for identifying dirty dishes as they exit the
   dishwasher has been designed and implemented. While this process appears to be
   acceptable, there is certainly room for improvement to insure any dirty dish to be
   detected as such. The consequence of passing a stained dish as clean is minimal
   since a stain will not contaminate other dishes and can culled during later
   inspections prior to the next serving.

3. Threshold imaging control in real-time has been introduced to allow for changing
   system specific variables such as fluctuating power sources that may cause changes
   in the lighting intensity, or cameras that require a warm-up period to become stable
   in their CCD output.

4. For applications were shiny or specular objects, such as the dishes used in this
   research, are inspected for dirt particles or other surface defects, indirect ultra-violet
   lighting can be obtained using fluorescent lamps and UV band pass filters,
   providing good contrast between the surface of the object and the defect by
   effectively eliminating glint.

Future Research

Finding the optimum illumination technique for machine vision applications can be somewhat limited by time and resources. It is recommended that software and graphics be developed for personal computer simulation of CCD imaging that would allow the designer to specify light box dimensions, light source types and orientations, and possibly filters. The simulated output of a CCD camera displayed graphically as these parameters are changed would eliminate the need build an experimental set-up to evaluate each idea.

It is also recommended that means be found to efficiently illuminate dishes in the dishrack without the need for a light box. Banked arrangements of fiber optic point sources may have potential for this.

Further enhancing the contrast between food particle and dish surface may be achieved by staining the food particles with a color dye during the prescrub operation. This should be investigated.

An expansion of machine vision sorting and inspection in automated dishwashing should include the sorting and inspection of silverware, cups and trays. With the fundamentals supplied in this work, modifications could easily be made to accommodate the addition of these to the complete package of an automated dishroom.

REFERENCES

1. *AdeptVision AGS User's Guide.* (1990). San Jose, CA: Adept Technology, Inc.

2. *AdeptVision AGS-GV Programming Course Manual.* (1991). San Jose, CA: Adept Technology, Inc.

3. *AdeptVision Reference Guide* . (1990). San Jose, CA: Adept Technology, Inc.

4. Chen, J. and Tretiak, O. (1992). Fluorescence Imaging for Machine Vision, *Applied Optics*, **31** (11), 1871-1877.

5. Griswold, R. (1993, March). Personal communication. Hershey Foods Corp. Hershey, PA.

6. Harding, K. (1993). Optical Considerations for Machine Vision, *Lighting Engineering for Machine Vision: Techniques and Applications*, Ann Arbor, MI: Industrial Technology Institute.

7. Koelsch, J. (1991, October). Taking on Tough Burrs, *Manufacturing Engineering*, **107**, 35-38.

8. Langenfeld, M. (1992, June). Personal communication. VARTEC, Inc, Tulsa, OK.

9. Pletized Production Pleases People Packing Pepperoni Pizzas, (1992). *The Industrial Robot*, **19** (5), 39-40.

10. Phillip, M. and Michaeli, W. (1992, July). Using Machine Vision to Improve Quality and Process Control, *Modern Plastics*, **69**, 56-57.

11. Schaffer, G. (1986). *Machine Vision: A Sense for CIM.* Machine Vision: Capabilties for Industry, Dearborn, MI: Society of Manufacturing Engineers.

12. Schwind, G. (1992, July). Palletizing Robots: Flexible Backsavers, *Material Handling Engineering*, **47** 47-50.

13. Stephan, P. (1993). Robots in the Press Shop, *Robotics Today*, **6** (1), 3-4.

14. Tech Update, (1992, February). *Manufacturing Engineering*, **108** 20+.

15. Vaccari, J. (1992). Arc-welding Robots Reap Large Savings, *American Machinest*, **136** 43-44.

16. VanDommelen, C and Harding, K. (1990). Lighting Design Rules for Machine Vision, *RIA/ESD Robotics and Vision Conference*, Ann Arbor, MI: Industrial Technology Institute.

17. Wong, K. (1992). Machine Vision, Robot Vision, Computer Vision, and Close-range Photogrammetry, *Photogrammetric Engineering & Remote Sensing,* **58** 1197-1198.

18. Zuech, N, (1987). *Gaging With Vision Systems.* Dearborn, MI: Society of Manufacturing Engineers.

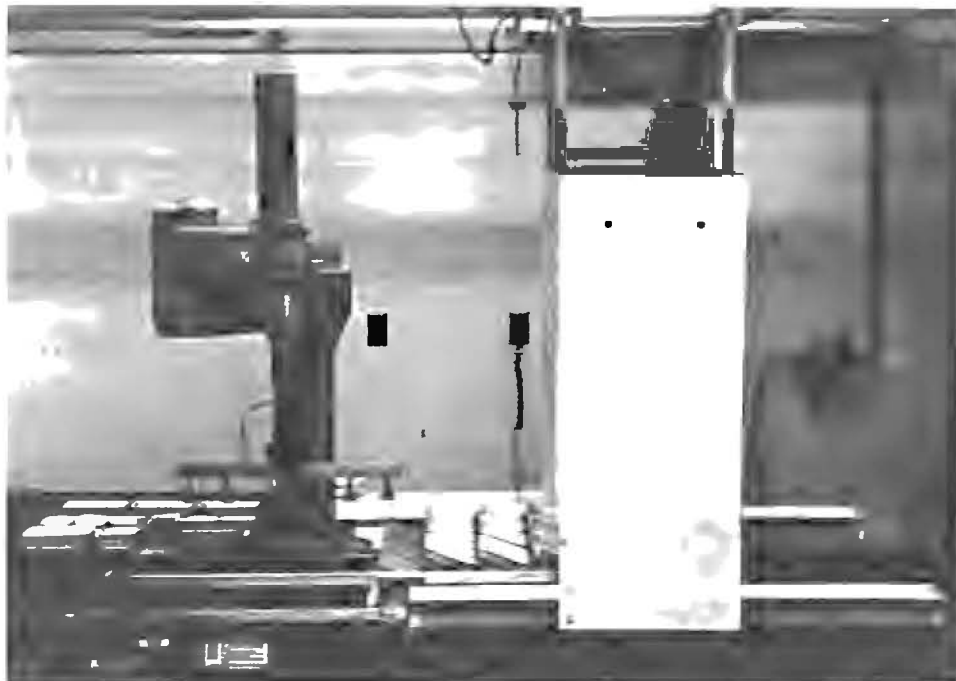APPENDIXES

APPENDIX A

PHOTOGRAPHS OF THE EXPERIMETAL SET-UP

Figure A.1 Experimental Set-up with AdeptOne SCARA Robot, Six-
point Vacuum-type Gripper, Modified Dishrack in the
Up Position Mounted to the Conveyor Belt, Light Box,
and Cameras

Figure A.2 Side View of Light Box with Side Panel
Removed, Cameras Overhead, and
Dishrack Showing Closed, Transition,
and Opened Positions.
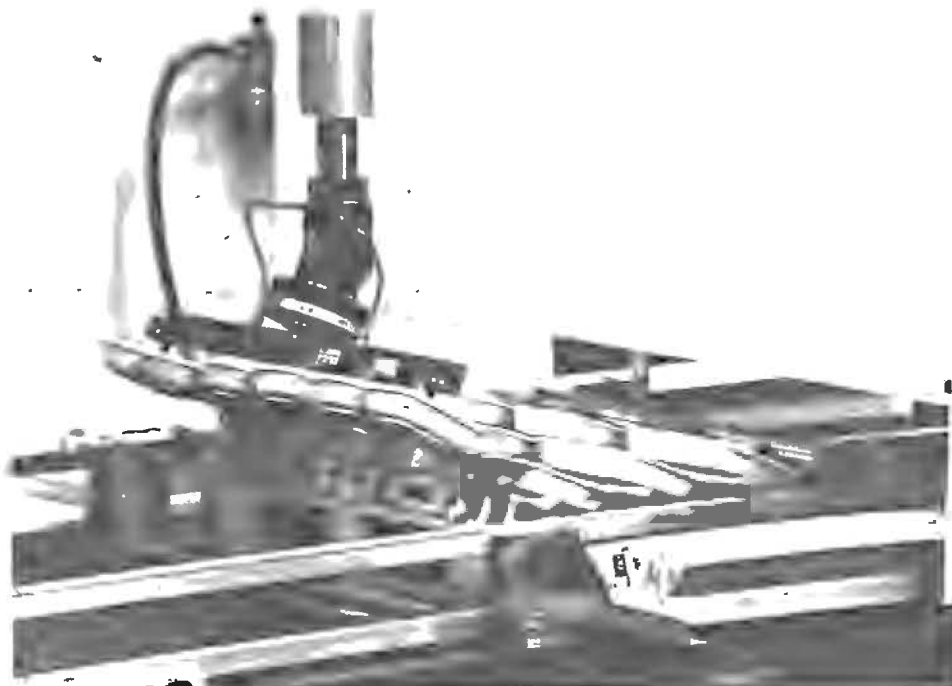
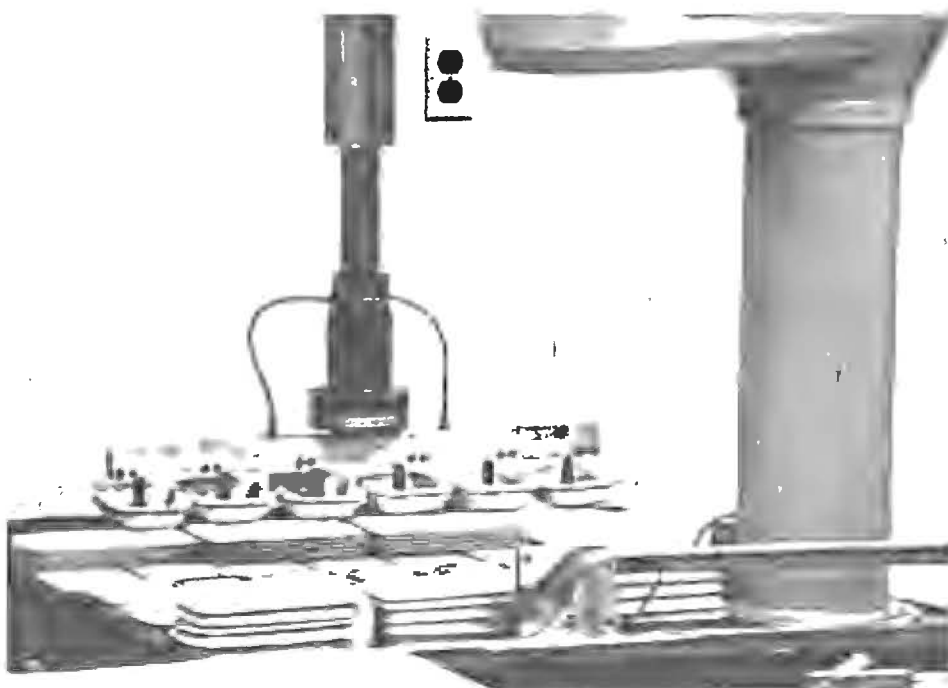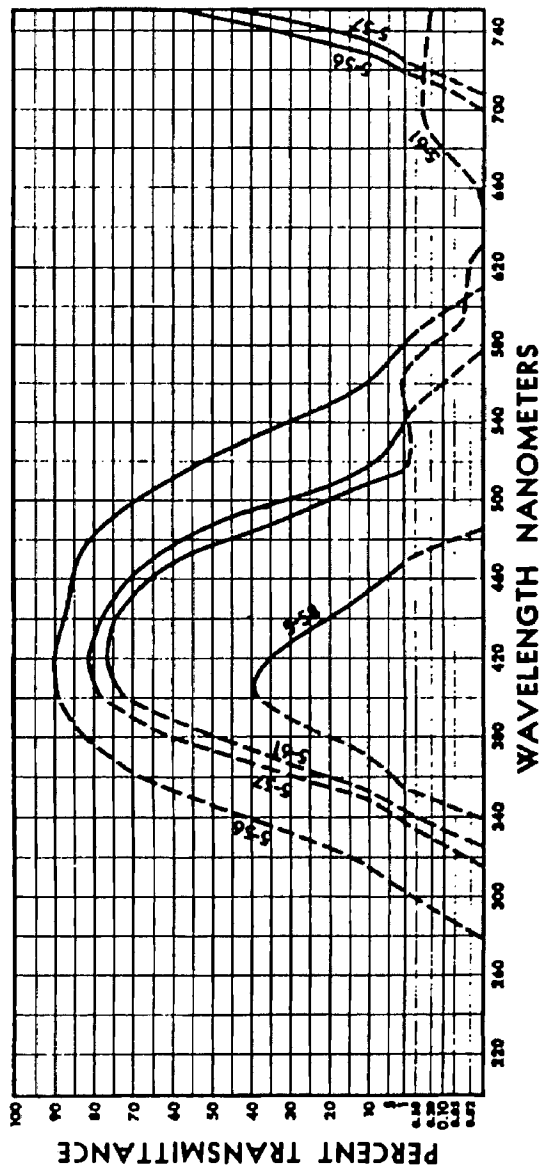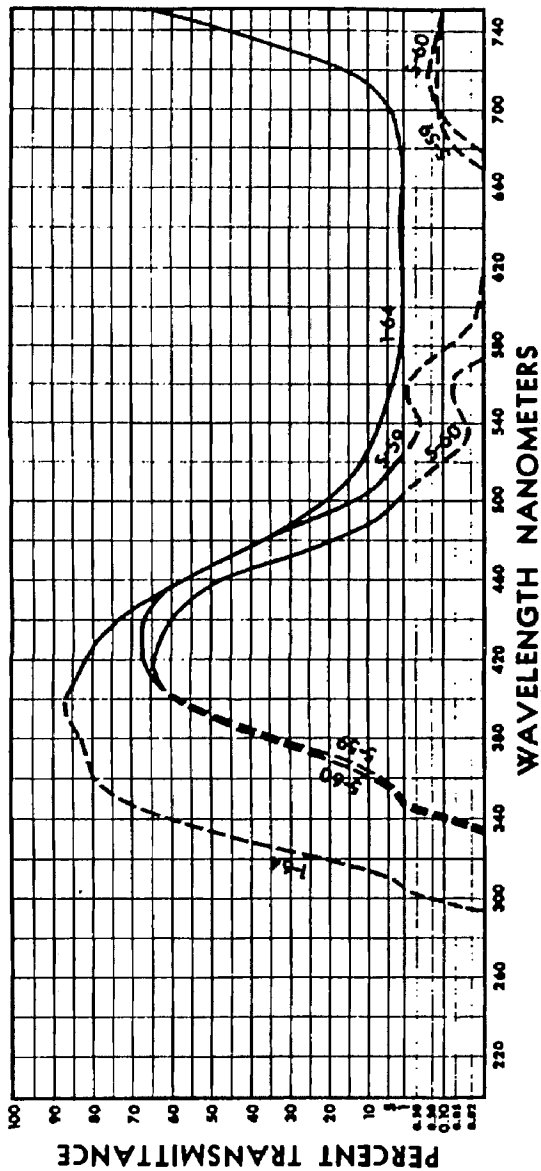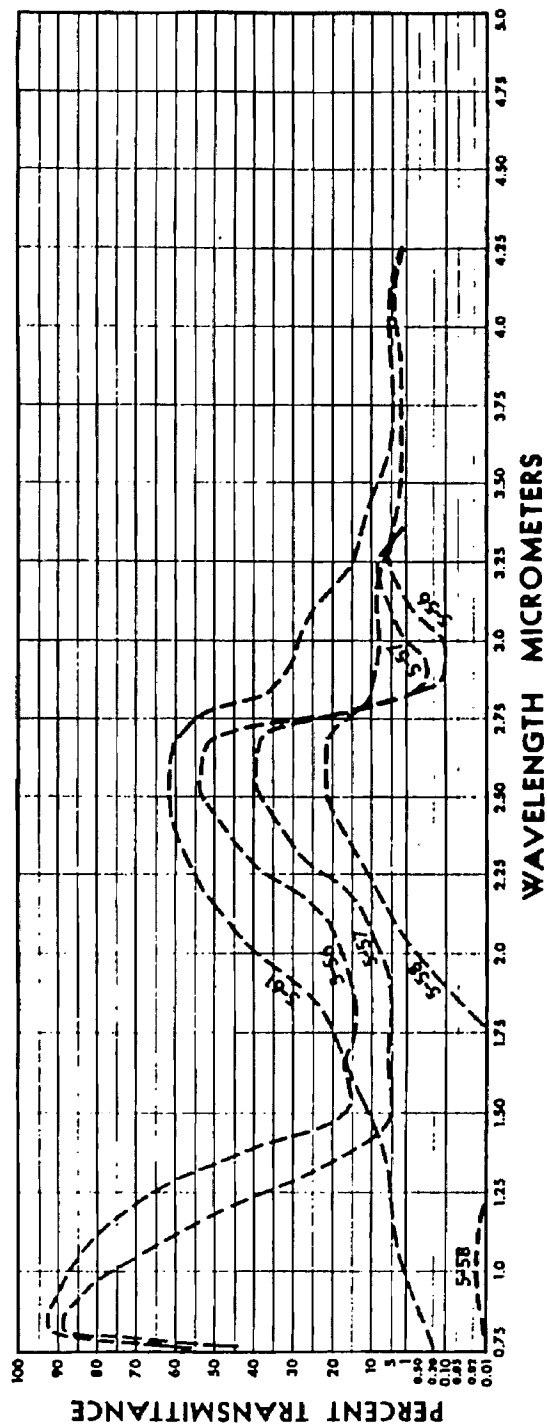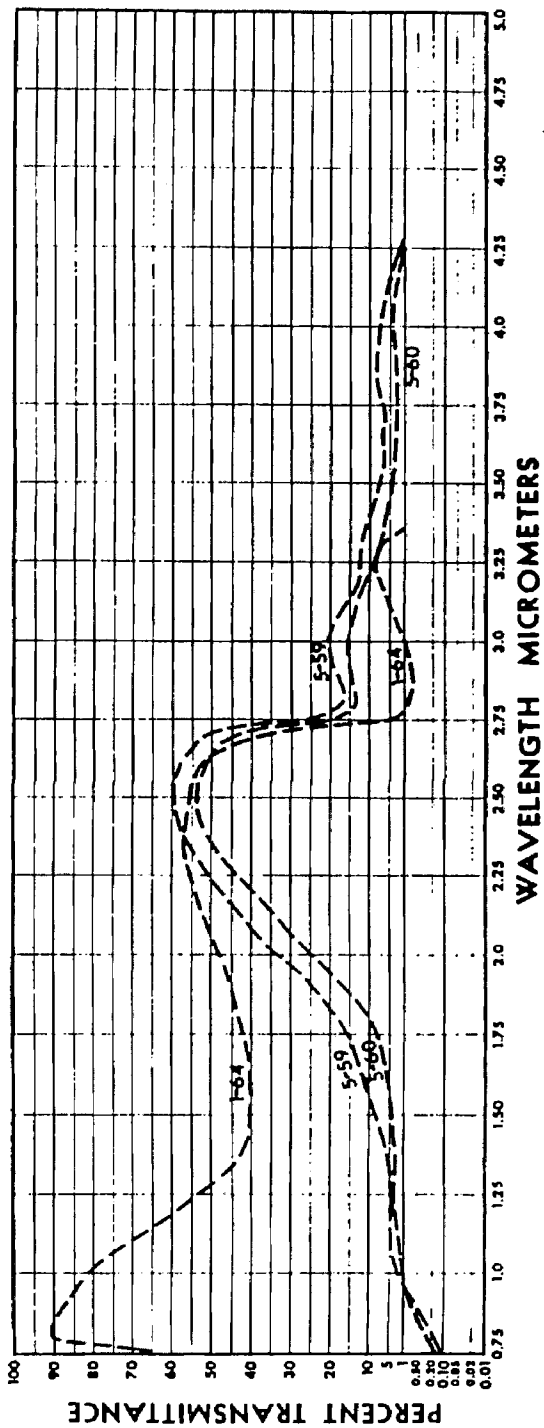Figure A.3 Six-Point End-Effector Removing Dishes from the Dishrack

Figure A.4 Dishes Being Placed in Proper Stacks by the Robot End-Effector

APPENDIX B

TRANSMITTANCE CURVES FOR ULTRA-VIOLET

FILTERS 1-64 AND 5-59

VITA

Anthony K. Johnson

Candidate for the Degree of

Master of Science

Thesis: MACHINE VISION SORTING AND INSPECTION IN COMMERCIAL AUTOMATIC DISHWASHING

Major Field: Mechanical Engineering

Biographical:

Personal Data: Born in Lincoln, Nebraska, January 6, 1967, the son of Galen K. and Marie K. Johnson.

Education: Graduated from Ardmore Senior High School, Ardmore, Oklahoma, in May, 1985; received Bachelor of Science Degree in Mechanical Engineering from Oklahoma State University in May, 1991; completed requirements for the Master of Science Degree at Oklahoma State University in July, 1993.

Professional Experience: Teaching Assistant, Department of Mechanical Engineering, Oklahoma State University, from January, 1990, to May, 1990; Junior Project Engineer, Oryx Energy, Breckenridge, Texas, from May, 1990, to August, 1990; Teaching Assistant, Department of Mechanical Engineering, Oklahoma State University, from August, 1990, to May 1991; Facilities Project Engineer, Chevron U.S.A. Inc., from May 1991, to August, 1991; Graduate Research Assistant, Department of Mechanical Engineering, Oklahoma State University, from August, 1991, to May, 1993.

Affiliations: ASME, SME, MVA of SME, Pi Tau Sigma, Tau Beta Pi