

Robotic Assistants for Aircraft Inspectors

Mel Siegel, Priyan Gunatilake, and Gregg Podnar
Intelligent Sensor, Measurement, and Control Laboratory
The Robotics Institute -- School of Computer Science -- Carnegie Mellon University
Pittsburgh, Pennsylvania 15213-3891

ABSTRACT

Aircraft flight pressurization/depressurization cycling causes the skin to inflate and deflate, stressing it around the rivets that fasten it to the airframe. The resulting strain, exacerbated by corrosion, drives the growth of initially microscopic cracks. To avoid catastrophe, aircraft are inspected periodically for cracks and corrosion. The inspection technology employed is ~90% naked-eye vision. We have developed and demonstrated robotic deployment of both remote enhanced 3D-stereoscopic video instrumentation for visual inspection and remote eddy current probes for instrumented inspection. This article describes the aircraft skin inspection application, how robotic deployment may alleviate human performance problems and workplace hazards during inspection, practical robotic deployment systems, their instrumentation packages, and our progress toward developing image enhancement and understanding techniques that could help aircraft inspectors to find cracks, corrosion, and other visually detectable damage.

KEYWORDS: automated robot visual NDI inspection enhanced remote stereoscopic multiresolution

1. INTRODUCTION

Pressurization and de-pressurization of an airplane during each takeoff and landing cycle causes its body to expand and contract like an inflating and deflating balloon. Cycling thus induces stress fatigue at the rivets that hold the aircraft surface skin to its frame, resulting in the growth of radial cracks. On April 28, 1988 Aloha Airlines flight 243, flying at 24,000 feet near Maui, abruptly lost a section of its upper fuselage skin from the floorline on the left side to the windowline on the right side, from just behind the passenger door to just in front of the wing, a distance of about 6 m [18 ft]. Although there were numerous injuries and one death (a cabin attendant was blown out during the initial decompression) the airplane landed intact except for the initial damage. The NTSB investigation¹ attributed the catastrophe to a previously unanticipated failure mode now called “multiple site damage” or MSD. In plain English, MSD means “two small cracks that joined to make one big crack”; since crack growth rate is essentially exponential in crack length, one big crack is much worse than two small cracks.

In the aftermath of “Aloha’88” Congress mandated an FAA Aging Aircraft Research Program, with a significant focus on inspection technology. By 1995 the specific aging aircraft infrastructure issues that led to Aloha’88 were declared “behind us”, and the nominal emphasis shifted from “aging aircraft” to “continued airworthiness assurance”. Nevertheless, in the November 3, 1997 issue of Aviation Week & Space Technology, we read that on October 27 the FAA “ordered stringent inspections and modifications to the oldest Boeing 737s flying” in response to “receiving reports from Boeing and airlines of widespread cracking in adjacent rivet holes in lap joints on the fuselage *lower* lobes of older aircraft”². We emphasize *lower* in this quote because until now interest in skin cracking has focused on the upper portions of the fuselage, where there is less strengthening substructure. It is thus apparent that, name change notwithstanding, a need still exists for research on aging aircraft inspection technology.

Early in the Aging Aircraft Research Program, CMU won an FAA grant to generate and analyze scenarios for robotic deployment of instrumentation to inspect for cracks of the sort that led to the Aloha’88 incident, to design a robotic system that implemented the most promising scenario, and to prototype an initial system to demonstrate this system’s feasibility. The ANDI (Automated NonDestructive Inspector) system, discussed in detail in Section 4.1 emerged from this effort. In addition, at least four more robots for aircraft skin inspection have been built in the US: one funded by the FAA at Wichita State University (ROSTAM), one funded primarily by the Ben Franklin Technology

Center of Western Pennsylvania at CMU (CIMP), one funded by the Air Force at NASA JPL (MACS), and one privately funded by AutoCrawler LLC (AutoCrawler). [See sidebar *other robots that crawl on airplanes* for photos and descriptions of ROSTAM, MACS, and AutoCrawler; the second CMU robot, CIMP, is described in detail in Section 4.2.] Outside the US, the Singapore Air Force is currently supporting a substantial effort for robotic underwing inspection of F-5 aircraft, and there are persistent rumors of one or more ongoing efforts, particularly in Japan, that are not being reported in the open literature.

In the aggregate, the five US prototypes have demonstrated all of the technical capabilities needed to implement a robotically assisted inspection system: *measurement, manipulation, mobility, and monitoring*. [See sidebar *the 4-Ms of robot assisted aircraft inspection* for a diagrammatic explanation of this hierarchy.] However no system has demonstrated all “4-Ms” at once. Probably only one robot, CIMP, has actually delivered to an aircraft inspector in the field anything that could with a straight face be called useful inspection data; however CIMP, the **Crown Inspection Mobile Platform**, succeeds in this goal (and the secondary goal of wireless operation) only by employing a design that restricts its mobility, as its name implies, to the aircraft crown, i.e., the upper fuselage, from cockpit to tail, from windowline to windowline.

2. CURRENT AND ALTERNATIVE AIRCRAFT INSPECTION METHODS

Visual inspection, augmented by instrumented inspection of known problem areas and instrumented confirmation of visually detected apparent problems, is required to ensure the structural integrity of an aircraft’s skin and its supporting substructure. A “major” inspection might be carried out on a commercial aircraft every 6 years, 24,000 flying hours, or 12,000 takeoff and landing cycles; in these 6 years, there might also be 2 cycles of alternating “light” and “heavy” checks. Additional checks, primarily visual, are made during “line maintenance”, which occurs each night; large aircraft that fly long routes, e.g., 747s, may also undergo a “turnaround check” each cycle. Details vary greatly from airline to airline, and within any one airline the fine structures of policies and schedules are likely to change with every upper management shuffle, but these typical figures accurately convey the spirit of how inspection and maintenance are scheduled.

Skin inspections are currently about 90% visual and 10% nondestructive inspection (NDI, meaning inspection using electronic instruments). Eddy current probes are the main NDI technology used for skin inspection. Although eddy current sensing is a near-surface sensitive technology, some depth resolution can be achieved by varying the excitation and detection frequency: the highest frequencies probe a thin layer just inside the skin surface, whereas lower frequencies and skillfully defined protocols can detect defects in second and third skin layers, and often in the ribs to which the skin is affixed.

Without exception, these inspections now require putting a human inspector on the body of the aircraft to closely examine its surface both for known structural problems and for more widespread defects, such as cracks, corrosion, damaged rivets, and lightning strikes. The practice raises safety issues for the inspector, is time consuming, and suffers at times from being ineffective due to inspector fatigue or boredom. Visual

and NDI inspection practices are detailed in two sidebars, *visual inspection: how it works*, and *instrumented (NDI) inspection: how it works*, and illustrated by Figure 1.



Figure 1: Aircraft inspector (with flashlight and 10x magnifier) and author P.G. inspecting the crown of a Boeing 737 airplane. Note safety harnesses.

3. WHY USE A ROBOT FOR AIRCRAFT INSPECTION?

Numerous hypothesized advantages of computer controlled mobile remote deployment platforms (for short, “robots”) for aircraft inspection instruments and remote cameras for visual inspection have been expounded at length elsewhere, for example References 3, 4, and 5, and other references therein. In contrast to most of these theoretical arguments, real experience has led us to believe that robotic deployment of inspection equipment has clear operational and economic advantages over manual deployment for three really crucial reasons: *thoroughness*, *correctness*, and *recordability*. Early theoretical arguments, particularly those relating to increased bodily safety of the inspectors and other advantages of “getting the man off the airplane,” now seem less plausible as the likely deployment scenario (primarily during heavy maintenance) and the personalities of the inspectors (they tell us they enjoy being on the airplanes) have been clarified by years of probing discussions and actual field experience.

THOROUGHNESS: The robot will cover the programmed inspection path or area completely, with a uniformly high level of concentration.

CORRECTNESS: The robot will deploy the correctly set up inspection instrument using exactly the programmed deployment protocol.

RECORDABILITY: The robot will faultlessly remember the location and result of every measurement*.

* It will presumably not elude the astute reader that the legal departments of commercial airlines may regard recordability and recording as not entirely desirable features; in the military aviation sector, however, they are usually favorably regarded.

With robotic deployment the correct data will always be available for interpretation (by computer software or by human experts), the location on the airplane where the data were obtained will always be known exactly (enabling advanced “C-scan” image-accumulation-and-display whatever the sensor), and precise trend analysis over arbitrary time periods will become possible (enabling better understanding of the development and evolution of problems, and allowing the operator and the regulatory authorities to choose statistically appropriate inspection intervals).

4. ANDI AND CIMP

4.1. ANDI: AUTOMATED NONDESTRUCTIVE INSPECTOR

The FAA Aging Aircraft Research program sponsored the design, construction, and testing of the Automated NonDestructive Inspector (ANDI) at CMU in a joint project of the Carnegie Mellon Research Institute (CMRI), the university’s applied research arm, and our lab, the Intelligent Sensors, Measurement, and Control Lab, in the Robotics Institute of the School of Computer Science. ANDI’s design was dictated by the FAA’s state of mind and by the state of NDI technology around 1990, when the project was defined and begun. The state of mind at the time, still dominated by Aloha’88, was that large scale eddy current “fishing expeditions” are a desirable way to head off future “Alohas”, and that large scale instrumented inspection could be made palatable to commercial airline operators if there were an economically acceptable automated device to deploy the sensors. We considered three broad design families for the robot’s mobility module. We call these the *car wash*, the *cherry picker*, and the *skin crawler*. These families are described in the sidebar *design scenarios for robots for aircraft inspection*, in which we also discuss the circumstances of commercial airlines operations and economics that led us to pursue the *skin crawler option* for ANDI.

The state of technology for eddy current sensors at the time was mainly manually manipulated pencil probes and complex impedance plane oscilloscope displays of their signals. These circumstances led to a design for ANDI that maneuvers most gracefully along fore-aft lines of rivets, maintaining precision alignment with them so that an eddy current pencil probe scanned parallel to the line of motion would follow the desired scanning path with little or no need for finely tuned closed loop path control.

The design developed for this scenario^{4,6} is drawn on the left side of Figure 2; to the right of the drawing is a photograph of the near-final ANDI on a DC-9 nose section at the Aging Aircraft Nondestructive Testing Center (AANC, Sandia National Laboratories, Albuquerque NM). This design, a form that is known in the robotics literature as a “beam walker,” achieves mobility by suitable motions of the bridges (arms) relative to the spine as the suction cup groups on the spine and the bridges are alternately affixed and released. The eddy current probe is scanned by one of the bridges moving along the spine while the spine’s suction cups are affixed to the aircraft skin. ANDI is equipped with four cameras for navigation and alignment: one each fore and aft to align the spine with the rivet line, one adjacent to the eddy current probe to verify location and alignment, and one high mounted with a wide angle field of view for navigation, obstacle avoidance, and proprioception (“self awareness”). In contrast to CIMP, the second CMU aircraft inspection robot (Section 4.2), whose capability is focused on enhanced remote visual

inspection, ANDI's cameras were not intended to have sufficient resolution to be useful for visual inspection per se.

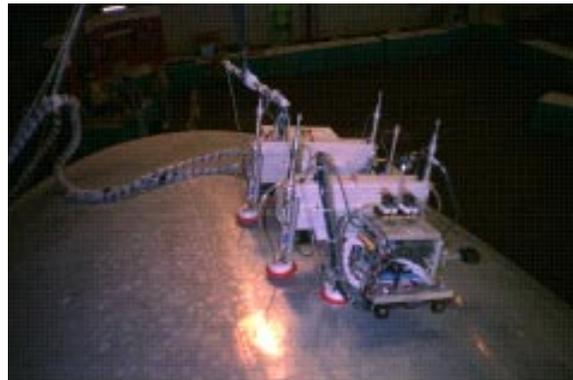
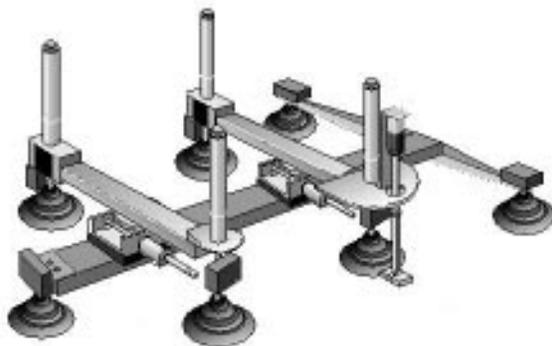


Figure 2: (left) ANDI as a design drawing, and (right) photographed on the nose section of a DC-9. In the drawing, the eddy current sensor is seen on the near end of the far bridge (arm). In the photo the small black box (on an outrigger on the far side) contains one of the two alignment cameras (see text for discussion of camera arrangement).

Despite essential successes in mobility (getting where it needed to be), automatic alignment (using a machine vision rivet line finding algorithm), manipulation (moving the eddy current probe smoothly along the desired path), and measurement (collecting and delivering eddy current sensor data to the ground), as well as the articulation of a comprehensive system architecture for integrating robotics and automation into aircraft maintenance and inspection, unforeseeable changes in the context for ANDI led to its early marginalization.

First among these factors was a de facto return to the pre-Aloha'88 model, which says that visual inspection should be the lion's share of skin inspection, with eddy current and other NDI technologies being used for backup, confirmation, and a relatively small number of directed inspections for specific flaws at specific problematic locations. A robot designed for large area eddy current inspection along rivet lines would have a hard time being economically competitive in an environment that views eddy current as a confirming technology for suspected visual flaws and as a survey technology only for a few specific fuselage locations, e.g., locations known from structural models or past experience to present specific cracking or corrosion patterns.

Simply stated, at least in the civilian sector, there is no economic interest in a robot that does the 10% of inspections that are instrumented; to make an impact with the commercial airline operators, a robotic inspection system will have to do the visual inspections that account for 90% of the inspection effort.

Another development that weighs substantially against the viability of ANDI is the recent advance in sensors and display systems for C-scan rendering of eddy current data. We now have linear and area arrays (or their equivalent in, e.g., magneto-optic imaging (MOI), mentioned in sidebar *instrumented (NDI) inspection: how it works*), and inspectors now expect to see false-color images rather than oscilloscope traces, making ANDI's mechanical optimization for point probes somewhat pointless.

4.2. CIMP: CROWN INSPECTION MOBILE PLATFORM

CIMP, built in our laboratory with support from the Ben Franklin Technology Center of Western Pennsylvania and our spin-off company Aircraft Diagnostics Corporation, is an aircraft inspection robot that is explicitly *not* a "wall crawler." Chastised by the two lessons of ANDI:

if you spend all your time working on the robot's mobility you'll never get any

inspection data

and

if you can't do visual inspection nobody will be interested in your robot,

we set out to demonstrate that a robot could generate data, first and foremost video data whose quality inspectors would gladly accept for routine visual inspection, and to deliver the data to an “inspector’s workstation” off the airplane. To allow us to concentrate on inspection data and not inspection equipment transportation, we designed⁷ an interim robot whose mobility is limited to the fuselage crown: CIMP, the Crown Inspection Mobile Platform, shown in Figure 3.



Figure 3: (left) CIMP on a 747 in a heavy maintenance bay at Northwest Airlines’ Minneapolis headquarters. The inspector, observed by a CMU staffer, is performing an eddy current check of a visual anomaly detected using the remote vision system shown in sidebar *remote 3d-stereoscopic visual inspection*. Future CIMP models would incorporate remotely operated eddy current sensing. (right) CIMP showing mobility (differentially driven wheels), sensor pod mounted off circumference-scanning carriage, and wide angle cameras for proprioception and navigation (upper right). The vertical stalk and the sensor pod rotate to change the camera viewing azimuth.

Because CIMP works with gravity instead of against it, it does not need a tether. It was designed for the curvature of a DC-9, and for window-line to window-line mobility on that aircraft type; however it turned out to be more convenient to test the prototype on a 747, on which it ran with no difficulty despite having the “wrong” curvature.

Because its power requirements are tiny compared to a robot that has to adhere to the fuselage in arbitrary orientations, CIMP does not need an umbilicus. It runs for several hours on its internal batteries; exactly how long depends on the variable demands of mobility, manipulation, illumination, etc. Control signals are transmitted to CIMP wirelessly using off-the-shelf model airplane transmitter technology. Video data are returned wirelessly using micropower radiofrequency channels; in the prototype these are off-the-shelf 2.4 GHz cable eliminators sold in the consumer market to wirelessly connect a home VCR and TV. In a commercial version of CIMP somewhat more sophisticated (and costly) channel options would be appropriate to avoid signal degradation due to multipath effects.

CIMP has been successfully operated on a 747 at Northwest (as shown in the accompanying figures) and on a DC-9 at US Airways. Working aircraft inspectors have been uniformly enthusiastic about the quality and utility of the imagery that the CIMP remote 3D-stereoscopic video system delivers. However many are skeptical about the economic benefits that might reasonably be expected from robotic deployment of inspection equipment. Some also question whether the introduction of robotic deployment equipment would enhance their job satisfaction; despite our best intentions to make the inspector's job easier, safer, etc., by "getting the man off the airplane", several inspectors have told us that they like their jobs in large part because they like being on the airplane.

4.2.1 CIMP'S "INSPECTOR'S CONSOLE"

CIMP's user interface, the "inspection console", provides for remote 3D-stereoscopic visual inspection of the aircraft surface, and for remote control and navigation of CIMP on the aircraft crown. The current prototype inspection console consists of two primary displays and their supporting equipment and a radio transmitter (of the type used to control model vehicles) that controls forward and backward motion, left right steering, camera position and orientation, and lighting selection and orientation. The first display is a monitor that delivers live video at NTSC spatial and temporal resolution to eye, providing stereoscopic imagery of either the inspection or the navigational camera pair. The second display is a Silicon Graphics Indy workstation with a GUI (graphical user interface) that we call the Intelligent Inspection Window (IIW). The IIW fulfills a variety of requirements: it displays the live monoscopic or still stereoscopic imagery; it is the operational interface and output display unit to the image enhancement and understanding algorithms (activated by the menus and buttons of the IIW); and, in the future, it will contain facilities for creating, editing, storing and retrieving multimedia records of surface defects. 3D-stereoscopic image pairs showing different aircraft skin features, the sensor pod containing the 3D-stereoscopic cameras, and an inspector at the console are shown, along with additional technical details, in the sidebar *remote 3d-stereoscopic visual inspection*.

5. SKIN DEFECT DETECTION AND CLASSIFICATION

5.1. BACKGROUND

To our surprise and delight, aircraft inspectors have been spontaneous and enthusiastic advocates for using computer image enhancement and automated image understanding for flaw detection; they are, however, skeptical about the likelihood that we will succeed at the latter.

The goal of an image understanding algorithm for aircraft inspection is to recognize and classify certain surface flaws that might appear in the live imagery. The recognition capability of an algorithm is achieved by correlating features of the imagery with prior or learned knowledge of the surface flaw types. However, developing a successful image understanding algorithm remains a non-trivial challenge, due primarily to the difficulty of generalizing and encoding in an algorithm the notions that humans use to discriminate normal from defective, the limited resolution and dynamic range of practical imaging systems, and the confounding effects of environment factors such as illumination.

Given these limitations, an attractive scenario for application of image understanding algorithms in remote visual inspection is screening large volumes of image data. The image understanding algorithm can conservatively label all plausible defects, as a coarse prefilter for the inspector who does the fine screening. Another scenario is the interactive use of these algorithms by an inspector to obtain a second

opinion about a particular suspicious flaw. This possibility is most attractive when the real-time inspector is relatively inexperienced, in general or with respect to a specific problem, compared to the inspectors whose expertise has been incorporated (explicitly or implicitly) in the algorithm; in this scenario the computer fulfills a training role in addition to its direct inspection role.

In the spirit of these anticipated operating scenarios, we have developed prototype algorithms⁷ that detect surface cracks, surface corrosion, and subsurface corrosion evidenced by surface pillowing.

5.2. CRACK DETECTION ALGORITHM

The crack detection algorithm that we have developed is modeled closely on the inspectors' practice of using grazing angle directional lighting to enhance crack visibility. We emulate the directional lighting produced by an inspector's flashlight by employing a remotely controlled rotatable directional light source on CIMP. The remote inspector can rotate the light source around a rivet location and examine the resulting live monoscopic or stereoscopic imagery of the rivet and its neighborhood for cracks. In addition, the inspector can run the crack detection algorithm on these images for detection or verification of cracks in the live imagery. The stereoscopic imagery can also be recorded (at slightly reduced resolution) on a standard S-VHS recorder for future examination or computer processing.

The first module of our crack detection pipeline finds rivets in the image presented to it. Since cracks appear in the neighborhood of rivets, finding rivets in the first stage permits subsequent states to focus on the areas ("regions of interest", ROIs) that are most likely to contain cracks. Rivets are identified by detecting the circular arc edges made by their heads; a region of interest is then defined as a square centered on a rivet. Figure 4 shows a section of an aircraft surface with three simulated cracks appearing as dark lines emanating horizontally from the two rivets, and the two ROIs found by the algorithm.

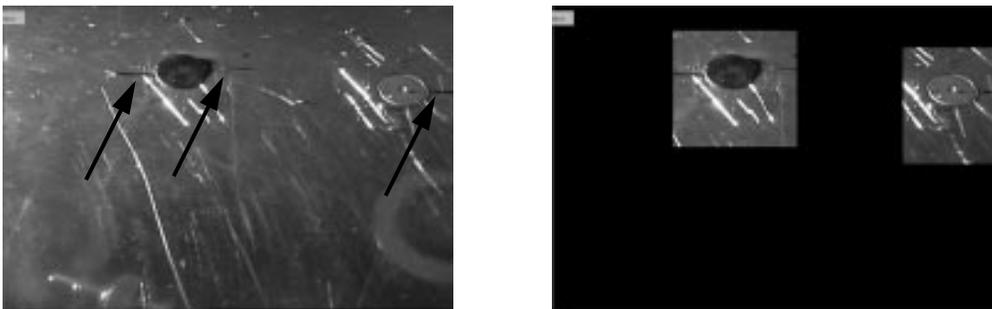


Figure 4: Rivet "region of interest" (ROI) finding algorithm. (left) Aircraft skin showing two rivets. Arrows point to three short black horizontal lines (hairs from a fine brush) that simulate narrow cracks. (right) ROIs found by the algorithm. Notice reasonable behavior even for a rivet near the edge of the frame.

Even the relatively small areas defined as ROIs still typically contain many edge-like features, most of which are due to rivet edges, scratches, dirt marks, lap joints, metal repair plates, and a few of which are due to real cracks. Thus for the second stage in the crack detection pipeline we need an algorithm that finds and attaches a quantitative description to each edge, and for the third stage an algorithm that discriminates the small fraction of edges that are cracks from the many edges that are of no interest.

A crack is typically very small compared to most visible features of the aircraft surface; this motivates us to develop, for the second stage, a multiscale (also called multiresolution) algorithm to detect edges based on quantitative spatial features of edges. For this stage we use wavelet filters (discussed in more detail in Section 5.3 in the context of surface corrosion detection) to project an ROI to different resolutions, and estimations of intensity variation at each resolution for multiscale edge detection. Edges of the same object or feature are usually present in more than one scale. We use a coarse-to-fine edge linking algorithm to attach a *propagation depth* to each edge appearing at the coarsest scale. The propagation

depth, and several additional features of the edge, are taken to comprise a feature vector that can be used to classify the edge to be a crack or to be “crack-like but non-crack”.

The classification is done in the third stage by a simple neural network trained on a set of feature vectors that had been identified by a person as corresponding to either a crack or a non-crack. The result of applying this pipeline to the area around a badly damaged rivet hole is illustrated in Figure 5.

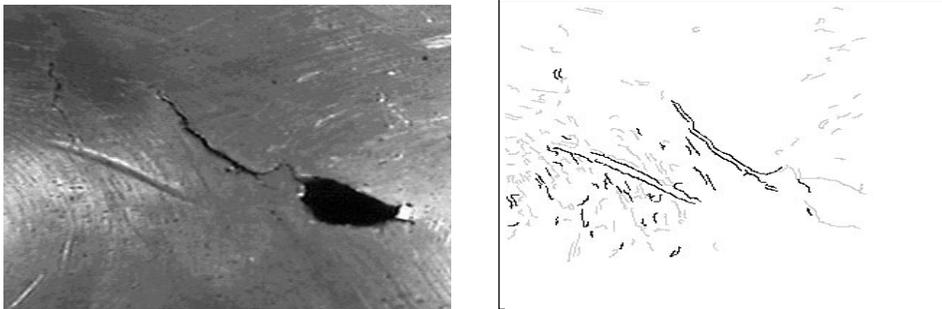


Figure 5: (left) Photograph of a badly damaged rivet hole with a large crack and a deep, crack-like dent. (right) Output of the crack detection algorithm, with crack-like edge features shown in dark gray and non-crack-line edge features shown in light gray. [This is a very severe example; really interesting cracks are so small, in length or opening or both, that they would be invisible when printed at this resolution.]

5.3. SURFACE CORROSION DETECTION ALGORITHM

A comprehensive corrosion detection algorithm needs to detect both surface and subsurface corrosion. Surface corrosion is detected by color and texture visually suggestive of corrosion, whereas subsurface corrosion is detected by distortion (“pillowing”) of the surface. Thus a comprehensive algorithm requires an image and a shape profile of the inspection surface to detect both types of corrosion. In this section, we describe the algorithm we have developed to detect surface corrosion; in the next section we describe our subsurface corrosion detection apparatus and algorithm.

We detect surface corrosion by a binary segmentation of the image into regions of texture suggestive of corrosion and regions of texture suggestive of freedom from corrosion. We implement the segmentation via a surface description that uses multiresolution-multiorientation wavelet functions to compactly represent the scale and orientation of texture in small surface patches. Since wavelets have good spatial and frequency localization, the wavelet coefficients provide the compact representation of texture that we seek.

In our most recent and refined implementation of this approach, the raw 24-bit RGB images are, in the first module, transformed into corresponding YIQ representation. Then we perform a three-level discrete wavelet transform (DWT) without sub-sampling of the Y-channel (luminance), and a two-level DWT with sub-sampling separately of the I- and Q-channels (chrominance components). The sub-sampling of the DWT of the Y-channel forces the Y-channel decomposition to be translationally invariant, which is a desirable feature for texture recognition. The DWTs decomposes the Y-channel into ten sub-bands and the I- and Q-channels each into seven sub-bands. Each sub-band corresponds to a particular texture scale

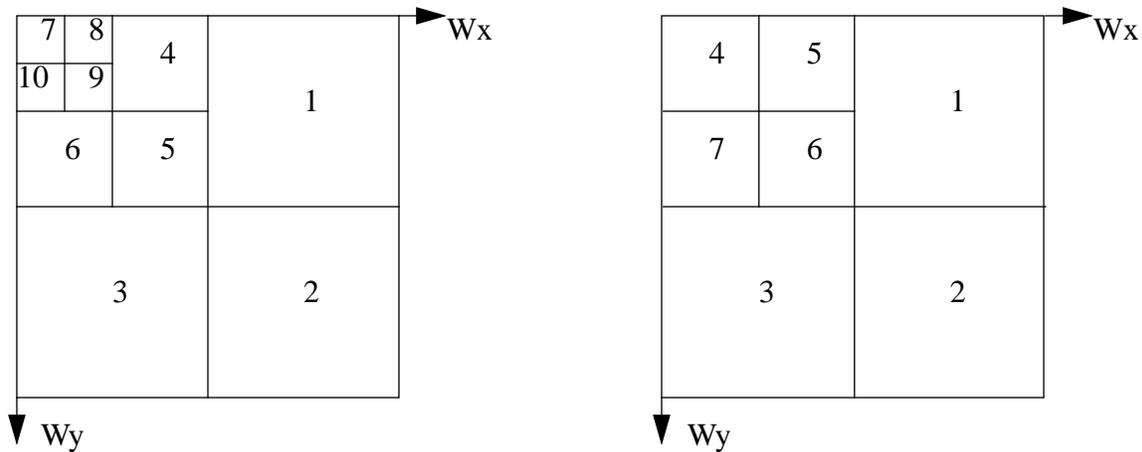


Figure 6: Sub-band decomposition of images based on texture scale and orientation. Smaller squares correspond to smaller spatial frequencies; the square adjacent to the origin is the DC component. Different squares of the same size correspond to different orientations at the same scale. (left) Y-channel contains the luminance information; 3-level DWT gives 10 sub-bands. (right) I- and Q-channels contain the chrominance information; 2-level DWT gives 7 sub-bands

and orientation, as shown shown in Figure 6. The actual wavelets we use are derived from the Battle-Lemaire 32-tap asymmetric FIR filter⁹.

In the second module, feature extraction, we calculate for each 32-by-32 non-overlapping block in the YIQ image, a vector of ten features. These are a subset of seven features extracted from the Y-channel decomposition and three extracted from the I- and Q-channel decompositions. Included in the Y-channel feature set are three rotationally invariant energy measures and three rotationally invariant orientation measures, plus the DC sub-band energy, sub-band seven in Figure 6. [“Energy” is defined as the weighted sum squared of the wavelet coefficients belonging to that block.] The three features from the I- and Q-channels are the I-to-Q ratios of the total energy in sub-bands 1, 2, and 3, in sub-band groups 5, 6, and 7, and in DC channel 4. Summing over sub-bands of the same scale averages over orientations, i.e., it removes sensitivity to viewing azimuth.

In the last module, we use a three-layer feed-forward neural network (10 inputs, 30 hidden layer neurons, 2 outputs) to classify the preceding feature vector as either corresponding to a region of corrosion and or corresponding to a region non-corrosion. The neural network was trained using a training set hand-classified vectors.

Figure 7 displays, on the left, an image that includes both corroded and uncorroded aircraft skin removed from an airplane as part of a maintenance program, and on the right, the output of our surface corrosion

detection. Bright and dark areas respectively indicate our algorithm's identification of the corrosion and corrosion free areas in this image.



Figure 7: (left) Corrosion surrounding rivet holes in a disassembled lap joint. (right) output of the corrosion detection algorithm.

5.4. SUBSURFACE CORROSION DETECTION SYSTEM

Subsurface corrosion may be visible externally because of the surface “pillowing” it induces. Pillowing is a change in skin surface shape rather than surface texture; it is detectable as an increase in skin surface height toward the center of each rivet row-and-column-bounded rectangle over a region suffering from subsurface corrosion. Stereoscopic cameras are well suited to creating surface altitude maps, but the low density of high contrast features on aircraft sheet metal expanses makes the critical step of identifying corresponding points in left and right images very difficult. We circumvent this difficulty by illuminating the surface with a laser that projects a square grid of 17x17 spots; by concentrating on these spots rather than on natural textural features of the surface, the correspondence problem is easily solved.

Figure 8 shows the result of applying this method to a slightly wrinkled piece of aircraft belly skin that is sloping downward from back to front (below).

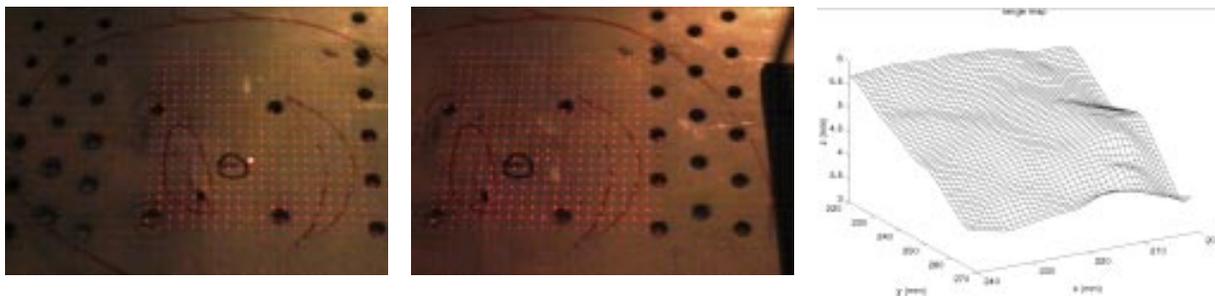


Figure 8: Method of detecting subsurface corrosion by elevation mapping of the visible surface. Three frames are left and right perspectives of the projected laser grid illuminating a sloping aircraft sheet metal surface, and the corresponding depth map. Elevation resolution is about 0.5 mm.

6. SUMMARY DISCUSSION AND FUTURE WORK

After the Aloha'88 “incident” several research groups embarked confidently on paths toward mobile robotic platforms and computer-based automatic control systems for deployment of NDI and visual inspection equipment on aircraft surfaces. The ambitious plans included agile vehicles, sophisticated deployment of sophisticated sensors, a high level of intelligent, sensible autonomy in task and path planning, navigation, and inspection, elegant and functional human-computer interfaces, hierarchical data and information displays matched to the needs of inspectors, supervisors, and management, totally automated networked integration of inspection with maintenance, engineering, operations, and management databases, and the emergence of a safer and more economical inspection and maintenance system based on massive analysis of massive quantities of data that would permit just-in-time, but never

earlier than necessary, predictive response to developing repair requirements. Yet there is still no working system that comes anywhere near these early expectations; there is not yet even a demonstration of a robot that is both agile (able to go 'anywhere' on the aircraft's skin) and functional (able to deliver data that inspectors want). Furthermore there is a sense of diminished confidence that there will be any such system any time soon, in large part because of how difficult it turns out to be to build a suction-cup based vehicle without a cumbersome umbilicus.

On the positive side, in fact all four key modules needed by a useful robotic inspection system -- **m**easurement, **m**anipulation, **m**obility, and **m**onitoring -- have been separately demonstrated. It has been hard to tie the modules together in a fully functional system in large part because the generic hard problems that must be faced in designing and building an aircraft skin mobility module -- adhering to arbitrarily curved and oriented surface regions, moving gracefully over lapjoints and buttonhead rivets, managing a safety tether and an energy-lifeline umbilicus -- have disproportionately diverted attention from the other three key modules.

However the diversion has actually paid off: we now have multiple examples of mobile platforms matched to various operational scenarios: ANDI for precision deployment of traditional point probe sensor types, AutoCrawler for manhandling large area survey instruments, and MACS for the anticipated next generation of light weight sensors, among others.

In our lab we set out with CIMP to demonstrate a complete system via the expedient of temporarily sidestepping the general mobility problem: we built only a simple (though wireless!) platform whose mobility is restricted to the fuselage crown. Thus we were able to concentrate the extremely limited resources that were available to us for this project on demonstrating the single most important capability of a robotic inspection system: its ability to deliver useful inspection data to the ground.

We successfully demonstrated CIMP's remote control and imaging capability to Northwest Airlines at their Minneapolis 747 maintenance and inspection facility and to US Airways at their Pittsburgh maintenance and inspection facility. Our demonstration showed that state-of-the-art 3D stereoscopic video technology implemented by us and operated by inspectors not specifically trained in its use, delivers imagery of sufficiently high visual quality that aircraft inspectors and NDI supervisors were willing to accept it (and sometimes prefer it) as an alternative to direct visual inspection.

We succeeded in delivering inspection-quality visual data to an inspector who was remotely driving the robot from a rudimentary but acceptable workstation. The mobility and manipulation components were comprehensive enough that the inspector could scan along a useful path, stop at a possible flaw, and inspect more closely by varying the camera's viewing angle and the character of the illumination (flood or spot), and the direction of spot illumination. In the lab, we have also made substantial progress toward useful image enhancement and image understanding algorithms for visually detectable cracks, surface corrosion, and subsurface corrosion.

These successes are, we believe, clear demonstrations that we are ready to respond to a well defined real-world application demand with a technically and economically justified system.

7. ACKNOWLEDGMENTS

The cooperation of the commercial airline operators, especially US Airways and Northwest Airlines, has been essential to this work. The ongoing support of Russell Jones and Roy Weatherbee of US Airways and Jeff Register of Northwest has been invaluable. The inspectors in both their organizations have always been extremely hospitable and open, and remarkably tolerant of our incessant silly questions and wild ideas. The ANDI project was supported by the FAA Aging Aircraft Research program. The CIMP project was supported by the Ben Franklin Technology Center of Western Pennsylvania and Aircraft Diagnostics Corporation. The staff of the ANDI project at CMRI included Bill Kaufman, Chris Alberts, Chris Carroll, the late Court Wolfe, and many others; Alan Guisewite and graduate student Ian Davis helped on the Robotics Institute side. The staff of the CIMP project at the Robotics Institute included Alan Guisewite, graduate student Huadong Wu, and undergraduate student Earl Crane. The basic

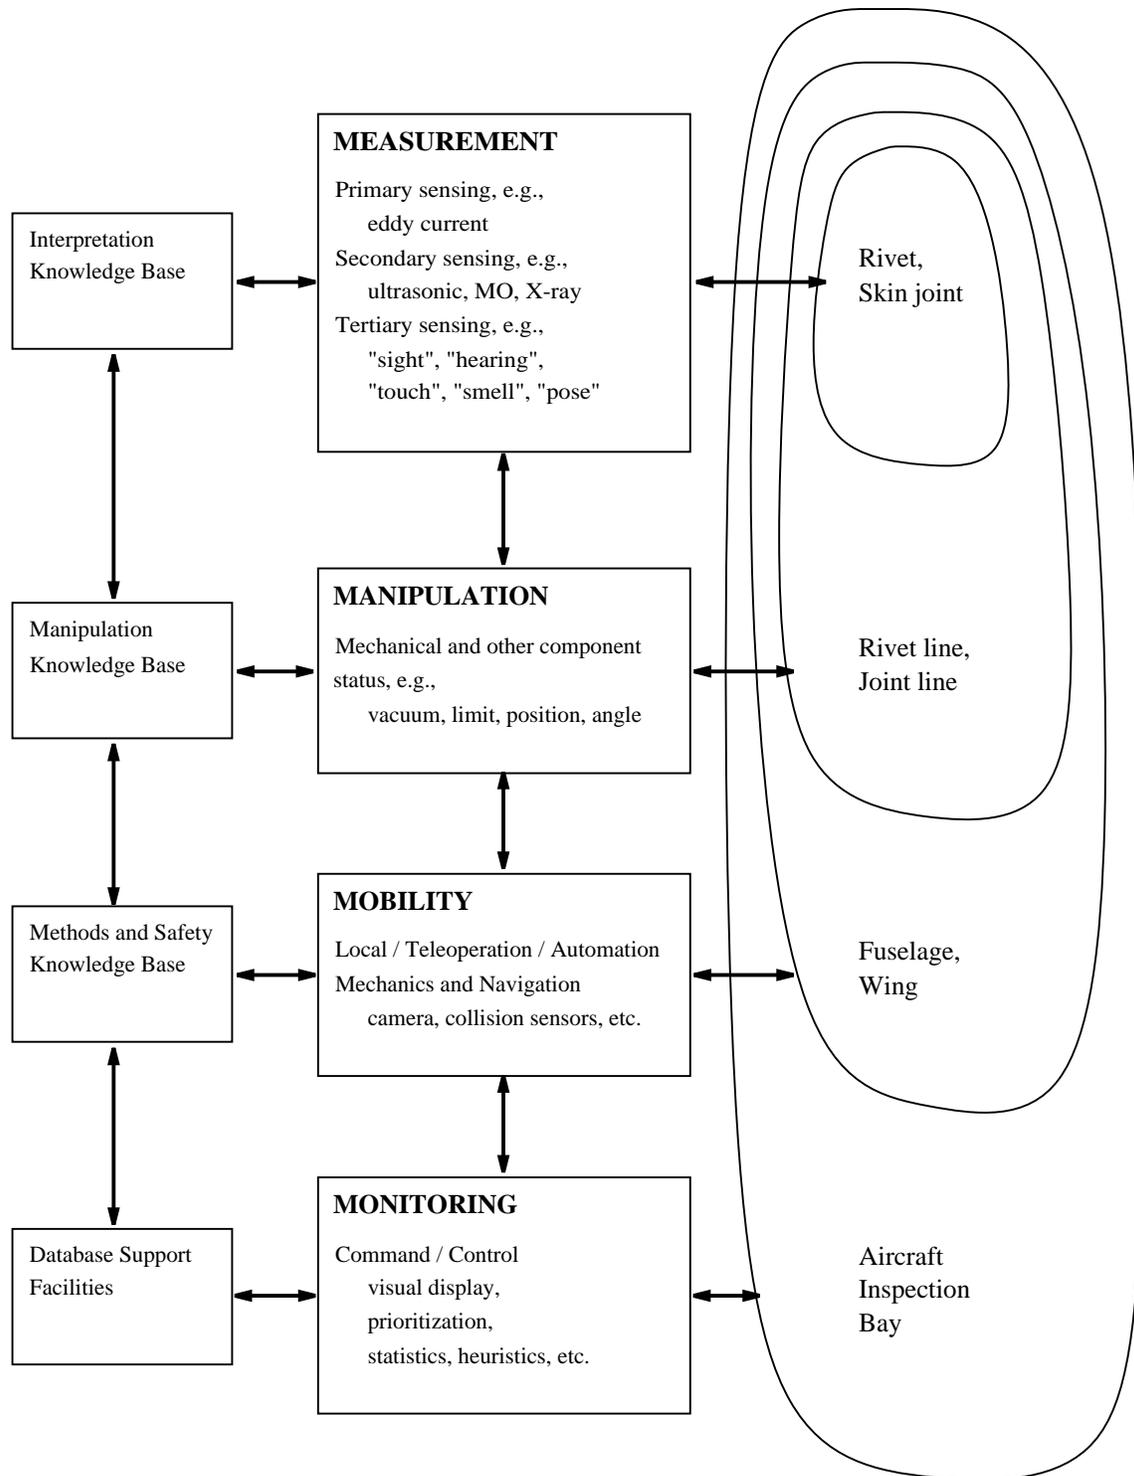
research on which the stereoscopic video system implementation is based was funded by DARPA's High Definition Systems program.

8. BIBLIOGRAPHY

1. Excerpts and photos from the NTSB report can be viewed at <http://www.aloha.net/~icarus/>
2. Aviation Week & Space Technology, November 3, 1997, "FAA Orders Inspections", p. 41.
3. C. Seher, "The national aging aircraft nondestructive inspection research and development plan", Proceedings of the International Workshop on Inspection and Evaluation of Aging Aircraft, May 21, 1992.
4. M. W. Siegel, W. M. Kaufman and C. J. Alberts, "Mobile Robots for Difficult Measurements in Difficult Environments: Application to Aging Aircraft Inspection", Robotics and Autonomous Systems, Vol. 11, pp 187 - 194, July 1993.
5. M. W. Siegel, "Automation for Nondestructive Inspection of Aircraft", Conference on Intelligent Robots in Field, Factory, Service and Space (CIRFFSS'94), Paul J Weitz (NASA/JSC), ed., AIAA, AIAA/NASA, pp. 367 - 377, Houston TX, March 1994.
6. Wolfe, C., M. W. Siegel, C. J. Alberts, "Robot with Cruciform Geometry", Patent US5429009, July 4, 1995.
7. Siegel, M. W., P. Gunatilake, "Remote Inspection Technologies of Aircraft Skin Inspection", Proceedings of the 1997 IEEE Workshop on Emergent Technologies and Virtual Systems for Instrumentation and Measurement, Niagara Falls, CANADA, pp 79-78.
8. Gunatilake, P., M. W. Siegel, A. G. Jordan, G. W. Podnar, "Image understanding algorithms for remote visual inspection of aircraft surfaces", Proceedings of the SPIE Conference on Machine Vision Applications in Industrial Inspection V, San Jose, February 1997, SPIE v. 3029, pp 2-13.
9. Wavelets and Filter Banks, G. Strang and T. Nguyen, Wellesley-Cambridge Press, 1996.

--- SIDEBAR ---

THE 4-MS OF ROBOT ASSISTED AIRCRAFT INSPECTION



--- SIDEBAR ---

DESIGN SCENARIOS FOR ROBOTS FOR AIRCRAFT INSPECTION

The *car wash* scenario imagines a central facility dedicated to inspection: aircraft are flown in specifically for inspection “with a fine tooth comb.” It can thus be carried out without interference from operations, maintenance, or anything else. The technically most excellent job can probably be accomplished by a *gantry robot* arrangement, like a huge automatic car wash from which precise deployment of a variety of inspection devices can be carried out unhurriedly and thoroughly. The conflict that this scenario presents for economical operation in the civilian sector (and perhaps for mission readiness in the military sector) probably makes it impractical despite its technical superiority over all alternatives.

The *cherry picker*, in contrast to the car wash, imagines bringing the inspection apparatus to the airplane rather than the reverse. A vehicle-mounted cherry picker, of the sort used for a variety of operations in typical maintenance and inspection hangars, is used to deploy inspection devices in much the same manner as in the car wash scenario: in both, mobility and manipulation use separate mechanisms and operate at substantially different scales. In the big picture, the cherry picker is less disruptive of normal operations than is the car wash. However discussions with airline maintenance and inspection managers uncover substantial objections, primarily on two grounds: first the fear that an automatically or teleoperated cherry picker will collide with and damage the aircraft under inspection, and second the complaint that the floor space around an airplane undergoing heavy maintenance and inspection is too busy and too cluttered to tolerate the routine intrusion of a cherry picker.

Given the substantial operational and economic objections to the car wash and the cherry picker, we are left with only the *skin crawler*: a small self-mobile device that adheres to the aircraft skin and maneuvers under some mix of teleoperation and autonomous control to carry out a sequence of inspections at a sequence of locations. An inspector affixes the crawler to the airplane at any convenient ground-accessible location, it crawls wherever it needs to go and does whatever it needs to do, then it returns to the original or another ground accessible location where the inspector removes it. The only problem is that building a crawler that will be practical in the aircraft inspection environment is easier said than done!

It is not easy to make a skin crawler because a crawler needs to adhere to the airplane, and the only practical way to make it both adherent and mobile is to use suction cups. Although passive suction cups are a possibility, operational and safety considerations demand active suction cups, i.e., suction cups that depend on a vacuum supply. Elementary analysis shows that the power required to obtain the necessary vacuum pumping speeds for a reasonable operating time exceeds what is available from any practical on-board energy storage system. So the only alternative is an umbilicus carrying a vacuum hose or, better, an air hose that can generate vacuum on board via venturi-effect “ejectors”. The problem is that the umbilicus gets in the way of the easy mobility that is a skin crawler’s biggest advantage. Even worse, managing the umbilicus becomes a frustrating, expensive, often simply intractable problem: the umbilicus literally becomes the tail that wags the dog.

--- SIDEBAR ---

OTHER ROBOTS THAT CRAWL ON AIRPLANES

ROSTAM I THROUGH IV: WICHITA STATE UNIVERSITY

Benham Bahr and his students at Wichita State University, Wichita KS may have been the first to describe a family of robots specifically conceived to carry NDI sensors and video cameras for aircraft skin inspection. The ROSTAM I - IV series is notable for a design that uses one very large diameter suction cup on its “belly” and a smaller suction cup on each “leg” (or “arm”). ROSTAM III is shown on a section of aircraft material (apparently wing) in the figure. Many theoretical aspects of the ROSTAM series design (suction cups, safety issues, sensory guidance) and their hypothetical inspection capability (automated crack monitoring using a vision system) have been reported in technical conferences and archival journals; however it does not appear that inspection data were ever delivered in field tests.

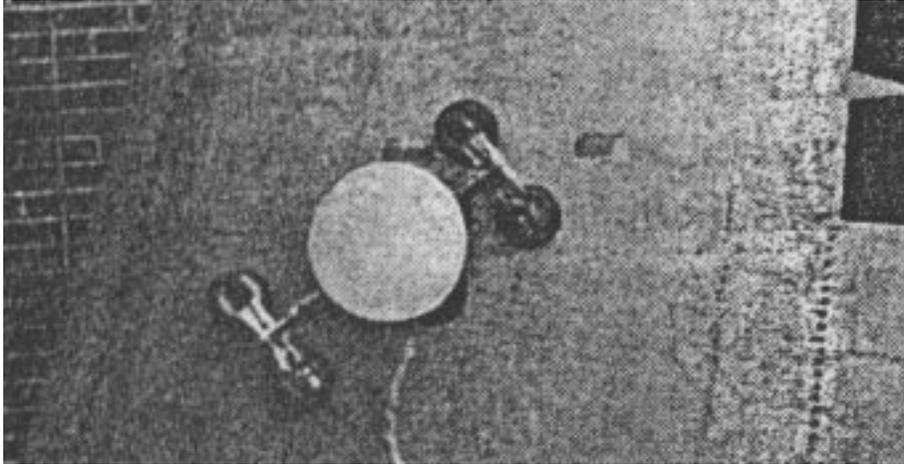


Figure 1: ROSTAM III on a section of aircraft material.

1. Bahr, B., “Automated Inspection for Aging Aircraft”, International Workshop on Inspection and Evaluation of Aging Aircraft, May 18-21, 1992, Albuquerque, NM, sponsored by the Federal Aviation Administration, hosted by Sandia National Laboratories.

AUTOCRAWLER: AUTOCRAWLER LLC

Henry R. Seemann's Seattle WA based company AutoCrawler LLC, with support from Boeing, has developed a tank-like multi-suction-cup-tracked vehicle, AutoCrawler, with a clever valving arrangement that applies vacuum only to those suction cups that are actually in contact with the surface. AutoCrawler is a behemoth of a mobile platform, capable of carrying enormous loads at very high speeds thanks to its powerful air motors and high capacity vacuum ejectors. On the other hand, it demands a large air compressor, and it is correspondingly noisy. Although suction cups of optimized material and shape have been custom designed and manufactured for the application, aluminum surface scuffing is still very evident in the AutoCrawler's wake. Boeing's experimental area array eddy current sensor and the PRI Magneto-Optic Imager have been carried by AutoCrawler on a 737 fuselage section with known defects, and images from the Boeing sensor have been exhibited in AutoCrawler's sales literature. AutoCrawler's mechanics are well suited to area-type sensors, as they do not require precise placement and scanning. The AutoCrawler has not been reported in technical conferences or archival journals.

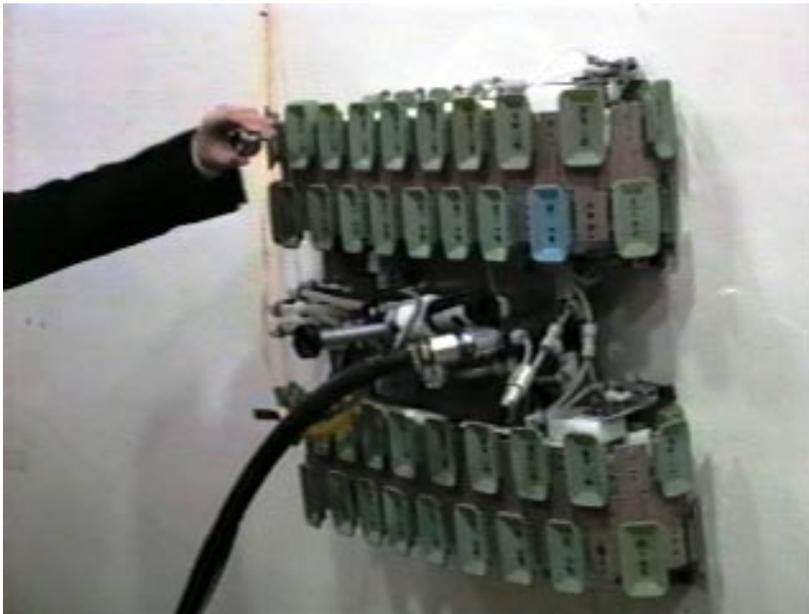


Figure 2: AutoCrawler on the side of an airplane. The hand is about to install, on the "periscope" at the center of the AutoCrawler, a retroreflector that is part of the laser tracking system used for locating the robot absolutely relative to the hangar floor.

2. Seemann, H. R., AutoCrawler Corporation, Seattle, WA. Personal communication and corporate sales literature, videos, etc. Mr. Seemann, CEO and President of AutoCrawler Inc., can be found at 206-367-8163, FAX 206-440-8893, email H579@aol.com.

MACS I THROUGH III: NASA JPL

The Air Force Robotics and Automation Center of Excellence (RACE) at Kelly Air Force Base, San Antonio TX, funded a group led by Paul Backes at NASA's JPL, Pasadena CA, to develop a series of mobile platforms, Multifunction Automated Crawling System (MACS) I through III, a family of small, light weight, high carrying capacity ratio mobile platforms that use suction cups for attachment and ultrasonic motors for motion. The MACS family walking paradigm of alternate attachment and detachment of half the suction cups is essentially the same as ANDI's. The JPL group reports that in the future a descendant of the MACS I through III series with increased on-board intelligence, tetherless operation, operation over the internet, and integration of multiple sensor payloads might be able to carry NDI sensors, e.g., new miniature cameras, tap testers, eddy current sensors, ultrasonic sensors, etc., on an aircraft surface; however there are no reports of any field tests that delivered useful inspection data.

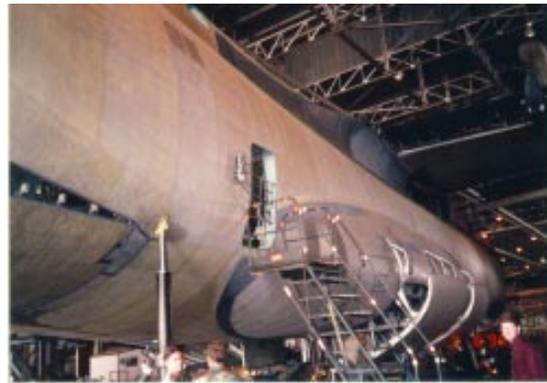


Figure 3: MACS on a piece of sheet metal in the lab (left) and on a C5 airplane (above, at shoulder height, slightly to the left of the doorway).

3. Backes, P. G., Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, CA 91109. Personal communication. See <http://robotics.jpl.nasa.gov/tasks/macs/homepage.html>.

4. Backes, P. G., Y. Bar-Cohen, B. Joffe, "The Multifunction Automated Crawling System (MACS)", Proceedings of the 1997 IEEE International Conference on Robotics and Automation, Albuquerque NM, April 1997, pp 335-340.

5. See <http://www.kelly-afb.org/links/orgs/race/race.htm>.

--- SIDEBAR ---

REMOTE 3D-STEREOSCOPIC VISUAL INSPECTION

Watching how visual inspectors work, we concluded that they use binocular disparity (the small differences between left and right eye perspectives) in several important ways. First, binocular disparity is the primary origin of stereopsis, the human perception of depth via the fusion of slightly different left and right eye images; depth perception is important for perceiving the difference between dents and lighting anomalies, bulges and depressions, etc. Second, aircraft inspectors routinely use dynamic lighting and grazing angles of illumination or observation to discern subtle textural anomalies even on essentially flat surfaces; these they apparently discriminate via the strong binocular disparity that originates in specular (vs. diffuse) reflection features. Thus we decided to provide CIMP with a 3D-stereoscopic video system that gives the inspectors remote *binocular* inspection capability.

The first figure shows two examples of the imagery returned by the inspection cameras. [These pictures have been through several stages of subsampling prior to recording, 8mm taping, digitizing, and MPEG-type data compression, so their quality is not indicative of what the inspector sees live; the live view actually gives each eye an independent NTSC resolution signal stream with very high perceived quality.]

The second figure shows a close-up of CIMP's sensor pod, which contains the 3D-stereoscopic cameras and remotely controlled dynamic lighting, and potentially a variety of other sensors, e.g., eddy current probes.

The third figure shows an inspector at the workstation, operating the wireless remote controller and observing the 3D-stereoscopic imagery.

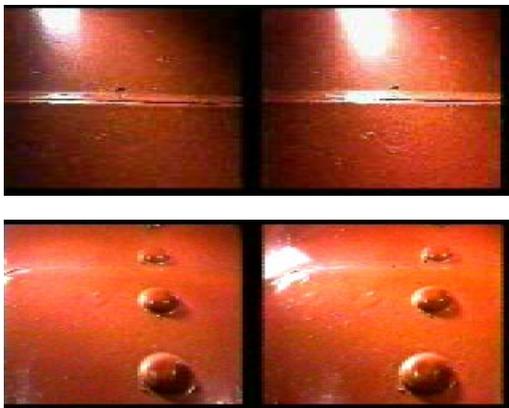


Figure 1: Two left and right eye 3D-stereoscopic views from the 4x 3D-stereoscopic camera in the sensor pod. Top is a lap joint, bottom is a row of buttonhead rivets. Note, in addition to the perspective differences between the two views (which stimulates depth perception), the distinct differences in specular reflections (which we believe stimulates texture recognition).

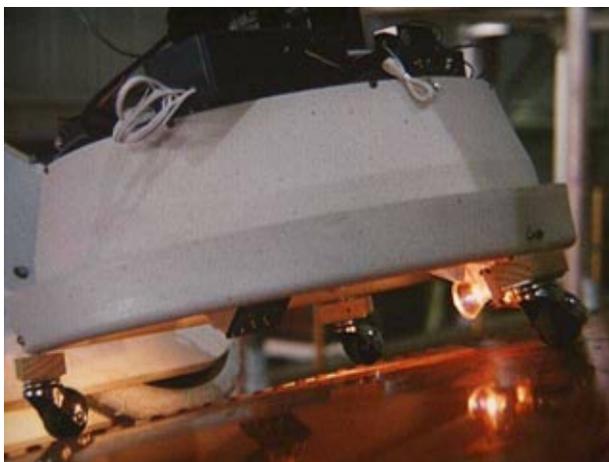


Figure 2: Sensor pod, showing 3D-stereoscopic camera (white box with black endcap left of and below center) and remotely movable low angle illuminator. Inspector can remotely swing this illuminator in a 300 degree arc centered on the forward viewing direction of the camera, reproducing the way he typically uses his flashlight to pick up highlights. The flood illuminator is not visible in this view.

Figure 3: Inspector at the prototype workstation. Small monitor at left shows one eye's view. Large monitor in front of inspector shows left and right eye views 3D-stereoscopically when viewed through the goggles seen. Inspector is driving robot, controlling lighting, cameras, etc., via the model radio controller joysticks, switches, and control knobs.



--- SIDEBAR ---

VISUAL INSPECTION: HOW IT WORKS

There are a wide variety of practices in visual inspection. In the commercial sector, visual is the main method of formal inspection, whereas in the military sector it may be regarded as something that just happens spontaneously, as a side effect of scheduled instrumented inspections, maintenance activities, and sheet metal repairs. Among some commercial airlines the visual inspection corps comprises an organizational unit that is separate from the instrumented inspection (called NDI, NDT, or NDE) corps; at other airlines the NDI inspectors are also responsible for visual inspections. Here we describe how it typically works in a commercial airline that has separate visual and NDI inspection units.

A visual inspector typically comes from the ranks of the maintenance organizational unit after accumulating ten to fifteen years of experience there; supervisors have noticed, over these years, that he (or she) has a natural talent for spotting potential problems, and a personality that insists on following up on these observations personally and administratively. The newly tapped inspector receives a few weeks of training in the administrative aspects of visual inspection, its paperwork and document flow, but no explicit training in the perceptual science of visual inspection or its enhancement by technological aids, e.g., lighting and optical magnification; these are learned on the job, by working initially alongside an experienced visual inspector.

Visual inspection has both a directed component and an undirected component. The former involves intensive close inspection, sometimes with specified lighting and magnification, over a small specified area in which problems are anticipated by previous experience or by the aircraft manufacturer's stress models. The undirected component is essentially a "fishing expedition", an overview intended to find problems that could appear anywhere: pilling indicative of substructure corrosion, cracking somewhere other than at the high stress locations that are the subjects of directed inspections, rivets pulling loose, dents caused by baggage and food service vehicles, lightning strikes (typically a front-to-rear trail of five to ten metal splatters, each 0.5 to 1 cm in diameter, spaced 15 to 25 cm apart), hardened and cracked sealants under antenna patches, fluids leaking from the hydraulic systems, the galleys and lavatory plumbing systems, etc. The "fishermen" are as serious and thorough about the undirected component as they are about the directed component, working slowly and methodically, but typically using only their naked eyes and a high intensity flashlight, bringing out magnifiers, dye penetrants, and requesting NDI (e.g., eddy current) confirmation only when a problem has been spotted at a distance. However, with sometimes hard-to-believe natural abilities, honed by years of experience, we have personally seen them spot, from 1 or 2 m away, millimeter length cracks so thin that we then have difficulty seeing up close even after they have been circled with a red china marker ("grease pencil").

The visual inspector's tools are thus mainly his (or her) eyes, a flashlight, and a 10x magnifier. The flashlight, used to illuminate the aircraft skin at a glancing angle from an arc of azimuths, helps high spots (e.g., loose rivets) and low spots (e.g., dents, scratches, cracks, and rivets pulled inward by stressed substructures) to stand out visually. The presence or absence of specular reflection distinguishes a scratch from a crack: a scratch has a shiny bottom, a crack is a "black hole". The visual inspector also carries a small pick, which he can use to attempt to enlarge an apparent crack, lift a paint chip that might be mistaken for a loose rivet, poke at a sealant bead, etc. Cases that are still questionable after these procedures are resolved first by the dye penetrant kit (which the visual inspector is authorized to borrow from his NDI unit colleagues), and if still ambiguous, by requesting an NDI inspector to perform an eddy current check.

In addition to cracking, corrosion is of interest, as corrosion is common due to the frequent exposure of the aircraft body to liquids such as aircraft operating fluids, galley spills, lavatory contents, moisture condensed from sea air, etc. Both surface and subsurface corrosion can be seen visually. Surface corrosion is recognized by its color and texture. Subsurface corrosion is recognized by the bulging of the affected surface region, called "pilling". Since corrosion results in a loss of structural material of the affected area, early detection is crucial. Corrosion is also known to induce cracking.

1. S.N. Bobo," Visual Nondestructive Inspection Advisory Circular", 1993.

--- SIDEBAR ---

INSTRUMENTED (NDI) INSPECTION: HOW IT WORKS

NDI (**NonDestructive Inspection**, also called NDT, for **Testing**, and NDE, for **Evaluation**) has three main components: (1) inspections done off-line in the NDI shop, e.g., ultrasonic inspection of turbine blades, (2) inspections done on assembled or partly disassembled engines, e.g., eddy current inspection of a particular weld, and (3) inspections done on the intact aircraft, e.g., x-ray of major structural members. We concentrate on the **instrumented** inspections done on the intact aircraft, and in particular, on inspections of the skin surface and the closely adjacent substructure that is accessible to NDI techniques such as eddy current and ultrasonic probing.

Unlike visual inspectors, most of whom move up from the maintenance corps, NDI inspectors often have a technical school background specifically in NDI; many have associate's degrees, and some have bachelor's degrees, in an engineering field. Most have some years of experience in maintenance as well. In any case, a new NDI inspector receives several months of training in the technical as well as the administrative aspects of NDI inspection. The technical training emphasizes primarily the particular airline's specific NDI equipment, procedures, protocols, and policies. Specialized training off-site at the location of an equipment manufacturer is sometimes provided.

Despite the fairly high level of training and background knowledge that NDI inspectors have, their jobs are usually tightly controlled by the airline's inspection procedure and policy documents. An inspection supervisor, working on a manufactured calibration specimen that simulates the sought flaw, develops and documents an inspection protocol. The protocol specifies the particular probe to be used, instrument settings, probe path, any templates or probe guides and how to align and follow them, and a (usually graphical or pictorial) description of the signals that are expected from good and bad samples. On the airplane, the inspector is expected to set up the instrument according to the supervisor's procedure document, confirm the probe's response to a good and a bad (usually manufactured) sample, and proceed with the inspection without changing anything.

Of the estimated 10% of inspection that is instrumented (the 90% majority being visual), an estimated 90% is eddy current, and the remainder mostly ultrasonic, for a few specific problems

Eddy current "point" probes (sensitive to the metal in a localized measuring volume) may be either "pencil" style, with one coil, or "reflectance" (also called "pitch-catch") style with separate excitation and detection coils. In either case, the display is typically an x-y oscilloscope trace in a half-plane that represents complex impedance with increasing resistance to the right, inductive reactance up, and capacitive reactance down. If the probed volume includes a crack, the normally predominantly inductive character of the material within the sensitive volume decreases, and its capacitive character increases. Alternative means of recording and displaying the signal are sometimes seen, e.g., plotting the real and imaginary components of the complex impedance vs. time on a "heart monitor" trace ("A-scan" mode).

In recent years, alternative, easier to deploy, eddy current probe variations have been developed. These include linear arrays (several raster scan lines are collected in parallel), area arrays (making mechanical scanning unnecessary, at least over small areas), and magneto-optic imaging, MOI (an electro-optically active sheet in contact with the surface under inspection renders the eddy current pattern directly visible over a square of about 10 cm on each side).