DOT/FAA/AR-97/69

Office of Aviation Research
Washington, D.C. 20591

Automated Inspection of Aircraft

April 1998

Final Report

U.S. Department of Transportation
Federal Aviation Administration

NOTICE

1. Report No. DOT/FAA/AR-97/69	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle AUTOMATED INSPECTION OF AIRCRAFT		5. Report Date April 1998
		6. Performing Organization Code
7. Author(s) C. J. Alberts, C. W. Carroll, W. M. Kaufman, C. J. Perlee, and M. W. Siegel		8. Performing Organization Report No.
9. Performing Organization Name and Address Carnegie Mellon Research Institute, A Division of Carnegie Mellon University 700 Technology Drive, P.O. Box 2950 Pittsburgh, PA 15230-2950		10. Work Unit No. (TRAIS)
		11. Contract or Grant No. FAA Grant No. 94-G-018
12. Sponsoring Agency Name and Address U.S. Department of Transportation Federal Aviation Administration Office of Aviation Research Washington, D. C. 20591		13. Type of Report and Period Covered Final Report
		14. Sponsoring Agency Code AAR-430

15. Supplementary Notes

Federal Aviation Administration, William J. Hughes Technical Center COTR: David Galella, AAR-433

16. Abstract

This report summarizes the development of a robotic system designed to assist aircraft inspectors by remotely deploying non-destructive inspection (NDI) sensors and acquiring, processing, and storing inspection data.

Carnegie Mellon University studied the task of aircraft inspection, compiled the functional requirements for an automated system to inspect skin fastener rows, and developed a conceptual design of an inspection robot. A prototype of the robotic inspection system (the Automated Nondestructive Inspector or ANDI) was developed. The first phase of system development resulted in a laboratory system that demonstrated the abilities to adhere to the surface of an aircraft panel and deploy a standard eddy-current sensor. The second phase of development included enhancing the mechanics, adding video cameras to the robot for navigation, and adding an on-board computer for low-level task sequencing. The second-phase system was subsequently demonstrated at the FAA's Aging Aircraft NDI Validation Center (AANC). During the final phase of development, emphasis was placed on the enhancement of the robot's navigational system through automated recognition of image features captured by the navigation cameras.

A significant development effort remains to be accomplished before this robotic inspection technology is suitable for operational deployment. Outstanding development issues include: (1) reducing the weight of the robot so that it is more comfortable to lift and position on the aircraft; (2) improving the mechanical reliability and speed of the system; (3) minimizing the scratching of the skin surface by the suction cups and eddy-current sensors; (4) reduction or elimination of the umbilical cable; and (5) automation of the manually controlled operations. To commercialize the technology, a new mechanical system would need to be designed and built incorporating the lessons of this work. Further integration of the software system with additional human-computer interaction features would also be required to simplify operation by the end user. However, before efforts are undertaken to commercialize this technology, a rigorous cost and operational benefit analysis should be accomplished to justify further development.

17. Key Words Robotics, NDI, Eddy Current, Data Management, Sensor Deployment, Navigation, Path Planning, Robot Control, Sensor Calibration, Inspection Robot, ANDI		18. Distribution Statement This document is available to the public through the National Technical Information Service (NTIS), Springfield, Virginia 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 85	22. Price

Form DOT F1700.7 (8-72)   Reproduction of completed page authorized

TABLE OF CONTENTS

Page

LIST OF ILLUSTRATIONS

LIST OF TABLES

EXECUTIVE SUMMARY

Carnegie Mellon University (CMU) conducted a research program for the Federal Aviation Administration's (FAA) National Aging Aircraft Research Program to investigate robotic tools to assist aircraft inspectors by automating the collection, archiving, and post-processing of inspection data. Potential benefits of robot-assisted inspection include improved flaw detection and inspection repeatability, reduced inspection time, electronic capture and archive of inspection data, and improved safety for inspectors.

The results of this program established the technical feasibility of using robotic nondestructive inspection (NDI) systems in major aircraft maintenance facilities. This program was funded through the William J. Hughes Technical Center and the U.S. Bureau of Mines. USAir supported the project by providing technical guidance from experienced aircraft inspectors and access to aircraft in its Pittsburgh, PA, maintenance facilities.

During a preliminary phase of work under this program, CMU studied the task of aircraft inspection, compiled the functional requirements for an automated system to inspect skin fastener rows, and developed a conceptual design of an inspection robot. The purpose of the robotic system was to automatically deploy conventional sensors used by aircraft inspectors. The system was designed to be sufficiently flexible to incorporate new sensor technologies, including those developed by other organizations participating in the FAA's National Aging Aircraft Research Program.

This report describes the multiphase design and development of a prototype robot (the Automated Nondestructive Inspector or ANDI) capable of walking on an aircraft fuselage and, using an eddy-current probe, inspecting the skin for cracks and corrosion.

The first phase of system development resulted in a laboratory system that demonstrated the ability to adhere to the surface of an aircraft panel in various orientations (e.g., vertical and inverted surfaces). It was able to deploy a standard reflectance mode eddy-current sensor and scan a portion of a rivet line for surface and subsurface flaws. The eddy-current instrument was calibrated using the same standard specimens used for manual inspections. The robot could walk over a horizontal surface as directed by an operator via computer interface. Results of the first phase are documented in report DOT/FAA/CT-94/23, June 1994.

The mechanical system was enhanced during the second phase. It was able to walk at any orientation around the circumference of the fuselage. A marking apparatus for indicating the locations of suspected flaws was mounted on the sensor platform. The control system was implemented using a distributed network of computers installed on the robot and on the ground. Intuitive, graphical operator interfaces were added to the system to simplify operator control of the robot. Video cameras were installed on the robot to give feedback about the robot and its surroundings and allow the operator to remotely align the robot to the rivet line. Automatic alignment algorithms were developed for the robot and independently tested in the laboratory. The robot was demonstrated on a DC-9 nose section during the 1994 Air Transport Association (ATA) NDT Forum hosted by the FAA's Aging Aircraft NDI Validation Center operated by

Sandia National Laboratories in Albuquerque, New Mexico. Results of the second phase are documented in report DOT/FAA/AR-95/48, June 1996.

The robot's control system was the focus of development during the program's third and final phase. The automatic alignment algorithms were applied to the robot vision system. A three-dimensional rendering of the robot was added to the operator workstation interface to assist the operator in visualizing the robot's status.

In its current state, the robot can be remotely operated by an inspector. Several operational limitations of the robotic inspection system were exposed during the laboratory testing and field demonstration. Specific problems include the prototype robot's speed, absence of a high level of operator control, and mechanical reliability. As a result of these experiments, insight has been gained into means for improving speed, ease of operation, and mechanical performance. Future development efforts should be directed towards addressing these mechanical issues, obtaining more autonomous robot operation by providing inspection path planning, higher level operator controls, automated signal interpretation for alarm generation, and an implementing an architecture for data management.

1. INTRODUCTION.

Carnegie Mellon University (CMU) has conducted a research program for the Federal Aviation Administration's (FAA) National Aging Aircraft Research Program to develop robotic tools to assist aircraft inspectors by automating the collection, archiving, and post-processing of inspection data. The project established the technical feasibility of using robotic systems in aircraft maintenance facilities. Although operational concerns were used in the selection of appropriate robotic technologies for use in maintenance hangars, detailed operational and economic issues associated with robotic inspection of aircraft are beyond the scope of this project.

Manual inspection techniques do not incorporate extensive data acquisition and archiving. Currently, inspection results are recorded on paper and, therefore, inspectors do not have easy access to previous inspection results. This information, if stored in an electronic database, could be accessed easily and used to identify problem areas on an aircraft or on a specific class of aircraft.

The benefits anticipated to be derived from the successful commercial deployment of an automated inspection system are:

- Improved Detection—The automated inspection system will consistently maintain a high probability of detection from start to finish.

- Improved Repeatability—The robotic system will perform the inspection in the same way every time.

- Reduced Aircraft Downtime—The automated system will help an inspector complete more inspections per shift.

- Electronically Retrievable Inspection Data—The system will keep a continuing record of the development of structural flaws for post-processing and analysis.

- Improved Safety for Inspectors—The use of robots to inspect hazardous areas will reduce risks to the health and safety of inspectors.

The emphasis of CMU's development effort was to provide inspectors with tools designed to help them do their jobs more efficiently. Inspection tasks are divided between the automated inspection system and the human inspectors. The automated system will deploy the sensors in a consistent manner and process sensor signals for any abnormal indications while the inspectors will monitor the system and will be required to make final judgments on unusual sensor readings. The inspector monitoring the system will be required to investigate any areas on the aircraft that produced the abnormal readings; the inspector will still be responsible to decide if a flaw exists at that location. Thus, the robot will deploy the sensors while the inspector will interpret the meaning of the data. A major benefit of the deployment of an automated inspection system is not only the ability to gather inspection data, but also the ability to save this data for future analysis.

A data management system can contain multiple inspection methods such as eddy-current and visual data, as well as other relevant inspection information.

The task of aircraft skin inspection was selected as the first application to be automated because it is a highly repetitive task that is suitable for robotic augmentation. During this inspection task, the aluminum skin around the fasteners that attach the skin to the aircraft frame is examined for cracks. The skin is also examined for corrosion that may have developed during the years that the aircraft has been in use.

The long-term objective of this project was to develop an inspector's assistant comprising automated robotic data acquisition, operator visual and graphical interfaces, and integrated data management. The system, called the Automated Nondestructive Inspector (ANDI), would alert inspectors as perceived abnormalities are encountered; the inspectors would then interpret the data and decide what action, if any, to take. The development of ANDI has demonstrated the feasibility of using robotic systems at the airlines' major inspection and repair facilities.

2. PROGRAM HIGHLIGHTS.

CMU approached the system development in phases. The product of each development phase would then be a system with a useful subset of the complete system's capabilities.

2.1 PRELIMINARY PHASE, CONCEPTUAL DESIGN.

Work on the development of an automated inspection system for aircraft skin inspection was initiated in May 1991. The initial work comprised reviewing airline inspection procedures and state-of-the-art robotics technology through visits to maintenance hangars, discussions with airline maintenance personnel, and participation in relevant conferences. A conceptual design of an automated inspection system was developed that addressed the requirements of the inspection process and the hangar environment. This effort was completed in December 1991 [1].

2.1.1 Mechanical System.

The mechanical device was envisioned as a cruciform robot that would adhere to the aircraft fuselage using suction cups. The robot was designed to operate upon the central areas of the fuselage and was not designed to inspect the high-curvature areas such as the nose of the aircraft. The cruciform geometry and range of precise motion was well matched to deploying eddy-current sensor to lines of rivets. An umbilical comprising air, electrical power, and electronic control lines would connect the robot to the control area. Air passing through aspirating ejectors would create a vacuum for the suction cups. The robot was designed to deploy a suite of sensors, and physically mark the locations where abnormalities were encountered.

2.1.2 Sensors.

The robot would deploy eddy-current sensors on the skin to test for surface and subsurface cracks and/or corrosion, which may have begun to form. Contact sensors would sense any obstacles directly in the path of the robot. Proximity sensors would also be used to sense for obstacles in the robot's path. In addition, small on-board cameras focusing on the fuselage surface would

provide the inspector with images of the skin. At least two cameras would be used, one focusing at short distances to allow inspectors to view the fasteners and a second focusing at long distances to allow inspectors to look at large areas of the fuselage.

2.1.3   Control System.

Based on the map of a specific aircraft contained in the database and the inspection to be performed, the control software would generate a path for the robot to follow. The path would avoid all known obstacles on the surface of the aircraft. Sensors would search for the skin seams on the aircraft to establish the location of the robot; the seams and seam intersections provide benchmarks to identify the robot's position on the fuselage surface. Between such benchmarks, dead reckoning would guide the robot. The software would also control the two basic motions of the robot: scanning and walking. During a scan the robot is stationary while it deploys the eddy-current sensor over a fastener line and takes readings as the sensor moves along the line. The robot walks as it physically moves from one location on the aircraft to another.

2.1.4   Data Management System.

Information on the types and locations of flaws would be stored in a database. The database would contain a map of the surface of each type of aircraft. Each specific aircraft would have a record of flaws that would describe the locations of repairs and flaws; a map and the record of flaws would give all of the information for a specific aircraft. Locations yielding sensor signal abnormalities that have not reached their respective thresholds for cracks or corrosion would be noted as well. This data would be shared with the entire maintenance facility as well as with all of the carrier's maintenance facilities nationwide. Inspectors could use this information to track the evolution of the growth of a crack or corrosion. Statistical analyses would be performed on the data to note trends for a specific aircraft, a class of aircraft, or for an operator's fleet. The database would also be used by the control system for path planning for the robot; based on the inspection being performed and the map of the aircraft, the system would generate the path that the robot must follow. Signal processing software would be provided to analyze the data acquired by the on-board sensors. Image processing software would be developed to process images from the video cameras. Any data concerning flaws would be stored in the appropriate format as part of the database and would be efficiently organized by aircraft model, series, modification, or individual aircraft.

2.1.5   Human-Machine Interface.

This would comprise the following elements: video monitors, teleoperation controls, and a workstation with graphics capabilities. The video monitors would provide operators with the visual images from the on-board cameras. The teleoperation controls would allow the users to manually control the motion of the robot from the operator's console for those cases when manual control is necessary. The workstation would be capable of providing complete process information to the users. Before an inspection would be performed, the monitoring system would generate the robot's path, and the users would be given the opportunity to edit the path as appropriate. Automatically generated path plans would be presented as graphics (perhaps overlaid on video); alternative plans could be presented for selection by the operator. During the

3

inspection, the workstation would provide graphical images of the robot's position and orientation to the users. The operators would also use the workstation to graphically view the locations of suspected defects on the aircraft. All relevant inspection data would be accessed via the workstation. In addition, all control functions for the robot would be run by the users from the workstation environment. The operator would indicate high level tasks for the inspection system which would be controlled at a lower level by the computer.

2.2   FIRST-PHASE DEVELOPMENT, EARLY PROTOTYPE SYSTEM.

Work building a prototype of the cruciform robot design and developing software to support its motion began May 1992. This was completed in January 1993 [2]. The general objective of this phase was to develop a lab prototype capable of both adhering to an aircraft panel and performing a multifrequency eddy-current inspection of a line of fasteners.

2.2.1   Mechanical System.

The first prototype of the robotic system was developed and tested. It weighed approximately 30 pounds (14 kg) and was able to adhere to the skin of an aircraft at any orientation (top, side, and bottom of the fuselage) and perform a scan. The operator could maneuver the robot into alignment with the site on the skin to be inspected and the sensor platform was then translated along one axis to perform the scan.

The robot was connected to the operator station via an umbilical cable consisting of multiple air, electrical, control, and sensor lines. Compressed air is required for the on-board pneumatic cylinders and for the creation of a vacuum by passing the air through on-board aspirating ejectors. The system's source of air was an air tank; although later versions of the system made use of air compressors as an air source.

2.2.2   Sensors.

A reflectance eddy-current sensor (also known as a pitch-catch sliding probe) was used in the prototype system. The specific probe was a Nortec SPO-1958. It is driven using an SE Systems, Inc. SmartEDDY™ 3100 eddy-current instrument. The eddy-current data acquisition and analysis software used in this phase of the project was the SmartEDDY™ 3.0 package provided by the instrument manufacturer. The sensor was run simultaneously at two frequencies searching for both surface and subsurface defects; the electronics effectively made this two independent sensors.

A method to ensure that the eddy-current sensors of the automated inspection system were properly calibrated was designed. To keep the automated inspection system consistent with airline procedures, the calibration standards and inspection frequencies used in manual inspection procedures were also used with the robotic inspection system. During calibration, the robot deployed the eddy-current sensor and scanned the two rows of fasteners on the standard. The parameters of the eddy-current software were adjusted to obtain the standard calibration curves. In this way, the robotic system was calibrated simultaneously at both a high and a low frequency.

Experiments were conducted with a miniature video camera to help determine the specifications necessary to include vision as a sensing means in future phases of the project. A video specification was developed outlining the recommended hardware and its placement on the robot. Prototype algorithms were investigated to enable machine identification of lap joints and rivets in video images, and to aggregate these rivets into rivet lines useful for navigating and positioning the robot on the surface.

2.2.3   Control System.

As shown in figure 1, one ground-based personal computer (PC) controlled the scanning and walking motions of the robot while a second ground-based PC controlled the data acquisition functions for the eddy-current data. There was no communication between these computers; data acquisition was initiated manually by the operator.
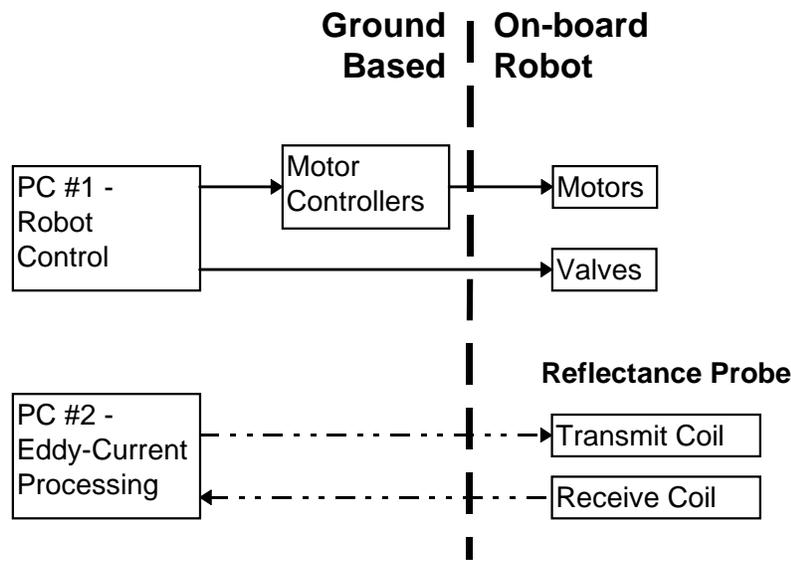
FIGURE 1. FIRST-PHASE SYSTEM ARCHITECTURE

The robot control and operator interface software was monolithic, with a single software module for both the operator interface and robot control. The eddy-current data acquisition was controlled by the stand-alone eddy-current system software discussed in section 2.2.2.

Software development was initiated for the guidance and control of the robot. The software was used to control the robot as it positioned and moved an eddy-current sensor along a length of skin seam.

2.2.4   Human-Machine Interface.

The operator interfaces for the robotic system were displayed on two separate color video monitors for control and inspection. A simple menu-driven interface was provided on the first monitor for use by the operator to control the robot. The output from the commercial eddy-current system was shown on the second monitor.

<u>2.2.5   Laboratory Testing</u>.

A simulated aircraft panel manufactured by Foster-Miller, Inc., was provided to the CMU Automated Inspection of Aircraft Project by the FAA Aging Aircraft Nondestructive Inspection Validation Center (AANC). The panel's curvature and fastener layout are similar to that of a Boeing 737. During the initial phase of development, the simulated aircraft panel was mounted horizontally in the lab; all testing on the panel during this phase was performed with the panel mounted in this position. Because the panel was a curved surface, the robot was also tested on a modest slope during this phase. The lab setup is shown in figure 2.

FIGURE 2. FIRST-PHASE ROBOT AND LAB SETUP

The robot was able to adhere to the skin of the horizontally-mounted aircraft panel. In addition, to confirm that the robot could hold to vertical surfaces, the robot was successfully affixed to the laboratory wall.

The walking motion of the robot was also successfully tested at orientations near horizontal or at modest slopes using the control software. By the end of the first development phase, the robot was able to scan a length of skin seam, take a step, and scan a second length of skin seam, etc. This ability to scan and step was a system capability that went beyond the objectives for this phase of the project. The walking motion was not tested while the robot was affixed to a wall.

This test was deferred until the next phase of work when a tether was added to prevent damage to the robot should it lose adherence to the skin.

Simultaneous high and low frequency scans were completed both on a calibration standard and on the simulated aircraft panel. Scans were completed by translating the sensor platform along one axis; the results of the eddy-current scans were displayed on a PC screen. Flaws were identified by examining the traces produced on the screen.

2.3   SECOND-PHASE DEVELOPMENT, REFINED SYSTEM.

Work improving the first-phase prototype inspection system began January 1993. This was completed in June 1995 [3]. At the conclusion of the second phase, the capabilities of the system had increased substantially. The robot possessed a degree of autonomous motion; it had some ability to manage data; and it was able to physically mark the surface when abnormalities were encountered. A six-week field test and demonstration of the robot on the nose of a DC-9 took place at the FAA AANC.

2.3.1   Mechanical System.

The mechanics were upgraded to allow the robot to walk in all orientations on a fuselage surface. A tether to secure the robot in the event of a fall and a device to physically mark the locations of defects (using washable stamps and markers) were incorporated into the system.

2.3.2   Sensors.

The commercial eddy-current software was replaced with custom software which was integrated with the control input and status display software. Acquisition of eddy-current data was synchronized with the robot's movements by the robot's control process. Inspection data was displayed to the operator or archived for later analysis.

Video cameras were added to the robot to enable the operator to navigate it over the surface of the fuselage and align it with the rivets to be inspected. Automatic algorithms to perform the alignment task were developed and tested under laboratory conditions; however, they were not integrated with the system at this time. The cameras also provided close-up images of the rivets being inspected and visual feedback of the robot's state to the operator.

Limit and vacuum switches were added to the robot to provide information about the robot's status to the operator and the control system.

2.3.3   Control System.

The electronic and computer system architecture was upgraded to provide a more flexible, modular architecture shown in figure 3. Multiple computers were employed to divide the processing load with communications between the individual computers using both point-to-point and network connections. A computer was installed on board the robot to control the low-level sequencing of mechanical operations and to monitor the status of the robot in real-time.
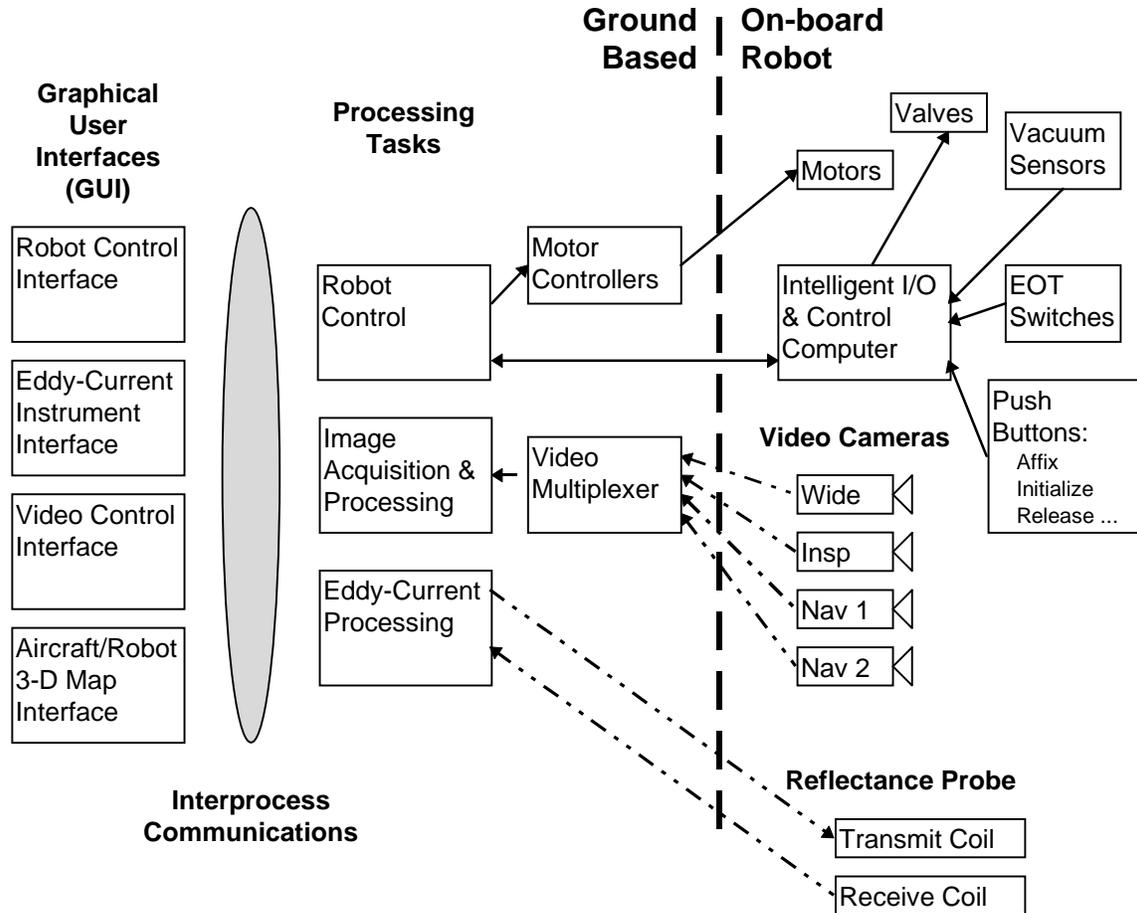
FIGURE 3. SECOND-PHASE SYSTEM ARCHITECTURE

The first-phase system's software was monolithic rather than modular, section 2.2.3. The software was redesigned to be modular during the second phase, incorporating many cooperating processes. By dividing the system into multiple processes and modules, it was possible to handle concurrent control activities and split the computational load over several computers. Multiple coordinated robot control and operator interface processes were created. Coordination of these software processes requires them to exchange information while executing. The interprocess communications entity shown in figure 3 is in control of all communications between any of the processing tasks or between the graphical user interfaces and the processing tasks. It routes text or binary messages between the software entities executing on the various computers forming the system.

By the end of the second phase, significant improvements had been made to the teleoperation of the robot. Software controlled walking enabled the robot to perform multiple steps in any one direction, allowing the robot to move several feet without operator intervention. A method to allow the operator to remotely align the robot with the rivet line using video images displayed to the operator was implemented. The acquisition and display of eddy-current data was integrated with the probe deployment and scanning motion of the sensor bridge.

2.3.4   Human-Machine Interface.

Graphical user interfaces were developed during the second phase to simplify the use of the system and the display of inspection data. These used a combination of pull-down menus, push buttons, text fields, and command lines to control the system. Feedback generated by using these controls were rapidly provided through visual indication of the robot's status, graphs representing eddy-current data, and wire-frame models of the robot and the aircraft being inspected.

The X-Windows interface was chosen for implementation of the robot system due to its portability to multiple types of computers and operating systems. The interface was implemented using a client/server model and is usable over a network. This allows multiple computers to transparently share inspection subtasks and results. The use of X-Windows also allows multiple interfaces to be simultaneously displayed in separate windows on a single computer monitor. This would enable a single operator to control and monitor several inspection robots simultaneously.

The U.S. Bureau of Mines developed a three-dimensional animation and rendering of the robot and aircraft surface in collaboration with the CMU team. This tool enabled the off-line visualization of the robot's motion and its interaction with a simulated aircraft. A simplified version of the three-dimensional model was added to the operator interface during the third phase of the project. Eventually, a complete model coordinated with the robot's status and position on the aircraft in real-time could provide an important tool to the operator for use of the inspection system.

2.3.5   Laboratory and Field Testing.

A fixture to support a simulated aircraft panel was designed and fabricated for testing the robot in the laboratory. This fixture allowed the panel to be rotated 360° and locked into any desired position. It could simulate the top, bottom, and sides of an aircraft. A trolley, for tethering the robot, ran above the panel. Figure 4 shows the panel mounted in the fixture (the robot is visible sitting on a cabinet in the foreground).

The second-phase robot was demonstrated at the Air Transport Association's Nondestructive Testing Forum on November 2, 1994, at the FAA Aging Aircraft NDI Validation Center (AANC). Prior to this demonstration, extensive tests were conducted in the laboratory and at the AANC hangar. A picture of the robot walking on the DC-9 nose section is shown in figure 5. The safety tether had a passive mobile follower that ran along a strap suspended above the nose section. The umbilical cable is supported to the right of the robot and had sufficient slack to accommodate a moderate range of motion.

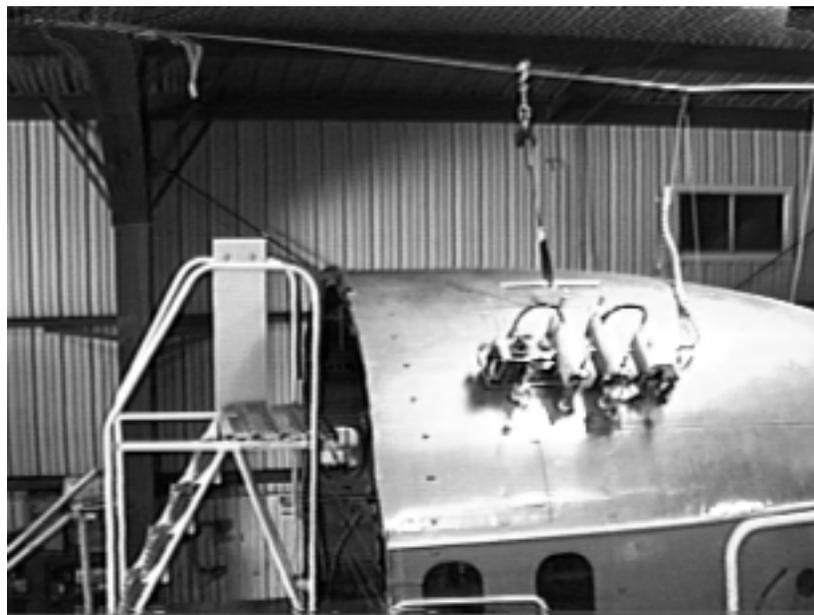FIGURE 4. LABORATORY TEST PANEL FIXTURE WITH TETHER

FIGURE 5. SECOND-PHASE ROBOT ON DC-9 NOSE SECTION

2.4   THIRD-PHASE DEVELOPMENT, NAVIGATIONAL AND OPERATIONAL ENHANCEMENTS.

The final phase of the project began June 1995 and was completed in February 1997. The emphasis of this phase of work was to improve the operational aspects of the robot to enhance reliability, autonomy, and ease of operation.

2.4.1 Mechanical System.

No major modifications were made to the basic mobility and manipulation components of the mechanical system. Mechanical improvements focused on the sensor and adhesion components. Opaque hoods were installed over the two navigation video cameras to improve the quality of the images for automatic rivet detection. The suction cup vacuum system was reworked to significantly reduce the robot's compressed air requirements.

2.4.2 Control System.

The linear motor control was improved to reduce strain on the robot while walking which had been observed during the field demonstration. Additional control links were made between the operator workstation PC and the video-processing PC to enable the video processing algorithms to be invoked by the control process and to present live video data to the operator through the computer interface.

2.4.3 Sensors.

The prototype algorithms for machine identification of rivets and the aggregation of these rivets into lines of rivets from video images were applied to the robot's machine vision system. The software was ported to the video-processing PC and the algorithms were modified based upon their performance on data from the robot on the Foster-Miller panel and a USAir DC-9.

2.4.4 Human-Machine Interface.

Components of the free standing robot animation and rendering software (section 2.3.4) were rewritten to enable their use on the operator workstation PC. A three-dimensional rendering of the robot can be displayed on the workstation monitor. Coupling this display capability to the robot control process and the aircraft map in future systems would enable the operator to monitor the inspection process from the viewpoint of a virtual camera in real-time.

2.4.5 Laboratory and Field Testing.

The robot was taken to a USAir maintenance hangar at the Pittsburgh International Airport to acquire realistic surface images using the new camera hoods for testing of the automatic rivet line detecting algorithms. Since the goal was to acquire test imagery under hangar conditions, not test the robot's mobility, the robot was manually positioned and held on the upper fuselage of a DC-9. This simplified the test setup and minimized the amount of equipment that needed to be fielded.

3. MECHANICAL SYSTEM DETAILS.

During the conceptual design phase, three potential approaches were considered for the automation of aircraft skin inspection: an overhead robot, a ground-based robot, and a surface-walking robot. The overhead-gantry approach would have consisted of a permanent structure through which aircraft would be towed for inspections; this is also called the "car wash" approach. This proved to be incompatible with airline maintenance procedures and not feasible

from an economic point of view. The ground-based-robot approach would have had a mobile platform that would carry an arm and probe manipulator to areas of aircraft that needed to be inspected. This type of robot would be difficult to incorporate into the crowded and cluttered hangar environments of commercial airlines. The third approach, the surface-walking robot, consisted of small mobile robots that would crawl over the fuselage surface. This is compatible with airline maintenance procedures and could easily be incorporated into the crowded hangar environments. The design of the surface-walking robot chosen by members of the Carnegie Mellon team was best suited for the application for which it was to be used.

A fuselage is generally a cylindrical surface with orthogonal rows of rivets. The rivets tend to run in two directions: circumferentially around the fuselage, which is called the vertical direction in this document, and approximately parallel to the horizontal axis of the aircraft from nose to tail, which is called the horizontal direction. An aircraft typically has more curvature and non-orthogonal rivet patterns at the nose and tail than at the main fuselage section; however, very few inspections are required to be performed at the nose and tail. One of the main goals of the Carnegie Mellon program was to develop useful inspection tools that have the potential to be commercialized quickly. Thus, the decision was made to design a tool that would perform most of the conventional inspections on the fuselage. Developing a prototype to inspect areas requiring few inspections was not considered in this program. Inspection personnel from USAir and other airlines clearly indicated that the robot must be able to perform large-area inspections on the main fuselage sections of commercial aircraft. In addition, there are more inspections performed along the horizontal lines of fasteners than along the vertical lines of fasteners, and there are repair patches and other obstructions scattered over any given aircraft. From this information, a decision regarding the minimum requirements for the robot was reached. It would be required to move gracefully in the horizontal direction, to move adequately in the vertical direction, and to possess some steering capability. A cruciform robot with some degree of compliance between the axes would satisfy these walking requirements.

The weight of the robot had to be kept to a minimum. Inspectors would be required to lift the robot when they affix it to the aircraft surface before initiating an inspection procedure and when they remove it from the surface after an inspection has been completed. The robot must be light enough for inspectors to lift. The ultimate goal was to produce a robot that weighed between 15 and 20 pounds (7 and 9 kg) and could both walk on the aircraft surface and deploy sensors. However, the weight goal of the prototype robot was relaxed to approximately 30 pounds (14 kg) to use commercially available pneumatic and mechanical devices to speed the development of the prototype.

Adding additional mechanical elements (e.g., a robotic arm) to deploy NDI sensors would have added significantly to the weight of the device. Weight was saved by using by the same system components to walk and to deploy sensors.

3.1   MECHANICAL DESIGN OF FIRST-PHASE ROBOT.

To affix the robot to the fuselage, active suction cups were used. Small on-board vacuum generators, one for each suction cup, provided the vacuum necessary to hold the robot to the fuselage surface. Thus, the cruciform robot was able to adhere to surfaces in all orientations,

deploy sensors, and to walk across a curved surface. The major mechanical components of the robot design is shown in figure 6; pneumatic and electrical connections are not shown in this figure.
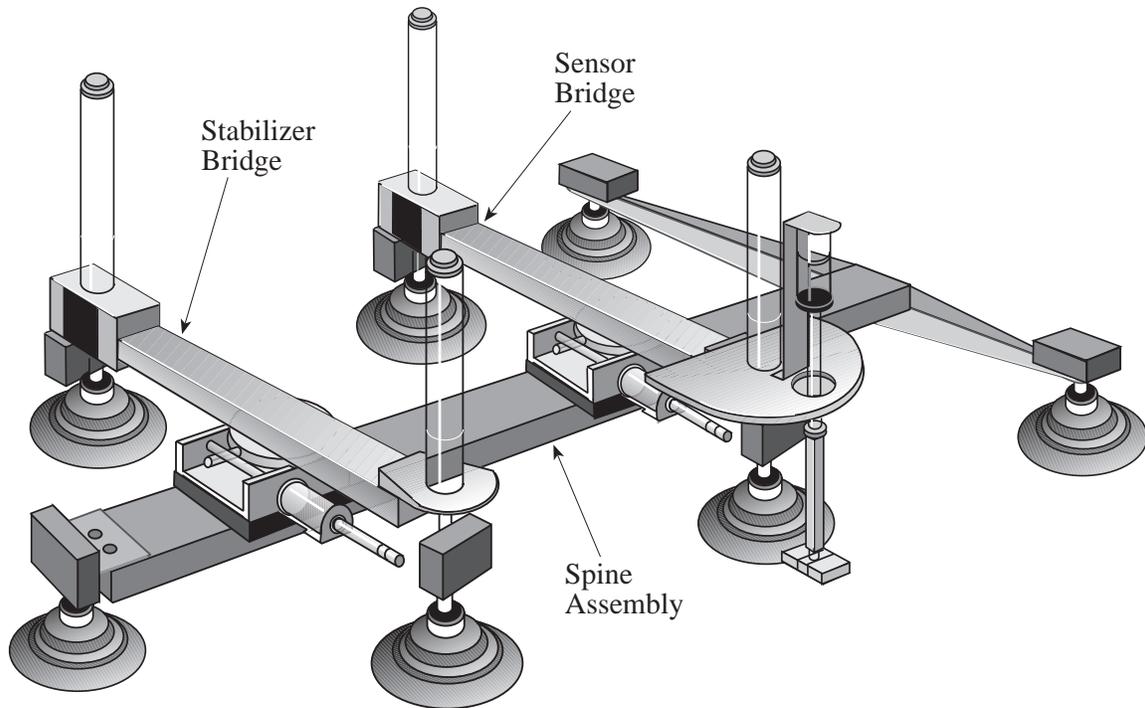
FIGURE 6. DRAWING OF THE CRUCIFORM ROBOT

The robot is comprised of a spine assembly and two bridges. The spine assembly is the main member upon which the bridges move. The sensor bridge possesses dual functionality; it is used to deploy the eddy-current sensor during scanning and is also used for support during walking. The stabilizer bridge is used only for support during walking, although it may be expanded to have an inspection role. Each of these structures and the basic robot motions will be discussed in more detail below. The relevant dimensions of the robot are shown in figures 7, 8, and 9 indicates the reference directions which will be used when describing the system components.
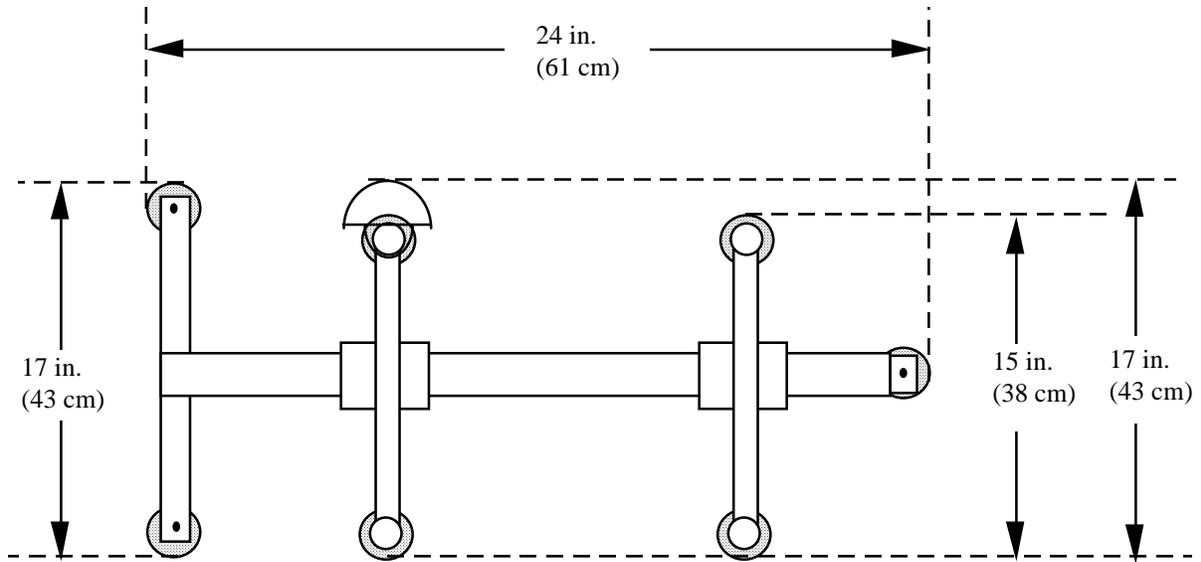
FIGURE 7. ROBOT DIMENSIONS (OVERHEAD VIEW)
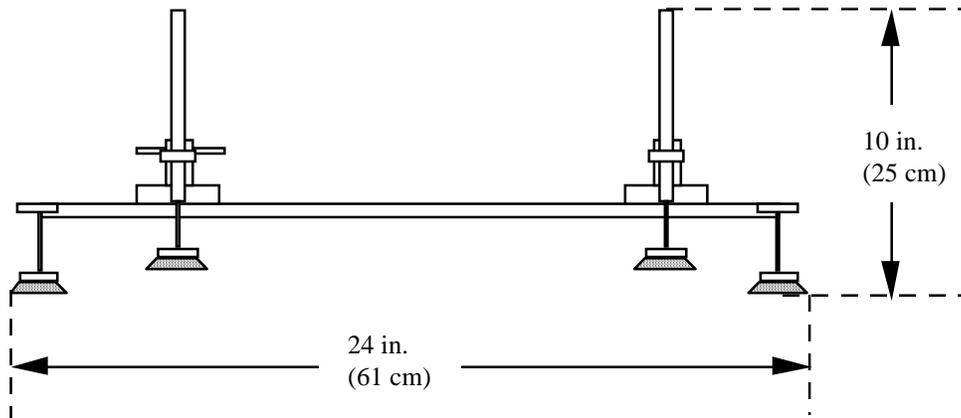
FIGURE 8. ROBOT DIMENSIONS (SIDE VIEW)

The robot design is symmetrical, permitting many possible configurations. The sensor bridge is the bridge closest to the tail end of the spine assembly. The sensor platform is mounted on the sensor bridge so that it is positioned on the left side of the robot. The stabilizer bridge is the bridge closest to the head end of the spine assembly.
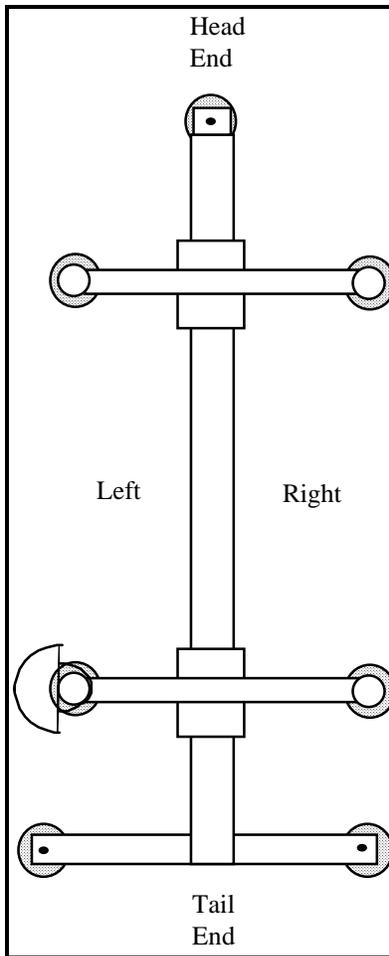
FIGURE 9. REFERENCE DIRECTIONS

3.1.1   Stabilizer Bridge.

The stabilizer bridge, shown in figure 10, is comprised of the following elements:

- Linear Motor—The positioning accuracy of the linear motor is ±0.003 inch (±0.076 mm), and the static force holding it to the platen (spine) is 180 pounds (801 N). It rides along the platen on an air cushion provided by an air bearing.

- Lead Screw Assembly—This assembly provides 6 inches (15 cm) of travel and is driven by a stepping motor.

- Lifters—These are double-acting pneumatic cylinders.

- Vacuum Pumps—Bernoulli-type vacuum pumps (one for each suction cup) generate the vacuum required to adhere to a surface. They require 1.12 cubic feet per minute (0.53 l/sec.) of air at a pressure of 73 pounds per square inch (0.5 MPa) to produce a vacuum of 28.1 inches of mercury (714 mm Hg).

- Suction Cups—These are 3 inches (8 cm) in diameter and are made of silicone rubber. At a vacuum of 28 inches of mercury (711 mm Hg), the lifting power for each suction cup on a smooth and clean surface is 44.0 pounds (196 N) in the direction normal to the surface and 24.2 pounds (108 N) in shear.

- Pivot Lock—This is a floating pneumatic-type pivoting mechanism that allows ±20 degrees of motion about its reference (locked) position. It consists of an air bearing support and a single-acting cylinder with a spring return.
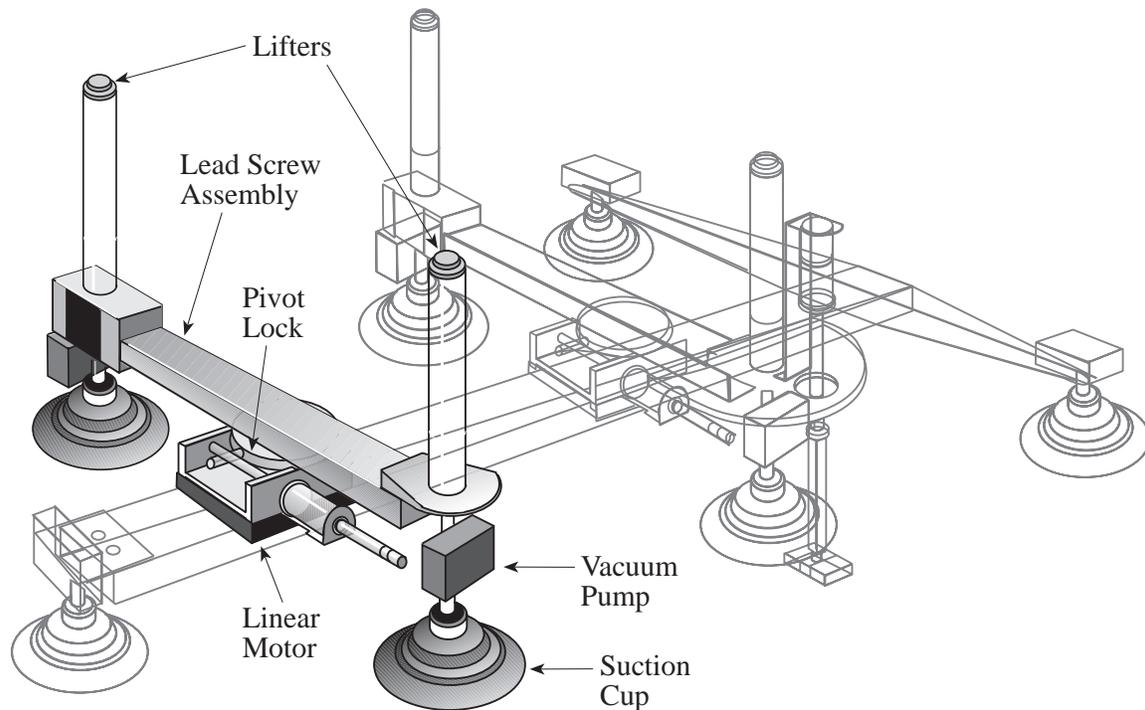
FIGURE 10. DRAWING OF THE STABILIZER BRIDGE

The stabilizer bridge can pivot with respect to the linear motor, can translate along the spine assembly, and can translate along its own axis. These motions are shown in figure 11. The stabilizer bridge can travel a maximum distance of 15 inches (38 cm) along the spine assembly, and the stroke of the bridge's lead screw assembly is 6 inches (15 cm). The pivot locks give ±20 degrees of motion as shown in figure 12. The locking mechanism consists of a single-acting/spring-return cylinder floating in an air bearing. In the position shown in figure 12, the cylinder is in its extended position, allowing the lead screw assembly to pivot with respect to the linear motor. When the cylinder is retracted, as shown in figure 13, the clamps lock the pivot into position securing the lead screw assembly perpendicular to the spine. This is the position used for the scanning and walking motions; the pivots are unlocked only when using the alignment motion. The details of these motions are described in section 4.
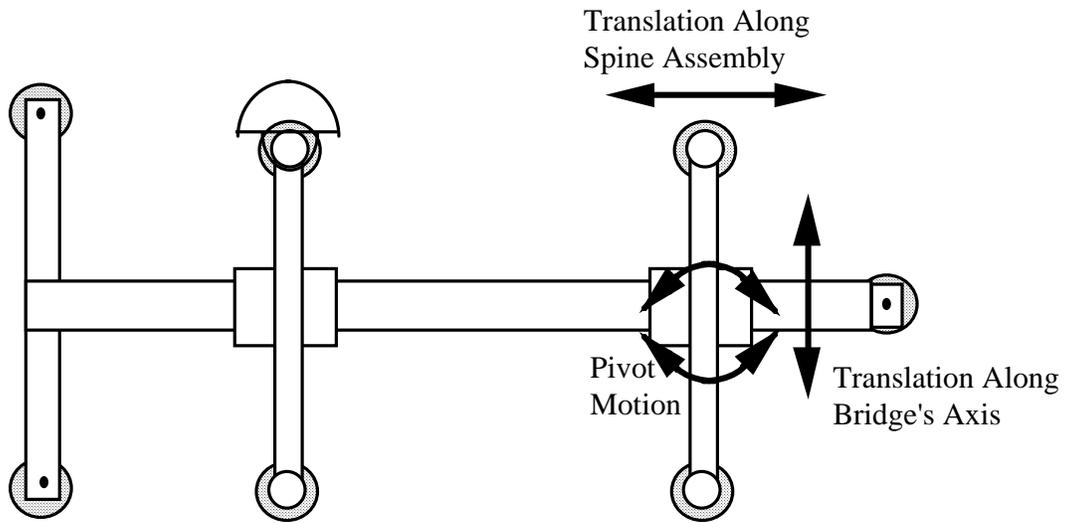
FIGURE 11. STABILIZER BRIDGE MOTIONS
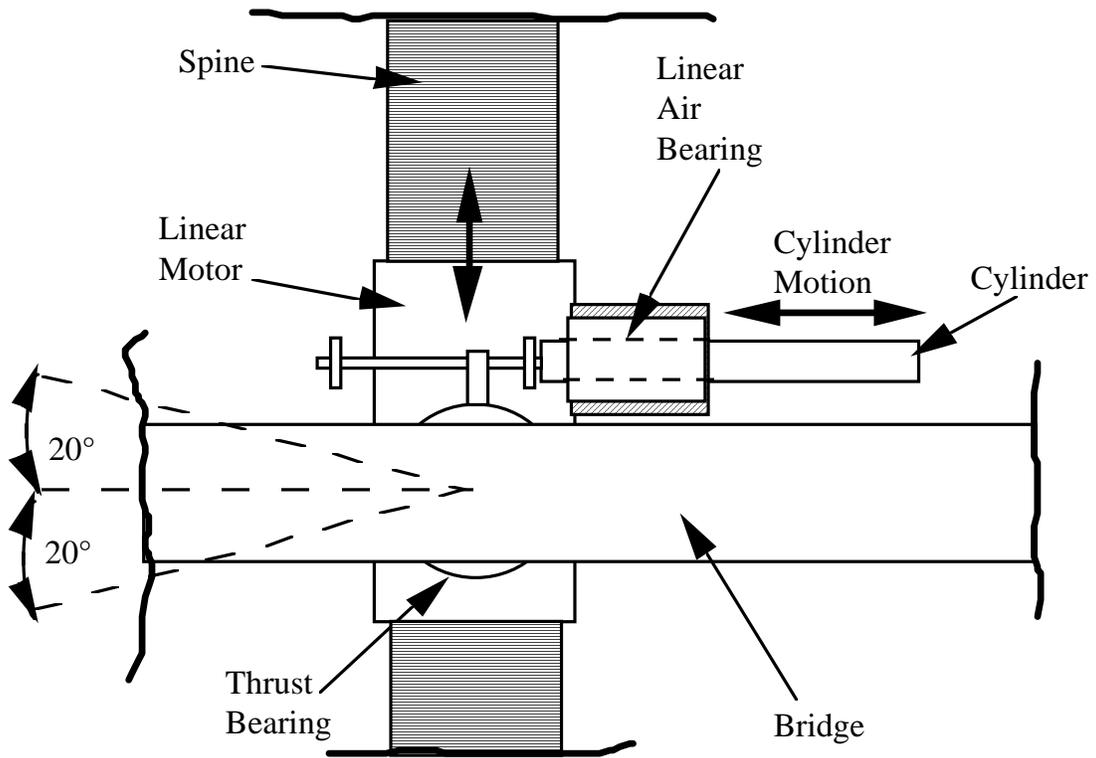
FIGURE 12. PIVOT LOCK SCHEMATIC (UNLOCKED)
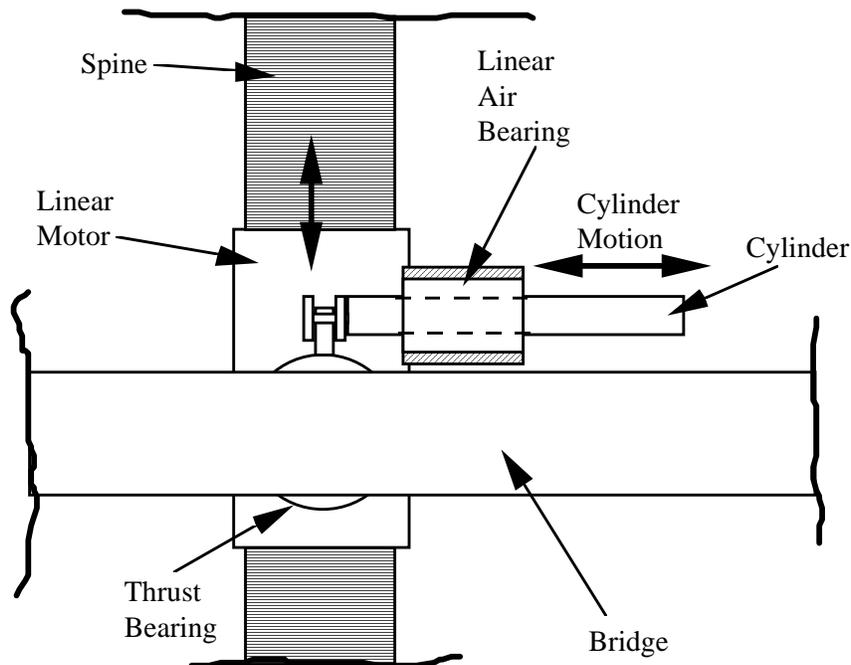
FIGURE 13. PIVOT LOCK SCHEMATIC (LOCKED)

3.1.2   Sensor Bridge.

The sensor bridge, shown in figure 14, is mechanically identical to the stabilizer bridge except that is contains the following additional items:

- Sensor Deployment Platform—A suite of sensors can be mounted on this platform. It can be indexed to deploy a specific sensor and locked into position. During this phase, one reflectance eddy-current sensor is mounted on the platform.

- Eddy-Current Sensor—A reflectance-type Nortec SPO-1958 eddy-current sensor is mounted on the sensor platform.

- Constant Force Deployment Assembly—A precision pneumatic cylinder is used to deploy the eddy-current sensor on the fuselage surface and to retract it from the surface. This assembly applies a constant force during sensor deployment which can be set from 1-3 pounds (4-13 N).

Because the sensor bridge is mechanically identical to the stabilizer bridge, its motions are the same. It can translate along the spine assembly a maximum distance of 15 inches (38 cm), translate along its own axis a maximum distance of 6 inches (15 cm), and pivot with respect to the linear motor ±20 degrees.
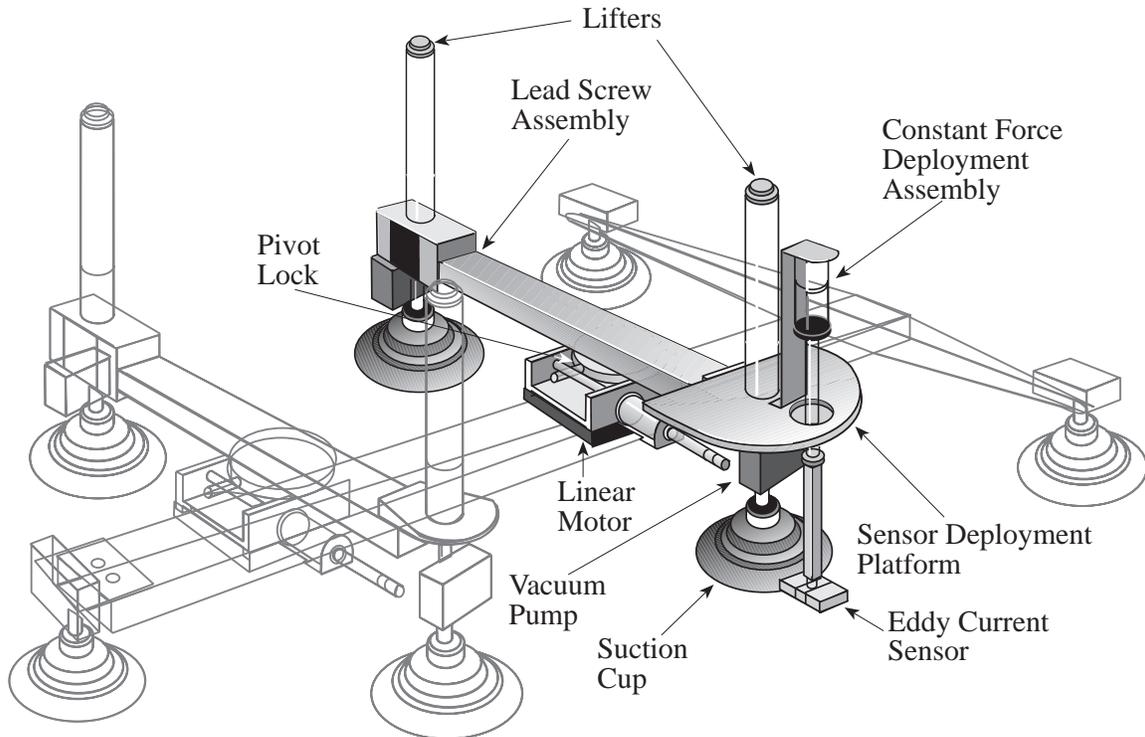
FIGURE 14. DRAWING OF THE SENSOR BRIDGE

3.1.3   Spine Assembly.

The spine assembly, shown in figure 15, consists of the following items:

- Spine—This is the platen upon which the linear motor travels.

- Tail Beam—This is a machined 6061-T6 aluminum T-section.

- Control Valves—These are three-way 24 volt direct current solenoids. The air flow through the valve is 0.5 cubic feet per minute (0.2 l/sec) at a pressure of 100 pounds per square inch (0.7 MPa).

The spine assembly can be thought of as the backbone of the robot. The bridges are mounted on the spine assembly and translate with respect to it during both scanning and walking. In addition, the control valves are mounted at the head and tail ends of the assembly.

Notably, all of the major parts used to fabricate the robot were standard off-the-shelf components; very few modifications were made to the parts. Using standard components shortened the fabrication process and minimized the fabrication costs. To complete the assembly, designing and fabricating custom bracketing were necessary; detailed designs of the custom bracketing are not provided in this report.

A photograph of the finished first-phase robot including the air and electrical supply lines is shown in figure 16.
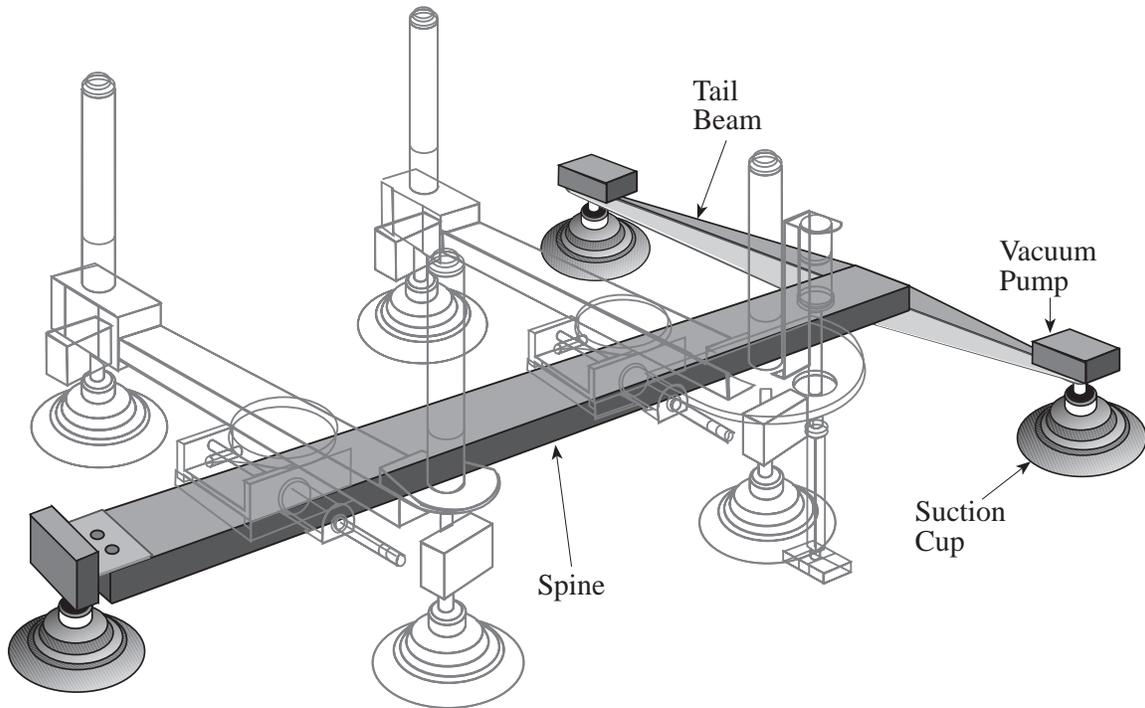
FIGURE 15. DRAWING OF THE SPINE ASSEMBLY
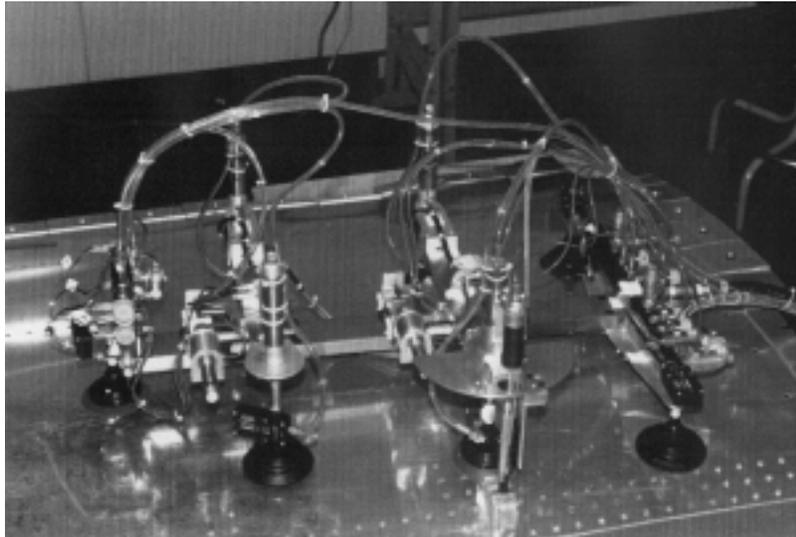
FIGURE 16. PHOTOGRAPH OF THE FIRST-PHASE ROBOT

3.1.4   First-Phase Robot Performance.

The first-phase robot was able to affix to flat horizontal surfaces such as tables, flat vertical surfaces such as walls, and curved surfaces such as the simulated aircraft panel. No limitations were observed with respect to the ability of the robot to adhere to surfaces, regardless of the orientations of the surfaces. The alignment, scanning, and walking motions of the robot were tested to ensure that they functioned as specified. The manually controlled alignment procedure

successfully aligned the robot to a line of fasteners, and the robot's scanning motion was used to deploy a reflectance eddy-current sensor and move it along a fastener line. The robot was able to walk along the aircraft panel, and a coordinated scan and walk motion was performed. Thus, the design of the robot satisfied the mobility and manipulation requirements of the skin inspection application.

The robot needed to be realigned to the fastener line after each step; the accumulated error was too large if this was not done. The robot strayed from its path approximately 0.125 inch (3.2 mm) during each step due to excessive flexibility of the system. The reflectance eddy-current sensor must align within ±0.05 inch (±1.3 mm) with the center of a fastener, thus, a realignment motion was necessary after every step with the first-phase robot. To reduce the problem of error accumulation during walking, the stiffness of the system needed to be increased. The bearing arrangement holding the linear motors to the spine and the stiffness of the lifters when fully extended were identified as areas for improvement.

3.2   SECOND-PHASE MECHANICAL IMPROVEMENTS.

The basic mechanical design of the second-phase robot was identical to that of the first-phase robot. The robot's mobility was enhanced by implementing the improvements identified during the first phase. It was capable of operation around the entire circumference of the fuselage. The robot was made modular in this phase, meaning that each of the three subassemblies contains all of the components (e.g., valves, manifolds, etc.) necessary for its operations. A photograph of the second-phase electromechanical assembly is shown in figure 17.
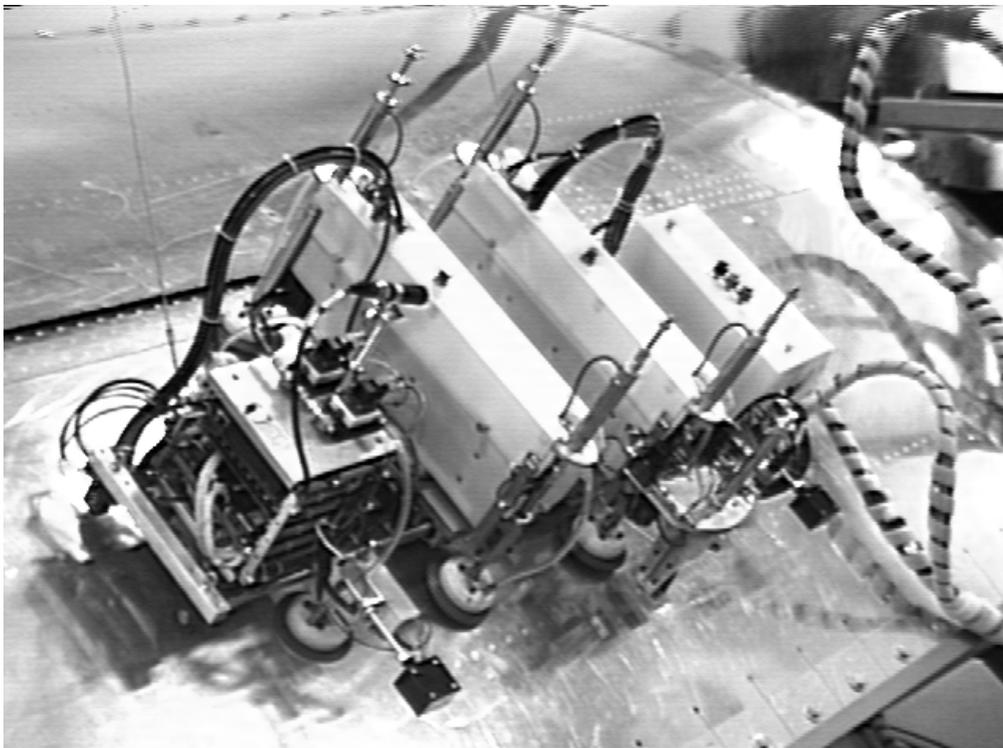
FIGURE 17. SECOND-PHASE ROBOT

21

3.2.1   Modular Design.

Because the robot was a research prototype, a more modular hardware design was used for the second phase to enable easier modifications for experimentation. The first-phase prototype of the robot was not modular, and it was difficult to make changes or additions to the mechanical system without a major overhaul of the entire robot. In a modular robot, each of the bridges and the spine would be independent mechanically, pneumatically, and electrically. Thus, one could remove the sensor bridge from the robot, modify it, test it, and after it was working properly, reattach it to the robot. Later in this section the addition of a set of rails to the spine assembly will be described. The modular design of the robot enabled this change to be made with a minimum of disruption to the rest of the robot and with a minimum of downtime.

3.2.2   Spine Assembly Modifications.

One area where significant weight was removed was the spine assembly. The spine assembly weighed roughly 12 pounds (5 kg) or about 40 percent of the total weight of the first-phase device, where the spine was a solid steel bar with a rectangular cross section. Experimentation with a test piece, which was machined into a U-shaped cross section, showed that up to two thirds of the weight could be removed from the spine without adversely affecting the functionality of the robot. The spine of the second-phase robot was machined into the U-shaped cross section, both lightening the robot and providing a channel for the electrical and pneumatic connections to be run between the robot's head and tail sections.

When the first-phase robot was tested on surfaces approaching vertical, the linear motors that drive the bridges along the spine stalled. The single suction cup at the front of the robot did not provide enough stability and the torque between the spine and the linear motors caused the air gap between the motor and spine to become uneven. To reduce the torsion about the front suction cup and its ball and socket joint, a head beam with two suction cups, similar to the tail beam, was added to the spine assembly. A pair of rails was also added to the robot to remove the load from the linear motor air bearings. The rails are anchored in the head and tail beams of the robot and run parallel to the spine. Each bridge was equipped with two linear pillow block bearing assemblies which ride along the rails. Thus, the forces caused by the bridges' weight act only upon the rails and their associated linear bearings. This relieved the torque about the linear motors and the air gap remains uniform, thereby solving the problem of linear motor stalling.

3.2.3   Leg Design.

It was desired that the robot operate in a variety of orientations, simulating various positions on an aircraft fuselage. For the robot to be able to withstand the forces present in all orientations, its legs were stiffened. Two enhancements over the previous leg designs were implemented: the suction cup fittings were redesigned to reduce the tipping moment and to reduce the effect of the maximum shear force; and a linear bearing assembly was added to withstand the maximum shear force. Figure 18 shows the second-phase leg assembly.

FIGURE 18. LEG ASSEMBLY

The suction cup from the first phase was a commercially available, off-the-shelf item. It consisted of a 3-inch-diameter silicone suction cup, an aluminum suction cup fitting, and a ball and socket joint. There were two problems with this suction cup: the length of the tipping moment arm from the surface to the ball and socket joint created a large tipping moment, and according to the manufacturer, the ball and socket joint was not designed to withstand a shear stress. These problems had to be addressed to prevent a failure during the tests. A new suction cup fitting was designed and fabricated. The new fittings contained an integral ball and socket joint and were machined from P.E.T., a lightweight, easily machinable plastic material. The joint was designed to withstand the expected forces, and by lowering it into the suction cup fitting, the length of the tipping moment arm was reduced, thereby reducing the tipping moment.

The leg assembly from the first phase was composed of the suction cup, suction cup fitting, the ball and socket joint discussed in the previous paragraph, and an air cylinder. Besides the problems with the suction cups, there was also a problem with attaching the air cylinder directly to the suction cup assembly. Exposing an air cylinder to shear loads over a prolonged period of time could break the seal around the rod, causing air leaks. To correct this problem, a linear pillow block bearing assembly was added to the leg configuration. The shaft of the bearing assembly was attached to the ball of the custom ball and socket joint previously described. The other end of the bearing shaft was connected to the cylinder rod via a linear coupler. Thus, in this design, all of the forces acted upon the bearing assembly, and the cylinder was used only for motion. The air cylinder was not exposed to shear loads, thus, protecting the seal around the rod.

The legs on the spine assembly are fixed in place. A standard 1/4-20 stainless steel stud is attached to the ball of the custom ball and socket joint. The studs are mounted through holes at the ends of the head and tail beams, and they are held in place using locking nuts.

23

<u>3.2.4 Tether Design</u>.

A tether to secure the robot from above was designed, built, and tested during the second phase of development. In operation, the tether could be secured to the safety lines already present for human inspectors; these safety lines are mounted on trolleys that run parallel to the horizontal axis of the aircraft from nose to tail. The tether would be payed out as necessary while the robot moves along the aircraft panel; if the robot loses adhesion and starts to fall, the tether must be locked in position to keep the robot from falling. A standard safety device known as a retractor was used in the tether. The retractor works much like an automobile's seat belt. With low acceleration, the cable is payed out as necessary; however, with a sudden acceleration the mechanism locks, preventing the cable from extending further. Custom bracketing was designed and fabricated to attach the tether to the head and tail beams of the robot. The tether was secured to a trolley that runs along a rail suspended above the area where the robot is deployed.

The tether arrangement for the robot is not effective when the robot is positioned on the underside of the aircraft.   The clearance underneath a typical mid-size jet aircraft (e.g., McDonnell Douglas DC-9 or Boeing 737) is about 40 inches (1 m), and the approximate radius of the fuselage is 75 inches (1.9 m). If the robot were positioned on the belly of the aircraft as shown in figure 19, the tether would follow the circumference of the fuselage to this point. If air is lost when the robot is in this position, the tether will not prevent the robot from hitting the ground since the tether would be too long. This is illustrated in figure 20. One of the designs for tethering the robot for this special case under consideration is shown in figure 21; in this design, when the robot is inspecting the belly of the aircraft, a second tether suspended from the opposite side of the fuselage would be attached to the robot in order to prevent the device from falling if the air supply is lost. This tether configuration has been used in the laboratory on the Foster-Miller panel to protect the robot while working inverted.
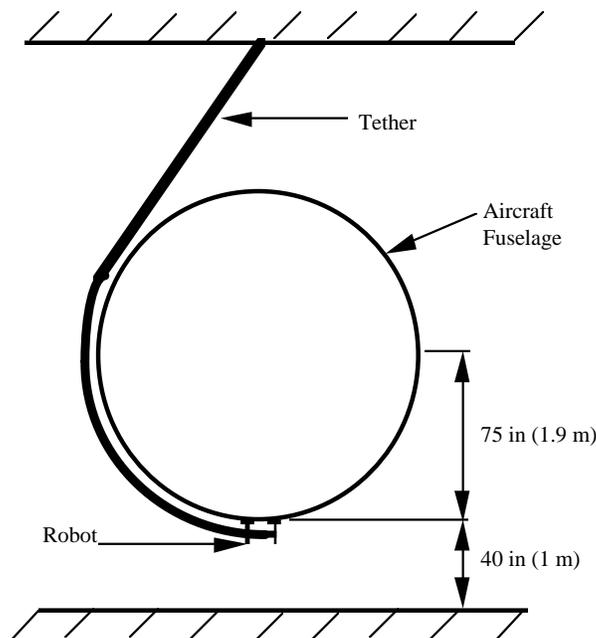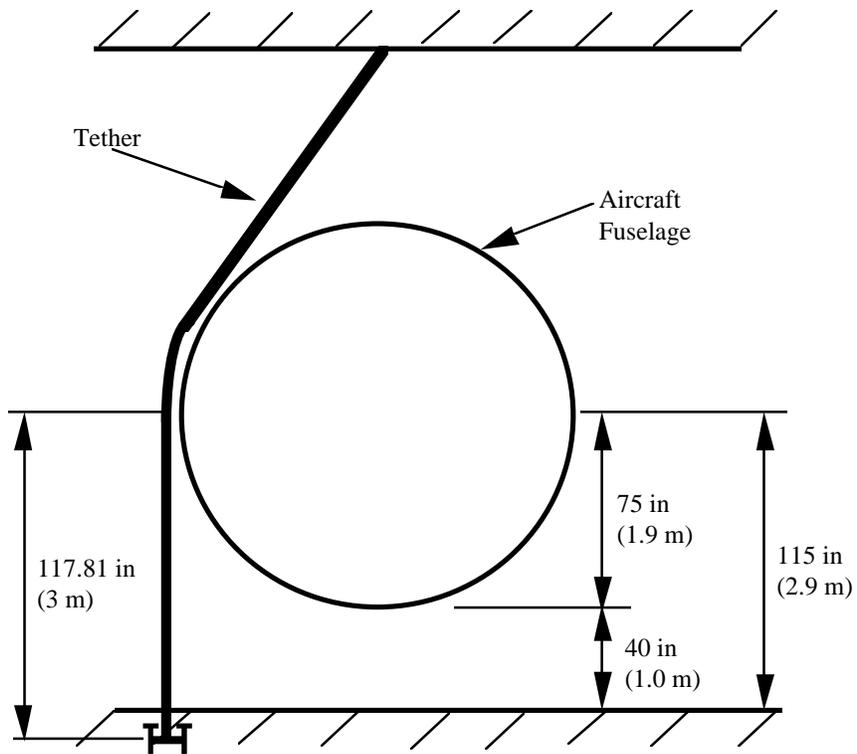

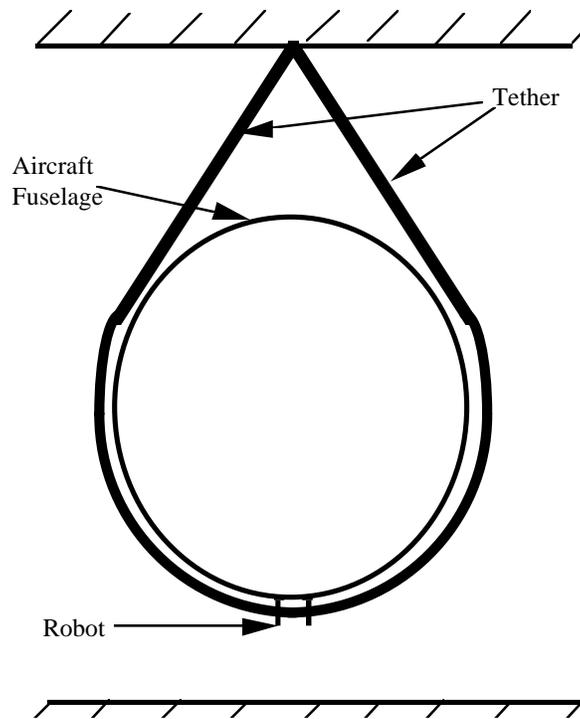FIGURE 19. ROBOT TETHERED TO BOTTOM OF AIRCRAFT

FIGURE 20. TETHER LIMITATION

FIGURE 21. SPECIAL CASE TETHER

25

3.2.5   Air Consumption.

The air consumption of the first-phase prototype robot varied depending on the number of suction cups activated at a given time and whether the eddy-current sensor was deployed or retracted. The maximum rate of air consumption for the robot was 7.2 standard cubic feet per minute (SCFM) of air. USAir inspectors had indicated that the robot should consume no more air than a standard pneumatic tool, about 5-10 SCFM of air. The first-phase robot fell within this range; however, many standard air compressors used to run pneumatic tools do not produce 7.2 SCFM of air. The Carnegie Mellon team decided to reduce the robot's air consumption to a maximum of 6 SCFM to run the robot from most standard air compressors and to minimize pressure loss in the air line between the robot and the ground. The laboratory compressor produced 6.2 SCFM at 100 psi.

The air consuming parts of the sensor deployment unit and the pivot lock air bearings were eliminated. The pivot lock air bearings were replaced with a dry bearing. The airpot of the sensor deployment unit, which consumed 0.6 SCFM of air, was replaced by an air cylinder, which did not require a constant draw of air.

Experiments measuring the leak rates at various positions on aircraft showed that vacuum pumps with one half of the air consumption of those on the first-phase robot could be used without sacrificing system performance. The performance of the reduced-flow vacuum pumps vis-a-vis leak rates on aircraft surfaces was nearly identical to the performance of the larger vacuum pumps.

The maximum rate of air consumption of the second-phase prototype robot is 5.6 SCFM of air. During normal operation, only four suction cups are affixed to the surface. The normal requirement of air is 3.2 SCFM (four suction cups and the linear motor air bearings). When the robot is walking, during the transition from the spine legs being affixed to the surface to the bridge legs being affixed (or vice versa), all eight cups will be affixed to the surface for a brief moment. The moment when all eight cups are affixed to the surface plus the constant draw of the linear motors gives the maximum air consumption of 5.6 SCFM.

Additional modifications were made to the vacuum circuits during the third phase of development to further decrease the air consumption. This was motivated by the desire to reduce the size of the compressed air line connecting the robot to the compressor. By using needle valves as flow restricting buffers, a single vacuum pump was able to support several suction cups. Adjusting the needle valves provided sufficient isolation between the cups so that a loss of seal on one cup would not spoil the vacuum to the remaining cups. Using this configuration the eight individual vacuum pumps were replaced by three vacuum circuits, one on the spine and each of the sensor bridges. The resting air consumption was reduced to approximately 2.0 SCFM with a maximum consumption of 2.6 SCFM.

3.2.6   Marking System.

A marking system was designed to physically mark the surface where abnormal sensor signals are encountered, allowing the inspector to locate and further inspect the marked areas at a later

time. The marking system contains a spring-return cylinder and either a china marker or a self inking stamp. A three-way valve controls the operation of this feature, and the air pressure and flow to the cylinder are regulated to ensure an appropriate supply of air.

3.3   SYSTEM PERFORMANCE.

The time that it takes to complete one scan and the time that it takes to inspect a line of fasteners are two parameters that help to define the inspection speed of the robot. During this phase, the scanning speed was measured, but because of the need for manual intervention in the alignment motion, the time required to automatically inspect a line of fasteners could not be defined.

When inspecting manually, inspectors move a reflectance eddy-current sensor across a line of fasteners at an average rate of 2.4 inches per second (6 cm/sec). This number was obtained by having an NDI supervisor measure the average time for eddy-current inspections performed at the USAir facilities in Pittsburgh, Pennsylvania. Although this is not an absolute number, it is a fairly good estimate of the manual-scanning rate needed for a rough comparison to the automated scanning speeds of the robot.

The robot's scanning rate can be set using the operator interface. During initial testing, the robot was able to move the sensor at 5 inches per second (13 cm/sec) and still produce high-quality traces on the PC monitor. This rate is considerably higher than the manual rate. However, this is not the upper limit for the scanning speed. The linear motors which control the rate of scanning have a limit of 100 inches per second (2.5 m/sec), and the SmartEDDY 3100 can sample at a rate of approximately 3,600 data points per second. Thus, the scanning rate for the automated system is potentially faster than the manually deployed sensors.

3.4   LONG-TERM MECHANICAL DESIGN ISSUES.

The robot was designed to inspect the main fuselage sections of aircraft the size of the McDonnell Douglas DC-9, Boeing 737, and larger aircraft. The design can be adapted to inspect highly curved areas such as the nose or tail, hard-to-reach areas such as those around windows or doors, and smaller aircraft. The robot designed by the Carnegie Mellon team can be scaled either larger or smaller; most components used to build the robot are commercially available in a variety of sizes. Thus, for smaller aircraft or for areas on larger aircraft with greater curvature, a smaller version of the prototype could be produced. This version would weigh less than the current version, but it would also have a much shorter maximum scanning distance. It would have to take several steps to inspect the same number of fasteners that the current prototype can inspect in one step.

When inspecting large aircraft such as the McDonnell Douglas DC-10 or the Boeing 747, it might be advantageous to build a large version of the robot to cover even greater distances with a single scan. Of course, the larger versions would weigh more, so there is a practical limit to the maximum size. In addition, scaling relations favor smaller devices when suction cups are used as the affixing means. Ultimately, the size of production versions will be driven by the requirements of aircraft inspectors. Several sizes of spine assemblies and bridges may be commercially produced, allowing end users to buy robots to fit their needs.

In addition, special attachments could be developed to enable the robot to inspect hard-to-reach areas. For example, inspecting areas around the windows of an aircraft could prove troublesome because these are areas on the fuselage on which maneuvering might be difficult. If some type of window inspection is required, a special attachment allowing access to the inspection area from an area above (or below) the windows might be used.

The design can also be easily extended to incorporate other NDI sensors. The prototype can deploy a reflectance eddy-current probe and several small video cameras. In the future, a number of NDI sensors could be mounted on the sensor platform. Of the probes commonly used by inspectors today, pencil eddy-current sensors and ultrasonic sensors are good candidates for inclusion in the sensor suite. New NDI sensors, such as the magneto-optic sensor developed by Physical Research, Inc. [4], and new techniques, such as shearography [5] can be evaluated and included as part of the sensor suite if appropriate.

Several issues related to the mechanical design were brought to light during the laboratory and field testing of the robot prototype. Any future development of the mechanical system will need to address these items.

3.4.1   Linear Motors.

Several problems associated with the linear motors were found during system testing. The linear motors tended to stall following 1/2 to 3/4 hours of use. The motors generate heat when they are used, and there is no efficient way to dissipate the heat. Because the platen (spine) was machined into a U-channel to save weight, there is very little material for heat conduction. The heat buildup most likely causes a thermal expansion of the platen, and over time, the platen expands enough to cause the motors to stall.

Severe wear was also noted on the linear motor platen. This may have two causes. First, the thermal expansion of the platen may cause wear on the platen before the motors stall. Also, because the platen was machined into a U-channel, it may not be perfectly flat, causing the motors to rub against the surface of the platen, creating wear. The stabilizer bridge did not move freely even when the system was cold, indicating that the spine may not be uniformly flat.

Finally, there was a loss of motor position calibration after either of the bridges ran into a hard stop or after the bridges ran into each other. If this happened, the motors had to be sent back to their home positions for proper initialization.

Enough problems with the linear motors were uncovered that if future versions of this robot are built it might be best to explore other options for creating the linear motion along the spine.

3.4.2   Lead Screw Assemblies.

The lead screw assemblies that provide linear translation in the direction perpendicular to the spine were slow, on the order of 10 inches per minute (4 mm/sec). When the motors were run at higher speeds, they did not provide enough torque for motion. The installation of nominally higher torque motors for the same assemblies provided only marginally greater torque at higher speeds, and speed was not usefully increased. This issue needs further investigation.

3.4.3   Umbilical and Tether Management.

The problems managing both a thick umbilical cable and a tether are operational issues that must be solved before surface-crawling robots can effectively be used in a commercial aircraft inspection facility. Although these issues were beyond the scope of this project, it was apparent from the field demonstration that tether and umbilical management are critical.

The umbilical cable is the line that extends from the control area to the robot. The umbilical cable is comprised of three 1/4-inch (6 mm) compressed air lines, two eddy-current cables, four video cables, four motor control cables, a power cable, and a computer communications cable. The diameter of this bundle is about 1 1/4-inches (40 mm).

During the demonstration, the robot walked over the crown of the nose section of a McDonnell Douglas DC-9 (the section that stretches from the nose of the aircraft to just past the second passenger window). For this, the length of the umbilical cable was 50 feet (15 m); it was very difficult to manage a cable of that thickness and length. To inspect an entire fuselage the robot would require a much longer cable. If the robot were to inspect a wide-body aircraft, several hundred feet of umbilical cable would be required. If robotics is to be practical in a commercial hangar environment, the umbilical must be either made more manageable, or if possible, eliminated altogether.

The tether for the demonstration ran along a safety line suspended above the DC-9; however, the tether did not move freely enough to gracefully move in concert with the robot. There were times when the tether became entangled with the bridges, thus interfering with the robot's motion. In the future, coordinated motion between the tether and the robot may be necessary. An active tether whose motion would be synchronized with the robot is an area of potential research for the future.

3.4.4   Navigation.

During the field trials, it was noted that navigation around seams was complicated by the spacing and size of the suction cups and the 3 1/2-inch (89 mm) sideways step size of the bridges. There are several potential solutions to this problem. First, clusters of suction cups could be added to each leg; this would have the disadvantage of increasing the air flow needed by the system. A second potential solution is to modify the dimensions of the bridges to make navigating around seams easier.

3.4.5   Walking Motion.

The walking motion of the robot created some operational problems. On inverted, angled surfaces, the extend and retract motion, in which the spine is raised from and lowered to the surface, was unreliable. This was primarily due to a combination of the weight of the robot and the sizes of the pneumatic lifting cylinders. Also, the extend and retract motion severely jarred the mechanical system, loosening and breaking electrical contacts and shifting camera views. A third problem with the walking motion was that the variation in the robot's position between the extended and retracted states made alignment difficult. Alignment was performed in the

extended state while scanning occurs in the retracted state. One apparent solution for these problems is to replace the fixed legs on the spine assembly with legs that can extend and retract. During walking, the spine would stay at a uniform height while the individual legs would be raised and lowered.

3.4.6   Surface Damage.

During field testing, it was noted that the suction cups left scuff marks on the aluminum skin; the scuff marks were much worse on the Foster-Miller panel than on the DC-9 fuselage surface. This problem may be connected with the dirt on the panel and fuselage surfaces. Even though the surfaces had been cleaned, a film of fine dirt particles remained on both surfaces. As the robot walked along either surface, the dirt particles collected on the suction cups. As the cups attached to and released from a surface, they tended to slide a little, especially when the spine assembly was lowered to the surface. The dirt trapped on the bottom surface of the suction cups scored the surface slightly as the cup slid. A potential solution, suggested by a representative of an aluminum producer, is the replacement of the soft silicone suction cup with a harder rubber that would be less likely to trap dirt.

In a related problem, the eddy-current probe did not initially ride flat on the skin surface, compromising scan data. To compensate for this, the pressure used to deploy the sensor was increased. As with the suction cups, dirt accumulated on the probe's surface, and when it was deployed, the probe also scored the skin surface.

The surface abrasion caused by the robot did not result in deep scratches to the surfaces and could easily be buffed away. However, this may not be acceptable to aircraft maintenance personnel and must be further investigated. The tendency to score an aluminum surface is a characteristic of all surface-walking robots. Dirty surfaces will be regularly encountered during routine inspections in a real environment, so this problem must be addressed before surface-walking robots can be routinely deployed in the field.

4.   ROBOT MOTIONS.

The cruciform robot was designed to move gracefully in the horizontal direction, to move adequately in the vertical direction, and to possess some degree of steering capability. Since the majority of skin inspections lie along fastener lines oriented in the horizontal direction, the first phase of testing was performed with the robot moving in that direction. Below is a description of the motions necessary for the robot to inspect horizontal rows of fasteners.

The robot has three basic motions: walking, alignment, and scanning. The act of the robot moving to another location on the aircraft surface is called walking. The alignment motion refers to the condition where the spine of the robot is aligned parallel to a line of fasteners; this is done so that the sensor may be deployed and information may be gathered. The robot performs the scanning motion when it deploys the eddy-current sensor and moves it along the fastener line. After the robot completes a scan, it is ready to move to another location on the aircraft to gather more data. Each of these motions will now be described in greater detail.

The initial position of the robot during walking is shown in figure 22, where both of the bridges are toward one end of the spine. This is the robot's position after it has performed a scan. If the bridges are not in this position (e.g., after alignment), the robot will move the bridges into the configuration shown in figure 22 when the inspector instructs the robot to take a step.
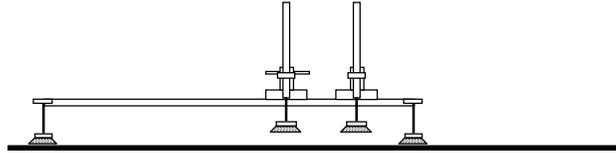
FIGURE 22. INITIAL CONFIGURATION FOR WALKING

Next, the lifters extend until the bridges' suction cups touch the surface. The suction cups on the bridges are activated, and all eight suction cups are affixed to the surface. The suction cups on the spine assembly are then deactivated, allowing the lifters to extend completely (see figure 23, lift exaggerated). The spine assembly is now lifted off of the fuselage surface.
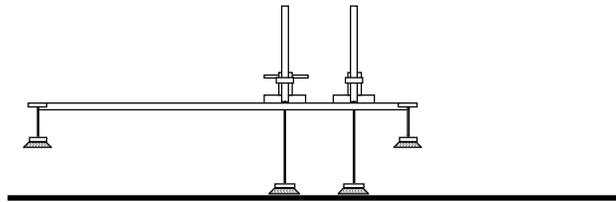
FIGURE 23. WALKING MOTION: SPINE ASSEMBLY RAISED

Both linear motors run simultaneously (and at an identical rate), moving the spine assembly with respect to the bridges (which are affixed to the surface). When the spine assembly has moved a preset distance, the robot will be in the configuration shown in figure 24.
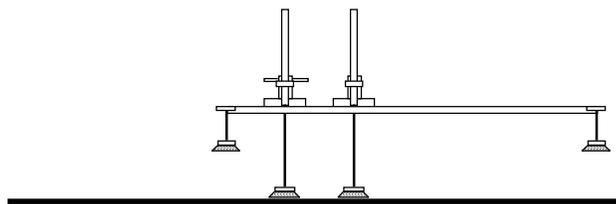
FIGURE 24. WALKING MOTION: SPINE ASSEMBLY MOTION

Next, the lifters retract until the suction cups of the spine assembly touch the surface. Those suction cups are activated, and all eight suction cups are affixed to the surface. The suction cups on the bridges are then deactivated, allowing the lifters to retract completely as shown in figure 25.
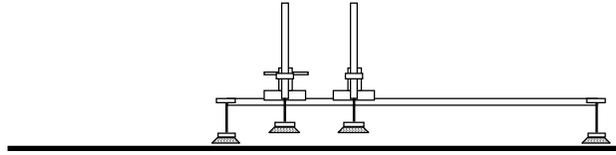
FIGURE 25. WALKING MOTION: SPINE ASSEMBLY LOWERED

Finally, the stabilizer bridge is moved to the far end of the spine as shown in figure 26. If the robot needs to be aligned, the alignment process is initiated. Once aligned, it is ready to scan a line of fasteners.
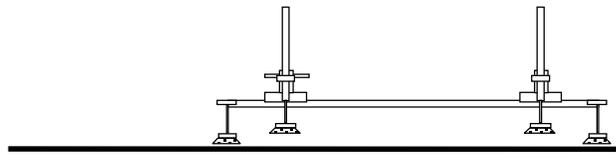
FIGURE 26. WALKING MOTION: STABILIZER BRIDGE MOVED INTO POSITION

Although the process of walking has been described for a horizontal step, a vertical step may be performed by substituting the bridge motions using the linear motors with analogous motions using the stepper motors. The vertical step is used when moving the robot from one row of rivets to another around the circumference of the fuselage.

4.2   ALIGNMENT MOTION.

The alignment motion is used in two basic circumstances: when the operator initially affixes the robot to the aircraft surface and, if necessary, after the robot takes a step. As shown in figure 27, two interim alignment posts were attached to each end of the first-phase robot's spine. During the second phase, these posts were replaced with a pair of video cameras. When its spine is not parallel to the fastener row and the tips of the alignment posts are not in line with the rivet row, the robot needs to be aligned; this is the situation shown in figure 27.
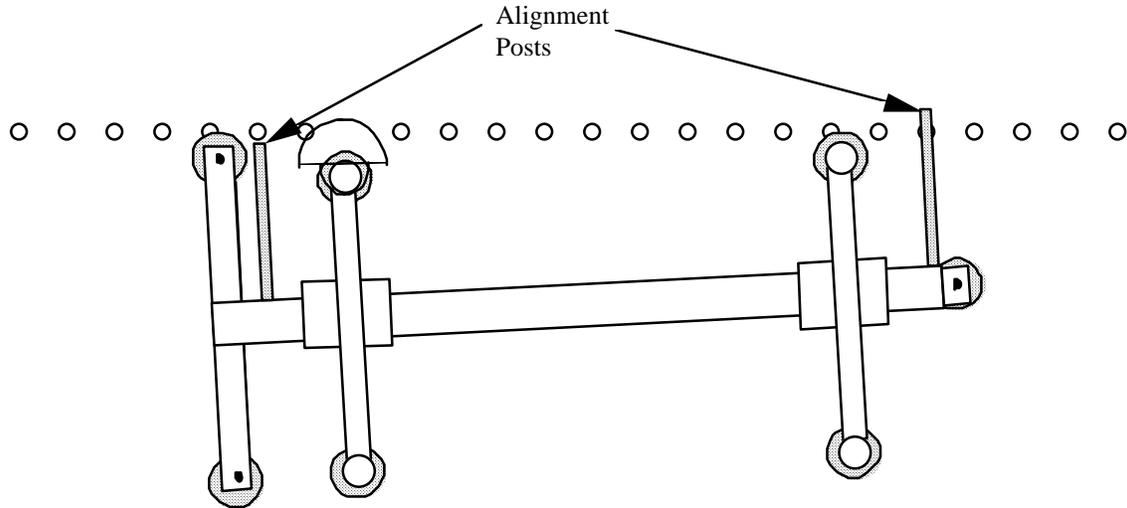
FIGURE 27. UNALIGNED ROBOT (OVERHEAD VIEW)

After the operator initially affixes the robot to the aircraft surface or after a step has been taken, the three suction cups on the spine assembly are activated and hold the robot to the aircraft surface. A profile of the robot in this position is shown in figure 28.
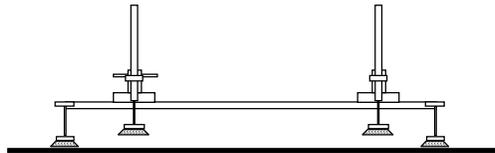
FIGURE 28. UNALIGNED ROBOT (SIDE VIEW)

The sensor and stabilizer bridges are at opposite ends of the spine. The suction cups on both bridges are deactivated, and the lifters are retracted. The sensor is also lifted from the surface (the sensor is not shown in figure 28).

To start the alignment procedure, the lifters are extended, the suction cups on the bridges are activated, and the suction cups on the spine assembly are deactivated; these actions must be performed in this order. A profile of the robot in this position is shown in figure 29 (the lift is exaggerated for clarity).
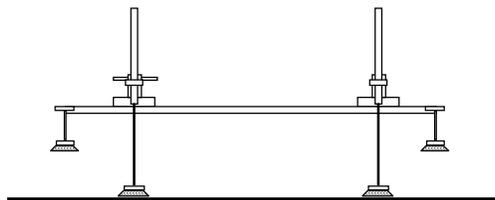
FIGURE 29. STEP 1 OF THE ALIGNMENT PROCESS

Next, the pivot locks on both bridges are unlocked. Following this, the stepping motors on the lead-screw assemblies drive the bridges to adjust the position of the spine assembly. When the

spine is parallel to the rivet row and the tips of the alignment posts are in line with the fasteners as shown in figure 30, the robot's spine is properly aligned to perform a scan.
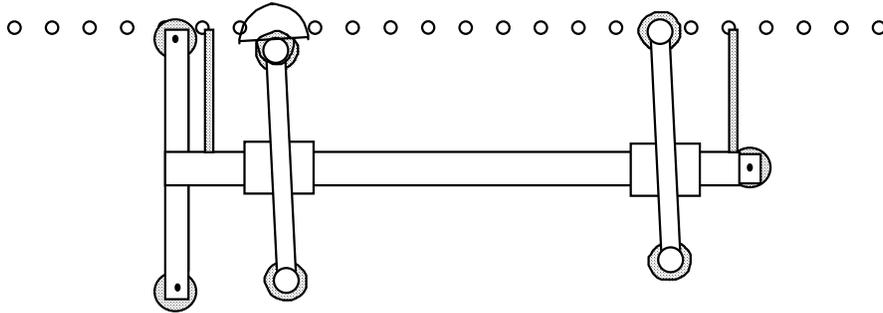
FIGURE 30. STEP 2 OF THE ALIGNMENT PROCESS

The lifters on both bridges are retracted, the suction cups on the spine assembly are activated, and the suction cups on the bridges are deactivated; as before, these actions must be performed in this order. Finally, the pivot locks on both bridges are engaged, locking the bridges into position perpendicular to the spine. The alignment has been completed and the robot is oriented as shown in figure 31. The sensor is moved to its starting position, and the robot is ready to scan a line of fasteners.
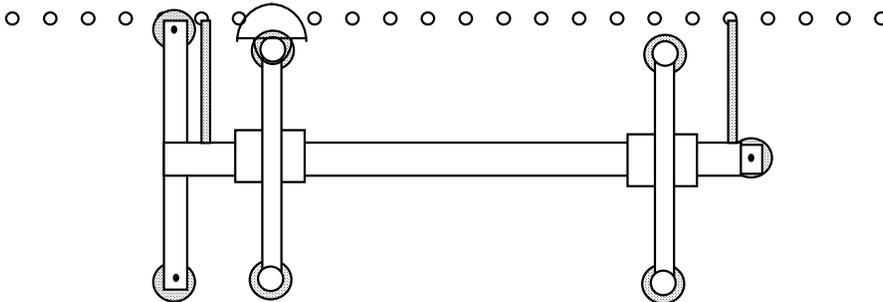
FIGURE 31. ALIGNED ROBOT (OVERHEAD VIEW)

It is assumed that the position of the robot relative to the line of rivets does not change during the retraction process. In practice, small deviations from this assumption on non-horizontal portions of the fuselage made alignment difficult and should be addressed in future designs.

The alignment motion can also be used to provide a degree of steering for the robot. For example, if obstacles are encountered, the direction of motion can be changed. Figure 32 illustrates how the alignment motion is used to steer the robot.
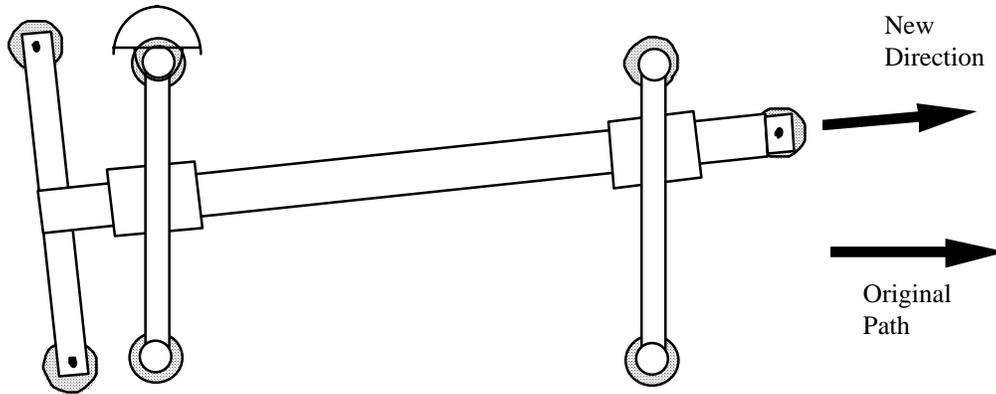
FIGURE 32. ROBOT STEERING

4.3   SCANNING MOTION.

After the inspector instructs the robot to perform a scan, the robot deploys the eddy-current sensor on the surface of the aircraft; the sensor is deployed with a light constant pressure simulating that provided during manual scanning. The eddy-current system then begins acquiring data. The robot moves its sensor bridge a preset distance in the direction indicated in figure 33 until it reaches the position shown in figure 34. Acquisition of eddy-current data is then halted.
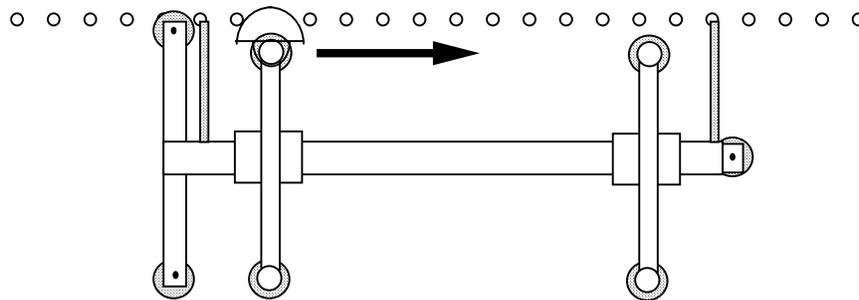
FIGURE 33. DIRECTION OF SCANNING MOTION
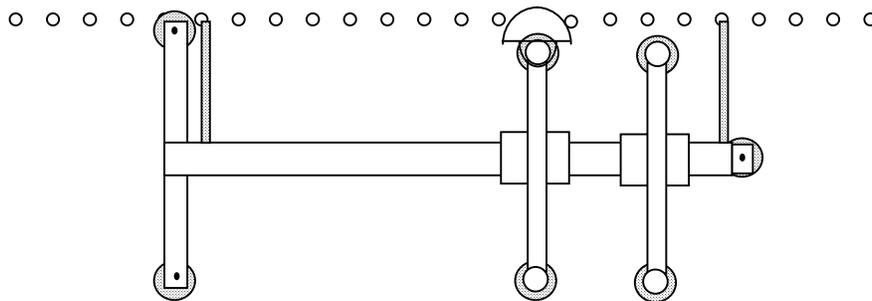
FIGURE 34. ROBOT'S CONFIGURATION AFTER SCANNING MOTION

35

5. CONTROL SYSTEM DETAILS.

5.1 ON-BOARD ROBOT ELECTRONIC SYSTEMS.

A block diagram of the electronic systems carried on board the robot is shown in figure 35. A major improvement to the second-phase robot was the addition of a dedicated computer carried on the robot. This processor manages all of the low-level robot control, feedback, and sequencing. A commercially available single-board industrial computer is used for the on-board computer. This processor card provides integrated input/output capabilities over a -40 to 85°C operating range. Mechanical limit switches and vacuum switches provide TTL level signals to the CPU card for monitoring the robot's status. A reed relay expansion card provides the 12 VDC outputs used to drive the solenoid valves for operation of the pneumatic systems. The circuit boards for the on-board computer are housed in a card cage mounted on the head end of the robot, as shown in figure 36. Commands are received, and status information is sent via an RS-232 serial communication link with the ground-based operator workstation. The eddy-current and video signals, as well as the motor drive signals, are managed directly from the ground. A DC/DC power converter is used to generate the 5 VDC power required for the computer cards from the robot's normal 12 VDC supply.
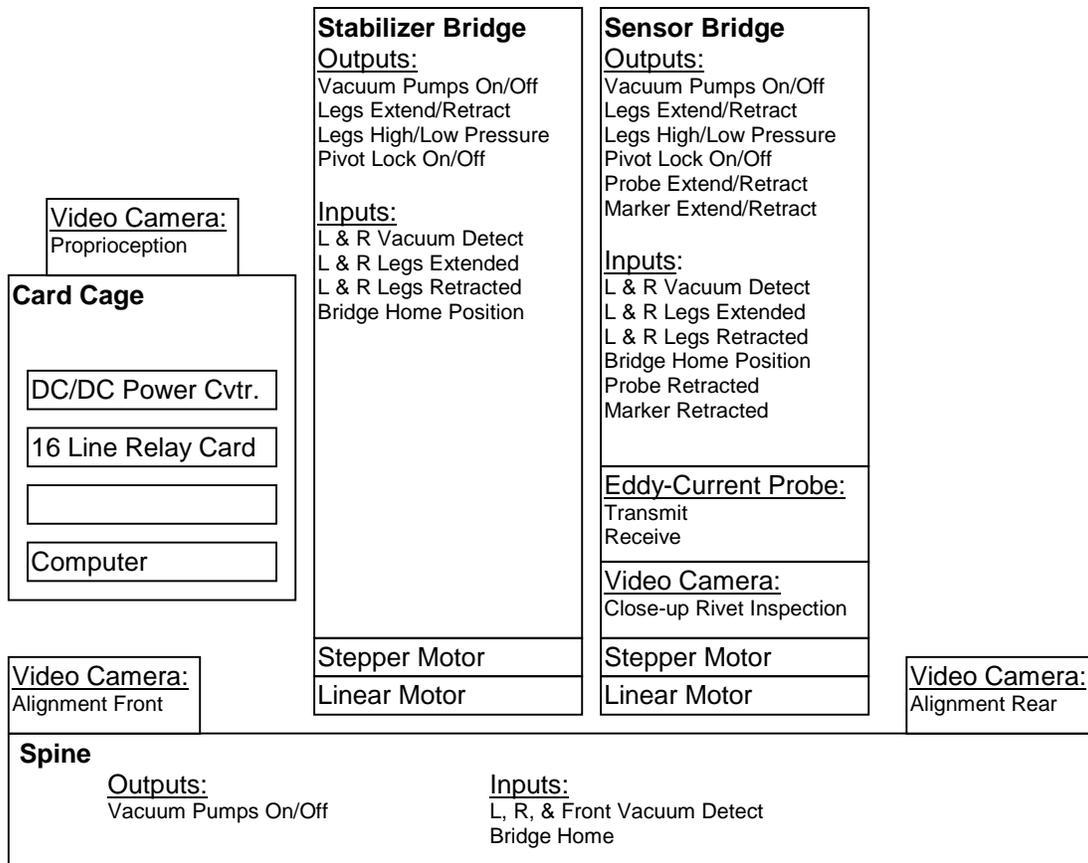
FIGURE 35. ON-BOARD COMPUTER AND ROBOT ELECTRONICS
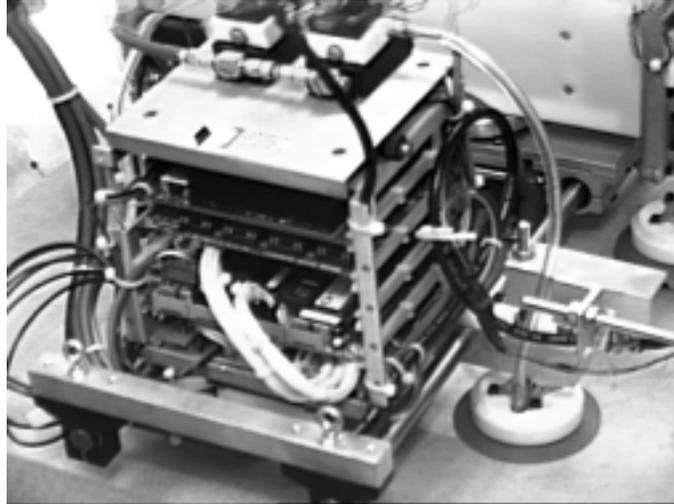
FIGURE 36. ON-BOARD PROCESSOR

The on-board computer executes a multitasking program enabling complex sequences of actions to be simply controlled from the ground. A sample task called "SmplExtendV," short for "sample: extend robot legs with vacuum," would be used during a walking sequence when performing the transition from having the spine attached to the airplane's surface to having the bridge legs attached and extended. The steps of this task break down to the operations:

1.   (a, b) set the leg cylinder pressure to low
2.   (a, b) turn on the bridge legs extension valves
3.   (a, b) turn on the bridge vacuum pumps
4.   (a, b, c, d) wait until all four bridge suction cups have achieved vacuum
5.   turn off the spine vacuum pumps
6.   (a, b, c, d) wait until all four spine suction cups have lost vacuum
7.   (a, b) complete extension by setting leg cylinder pressure to high

To perform this task the computer would execute the following instructions (italicized numbers in brackets refer to the previous list):

```
BB NewTask SmplExtendV
     WritePort SensHPBypass 0 EndInstruct              [1a]
     WritePort StabHPBypass 0 EndInstruct              [1b]
     WritePort SensExtend 1 EndInstruct                [2a]
     WritePort StabExtend 1 EndInstruct                [2b]
     WritePort SensVacuum 1 EndInstruct                [3a]
     WritePort StabVacuum 1 EndInstruct                [3b]
     SendMsg SmplExtendV:   Waiting for bridge vacuum make
EndInstruct
     ReadPortUntil SensVacDetR 0 EndInstruct           [4a]
     ReadPortUntil SensVacDetL 0 EndInstruct           [4b]
     ReadPortUntil StabVacDetR 0 EndInstruct           [4c]
     ReadPortUntil StabVacDetL 0 EndInstruct           [4d]
```

```
        WritePort SpineVacuum 0 EndInstruct                    [5]
        SendMsg SmplExtendV:  Waiting for spine vacuum break
    EndInstruct
        ReadPortUntil SpineVacDetR 1 EndInstruct               [6a]
        ReadPortUntil SpineVacDetL 1 EndInstruct               [6b]
        ReadPortUntil SpineVacDetFR 1 EndInstruct              [6c]
        ReadPortUntil SpineVacDetFL 1 EndInstruct              [6d]
        WritePort SensHPBypass 1 EndInstruct                   [7a]
        WritePort StabHPBypass 1 EndInstruct                   [7b]
        SendMsg SmplExtendV:  Done EndInstruct
        TaskDone EndInstruct
    EE
```

Additional task capabilities include timed waits, other conditional and looping constructs, and the ability to stop or start other tasks.

Three push-button switches are currently mounted on the robot's tail beam to allow the operator to begin and end the inspection process. The buttons perform the following tasks:

- Initialize—This button returns the robot to its initial state where all suction cups are turned off, the legs are retracted, and the pivot locks are engaged. It is generally used before the robot is placed on the aircraft.

- Affix—Pressing this button will activate the four suction cups on the spine assembly, allowing the operator to affix the robot to an aircraft.

- Remove—When the operator presses this button, all suction cups are deactivated, allowing the operator to remove the robot from the aircraft surface.

5.2   GROUND-BASED ELECTRONIC SYSTEMS.

The ground-based electronic robot control equipment consists of the operator workstation computer and the satellite equipment enclosure. Their contents will be discussed in the next two sections.

5.2.1   Operator Workstation Computer.

The operator workstation computer is the main control point for the robotic system. It provides the display and interface resources required by the operator to control and monitor the robot. These interfaces are described in further detail in section 5.3. All high-level robot control and coordination are performed by this machine. Communication channels are provided to the motor controllers, the on-board computer, and the video-processing computers.   The eddy-current instrument is contained in and controlled by this machine. A photograph of the operator workstation is presented in figure 37.

FIGURE 37. OPERATOR WORKSTATION AND VIDEO MONITOR/RECORDER

5.2.2   Satellite Equipment Enclosure.

The satellite equipment enclosure is a small, mobile, 19-inch equipment rack used to house much of the ancillary support equipment for the system such as power supplies, motor controllers, and the camera control units.

All electrical lines between the robot and the ground terminate at this enclosure. Communication and eddy-current signal lines are then routed to the operator workstation, while video lines go to the video processor. A photograph of the satellite equipment enclosure is presented in figure 38.
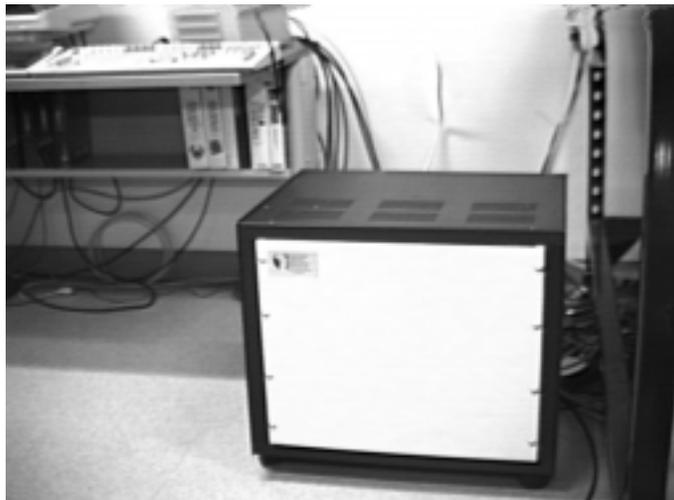

FIGURE 38. SATELLITE EQUIPMENT ENCLOSURE

The ground-based PC communicates with both the motor controllers and the on-board computer over four RS-232 serial connections as shown schematically in figure 39. While this serial communication configuration is sufficient for use in the lab, it is expected that a combination of

RS-422 and optical isolation will be used in the hangar to support longer line lengths and to provide greater noise immunity. A four-line serial interface is used to provide the additional communication lines from the PC. Software drivers enable these additional lines to be accessed in the same manner as standard PC communication (com) ports. The original 16450 UART ICs on the communication board were replaced with 16550 UARTs (with on chip character buffers) to provide more reliable high-speed serial communication links in the multitasking environment.
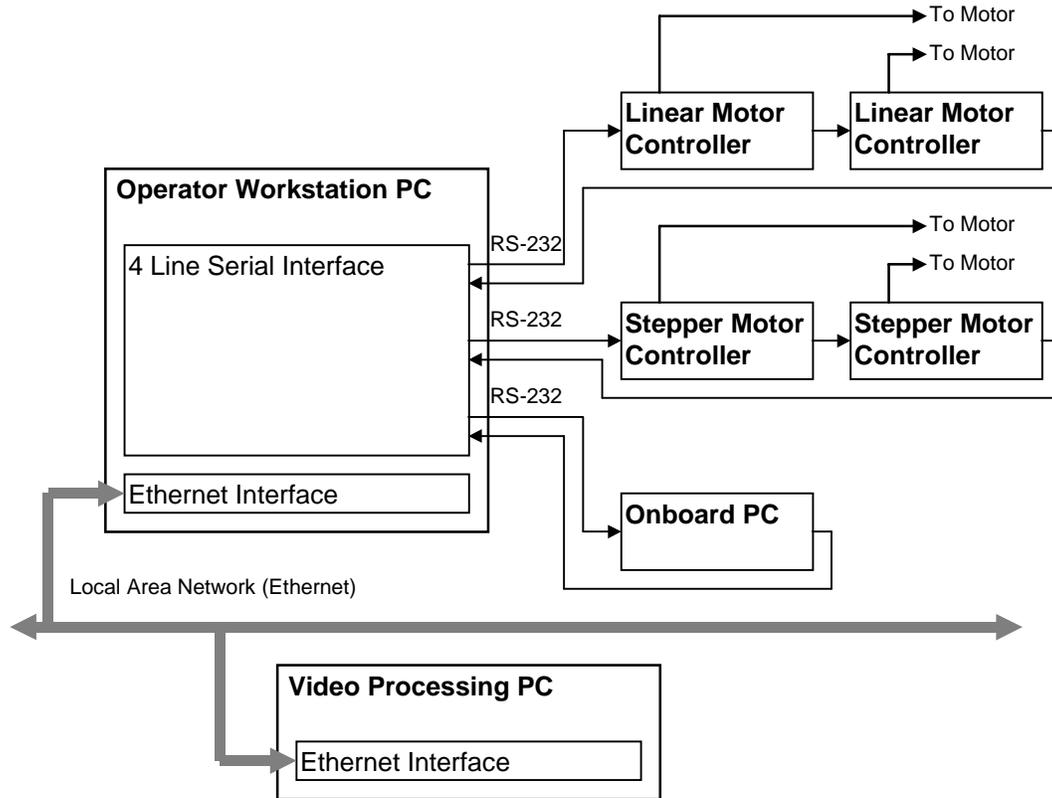
FIGURE 39. COMMUNICATIONS SUBSYSTEM

The stepper motor controllers drive the lead screw motors which are used to translate the bridges from side to side. The linear motor controllers drive the forcers which are used to move the bridges along the spine. Each of these motor controller systems is configured in a daisy chain, allowing multiple controllers to share a single communication line. Commands can be directed to specific controllers (which have separate addresses), allowing independent motion control and status reporting for each bridge.

The operator workstation and video-processing PCs are also connected to an Ethernet local area network. Communications between these two machines use standard TCP/IP protocols for high-speed command and image data exchange. This connection also provides access to resources on other computers such as disk storage, printers, and tape drives. Software development can be carried out on multiple computers and then easily transferred to the operator workstation as completed. This data path can also be used to provide access to the X-Windows user interfaces for control and monitoring of the robot through remote computers.

5.3   CONTROL SYSTEM SOFTWARE.

During the first-phase development, the main objectives for the control system were to develop a strategy that would control the walking and scanning motions of the robot, to implement the hardware and software necessary to perform the scanning motion, and to develop a prototype menu-driven operator interface for the system. The operator interface for the system consisted of character-cell based menu bars and pull-down menu selections. An integrated control system was designed and developed, and the software was able to control both the walking and the scanning motions of the robot; this was beyond the requirements of phase one, which specified only that the robot be able to scan a length of skin seam.

During the second phase of development the control software was split into a control process and a more intuitive, graphical user interface. The robot control process controls all higher level robot functionality. It supports the following operations:

- Walk (forward, back, left, right)
- Scan (forward, back)
- Home (calibrates motor position by sending bridges to a known location)

The robot control process provides the command interface to the motor controllers and the on-board PC and maintains a copy of the current robot status. The inspector interacts with the robot control process through the control/debug, router, and message interfaces.

The control/debug interface, figure 40, provides a schematic representation of the physical robot hardware for operating the robot. The spine and both bridges are presented. Individual control points are accessible via graphical on/off or arrow push buttons. Selecting these buttons with the mouse sends commands to the on-board computer to turn pneumatic valves on or off, or to move the motors. A key to the functions of the various control points is provided in table 1.

TABLE 1. ROBOT CONTROL INTERFACE LABELS

Label	Function
BRK	Turn leg brake on/off, (no longer used)
BYP	Turn leg high-pressure bypass on/off
LEG	Extend/retract bridge legs
MLK	Lock (synchronize)/unlock the sensor and stabilizer bridge motors
MRK	Extend/retract the flaw marking device
PRB	Extend/retract eddy-current sensor probe
VAC	Turn suction cup vacuum on/off

Status information obtained from the robot is indicated through color changes in the display:

- control points (On=White, Off=Gray)
- limit switches (On=Yellow, Off=Gray)
- suction cups (Vacuum=Green, Ambient pressure=Red)
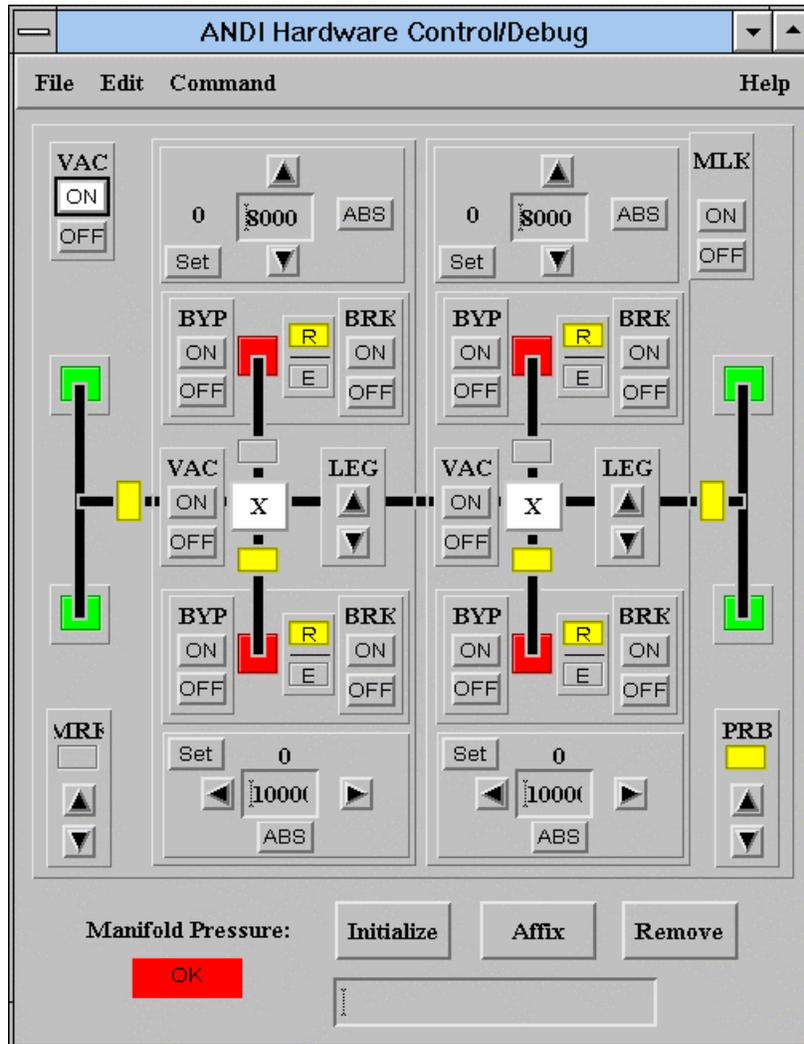
FIGURE 40. ROBOT CONTROL INTERFACE

The ANDI command interface, shown in figure 41, sends textual commands to any of the cooperating interfaces or processes through the router process. This provides support for both system debugging as well as more complicated motion sequences run from the control process. These moves include walking, scanning, and homing the robot's motors.

For system debugging purposes, low-level messages from the on-board computer are displayed in the communication screen, figure 42. The current status of each of the robot's inputs and output are divided into groups of eight bits, forming data bytes which are displayed using hexadecimal notation. There are currently five input and five output bytes used to reflect the robot's status.

Output transitions, text messages generated by tasks, and text responses to debugging commands are also displayed here.
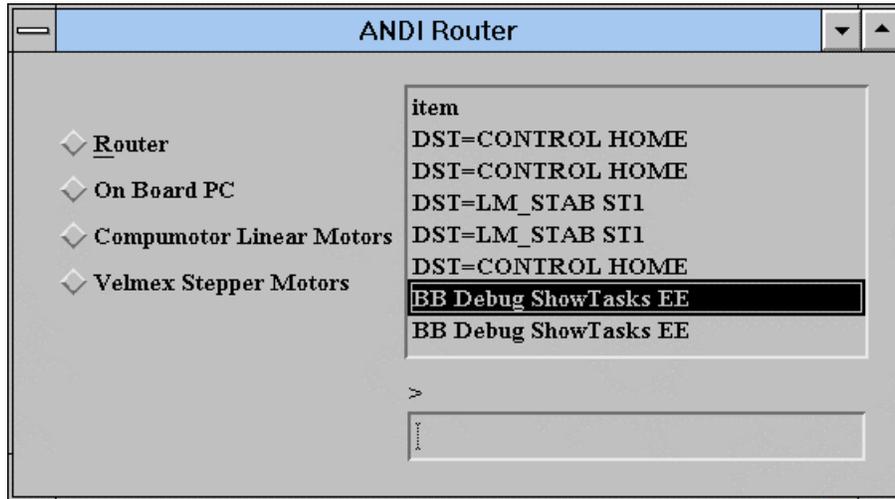
FIGURE 41. COMMAND INTERFACE
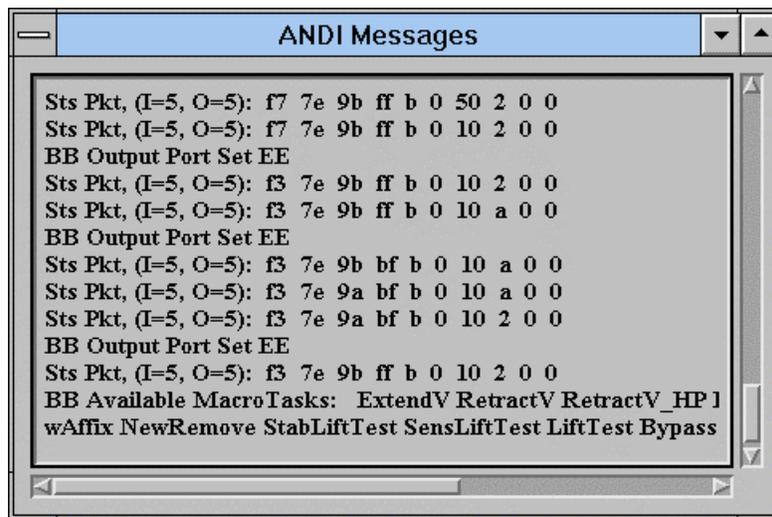
FIGURE 42. COMMUNICATION DISPLAY

5.4   NAVIGATION AND VISUALIZATION.

A prototype graphical display of the robot's inspection path was created using a three-dimensional, wire-frame representation of the surface of the airframe under inspection. While CAD data obtained directly from the aircraft manufacturer or the airlines would be preferred for use as the basis for aircraft maps, it is typically not available for the older generation aircraft now being inspected. An alternative is the use of models of various aircraft which are available from computer animation companies. While these models are optimized for their visual representation and do not generally include structural information, they do provide a sufficient framework to

provide navigation information at the resolution required by the robotic system. The frame and stringer or longeron intersections serve as navigation points whose approximate location can be obtained from service documentation and measurements from the aircraft. These intersection points define the lines where rivets are found and can be manually fit to the aircraft model using computer aided drafting tools. Individual rivets are not modeled. In addition to the ability to model aircraft, information concerning test fixtures, such as the Foster-Miller laboratory panel, are easily input. A sample display of the lines connecting the frame and stringer intersections on the Foster-Miller test panel used in the laboratory is provided in figure 43.

The operator has the ability to import various map data files into the viewport and select the point of view. Rivet lines are indicated in a contrasting color. In the envisioned system, the operator would be able to define, via an input device such as a mouse, the robot path by selecting appropriate lines of rivets for inspection. Areas that are selected for inspection or have been inspected or where anomalies have been found may all be indicated through the color used to display the segment. It is also possible to overlay an image of the robot on this display to monitor the position and configuration of the robot on the aircraft being inspected.
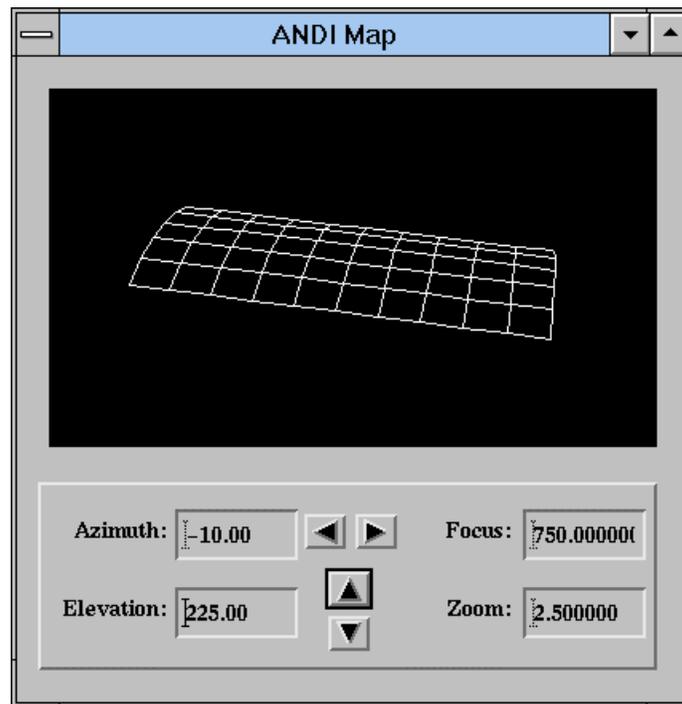
FIGURE 43. AIRCRAFT MAP DISPLAY (FOSTER-MILLER PANEL)

The U.S. Bureau of Mines has developed a three-dimensional animation and rendering of the robot and aircraft surface. This tool, which runs on a Silicon Graphics workstation, enables the visualization of the robot's motion and its interaction with the surface being inspected. The animation of the robot can be controlled using either predefined scripts of motion commands or via the X-Windows robot control interface described in section 5.3.

Components of the robot rendering software were rewritten during the third phase of effort to allow them to execute on the operator workstation computer. A view of the three-dimensional rendering of the robot is presented in figure 44. By coupling this display capability to the robot control system and the map of the aircraft, a virtual inspection environment that would be much more realistic than the wire-frame display shown in figure 43 could be presented to the operator.
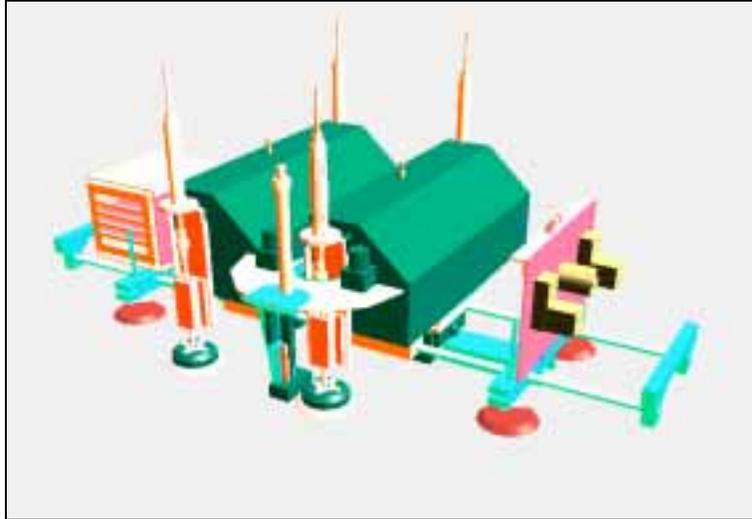
FIGURE 44. THREE-DIMENSIONAL RENDERING OF THE ROBOT

6.   SENSOR SYSTEM DETAILS.

6.1   EDDY-CURRENT SENSORS.

The majority of the nonvisual inspections are currently performed using eddy-current instruments. Thus, the major objective in selecting the sensors for the robotic inspection system was to be able to provide for both visual and eddy-current inspections. The automated deployment of eddy-current sensors was explored first because it is much easier to implement than is an automated visual inspection system. More specifically, the objective used to develop the sensor system was to initially investigate at least three NDI sensors for inclusion on the sensor platform: a high-frequency eddy-current sensor to find surface flaws, a low-frequency eddy-current sensor to find subsurface flaws, and video cameras for visual examination of the surface and for navigation.

6.1.1   Eddy-Current Inspection Hardware.

There are many types of eddy-current sensors used by inspectors to examine the fuselages of aircraft. Sensors were chosen that are both FAA-approved and operator-accepted. A comparison of results can then be made between the sensors that are deployed manually and those that are deployed robotically. One of the most popular types of sensors is the reflectance probe, also known as the sliding probe. Reflectance probes (pitch-catch) have separate driver and pickup coils and are deployed over a row of fasteners as shown in figure 45. They are scanned over the fastener heads, and the differences in the complex impedance-plane displays for good fastener

holes and those with cracks are noted. Because of its universal acceptance, the reflectance probe is a good choice to be included in the robotic system.
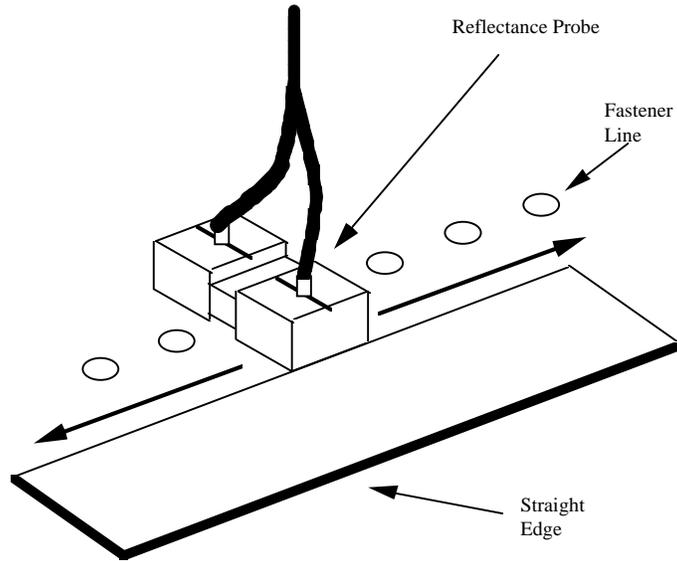
FIGURE 45. REFLECTANCE EDDY-CURRENT SENSOR

The Nortec SPO-1958 eddy-current sensor is the reflectance-type sensor chosen for the robotic system. This probe can operate at frequencies between 100 Hz and 40 kHz, allowing for both subsurface crack detection at low frequencies and surface crack detection at high frequencies. At its lowest frequencies, the SPO-1958 produces eddy-currents that have a working depth of 0.50 inch (13 mm). A photograph of the first-phase robot deploying the eddy-current sensor on the simulated aircraft panel is shown in figure 46.
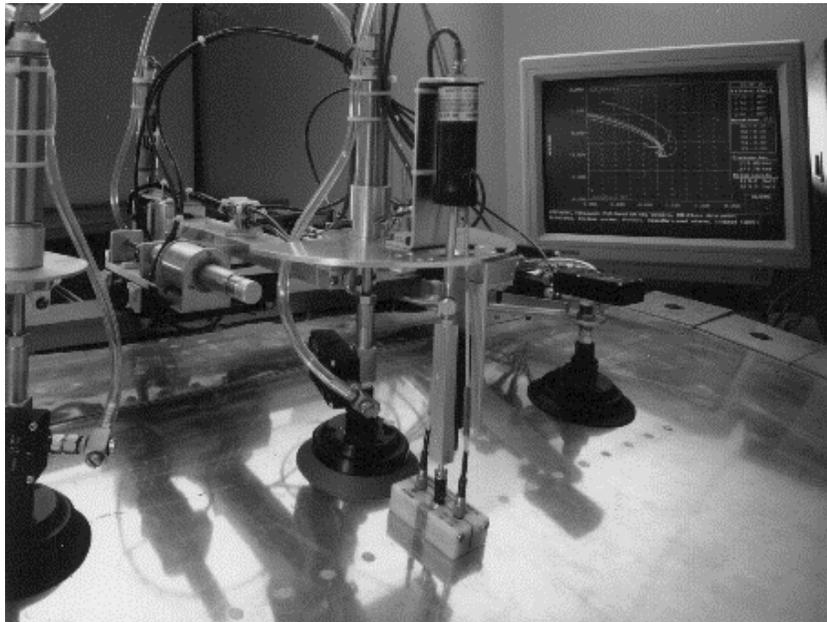
FIGURE 46. EDDY-CURRENT SENSOR DEPLOYED BY FIRST-PHASE ROBOT

Various eddy-current systems were evaluated for use in this project. SE Systems SmartEDDY 3100 system was chosen because it operated on a PC platform and could easily be integrated into the robotic system. The SmartEDDY 31000 has a dual frequency feature, meaning that one probe can be run simultaneously at two different frequencies. Thus, for the robotic system, one SPO-1958 can simultaneously search for both surface and subsurface defects; the electronics effectively makes this two independent sensors.

6.1.2   Eddy-Current Sensor Calibration.

During the first phase of development, the main objective in the calibration of the robotic system was to ensure that the operator sees a clear and distinct separation of signals between those produced by rivet holes on an aircraft containing cracks and those produced by crack-free holes.

The method of calibrating the eddy-current sensor of the robotic system was based upon the accepted methods of calibrating eddy-current sensors for manual inspections. Those methods require eddy-current sensors to be initially deployed on a calibration standard before they are deployed on an aircraft. A calibration standard contains both areas without flaws which result in a consistent trace on eddy-current instruments and areas containing flaws which result in traces that differ from those produced by the flaw-free areas. A standard is shown in figures 47 and 48, depicting the front and back of the standard, respectively. The front of the standard represents the outside skin which is normally accessible to inspectors while the back represents the substructure which is normally inaccessible to inspectors.

Manufactured Surface Crack
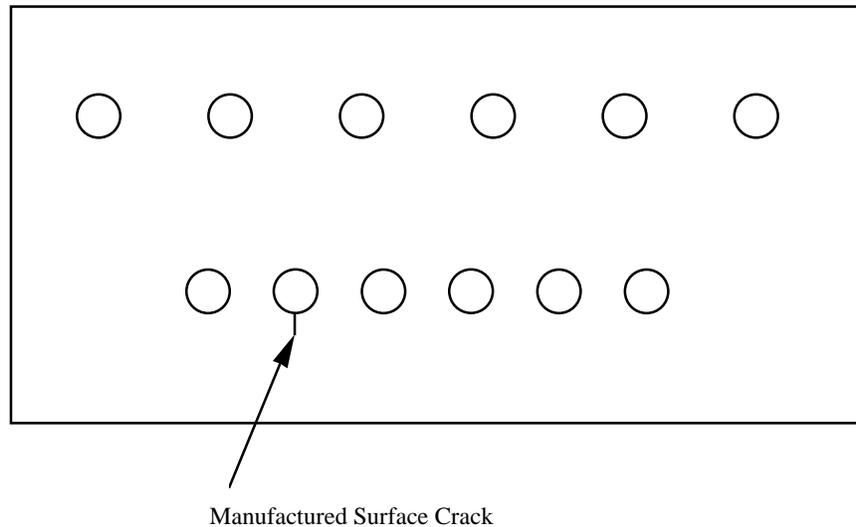
FIGURE 47. CALIBRATION STANDARD (FRONT)

Operators can see the differences between signals produced by flaw-free areas and those produced by areas with flaws by deploying an eddy-current sensor on a standard. After the sensors are manually deployed on a standard and the differences in traces are noted by the operator, the sensors are considered to be calibrated. When an inspection is executed on an aircraft fuselage, the unknown areas are classified with respect to the good and bad signals previously generated from the calibration procedure. Similarly, before the robot deploys its

47

eddy-current sensor on the simulated aircraft panel, the sensor is first deployed on a calibration standard. To keep the automated inspection system consistent with airline procedures, the inspection frequencies and calibration standards used in manual inspection procedures were also used by the robot.
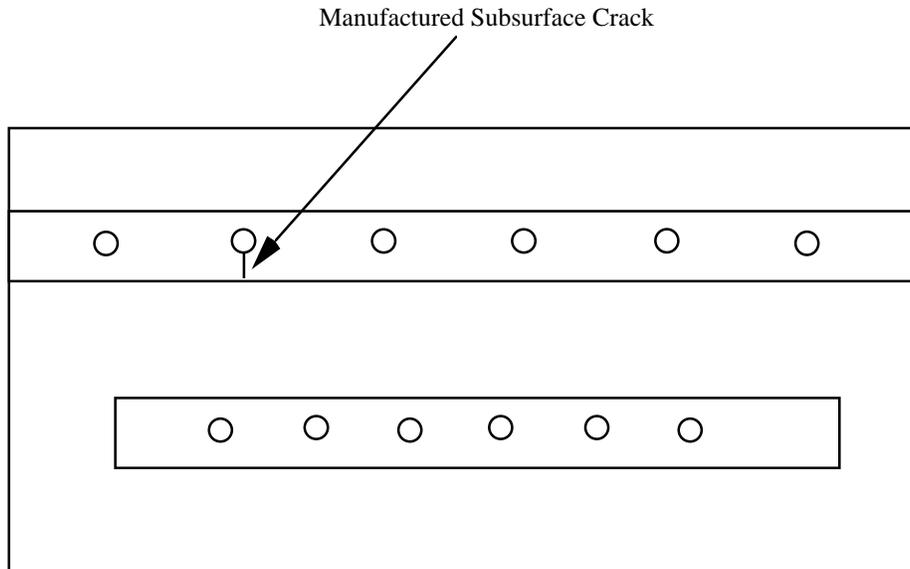
FIGURE 48. CALIBRATION STANDARD (BACK)
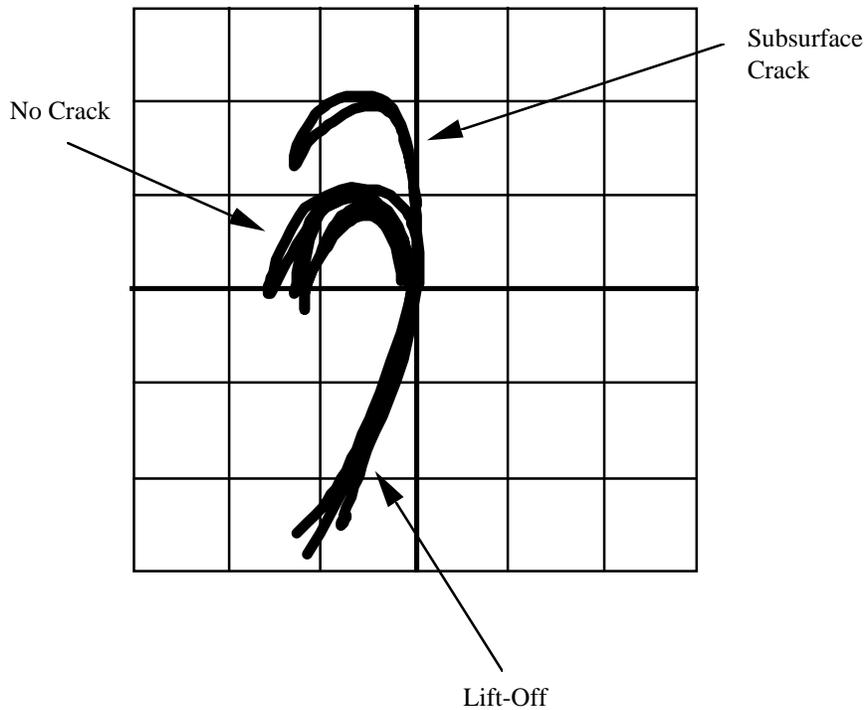
FIGURE 49. EDDY-CURRENT INSTRUMENT DISPLAY FOR SUBSURFACE CRACK

The calibration standard shown in figures 47 and 48 is the type used to calibrate the robotic system; it contains both surface and subsurface flaws. Because it contains both types of flaws, it is an appropriate standard for the robotic system. When it is being calibrated, the eddy-current data acquisition and analysis software is set for two frequencies: low frequency to detect subsurface flaws and high frequency to detect surface flaws.

The low-frequency setting is calibrated when the robot deploys the sensor and scans the top row of rivets in the standard. The system parameters for the low frequency calibration are set such that the eddy-current instrument display resembles figure 49, which is a drawing representing the oscilloscope output that an operator sees. An x-y position on the display represents the complex number that characterizes the impedance of the metal under the probe; the curves in figure 49 represent the variation of impedance as the probe is moved over several good fasteners and one containing a subsurface crack.

The signal produced by the subsurface crack is easily separated from the signals produced by the flaw-free areas. When this type of display is produced after the system has deployed the sensor over the top row of fasteners, the system has been calibrated to find subsurface cracks similar to those found in the standard.

Similarly, the high-frequency setting is calibrated when the robot deploys the sensor and scans the bottom row of rivets in the standard. The system parameters for the high-frequency calibration are set such that the eddy-current instrument display resembles figure 50.

FIGURE 50. EDDY-CURRENT INSTRUMENT DISPLAY FOR SURFACE CRACK

The signal produced by the surface crack is easily separated from the signals produced by the flaw-free areas. When this type of display is produced after the system has deployed the sensor over the bottom row of fasteners, the system has been calibrated to find surface cracks similar to those found in the standard. After both the high- and the low-frequency settings have been properly calibrated, the sensors can be deployed on the simulated aircraft panel to inspect for flaws.

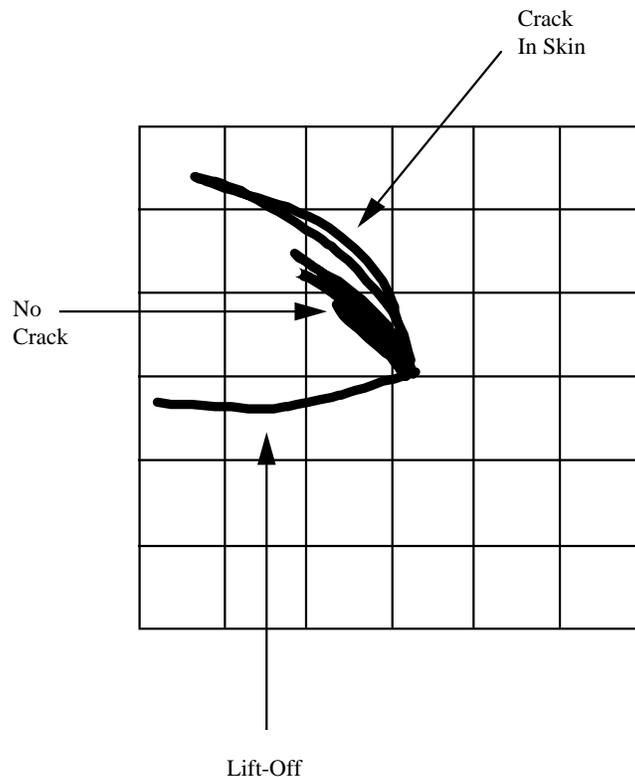The method described above to calibrate the system proved to work well. The system parameters for both high and low frequencies were adjusted to produce signals comparable to those shown in figures 49 and 50, respectively, providing a clear separation of signals for the operator. Thus, the eddy-current sensor on the robotic system was properly calibrated to find both surface and subsurface cracks simultaneously. The results were identical to those resulting from manually deployed eddy-current sensors.

6.1.3   Eddy-Current Software.

While a self-contained eddy-current inspection software package was included with the SE Systems, Inc. instrument, it was structured as a dedicated application. It could not exchange information with other parts of the control program, and it required the entire capacity of the computer. To integrate the eddy-current instrument with the robotic system, a custom software interface was written. The eddy-current process manages the instrument inspection mode, signal level, signal frequency, and sampling rate. It also performs the actual acquisition and buffering of the inspection data. The system can support sampling the eddy-current signal at rates of up to 1 kHz.

The eddy-current control interface is used by the operator to set the inspection parameters of the eddy-current instrument, as well as to display the acquired data. The inspection parameters that can be set in the eddy-current control screen in figure 51 include the A and B channel oscillator frequencies and levels, the reflectance or impedance inspection mode via the signal selection relays, and the data sampling rate. Although the instrument supports impedance mode operation, reflectance mode was required for the sample inspection tasks that were chosen for this project.
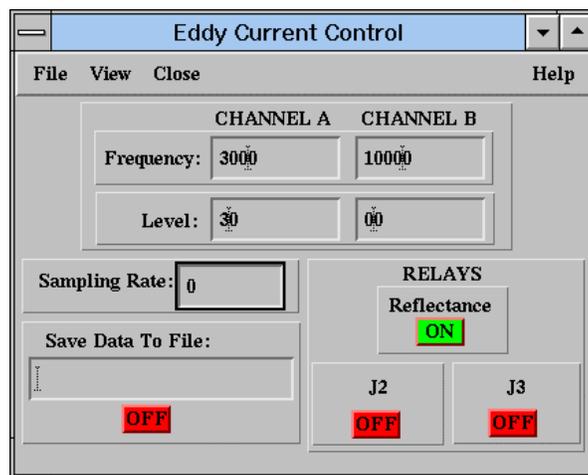
FIGURE 51. EDDY-CURRENT INSTRUMENT INTERFACE

Acquired data may be displayed to the operator or also archived to a file for later analysis. The data can be presented in both textual form (figure 52), and two graphical forms (figures 53 and 54). The graphical displays can plot the data as in-phase versus quadrature signals, figure 53, or as separate in-phase versus time and quadrature versus time plots, figure 54. Each frequency channel is displayed in a different color and the operator has the ability to zoom in on selected sections of the plots to view the data in greater detail. The data for these figures were obtained by scanning a calibration standard and include manufactured flaws on both the top surface and subsurface.

FIGURE 52. EDDY-CURRENT TEXTUAL DATA DISPLAY
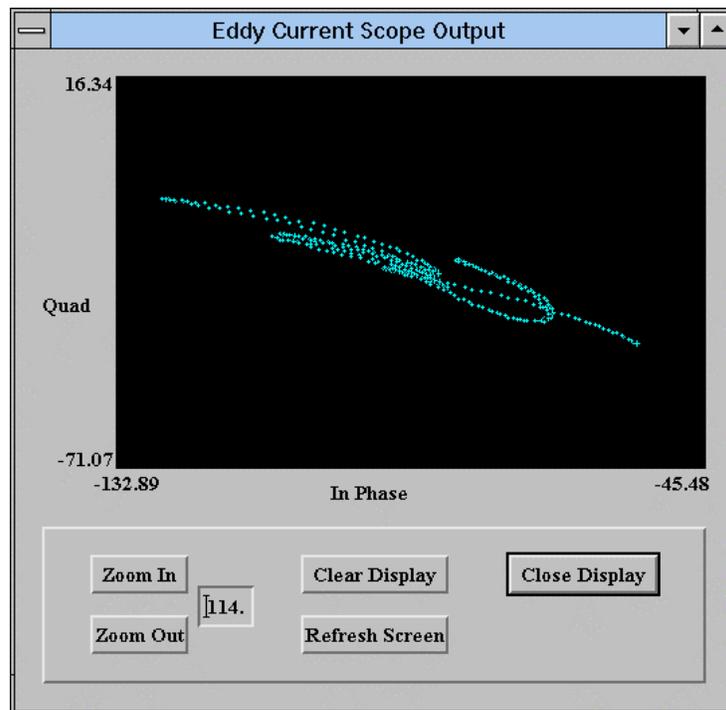
FIGURE 53. EDDY-CURRENT IMPEDANCE PLANE DISPLAY
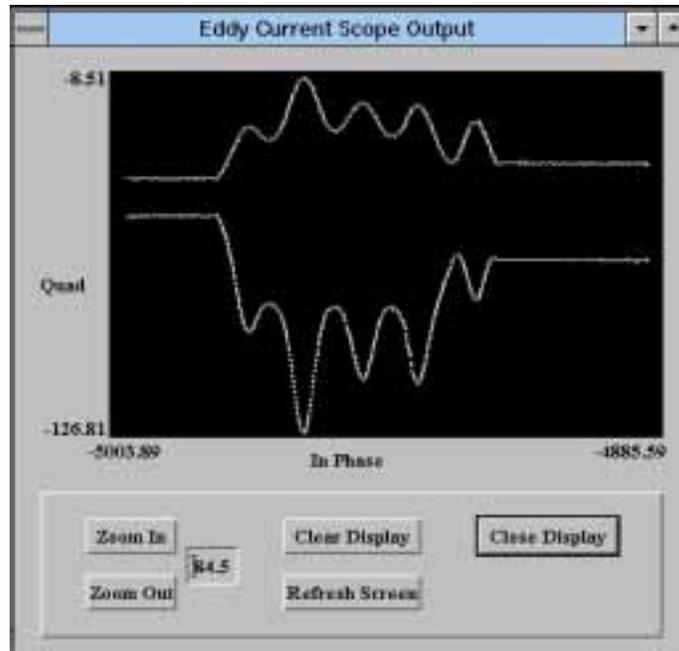
FIGURE 54. EDDY-CURRENT IN-PHASE/QUADRATURE VERSUS TIME DISPLAY

To aid operators in identifying abnormal signals, manufacturers of eddy-current instruments provide alarm thresholds. Such features are available in the commercial SmartEDDY software. In practice, these alarm features are seldom used during inspections since the oscilloscope display provides a richer kind of information, helping the operators to make subtle decisions. There is currently no automatic flaw detection capability in the system's eddy-current software. Simple, threshold-based flaw detection algorithms may be incorporated to more fully demonstrate the autonomous operation of the inspection system. An operator could be alerted to the existence of unusual signals as they are encountered.   The data will subsequently be examined by the inspector who will decide whether the signal was caused by a flaw or not.

6.1.4   Eddy-Current Flaw Detection.

Discussions with USAir NDI inspectors clearly showed that they classify eddy-current signals based on samples of good and bad signals generated when they deploy an eddy-current sensor on a calibration standard. They do not use quantitative means when classifying signals, but instead, base their decisions on how a signal compares to known patterns. While a person can easily distinguish between the signals produced by normal and flawed rivets, the signals are not unidimensional, and, in many cases, simple thresholding may not be definitive.   Pattern recognition methods, which capture the discrimination criteria being used by human inspectors without explicitly requiring the inspectors to articulate (in words) the criteria, are most suitable for automating this kind of decision making. In the future, pattern recognition techniques could be selected to aid in the classification of abnormal signals. The robot would mark the area producing an abnormal signal, and the operator would be responsible for examining the site and making the final judgment as to whether the abnormality was caused by a flaw or by some other anomaly.

52

As mentioned previously, during the first phase of system development, the only data analysis functions available to operators were those features found in commercial eddy-current systems. For example, to aid operators in identifying abnormal signals, manufacturers of most eddy-current instruments provide features which allow operators to set alarm regions. The SmartEDDY 3.0 eddy-current software package used as part of the robotic inspection system allows inspectors to set such thresholds. As described previously in the Calibration section, to calibrate eddy-current systems, inspectors deploy an eddy-current sensor on a calibration standard and set the eddy-current system's parameters to produce a clear separation between signals produced by fasteners without flaws and those produced by fasteners with flaws. Inspectors can also define a rectangular area on the eddy-current system's screen which will produce an audible and/or visual alarm if the impedance-plane signal enters this region. A drawing depicting this situation is shown in figure 55. In practice, these alarm features are seldom used during inspections; the oscilloscope traces provide better qualitative information for the operators to make intelligent decisions.
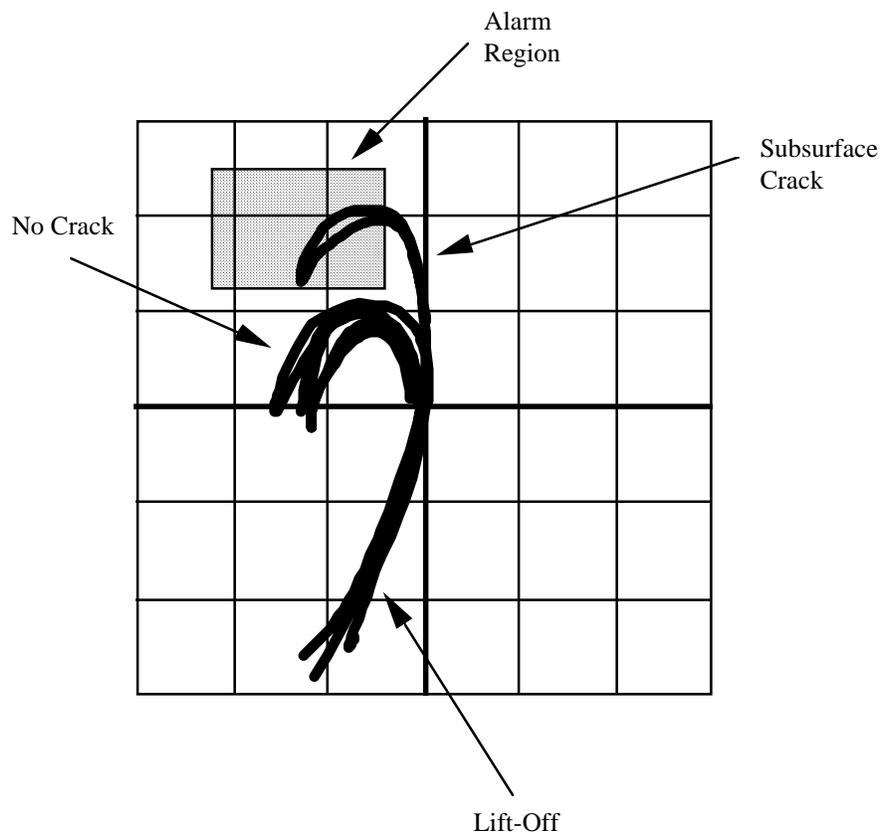
FIGURE 55. ALARM REGION ON AN IMPEDANCE-PLANE DISPLAY

6.2   VISUAL SENSORS AND VIDEO-BASED RIVET LOCATION ALGORITHMS.

Visual inspections currently account for a large portion of the inspections performed on U.S. commercial aircraft. Providing a means of acquiring images of the fuselage surface for inspectors is very important. Aircraft inspectors note that, even when they are performing nonvisual nondestructive inspections on a fuselage, they tend to notice other defects as well. For

example, when inspectors use eddy-current sensors to examine lap joints for cracks around the fasteners, they are also looking at the area for other types of flaws and damage, such as dished rivet heads and lightning strikes. Video imagery also provides information required to align the robot for taking eddy-current measurements and improves the teleoperation capability of the robot. A total of four video cameras were installed on the robot to support navigation and alignment, close-up visual rivet inspection, monitoring of the robot, and large area visual inspection. Although computer analysis of the imagery to automatically detect surface flaws and damage would further automate the inspection process, its development is beyond the scope of this project.

Two Chinon CX-060 miniature black and white video cameras were used for navigation and alignment of the probe with respect to the rivets on the fuselage. One camera is placed at each end of the robot's spine assembly looking down at the aircraft's surface. Each of these cameras has a field of view of approximately 4 by 6 inches (10 by 15 cm). These cameras are enclosed in opaque plastic hoods with controlled, low-angle illumination to provide high-contrast images of the rivets against the skin. The images obtained from these cameras are used to align the spine of the robot parallel to the row of rivets to be inspected. Algorithms have been developed for finding rows of rivets to provide feedback for automatically aligning the robot.

To provide additional operational feedback to the operator, two Elmo MN401E 1/2-inch CCD, (768H x 494V pixels), color cameras were installed on the robot. These cameras each require an auxiliary control unit, located in the satellite equipment enclosure, and a multiconductor control cable. The first camera is the "proprioception camera" which is mounted at the front end of the robot on top of the computer card cage. This camera has a wide field of view and is aimed back over the robot so that the operator may view the robot and its immediate surroundings. It is used for obstacle avoidance, long-range navigation, robot positioning and status evaluation, and wide-area visual inspection. The future use of a pan and tilt base with the camera, as well as additional illumination, could provide more flexibility in the application of this camera. The second color camera is mounted on the sensor platform and is aimed at the eddy-current probe. This camera provides a detailed visual view of the rivets that are being inspected. The camera is used to confirm probe location and to give the inspector a close-up view of selected rivets. Although this camera uses ambient Illumination, a ring light could also be mounted on the system.

A block diagram of the video subsystem is provided in figure 56. The selection of the currently active video input from the available video sources is done by the computer using the video multiplexer which has twelve video inputs and seven video outputs. The video signal is digitized by the EPIX 4MEG VIDEO™ Model 12 video digitizer/processor. This card has 4Mb of video memory, a 14.3 MHz pixel clock, and a dedicated TMS320C25 12MIP digital signal processing chip. It is capable of driving an RGB video monitor with either live or processed video imagery. Both of these cards are installed in the video-processing PC.
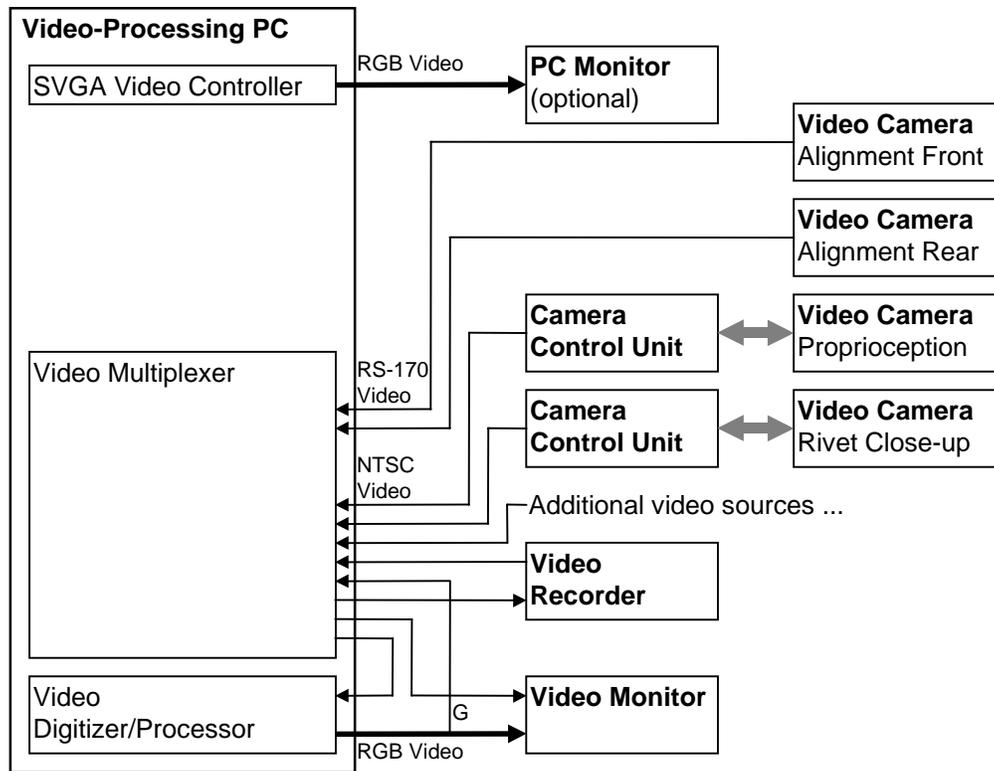
FIGURE 56. VIDEO SUBSYSTEM

A video monitor is normally used to display the video data to the operator. However, the capability exists to display the video data in a subwindow within the computer monitor. A S-VHS video recorder is provided to archive the video data from the cameras as well as to feedback previously recorded data into the system for the testing and validation of image analysis algorithms.

The video-processing computer is used to perform all of the video switching and processing functions required by the system. It is configured as a separate machine because of the large amount of processing that is required to perform computer video analysis. Thus, the processor resources used for image analysis are independent of the robot control, eddy-current data acquisition, and the user interface. A photograph of the video-processing computer is presented in figure 57.

The video-processing application runs as a TCP/IP network server on the video-processing PC. It sets up the video digitizer and multiplexer hardware and then listens on the local area network for commands from the operator workstation PC. Video data may be acquired and displayed from multiple sources, saved to or restored from disk, and overlaid with an alignment cursor.

FIGURE 57. VIDEO-PROCESSING COMPUTER

The video control interface is used by the operator to manually adjust the robot alignment prior to eddy-current inspection. It consists of both text and graphically-based commands and is shown in figure 58. Data from any of the cameras on the robot may be selected through the "Video Source" menu. Single frames of video are digitized and displayed each time the "Digitize" button is pressed. The look-up table used for displaying the digitized image data is presented as a horizontal bar at the bottom of the image area.
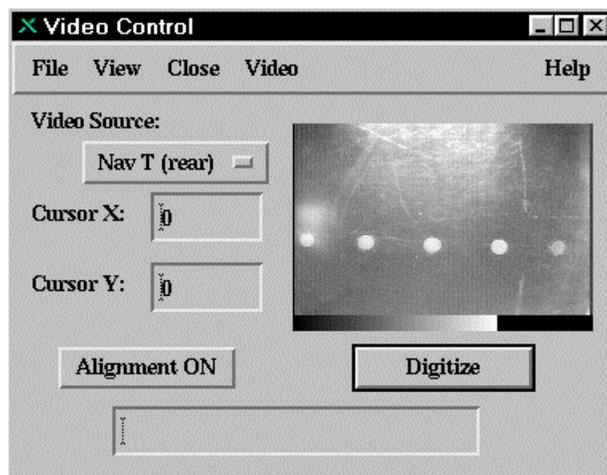
FIGURE 58. VIDEO CONTROL INTERFACE

The "Alignment ON" button is used to set the robot in the alignment mode with legs extended and pivot locks off. Electronically overlaying the video from the front and rear alignment cameras with a calibrated cursor enables the operator to manually align the robot with the rivet rows. The operator can then rotate and jog the robot to bring the eddy-current sensor into the

correct position for scanning by depressing the mouse button and dragging the mouse while the cursor is in the image display region of the window. Once properly aligned, the robot is returned to the scanning mode with legs retracted and pivot locks on. Text-based commands may be sent to the video-processing PC to set the position of the video cursor for the front and rear alignment cameras.

To allow the robot to align the eddy-current probe to the rivets undergoing inspection as well as to navigate autonomously over the aircraft's surface requires algorithms to detect the presence and measure the location of rivets in the video imagery from the navigation cameras. Once isolated rivets have been detected, they can be consolidated into lines of rivets which form a natural reference grid for navigation over the surface of the fuselage and also provide the feedback required for the deployment of the eddy-current inspection probe. Algorithms to do this rivet detection and rivet line abstraction were developed at the CMU Robotics Institute [6] and then extended and applied to the robot at Carnegie Mellon Research Institute.

Figure 59 depicts the role of rivet segmentation and line-fitting in the navigational control of the robot. The process begins when a pair of images of the aircraft surface is captured by the head and tail navigational cameras. The images are digitized, then processed by the video-processing computer to segment the rivets from the image background. A line-fitting algorithm fits a straight line through the rivets in each image. The results from the two images are combined to indicate the path that will position the eddy-current sensor directly over the rivets.

1A	1B	1C	1D	1E
camera (head) / camera (tail)	digitizer	rivet segmentation algorithms	line fitting algorithm	robot navigation control
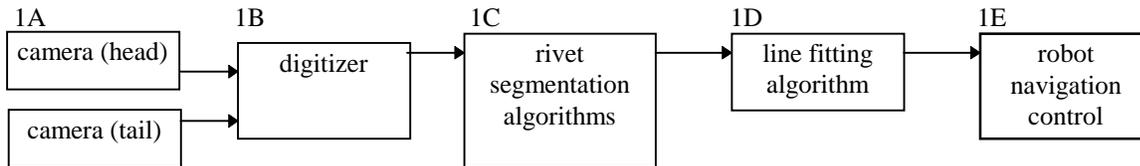
FIGURE 59. RIVET SEGMENTATION AND LINE-FITTING IN ROBOT NAVIGATION

Two types of segmentation algorithms were studied (block 1C in figure 59), neural networks and conventional methods. These two techniques present different approaches to image processing. Neural networks are empirical by nature, deriving a solution from observations of data, in this case images of the aircraft surface. Conventional techniques are parametric, requiring a mathematical description of the solution. Each approach was evaluated for computational speed, segmentation accuracy, and the complexity of its implementation. In the end, the conventional techniques outperformed the neural networks, providing the fastest computation time and simplest implementation while providing comparable accuracy.

Neural networks (NN) were initially investigated because of their unique ability to learn solutions. Unlike the conventional techniques discussed below, where the criteria for a classification match are mathematically specified, a NN gradually develops its solution by repeatedly viewing images along with the corresponding correct classification. Over repeated sessions, the NN learns to recognize which image features enable it to distinguish between classifications. For rivet segmentation, it was hoped the NNs would learn which combination of features distinguish rivets from other, perhaps similar looking surface marks. The NN would thus learn features that discriminate between the classes of rivet and background.

One inherent drawback to NNs is the need for significant quantities of training data (i.e., images of rivet and background classes) large enough to represent the problem space. This can be a labor intensive task, and if the problem space changes, such as when a new classification group is added or a different aircraft is used, new data must be acquired and the NN must be retrained. On the other hand, one of the purported advantages of NNs is that the burden of finding discriminating features and developing mathematical descriptions of the classes is left to the NN which should learn these by itself.

Conventional techniques refer to a commonly used set of image processing algorithms for image enhancement, restoration, and segmentation. They include such techniques as histogram equalization, edge detection, high-pass or low-pass filtering, and many others. Unlike NNs, there is no need to create large quantities of training data. However, they are generally most successful if the investigator has a good understanding of the problem and can make judicious choices between the available techniques.

6.2.1   Robotics Institute Video Rivet Location Algorithms.

If the rivets that hold an aircraft's skin to its frame all looked the same, then it would be easy to construct a dedicated algorithm that would reliably find and mark all the rivets in an image. However there are many types of fasteners in use on new aircraft, many more types in use on in-service aircraft, and even identical fasteners on the same aircraft take on different appearances depending on the details of their installation, location, weathering, buffing, coating, etc. To accommodate this variation range, an open ended approach was adopted using a neural network architecture that learns from visual examples presented to it which have been classified by a human trainer as containing or not containing rivets. The visual examples are small square windows on the full image; a window's side is about twice the diameter of the largest rivet head that will be encountered. Thus the window may contain all or part or none of a rivet head, but it will rarely contain more than one.

In operation, the trained neural network generates a numerical output, that we call "rivetness," that represents its estimate of the probability that the window it is being shown contains a rivet. The window is scanned over an image resulting in a rivetness map that is bright in areas where rivets are probable, dark in areas where they are improbable, and gray in ambiguous areas. A sample of the rivetness image is provided in figure 60. Via several subsequent conventional image processing steps the rivetness image is converted to a black-and-white (binary) image in which rivets are white, non-rivets are black. The binary image obtained from figure 60 is shown in figure 61. The binarization process completes the rivet finding step.
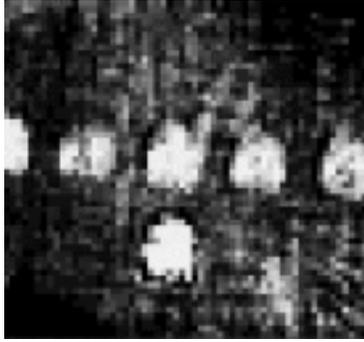
FIGURE 60. RIVETNESS IMAGE

FIGURE 61. BINARY IMAGE WITH LINE

Next, it is necessary for the computer to abstract isolated rivets into rivet lines and a rivet line grid; this kind of abstraction, which is the natural thing for the human brain to do when it is presented with isolated dots that happen to fall on lines and rectilinear grids, is a potentially difficult task for a computer. It is a simple and very common procedure to fit a line to a set of data points in some coordinate system, e.g., by the least squares method, but exactly where this line falls can be very strongly influenced by a small number of points that are far from the main distribution; that is, these methods are inappropriately sensitive to outliers. Any image containing a line of rivets will probably also contain a few extraneous rivets that would be effortlessly rejected as irrelevant by any human but that would not be recognized as irrelevant by conventional line-fitting algorithms. Fortunately, modern statisticians have recognized this problem, and in response they have developed robust line-fitting methods that do an excellent job of focusing on the main stream and ignoring the outliers. These algorithms have been extremely successful in identifying rivet lines. The line in figure 61 is the output of the robust line-fitting algorithm.

An existing Robotics Institute laboratory robot that was built as a development platform for enhanced visual inspection sensors (not part of the FAA program) was used to test the rivet detection and line-fitting algorithms. The "test inspection platform" (TIP) robot is capable of pure rotation (with negligible forward/backward or left/right motion), so it could follow the requested path precisely despite the path's right angle corners. It was guided around a rectangular path over a grid of rivet holes in a sample of DC-9 belly skin under the closed loop control of the algorithms. To provide a rapid demonstration of principle, a multicomputer pipeline was assembled in which the tasks of image collection, image transmission, rivet detection and rivet line abstraction, robot control strategy generation, robot control command generation, and robot control command execution were done on several different workstations, in several different locations, and communicating over the campus Ethernet LAN. Although the decision making process was slow, primarily because of the need to move large image files over the Ethernet, the navigation sequence was executed without error, the path following accuracy was excellent, and the TIP robot always discovered and corrected random orientation errors that the top-level program introduced to challenge the lower-level programs.

6.2.2   CMRI Video Rivet Location Algorithms.

The transfer of the initial algorithms into the robotic system was conducted in three steps, beginning with the neural network (NN) investigation. Despite several modifications to the network architecture, and numerous image enhancement efforts, the NN segmentation accuracy was unacceptable in this first stage of work. It was concluded that the images contained too much competing background clutter in the form of scratches, surface marks, and reflections for the NN to accurately discriminate between rivets and these other objects. In the second step, conventional segmentation algorithms were implemented on these same images, but with similar, unsatisfactory results. The video data obtained from the Foster-Miller panel exhibited much less contrast between the rivets and the surrounding skin and had many more surface scratches and dents than the imagery used in the previous algorithm development. Since it was evident that poor image quality was responsible for the poor segmentation in both studies, the third step in the investigation focused on improving the images by modifying the camera system to improve image quality. A camera hood was developed that greatly reduced surface reflections and consequently enhanced the contrast between rivets and background. The NN and conventional segmentation techniques were reevaluated on the new images. Results improved significantly. The conventional algorithms, however, demonstrated superior computation time over the neural networks while providing comparable segmentation accuracy.

Figure 62(a) shows a typical image from the head navigational camera, without the camera hood. This is the surface of the Foster-Miller panel (oriented roughly horizontally in its fixture) imaged under ambient lighting conditions. No image enhancement is shown. There are five rivets positioned in a horizontal line just above center in the image. Poor contrast makes the two leftmost rivets barely visible.

Ideally, segmentation of figure 62(a) will resemble that of figure 62(b) where rivets, displayed in white, have been segmented from the rest of the background, shown in black. Figure 62(b) was created manually to illustrate the goal. The following sections describe the efforts to develop segmentation algorithms to achieve performance as close as possible to the goal of figure 62(b).
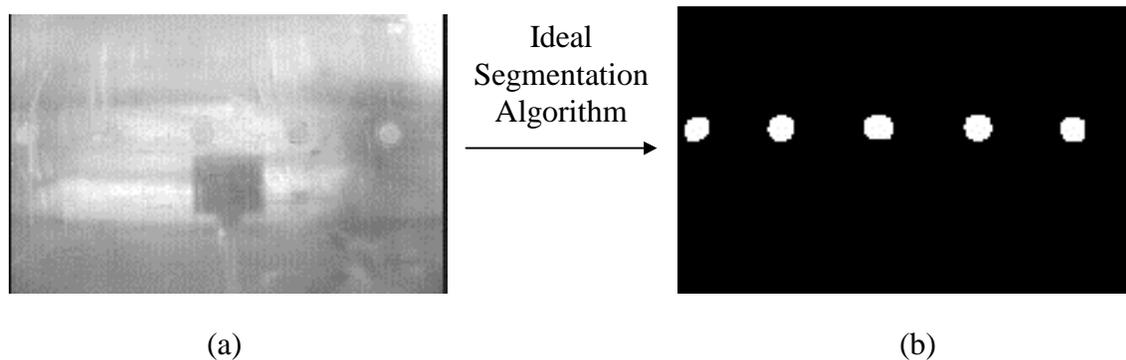

Ideal
Segmentation
Algorithm

(a)   (b)

FIGURE 62. (a) AN IMAGE FROM A NAVIGATIONAL CAMERA, (b) AN IDEAL
SEGMENTATION

To determine if NNs could provide the segmentation quality shown in figure 62(b), the first task was to create training and testing data for teaching the network. The data was designed to be representative of the problem space including images of rivets and a variety of surface markings such as scratches and reflections. Fifty-six images of the Foster-Miller panel were acquired from the navigational cameras under ambient lighting conditions. They were equally divided (14 images each) between the following conditions:

- head camera, legs retracted
- head camera, legs extended
- tail camera, legs retracted
- tail camera, legs extended

Examples are shown in figure 63. The camera lens is approximately 4 inches (100 mm) above the panel, however this varies depending upon the position of the bridges on the lead screws and whether the legs are extended or retracted. The field of view is slightly larger with legs extended (figure 63(b), (d)). In these examples, the rivets are positioned along a single horizontal line just above center in the images. The dark, square object at the image center is the shadow of the camera. The images are of size 188 x 120 pixels. The digitizer generates 752 x 480 pixel images. By reducing the images by a factor of two or four, to size 376 x 240 or 188 x 120, respectively, processing speed is greatly increased without sacrificing segmentation accuracy.
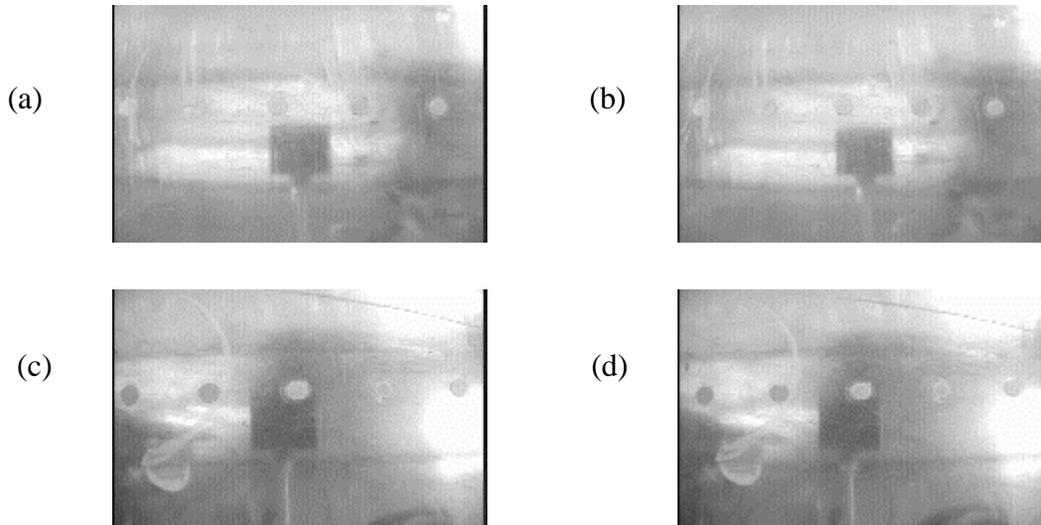
FIGURE 63. EXAMPLES OF IMAGES FROM (a) HEAD CAMERA, LEGS RETRACTED; (b) HEAD CAMERA, LEGS EXTENDED; (c) TAIL CAMERA, LEGS RETRACTED; (d) TAIL CAMERA, LEGS EXTENDED

Training and testing data were created as pairs of NN input/output responses. Each NN input, called a retina, R_i, was an m x m pixel subimage taken from the reduced image. Each retina R_i, centered at pixel location x_i, y_i, was a few pixels larger in each dimension than the diameter of a

rivet. For a 188 x 120 image, the retina was 13 x 13 pixels since the rivets are approximately 9 pixels in diameter.

The corresponding output paired with each retina was a scalar value, 1 or -1, depending on whether the retina contained a rivet or not. If a retina contained a full rivet it was classified as 1, i.e., a rivet. A retina with no rivet or only a partial rivet was classified as -1, i.e., background. Thus, the paired input/output response data for training and testing consisted of thousands of records, each record comprised of a 13 x 13 input retina (for a reduced-by-4 image), and a scalar output of 1 or -1. Figure 64 shows examples of four rivet retinas, and figure 65 shows four background retinas. The retinas have been enlarged for display purposes.
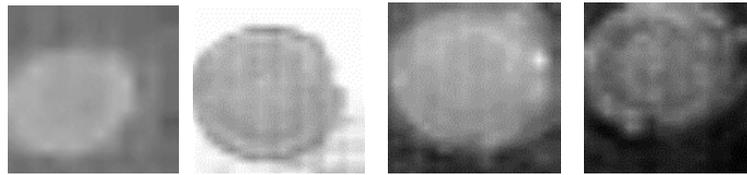
FIGURE 64. FOUR RETINAS OF RIVETS
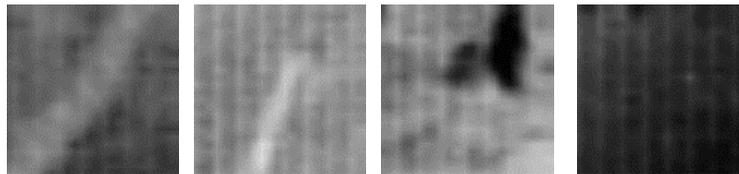
FIGURE 65. FOUR RETINAS OF BACKGROUND

Retina classifications for the training/testing data were established manually. Each image was displayed on a PC while an operator used a computer mouse to highlight the boundaries of each rivet. Software then filled in these regions and labeled all the filled pixels as class "rivet body." The pixels along the rivet edge were classified as type "rivet edge." A few images contained raised head rivets. These were manually segmented and labeled as type "raised rivet." All other pixels were classified as type "background." Figure 66 shows a manually segmented image with rivet edge pixels in white, rivet body pixels in gray, and background in black. Raised head rivets were omitted from all training data (there are none present in figure 66).

Rivet body and rivet edge pixels were labeled separately to provide more flexibility in the tests. Experiments were conducted to see if the NN was better able to identify a whole rivet body or just its edge. Results proved no significant advantage in either scenario, thus only rivet bodies were considered for the second half of the studies.
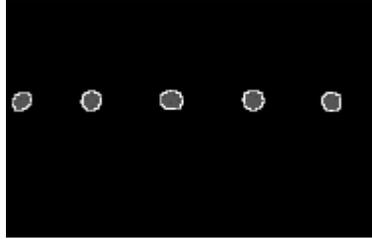
FIGURE 66. MANUALLY SEGMENTED IMAGE USED TO CREATE TRAINING DATA
WITH THREE LEVELS OF GRAY SCALE

Once all 56 images were manually segmented, a retina/classification pair was created for every pixel in the images. From this large set, thousands of pairs, called records, were randomly selected to create training and testing data. The NN learned on the training set, repeatedly viewing the data and modifying its internal connection weights each time it generated an incorrect output. Test set data were different records from those in the training set. Test data were used to monitor the learning process and stop the learning when test results no longer improved. At the start of the NN study, the training data contained an equal number of rivet records and background records. This is a good initial rule of thumb if each classification is equally important. However, later in the study, a distribution of twice as many background retinas as rivet retinas was shown to improve the segmentation accuracy by a few percentage points. This may be because most of the rivets are fairly similar in appearance, but the background retinas exhibit more variation in gray level and shape. Therefore, more retinas were needed to represent the background.

Neural networks are defined by a topology that describes the number of processing nodes and their connections. A learning algorithm defines how the connection weights are adjusted during learning. In this study, the topology was a three-layer network (input, hidden, and output layers) with 169 input nodes (one for each pixel in the 13 x 13 retina), a variable number of hidden nodes, and one output node. The output node generated a continuous value between -1 and 1. The closer the value was to 1, the more the NN thought the input retina resembled a rivet. For values closer to -1, the retina was thought to be a background retina. Back-propagation (BP) was selected as the learning algorithm because it is commonly used in complex classification problems.

Each processing node sums multiple inputs from the previous layer of nodes and operates on this sum with a nonlinear transfer function. The most commonly used transfer functions are the sigmoid and hyperbolic tangent (tanh). Previous experience in developing classifiers has shown that a tanh network learns faster than a sigmoid. Consequently, tanh was used here. Tanh generates an output value between -1.0 and +1.0 for any given input. Also, scaling the NN input values, i.e., the retina, to the range of -0.8 and 0.8 when using tanh helps prevent saturation of the nodes, which can stall learning. Therefore, while the retinas were originally gray scale 0 to 255, they were scaled to -0.8 and 0.8 before being input to the NN.

The number of hidden nodes, and two parameters known as the momentum and learning coefficient, were varied to optimize performance. Momentum and the learning coefficient affect the rate of learning. The number of hidden nodes also affects the learning speed, but more

importantly, it affects the network's ability to learn relationships between input and output data. Too many hidden nodes with respect to the number of training records allows the network to memorize the training data. Too few hidden nodes and the network may not learn a good solution. The selection of the number of hidden nodes is most often a trial and error process. There are several so-called automated techniques for selecting the number of hidden nodes; however, manually guided trial and error has been found to be the most reliable. One automated method, Singular Value Decomposition, was used to give a rough estimate, which was later confirmed by trial and error.

During learning, the number of hidden nodes, the momentum, and the learning coefficient were all varied in repeated tests to optimize the NN performance. Images were reduced in size from their original 752 x 480 by a factor 2 or 4 in both dimensions to reduce processing time. Larger images necessitate larger retinas and therefore more input nodes, which translates to more connections, which requires more training data, and the larger network takes longer to train. For these reasons, smaller images were preferred.

When each NN was fully trained, its segmentation accuracy was evaluated. The fully trained network processed each image one retina at a time, generating the segmented image one pixel at a time from the NN output node value. Each image to be segmented was dissected into retinas at every pixel location and input to the NN one by one. The NN output was the pixel value in the segmented image at the retina's center location. Since the NN output is a continuous value in the range of -1 and 1, the segmented image is gray scale.

Initial results were poor using both rivet body and rivet edge approaches. Although the NN recognized rivets well enough, there were too many background regions that were also classified as rivets. These incorrectly identified rivet areas would confuse the subsequent line-fitting algorithms. The presence of these areas was attributable to the cluttered quality of the original images. Numerous enhancement schemes were used to improve the quality of the input images. Histogram equalization, edge detection, median filtering, smoothing and others were attempted to improve the NN segmentation accuracy. Modest improvements were visible, but were not significant enough to make the NN a viable approach. Figure 69 shows an example of a reduced-by-2 image after histogram equalization and the resulting NN segmented image. Note that while rivets are visible, there is also considerable clutter that would impede any line-fitting algorithm. Reduced-by-4 image results were similar.
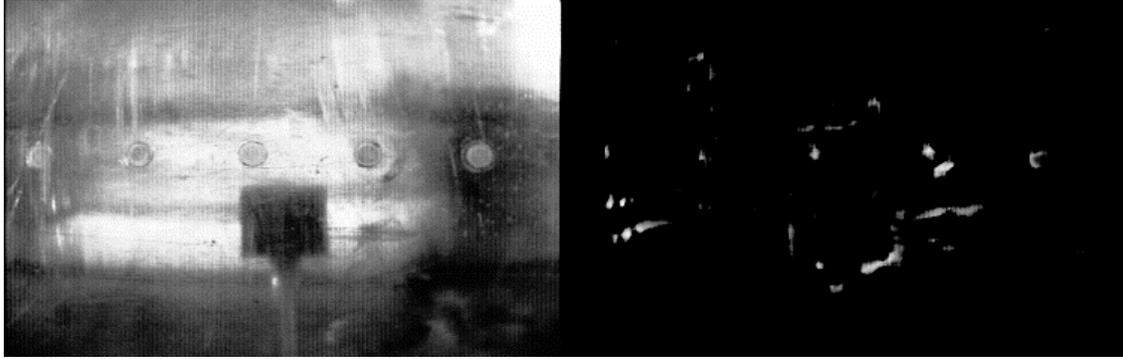
FIGURE 67. LEFT: REDUCED-BY-2 IMAGE, ENHANCED BY HISTOGRAM
EQUALIZATION, RIGHT: NN SEGMENTED OUTPUT

6.2.2.2   Conventional Segmentation Results.

The next step was to investigate more conventional algorithms of segmentation. Some of the image enhancement techniques that were started in the NN study were continued here. Camera images were enhanced to reduce noise and clutter, then a classifier was applied to recognize the rivets. Many enhancement algorithms were tested both alone and in combination. The best enhancement approach was found to be a combination of histogram equalization, median filtering, Sobel edge detection, and thresholding.

The first step, histogram equalization, improved contrast by evenly redistributing the gray scale values over the range 0 to 255. The next step, median filtering, smoothed the image slightly, reducing the visible vertical bands created by the digitizing process. Next, a Sobel edge detector is applied and its absolute value taken to extract only edges. The Sobel image is gray scale, and a threshold of 150 (out of possible 255) is applied to remove weak edges. Histogram equalization is then applied to improve image contrast. Figure 68 shows the results of this processing on a camera image. Note that while the rivets are visible in the final image, there is also considerable background clutter.
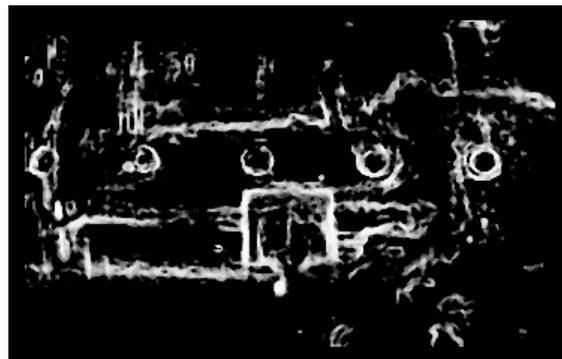

FIGURE 68. EFFECT OF IMAGE ENHANCEMENT ON GRAY SCALE CAMERA IMAGE

To complete the segmentation, a matched filter classifier was implemented to distinguish rivets from clutter. A matched filter correlates the image with an image of the object to be recognized,

65

in this case, a rivet. The image of the object is called the kernel and is shown in figure 69. This kernel is a rivet of average size, taken from one of the enhanced images that underwent the same processing as figure 68. Because of the curvature of the airplane panel, not all rivets are imaged to the same size. There is slight variation, and figure 69 shows an average size rivet in a 23 x 23 window.

FIGURE 69. EXAMPLE OF RIVET KERNEL

To perform a matched filter, the enhanced image was cross-correlated with the rivet kernel of figure 69. Ideally, the resulting cross-correlation image contains bright spots (highly correlated regions) at rivet locations and dark regions everywhere else. Figure 70 shows the output from the matched filter operation of figure 68 with 69. Results using smaller and larger rivets in the kernel were similar.
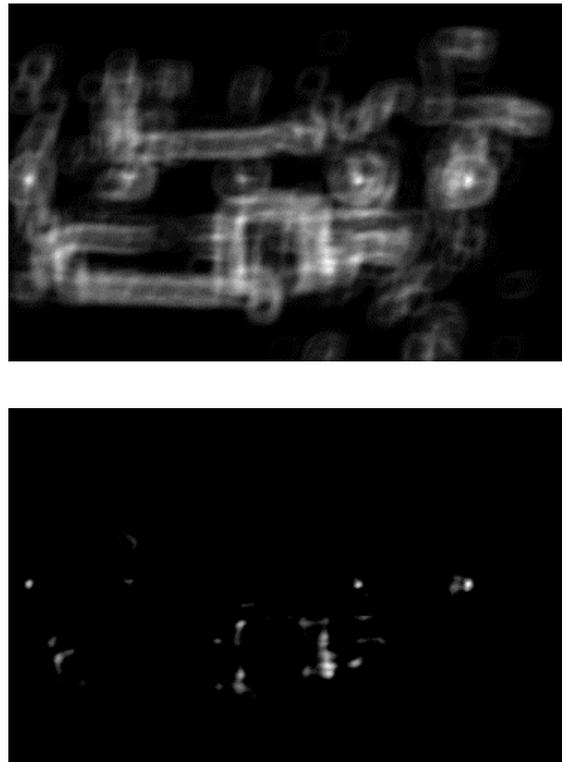
FIGURE 70. (TOP) MATCHED FILTER OUTPUT, (BOTTOM) AFTER THRESHOLD

A threshold of 150, out of a maximum 255, applied to the top image in figure 70 keeps only the most strongly correlated peaks and gives the final segmented image. Raising and lowering the threshold is a tradeoff in the final image quality. Raising the threshold will reduce clutter but eliminates some of the real rivets as well. Lowering the threshold may bring in more rivets but also introduces more clutter. For this example, a threshold of 150 was approximately optimal. Three of the five rivets were segmented. The camera shadow introduced a significant amount of

clutter around its edges. This can be removed since the position of the shadow is known. However, tests showed that the camera shadow was present in only about 50% of the images, and it often overlapped rivets. Therefore, removing the shadow also removed rivets. Good, segmented rivets were too rare to risk losing more by removing the camera shadow. Even when the camera shadow was removed, too many other false alarms remained. Numerous changes were made to the rivet kernel by changing its size and thickness to attempt to reduce the clutter, but improvements were slight at best. As in the case of the neural network investigation, the quality of the initial image was too poor in terms of contrast and clutter for the enhancement and classifier algorithms to produce good, consistent results.

6.2.2.3 Third Trial: Improved Camera System.

As the last step in the image segmentation study, focus returned to the image acquisition setup and how that could be modified to generate a less cluttered image with better contrast between rivets and background. Tests were conducted in the lab using the robot on the Foster-Miller panel using both natural and artificial lighting. Various materials were tested for construction of a hood over the cameras that would improve contrast by reducing surface reflections. The best combination, based on visual inspection of the resulting images, was a black-matte hood with an artificial light source. This was then implemented as two permanent hood-with-light source fixtures on the robot. The hood is constructed of plastic and painted on the inside with a black matte finish. The light source is an array of infrared (IR) light emitting diodes (LEDs). The cameras used for navigation are IR sensitive and IR LEDs have a high current to light power efficiency.

To fully test this new arrangement, two data acquisition experiments were conducted. The first collected data in the lab using the Foster-Miller panel. The second took the robot to the USAir hangar at the Pittsburgh International Airport and acquired images off of a DC-9. The latter experiment gave an opportunity to test the hood and the segmentation algorithms for the first time on images that were not acquired from the Foster-Miller panel and therefore on a surface different from the one on which the segmentation algorithms had been designed.

The first set of tests used images acquired from the Foster-Miller panel in the lab. Sixty images were acquired from various locations on the panel, fifteen images each from the front and rear cameras (each equipped with a hood and light source) with the robot legs extended and retracted, respectively. First, the conventional segmentation algorithms were revisited. The initial image, directly from the cameras and reduced by a factor of 4 in size, was noticeably improved over the prior images taken with no hood in place. Subsequent testing with the conventional algorithms used in the previous work resulted in improvements with these new images. The resulting segmentation algorithm that provided the best results was as follows: reduce image size by factor 4, apply 3 x 3 median filter to reduce digitizing noise, apply Sobel edge detection to bring out edges, threshold to eliminate weak edges (usually due to noise, not rivets), remove small and large blobs (those not of size similar to rivets), apply skeletonization technique, apply matched filter, and threshold to keep only brightest 50% of pixels. Figure 71 shows a sample segmentation of one of the images taken from the front camera with the robot legs retracted.
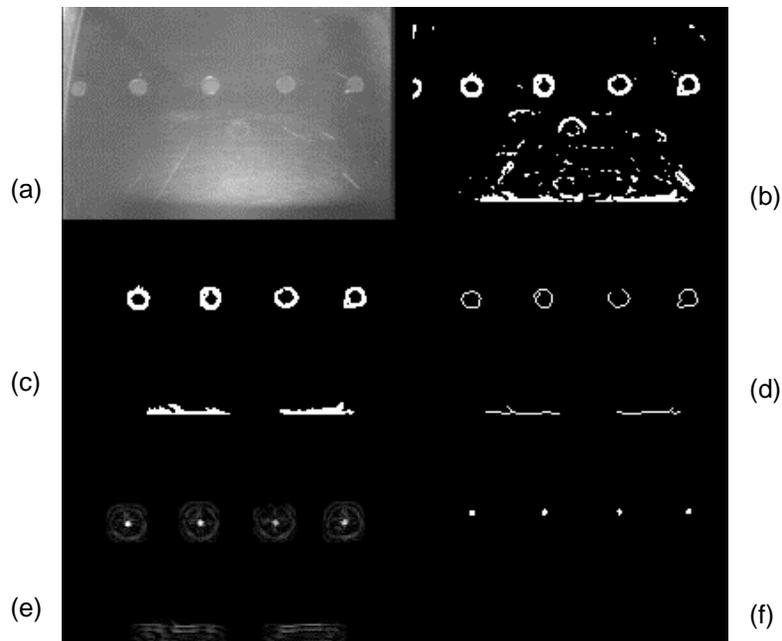
FIGURE 71. SEGMENTATION OF IMAGE TAKEN FROM FOSTER-MILLER PANEL
USING CAMERA HOOD AND LIGHT SOURCE

Figure 71(a) shows the gray scale image, reduced by factor 4 from the digitizer image. No additional processing had been done at this point. Figure 71(b) shows the image after median filtering, Sobel edge enhancement, and removing the weaker edges (i.e., edge values less than approximately 15% of the maximum edge values were removed). Note that at this point the rivet edges are clearly visible, and only a moderate level of noise is present. Much of the highly structured noise is due to either ambient illumination that has leaked beneath the edges of the hood or non-uniformity of the IR-LED light source. Both of these can be suppressed by its size, shape, or location in the image. Even with these problems, this image is already an improvement over the quality of the images before the camera hood was implemented. Next, blob filtering is applied by removing all blobs significantly smaller or larger than a rivet. Blob size is determined by the number of pixels in the blob. No shape information, such as roundness, is used. The result is shown in figure 71(c). Notice the significant reduction in clutter from figures 71(b) to 71(c). Next, a skeletonization algorithm is applied to thin out the remaining objects. This has a normalizing effect on the image, where just the shape of the objects are retained, not their thickness. The result is figure 71(d). The matched filter discussed earlier is applied to this image, resulting in figure 71(e). Finally, a threshold of 50% of the image's maximum pixel gray level is applied to the matched filter result, and the image is binarized, resulting in the image of figure 71(f). Note that all four of the whole rivets visible in figure 71(b) were successfully segmented in figure 71(f). This was not always the case, but over all the 60 images, the segmentation algorithm correctly segmented 70% of the rivets, with 34% clutter. Clutter was measured as the ratio of the number of nonrivet blobs to the number of rivets in the final image. It may not be necessary to remove all the clutter from the final image. With sufficient intelligence incorporated into the line-fitting algorithm, the nonrivet blobs may be recognized as such and eliminated from the final image used to navigate the robot. This is discussed more in section 6.2.3.

The processing time for all the steps in figure 71 averaged between 3 to 4 seconds on a Pentium 66 MHz PC. A neural network was implemented, resulting in similar accuracy, but its execution time was ten times that of the conventional algorithms. Thus, the conventional segmentation algorithms are preferred.

To further test the hood and segmentation algorithms, the robot was taken to a USAir maintenance hangar at the Pittsburgh International Airport, where more images were acquired on a DC-9. Although USAir did not paint their fleet, decals were used on the aluminum surface. Sixty images were acquired from various positions on the airplane body. Thirteen of the images were taken over the locations were the USAir decals had been removed. These proved much more difficult to accurately segment. The contrast of the images was not as good as the images taken from the Foster-Miller panel but still a significant improvement over the images taken earlier in the project without the camera hood. The first test of the segmentation algorithms was done on images reduced by a factor of two. Later, tests were done on images reduced by a factor of four. Overall, the reduced images were segmented to 74% accuracy. However, it is more enlightening to look at the results broken down into clear images and decal images. The clear images were acquired over positions on the airplane body that were clear of decals and crayon markings (put there by the inspector) while the decal images had remnants of removed decals in the field of view. Other body markings such as scratches or dents were sometimes evident in both types of images. The results are presented in table 2.

TABLE 2. RIVET DETECTION ALGORITHM PERFORMANCE

Image Type	Rivet Detection Rate	Clutter Rate
Clear	86%	50%
Decal (removed)	67%	357%

It is obvious from these results that even when removed, the decals present a particular problem to rivet segmentation. While a reasonable percentage of the rivets were segmented, they are overwhelmed by clutter. Even a sophisticated line-fitting algorithm would not be likely to distinguish rivets from clutter in these cases.

In the clear images, the segmentation accuracy was good (86%) but the clutter was higher than desired. However, almost half of the clutter resulted from bright edges in the image where the camera hood didn't completely cover the field of view. Modifications to the camera hood structure could be made to overcome this problem thereby reducing clutter to approximately 25%, a much more manageable range for the line-fitting algorithm.

The computation time of the segmentation algorithms for the reduced images was approximately 23 seconds on the Pentium 66 MHz PC. Images that were reduced by four were also tested because of the further improvement in execution time; however, the contrast reduction as a result of the size reduction was evident. Segmentation accuracy dropped to less than 50% and clutter was almost double the number of rivets. It is desirable to have less clutter than rivets in the final image for the line-fitting algorithm to be successful. No neural network processing was

implemented since previous tests showed comparable segmentation accuracy to the conventional algorithms, at the expense of greater computation time.

The addition of the camera hood and light source greatly improved the quality of the images from the robot, thus enabling robust, conventional segmentation algorithms to be developed. Many designs and tests of the algorithms resulted in a sequence of eight processing steps combining edge enhancement, thresholding, smoothing, and matched filtering to extract the rivets. Tests in the lab using a Foster-Miller panel and at the USAir hangar in Pittsburgh provided confirmation of the segmentation capabilities and highlighted that certain surface markings such as decals and crayon remain problematic.

6.2.3   Rivet Line-Fitting Algorithm.

Alignment of the robot with the seam under inspection requires identification of rows of rivets in each navigation camera's field of view. Image processing techniques are used to identify a set of suspected rivets. The resulting suspected rivets are represented by their pixel coordinates and relative mass (gray level intensity). These data are input to a line-fitting algorithm to identify likely lines of rivets.

Any field of view may contain multiple lines of rivets. This results in diagonal lines of rivets that are artifacts. The set of suspected rivets may also include objects that have been mistakenly identified as rivets. The line-fitting algorithm must be robust enough to recognize multiple rivet lines, ignore diagonal artifacts, and to ignore outlying objects. In addition to being robust in a statistical sense, heuristic knowledge can also be applied to improve the detection performance. Such knowledge could include the expected location of rivet lines given the previous location of the robot and its last move, the orientation of real rivet lines, and the spacing of rivets along real rivet lines.

The robot is initially aligned with the longeron rivet lines on the aircraft surface. Since the robot actively maintains this alignment, it is assumed that the rivet lines in the navigation cameras' fields of view will be within a small angle of horizontal or vertical. The line-fitting algorithm thus uses angle tolerance and distance tolerance parameters to identify lines of rivets.

Each pair-wise combination of two rivets defines a line segment. A table is used to track which segments have been assigned to a line. Each segment is investigated to see if it is part of a line of rivets. If the segment has not been assigned to a line, and the slope of the segment is within a specified angle tolerance of horizontal or vertical, then the equation of the line defined by the two points is found. The equation of the line is expressed in terms of horizontal pixels (x) and vertical pixels (y). To improve computational accuracy, the equations of horizontal lines are expressed in terms of x, and vertical lines are expressed in terms of y. The equations have the form:

Horizontal lines: $y(x) = x\dfrac{dy}{dx} + y(0)$

Vertical lines: $x(y) = y\dfrac{dx}{dy} + x(0)$

The distance to the line is found for each rivet in the set that has not been assigned to the line. The first rivet in the set that is within the distance threshold of the line is added to the line. A least squares fit is used to find a new equation for the rivet line. Since the equation of the line is modified, the process of finding rivets that fall within the distance threshold must be restarted. This continues until no more rivets are within the distance threshold to the line. If the resulting line contains more than two rivets, then the line is added to a list of lines. The segments linking the rivets in the line are marked as used in the segment table. If the segment's slope did not fall within the threshold, or the resulting line only contains two points, then these segments are also marked in the segment table. The process continues by investigating the remaining segments. When no more segments are left, a set of line equations is returned.

If matching horizontal segments are identified in each of the head and tail images, a corrective move can be calculated to bring the robot into correct alignment to perform the scan. Vertical lines of rivets could be combined with a map of the aircraft structure to keep track of the position of the robot along the longeron.

6.3   FEEDBACK SWITCHES.

To give the robot a sense of self awareness, switches were added to the second-phase robot to provide feedback when actions were completed.

When an action such as the extension of a leg occurs, an end-of-travel switch provides a signal to the control software to indicate that the action has been completed. Once the software has confirmation that an action has been completed, the next action in the sequence can be executed. If no signal is sent to the control software, the next action is not executed. This is in contrast to the first phase where timing loops were used to allow enough time for the robot to complete an action, but no feedback was provided. Even if an action was not completed, the next action in the sequence was executed, and the operator was required to notice the problem and stop the robot. With the limit switches, the robot's gait became much smoother.

Vacuum detection switches were also added to the robot during the second phase. They are used to detect if the suction cups are holding to a surface. The vacuum switches have an adjustable threshold; if, for some reason, there were a vacuum leak around a suction cup and if the threshold were not reached, the suction cups already adhering to the surface would not be released. These switches are used to minimize the possibility that the robot will lose adhesion to the surface. The status of all switches is shown symbolically on the debug interface screen described in section 5.3.

7.   CONCLUDING REMARKS.

Through laboratory testing at CMU and field testing at the AANC in Albuquerque, New Mexico, the technical feasibility of the robot inspection system has been demonstrated. Specifically, the robot was able to achieve the following goals:

- adhering to and walking over an aircraft fuselage regardless of the orientations of those surfaces,

- acquiring eddy-current inspection data that appeared identical to manually deployed eddy-current sensors, and

- being remotely operated using video and sensor feedback provided to the inspector.

Video camera-based automatic alignment and navigation were demonstrated in an auxiliary experiment using a separate robot moving on an aircraft skin surface. Through the addition of appropriate shrouding and illumination, these techniques were applied to the robot inspection system.

A significant development effort remains before this robotic inspection technology is suitable for operational deployment. Outstanding development issues include:

- reducing the weight of the robot so that it is more comfortable to lift and position on the fuselage,

- improving the mechanical reliability and speed of the system,

- minimizing the scratching of the skin surface by the suction cups and eddy-current sensor,

- reduction or elimination of the umbilical cable, and

- automation of many of the manually controlled operations.

To commercialize the technology, a new mechanical system would need to be designed and built incorporating the lessons of this work. Further integration of the software system by incorporating additional human-computer interaction features to simplify operation by aircraft inspectors would also be required. However, additional economic and operational benefit analyses may be required to justify further efforts to develop this technology.

8.   REFERENCES.

1.   Carnegie Mellon University, "Feasibility of Navigation (Guidance) of a Measurement Unit Over the Surface of an Aircraft," Federal Aviation Administration, March 1992.

2.   Alberts, C.J., W.M. Kaufman, and M.W. Siegel, "Development of an Automated Nondestructive Inspection (ANDI) System for Commercial Aircraft: Phase I," Federal Aviation Administration, DOT/FAA/CT-94/23, June 1994.

3.   Alberts, C.J., C.W. Carroll, W.M. Kaufman, and M.W. Siegel, "Automated Inspection of Aircraft," Federal Aviation Administration, DOT/FAA/AR-95/48, July 1996.

4.   S. Bobo, "Magneto-Optic Imaging Inspection of Selected Corrosion Specimens," Federal Aviation Administration DOT/FAA/CT-TN92/20, July 1992.

5.   S. Bobo, "Shearographic Inspection of a Boeing 737," Federal Aviation Administration DOT/FAA/CT-TN92/19, July 1992.

6.   Davis, I.L. and M.W. Siegel, "Vision Algorithms for Guiding the Automated Nondestructive Inspector of Aging Aircraft," *Proceedings of the 1993 International Symposium on Optics, Imaging, and Instrumentation*, SPIE The International Society for Optical Engineering, San Diego, CA, July 1993.