Eric Byler, Wendell Chun, William Hoff, Dan Layne

Autonomous Hazardous Waste Drum Inspection Vehicle

Automated inspection of radioactive and hazardous waste storage containers will reduce the exposure to personnel and create accurate, high-quality inspection reports to ensure regulatory compliance at storage facilities.

Hundreds of thousands of hazardous, radioactive, and mixed waste drums are being stored throughout the world, and the anticipated decommissioning of facilities will generate many more drums. Currently, in compliance with federal regulations, waste storage facilities at U.S. Department of Energy (DOE) sites are inspected manually for degradation and to verify inventories.

An Intelligent Mobile Sensing System (IMSS) has been developed for the automated inspection of radioactive and hazardous waste storage containers in warehouse facilities at DOE sites. The IMSS will reduce the risk of exposure to personnel and create accurate, high-quality inspection reports to ensure regulatory compliance.

The IMSS includes an autonomous robotic device with enhanced intelligence and maneuverability, capable of conducting routine multisensor inspection of stored waste drums.

INTRODUCTION

The purpose of the IMSS program is to create a system to automate monitoring and inspection of stored hazardous, radioactive, and mixed waste drums. The DOE has hundreds of thousands of storage drums stored in multiple facilities located at several sites in the United States (Figure 1). The Environmental Protection Agency (EPA) requires positive weekly inspection of each storage drum in a storage facility. This inspection process is time-consuming and presents inherent health hazards. The IMSS will automate the inspection process, lowering costs and providing safer, more accurate and more consistent inspections.

System functional requirements were developed from EPA requirements and discussions with waste operations personnel at four DOE sites (Oak Ridge National Laboratory, Hanford Engineering Laboratory, Idaho National Engineering Laboratory, and Rocky Flats Plant).

Problem

Most waste storage facilities contain 5,000 to 20,000 barrels per building. Barrel sizes include 35-, 55-, 87-, 93-, and 110-gallon drums; colors include white, yellow, silver, gray, and black. Storage facilities typically have barrels stored four to a pallet, with pallets arranged in single rows. Observed stacking heights for pallets vary from two to five pallets with an average of three high. Aisle widths vary from facility to facility, ranging from 26 in. to 30 in. to 36 in. Aisle lengths also vary from 20 feet to hundreds of feet. In general, space is left between the last pallet in a row and the adjacent wall. Positive inspection of each barrel is required and operator response to flagged barrels is required within 24 hours. Current methods used to inspect and monitor stored wastes are based on passive detectors or humans walking through the storage area with instruments. Passive monitoring relies on fixed sensors dispersed within the containment building, e.g., radiation or gas detectors. When an increase in radiation is measured, operators must enter the storage site and locate the leaking container. Walking inspections might include radiation

Eric Byler, Wendell Chun, William Hoff, Dan Layne: Martin Marietta, Astronautics, P.O. Box 179, Denver, CO 80201

 1070-9932/95/$4.00©1995 IEEE

Figure 1. Typical DOE storage facility (Courtesy of Hanford Engineering Laboratory).

detectors and gas detectors, but usually are visual inspections. Visual inspection of drums is required to detect dented, bulging, or rusting drums. However, visual methods are a function of operator acuity and fatigue and may vary between operators and even between individual drums. Operators may receive varying radiation doses during their inspections and must be examined for contamination before site exit. Required drum inspection frequency and operator lifetime radiation limits raise the total cost of this monitoring process and introduce health and safety risks.

In performing a visual inspection for mixed waste storage, a human operator evaluates the exterior condition of the exposed face of the drums to determine the integrity of liquid containment. Professional judgment is used to identify any negative conditions that may result in the escape of any liquids, or in the case of radioactive waste, any drum condition that may result in a release of airborne contamination (e.g., alpha particles). With these qualifiers in mind, the following extracted requirements for drum inspection were compiled in Table 1.

Quantitative numbers and/or conditions were selected to provide a decision-making basis whether a barrel should be considered defective. These numbers and the system performance are discussed later. Visual anomalies NOT to be mistaken or confused as a defect include:

1) Accumulations of dust or dirt on ridges, rims, or seams;
2) Condensation streaks in dust or dirt;
3) Symbols or other labels that are not bar codes;
4) Drum seams.

The inspection requirements are not standardized in general because of differences of the various state and federal regulating agencies and various DOE facility policies.

In operation, the mobile robot departs on an assigned mission, navigating through narrow aisles between rows of drums stacked on pallets, and avoiding obstacles along the way. The robot acquires and correlates data on each drum in each aisle. After completing an inspection mission, the IMSS vehicle returns to its home base, docks with a battery recharger, and transfers collected data across an Ethernet communication link to the operator's console located at the central control station. The control station supervises vehicles in several buildings from a remote facility via Ethernet commands to the docking station. The report forms created by the IMSS are the same as those created by inspectors. In general, inspection reports will be completed automatically once per week per area. A map of the area is included in the report on which the defective barrel is identified to aid operators in their responses.

IMSS Program

The IMSS program is a three-phase effort to develop an autonomous monitoring and inspection system/technology. The objective of the first phase was to demonstrate an integrated system performing all required functions; the objective of the second phase is to efficiently package the software and components and perform a hot demonstration in a full-scale storage facility; the objective of the third phase is to develop

Table 1. Barrel inspection requirements.

Sharp or Pointed Dents	No Depth Greater Than 1 in., Width or Length Not Critical
Rounded Dents	Ignore Unless Stability of Drum Is In Question
Superficial Rust (Paint Corrosion)	Track Diameter Size; If Rust Is Increasing, Identify
Streaks of Rust	Identify Streaks; Discriminate Between Streaks Of Rust & Streaks Of Condensation in Dirt Or Dust; Quantify By Length, Width & Position
Nonsuperficial Rust (Metal Corrosion)	Identify By Diameter
Tilted (Bulging) Drums	If Drums Are Banded, Identify If Base Of Drum Is Touching Bottom Storage Surface (Pallet, Plywood, Or Floor); If Drums Are Not Banded, Identify If Tilted (Any Angle Greater Than 20); Identify If Ribs Of Drum Cannot Be Distinguished
Stacking Levels	For Specific Storage Area, Identify If Stacking Level Is Exceeded
Condition of Pallets	Identify If Broken
Location Of Bar Codes	Upper Third of 55-gallon Drums or Top Half Of 35-Gallon Drums; Top Of Bar Code Not More Than 2 in. Below Drum Seal, Visible From Aisle; Note If Missing
Location of Hazardous Waste Labels	If Site Requires Hazardous Labels, Label Should Be Located In Center Third Of 55-gallon Drums Or Top Half of 35-gallon Drums

the product into a certified, commercially viable system.

The Phase 1 effort discussed in this paper assembled an integrated engineering demonstration model, including all components required to perform functional requirements. The vehicle was integrated from subsystems, some of which existed as part of a Mars rover prototype, including the motion platform, the sensor mast, and the operator's console. Other components that were added included the ultrasonic obstacle avoidance system, the sensor suite, omnidirectional wheels, and navigation software. The Phase 1 system was tested in a small-scale storage facility mockup to gather performance data for the Phase 2 design. There are two parallel efforts to the IMSS program: the Stored Waste Autonomous Mobile Inspector (SWAMI) being developed at the Savannah River Technology Center, and an intelligent inspection and survey robot being developed by the South Carolina Universities Research and Education Foundation (SCUREF).

The following portions of this paper discuss the design and performance of the engineering demonstration model beginning with a system overview and progressing through a discussion of the mobility base, mission sensors, and operator interface. The paper concludes with a presentation of measured system performance.

SYSTEM

The system functional architecture is shown in Figure 2. There are three major segments of interest: the vehicle, mission sensors, and the control station. These are coordinated by the executive that runs on the vehicle.

The vehicle is a self-contained entity that provides positioning services to mission sensors. Vehicle software is hierarchical. The architecture provides a structure that enhances modularity, both for portability and extendibility. The three major vehicle software components are the executive, navigation, and real-time operations (motion control, pointing, and obstacle avoidance). The motion base and sensor pointing system operate independently under coordination of the executive.

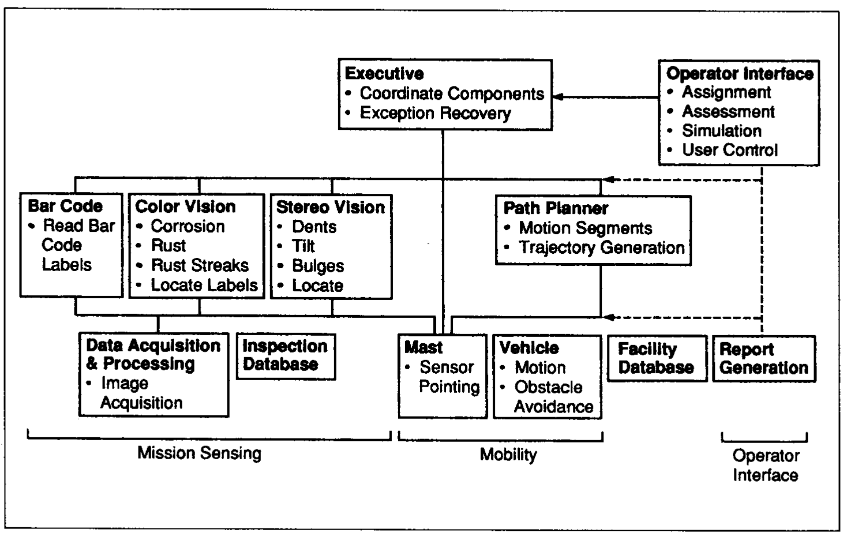

Figure 2. Functional architecture.

Mission sensing modules are each independent, parallel subsystems. This model is similar to DOE's Generic Intelligent System Controller (GISC) architecture [1] in terms of providing standard interfaces and calls to subsystem components. This structure separates control functions from sensing functions, but still allows the sensors to communicate as necessary through the world model or the executive for coordination.

The control station acts as the supervisor. Site operators can load and modify site databases; approve, modify, and initiate inspection cycles; and receive, analyze, and print inspection reports and data. All communications above the hardware control level were implemented using the Transfer Control Protocol/Internet Protocol (TCP/IP) on Ethernet. The network interfaces were implemented with the project-standard TCP/IP library. This library provides standard function calls for establishing a server, opening a connection, reading data, writing data, and testing for data availability of. This library was implemented on all the UNIX computer systems and the real-time VRTX embedded control system. This link is available as a radio link and as a hard link. No communications is required during operations.

Executive

The purpose of the mission executive is to direct the integrated action of various vehicle functions to achieve an efficient and complete site inspection. The functions (or agents) coordinated or controlled by this module include mobility and scan platforms, some parameter setting on the vehicle, the structured light 3-D vision module, the color vision module, the visual landmark navigation module, and the bar code reader. Because some modules can change the state of the vehicle, the mission executive also receives or requests relevant information pertaining to these state changes to be able to coordinate the remaining action effectively.

The Initial Plan Generation function is to generate, from the symbolic description obtained from the supervisor, a detailed, partially ordered action plan of consisting of integrated mobility and scan platform motions and various sensing operations. Efficient inspection is achieved in two ways: (1) inspection sequences that minimize changes in vehicle state and position will be efficient relative to those that do not, and (2) maximizing concurrency of task performance will tend to minimize inspection time. The plan is based on site information provided in the world model and costs and constraints associated with each potential action. Operations research and heuristic algorithms use this information to generate a complete, efficient plan for site inspection. Minor environmental perturbations can be accommodated by following local plan repair strategies, as discussed below, without significant impact on the overall efficiency of the plan. Strategies in initial plan generation are used to minimize state changes and to maximize concurrent performance of operations, thus maximizing the efficiency of the performed inspection sequence.

The Plan Dispatching and Monitoring function uses the plan constructed by the previous module. Commands are issued to each functional agent at the appropriate time, and the progress of each agent is monitored to maintain a coordinated inspection sequence. The most efficient vehicle operation entails concurrent performance of compatible operations. By overlapping actions where possible, significant time savings can be achieved in the inspection sequence. The most profitable of these are overlapping vehicle and scan platform motion, overlapping image processing with any other vehicle or sensor activity, and anticipatory state changes, particularly changing the attachment of the processing board that was used by two sensor systems. However, not all actions can occur concurrently, and ordering of actions-both concurrent and sequential-must be carefully controlled. The plan provided by the plan monitor sets up an ordering that allows the dispatcher to take advantage of the most profitable concurrencies. The dispatcher uses that plan to generate coordinated action sequences that achieve complete, efficient inspections. This is achieved through use of ordered command queues, monitoring of the state of the vehicle and the progress of each of the commanded modules, and checking of the required states for each of the pending commands.

There are three main command queues: pending, in progress, and complete. The pending command queue contains an ordered list of the next operations on the plan command list that have not yet been initiated. The in progress queue contains a list of all commands currently being executed and their status. The complete queue contains a list of all commands that have been fully executed and their final disposition, ordered by completion time.

Two types of state variables were tracked during execution for each of the vehicle and sensor components. One set tracked the actual state of the component (e.g., LAMP might be set to ON or OFF), and the other set tracked which component had the right to change state of a given component/module (e.g., the SATURN_ANDROX_BOARD control might be set to FREE_CONTROL or STEREO_CONTROL or LANDMARK_CONTROL). By carefully monitoring the current state and control of each variable, a determination can be made of what commands may and may not be concurrently executed. By setting control variables for each of the modules, ordering of commands is achieved when necessary, since any conflicting requirements are resolved in favor of the earliest item on the pending queue.

Although true excursions from the expected environment should occur only very rarely at these sites, when these anomalies do arise, the system must be fully competent to handle these anomalies and proceed with the inspection. The replanning function alters or refines the plan as the inspection sequence progresses to adjust to real-time inputs during the inspection. Under some sensor failure conditions, the vehicle simply continues the inspection sequence. However, under many conditions new actions must be planned and performed. Included in these replanning functions are procedures to return the vehicle to a safe and recoverable state/location in the case of serious vehicle or environment problems. This consists of the generation of a new, simple plan to return home given knowledge of current location and home and a strategy that allows the vehicle to search for home using local decision-making strategies.

MOBILITY SYSTEM

The mobility system will be discussed in terms of (1) vehicle hardware, (2) navigation, and (3) obstacle avoidance. The discussion on vehicle hardware includes the motion base, sensor pointing, electronics, and power.

Vehicle Hardware Motion Base—The Phase 1 vehicle (Figure 3) was originally designed as a testbed for planetary rover research [2]. It consists of two equipment bays, gearhead drives at each of the four wheels, a passive roll axis along the vehicle's longitudinal axis, and associated drive control electronics. Both bays are symmetrical and the vehicle is about 60-in. long by 20-in. wide by 24-in. tall. For simplicity, the vehicle was skid-steered. To adapt the wheeled rover to our inspection application, we outfitted the vehicle with four Mecanum omnidirectional wheels [3]. By independently controlling each wheel and the vehicle's inverse kinematics, we are able to achieve 3 degrees-of-freedom (DOF) Cartesian motion. As a result, simple Cartesian path planning may be used to position the vehicle precisely for measurements.

The wheel motors are driven by independent custom-built pulse-width modulation (PWM) amplifiers based on the Advanced Motion Controls' AMC-500 hybrid PWM amplifier. The amplifiers accept ±10-V differential control signals and produce +28-V pulse-width modulated drive for the motors. The current limit for each amplifier is 8 A. This limit is set well above its expected normal operating conditions.

Sensor Pointing—At the front of the vehicle is a two-axis gimbal. It provides pan-and-tilt motion for various mission sensors. This was originally designed to view only the ground surface for planetary navigation, which imposed limits on the height of barrels that could be examined. This will be modified in subsequent versions. The scan platform motors are driven by independent custom-built PWM amplifiers based on the linear integrated circuits' L292 switchmode driver. The amplifiers accept ±10-V differential control signals and produce +28-V pulse-width modulated drive for the motors. The current limit for each amplifier is 2 A.

Electronics—The VME bus card cage contains commercially available boards, including a Force SYS68K/CPU30 single-board computer, two Galil DMC530 motion control boards, a XYCOM XVME240 digital input/output (I/O) board, and a XYCOM XVME230 intelligent counter module. The Force SYS68K/CPU30 single-board computer is based on a 25-MHz 68030 with a 68882 math coprocessor. The CPU30 is fitted with Ready Systems' VRTX/Velocity programmable read-only memories (PROM). VRTX/Velocity is the real-time multitasking operating system used by the vehicle software. The CPU30 offers an Ethernet connection, serial ports, and 4 M of random access memory (RAM). Moti provide adjustable proportional integral derivative (PID) controllers for the wheels and the scan platform motors. The position sensing incremental opti-

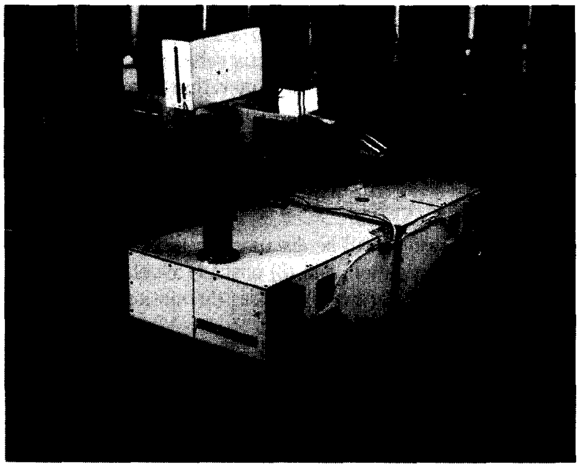

Figure 3. IMSS Phase 1 vehicle.

cal encoders are connected directly to the DMC530s. System digital discretes are controlled and monitored by the XVME240 digital I/O board. This board has a 64-signal capacity, configurable as 8-bit ports of inputs or outputs. The current configuration splits the ports, providing 32 inputs and 32 outputs. These signals are used for power state control for system components, monitoring abort conditions, and other system functions. The XVME230 intelligent counter module interfaces to the ultrasonic ranging system. One 16-bit up counter on the board is used as an event counter to acquire the distance to objects for each of the eight ranging channels.

Power—The primary power source is a set of lead acid batteries that supply up to 40 amperes. The +28-V bus provides power directly to the wheel, pan, and tilt motor PWM amplifiers. It also supplies the necessary input power for dc-dc converters, which provide the vehicle system with +5 and ±12 Vdc.

A secondary power source is two parallel 28-V, 12-A linear dc power supplies. The power supplies charge the batteries when the vehicle is off, and trickle-charge the batteries during intermittent operations (e.g., during debug and testing to avoid unnecessary battery cycling). Offboard supplies are connected to the vehicle +28-V bus by a 75-ft tether consisting of three 12-AWG power/ground pairs. This allows a maximum IR drop of 1.3 V. IR drop in the tether is significant because the supplies are configured to sense locally. If remote sensing is used with these supplies, IR drops of about 1 V cause instability, resulting in power supply damage.

Navigation

The vehicle navigation system is based on a combination of dead reckoning, landmark sightings, and an a priori map of the mockup storage facility. A block diagram of the navigation system is shown in Figure 4. This approach uses proven methods for a structured, indoor environment-continuous data availability from odometry and absolute position knowledge from landmarks. Odometry position estimation uncertainties that grow with travel distance are reset to zero by occasional sightings of known landmarks. Phase 2 batteries are sized for inspecting 12,000 drums per week.

Trajectory Generation—Position commands from the mission executive are used to build an acceleration and velocity limited velocity trajectory to move the vehicle from its present position as measured by odometry to the commanded new position in a smooth, orderly fashion. Positions are based on a priori map data or on information relative to drums or aisles. Maximum velocity and acceleration time to maximum velocity are programmable. The velocity trajectory is updated at the vehicle servo rate (0.1 sec) and used to calculate wheel speed commands using vehicle inverse kinematics.

Dead Reckoning—As we mentioned earlier, the IMSS vehicle has omnidirectional wheels designed to operate without translational slip. This allows estimation of vehicle position and orientation within the facility based on a time history of wheel angular velocities. The process involves sampling wheel positions at a fixed rate; differencing the positions to get the velocities; applying the kinematic velocity transformations between wheels, vehicle body, and facility reference frames; and integrating the resulting facility frame velocities to get the position and orientation. These estimated data are available at the basic servo-loop rate for coordination of vehicle and mission sensor scan platform motion. Error sources include kinematic modeling errors (e.g., wheel diameter) and wheel slippage on the floor. Error sources are compensated for by using landmark sightings.

Landmark-Based Pose Estimation—At regular intervals, landmarks are used to update the pose (position and orientation) of the vehicle in the facility, and reset accumulated errors from dead reckoning. In Phase 1, landmark sightings were commanded by the mission executive process after the vehicle had traveled about 5 meters or more or had turned a corner at the end of an aisle. In Phase 2, we plan to initiate landmark sightings at ends of aisles, during docking maneuvers, and possibly when a drum position-type defect is detected. Drum position-type defects include a missing drum or a drum not located in its expected position.

Landmark targets are fastened to the facility's walls and columns at regular intervals. Each target is a flat piece of metal or paper and is composed of a set of high contrast concentric circles in a predetermined pattern (Figure 5). These passive landmarks should normally require no maintenance, unless they are soiled or damaged. The choice of the concentric circle features and the layout pattern was motivated by our work in accurately locating objects for space robotic tasks [4, 5]. The concentric circles can be reliably extracted from images even in highly cluttered scenes, due to the distinctive arrangement of a white region co-centered with a black region. The small outer circles are used when the target is close enough so that they can be resolved (about 3 meters or less); otherwise the large inner circle is used. The distinctive pattern of three collinear circles is used to determine the correspondence of the small circles.

Because the main pattern of features (the concentric circles) is the same for all landmarks, an individual landmark is normally identified by its rough location in the facility. In other words, given a rough idea of the vehicle's pose, a landmark that is observed at a certain location can match only one of the known landmarks in the facility database. Another way of distinguishing landmarks from one another is by the

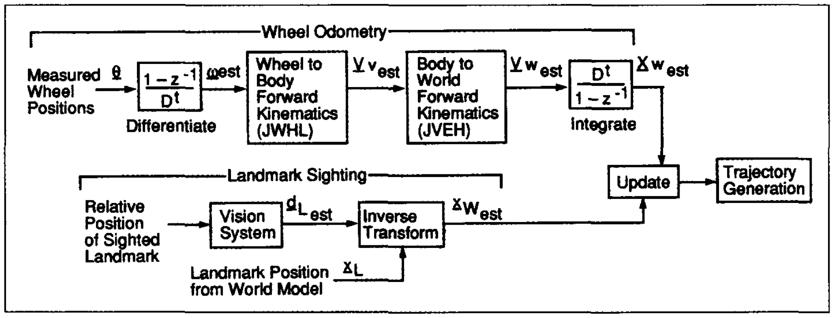

Figure 4. Block diagram of the navigation system.

unique bar code and text on each target. In this way, even if the vehicle is totally lost, it could locate itself by finding any landmark.

When the mission executive decides that it needs to update the vehicle's pose using landmarks, it sends a command to the landmark recognition system, along with a list of landmarks to locate. Using the approximately known pose of the vehicle, the landmark recognition system points a video camera in the direction of each landmark and grabs an image. Any video camera could be used, but in Phase 1 we used a dedicated navigation camera that had its focus and field-of-view optimized to locate distant targets.

The circle features are extracted from each of the images and combined into a list. The list contains for each of the circle features: (1) the unit vector direction to the feature in the vehicle coordinate frame, derived from the image location, and (2) the corresponding (x, y, z) location of the feature in the world coordinate frame, known from the facility database.

Next, an optimization algorithm is used to refine the pose of the vehicle in the facility. The optimization function is the least squared error between the observed directions to the features and the predicted directions to the features based on the current vehicle pose estimate. The "downhill simplex" algorithm was used [6] to refine the (x, y, q) vehicle pose estimate.

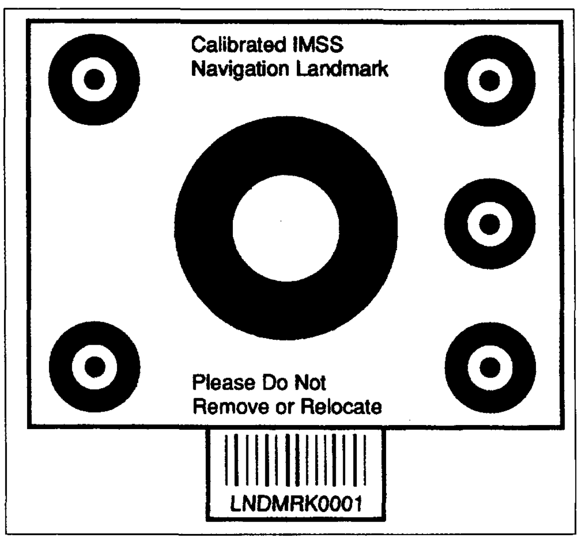

Figure 5. Representative landmark target.

The overall positioning accuracy of the vehicle was tested for various distances and angles traveled. Without a landmark update, the average accumulation of error was 0.008 meters per meter position error, and 0.25° per meter angular error. After landmark update, the average position error was 0.020 meter with a 0.24° angular error with respect to three different sightings.

Obstacle Avoidance

An obstacle avoidance mechanism is provided to prevent contact with unknown or out of position objects. Eight ultrasonic range sensors are located around the vehicle perimeter that provide range information indicating the proximity of objects in the vehicle environment. This range information is used to develop a virtual force potential (impedance) that wards off the vehicle from the object. When scaled properly, the force potential resulting from an imminent collision with an object may not be overcome by any motion command that would otherwise result in a collision with that object. A virtual force calculated from ultrasonic range information acts on the vehicle at each sensor location, resulting in a net force and torque at the vehicle center (vehicle reference coordinate frame origin). To apply the virtual force and torque to the vehicle through the controller requires generation of a velocity command using the force and torque. The relationship between force and velocity defines a mechanical impedance for which the simplest example involving mass is a second-order dynamic system. Using, for example, a mass-damper model to generate velocity commands from virtual forces and torques results in a well-behaved relation that allows adjustment of sensitivity and speed of response of the vehicle to obstacles.

The obstacle avoidance capability is not coupled to path planning. It operates independently for real-time obstacle avoidance to provide additional velocity inputs to the vehicle control system. There are two possible modes of operation: 3 DOF and 2 DOF. The 3-DOF mode implements the complete velocity response to obstacles as calculated from input forces and torques. The vehicle response includes rotations to move the vehicle from the obstacle in a least energy form. The 2-DOF mode is more appropriate when the vehicle is operating in narrow aisles where rotations of the vehicle may cause it to get "stuck."

The ultrasonic ranging system is capable of detecting objects at a minimum range adjustable from 0.7 to 1.3 ft and a maximum range of 32 ft. The stated typical absolute accuracy is 1% of the reading and the detection angle is approximately 10°. Each sensor module consists of a Polaroid 616342 electrostatic transducer and 615077 ranging board. These modules are connected to a common custom-built electronics package that accepts software-controlled digital discretes.

Assertion of a control signal to the electronics package initiates transmission by the corresponding transducer. At the same time, a 1-MHz square wave is fed into a 16-bit counter on the VME board. When the transducer detects an echo, the square wave is disabled. The value in the counter is then relatable to the distance of the object that caused the echo.

MISSION SENSING

The mobile robot shown in Figure 3 is equipped with an integrated sensor suite that gathers data to identify and report anomalous drum conditions. These defects include rust spots, rust streaks, corrosion, dents, tilted drums, drums missing or out of place, and missing bar-code identification labels. Subsystems involved include the geometric inspection system, the corrosion inspection system, and a bar-code reader.

Geometric Inspection System

The laser ranging sensing subsystem was used to inspect 3-D drum characteristics. Specifically, the tasks of initially locating the drums, detecting surface dents, and measuring drum tilt are essentially 3-D tasks (i.e., they require accurate measurement of 3-D points on the surface of the drums). A dense set of 3-D points is required so that small dents are not missed. The resulting dense set of 3-D points can be put into the form of a range image, in which integral positions of a 2-D array represent direction and the values stored in the array represent the range to points in the scene. Assuming the sampling intervals are consistent in the horizontal and vertical directions, the (i, j) position of an element in the array implicitly determines the 3-D direction vector to that point. The term "range image" is used because the array can be displayed on a video monitor and has the same form as an ordinary video image. The only difference is that the values in the range image represent distance instead of light intensity.

We elected to use structured light as the range sensing technique for our task of drum surface inspection. This selection was driven by our need for a low-cost, accurate system. The disadvantages of slow scanning time and limited depth-of-field were not perceived as sufficiently major to eliminate this technique from consideration. We developed our own structured light system consisting of a laser line projector, two CCD cameras, and a image processing workstation. The laser and cameras are mounted on the pan-tilt mechanism to allow scanning of the entire scene. This particular design for Phase 1 was motivated by the fact that all equipment, except the laser line projector, was available. In an operational system, image processors and framegrabbers will be present to perform rust and corrosion inspection.

During the creation of a range image, the pan-tilt mechanism was slowly panned across the drum surface, in two separate swaths, to image the complete drum. Images were grabbed from the two stereo cameras by an Androx ICS-400 image processing board with a Solbourne 5/501 host and processed at a rate of 7 frames per second. For each image, the location of the maximum intensity pixel along each row was determined (in the case of multiple pixels with the same maximum value, the midpoint of the group was determined). The resulting points, representing points along the projected light stripe, were saved to a file.

Next, a separate process converted the stripe image points to 3-D (x, y, z) points, based on a predefined calibration lookup table. Points from the two stereo cameras were combined at this time, using the following heuristic: if the corresponding points from the two cameras agreed in 3-D location within a threshold distance (2 cm), then the points were averaged. If not, the points were both thrown out. This two-camera approach helped eliminate false returns caused by interreflections and highlights. The complete set of remaining 3-D points was collected into a single range image. The characteristics of the structured light system are shown in Table 2. An example range image is shown in Figure 6.

Next, the drum was located within the range image. This was accomplished by a simple algorithm that used the fact that the drum was approximately the shape of a cylinder, with a known radius (known a priori in the facility database). It first computed the surface normal at each point in the range image, then projected inward from each surface point for a distance equal to the known radius of the drum (Figure 7). These points ideally all lie on the axis of the drum; although in reality, there is some displacement from the axis. The next step was to project all the "axis" points downward onto the floor and look for a large cluster of points (Figure 8). The location of the cluster gives the approximate location of the intersection of the drum axis with the floor.

The next step is to refine the drum's position and orientation, using raw measured surface points as data. This is accomplished with an optimization algorithm, which optimizes the least squared error distance between measured surface points and the distance to the predicted drum surface, based on the current drum axis pose estimate. The "downhill simplex" algorithm was used [6] to refine the (x, y, q, f) axis pose estimate. After the axis has been accurately found, the drum is tested to determine if it is tilted or out-of-place.

The final step in the processing is to find dents. Here, dents are defined to be deviations from the nominal drum surface of over 0.5-in. vertical depth and having a lateral surface area of 7.8 in.^{2}. Parameters were chosen empirically based on measured performance. We take the estimated pose of the drum and determine which points in the range image do not lie within a depth tolerance of the ideal drum surface. Points that lie at a radius less than the nominal drum radius are possible dent points and are clustered into regions. Those regions larger than the specified surface area are flagged as dent regions.

Corrosion Inspection System

Visual anomalies (rust, streaks, and corrosion) on storage drums are detected with a color camera and image analysis software. A 16mm lens is used, based on comparing various lenses for drum inspection requirements. Four to six images per drum are needed for medium-resolution inspection (to locate pea-sized rust spots) depending on barrel size. The camera is located at the center of the vehicle scan platform. Adjacent to the camera is a compact 20-W video (halogen) lamp that is turned on only during image acquisition. The lamp helps illuminate rust and streaks under variable lighting conditions. For inspection, the camera is nominally positioned about 1.0 m from the drum, using location informa-

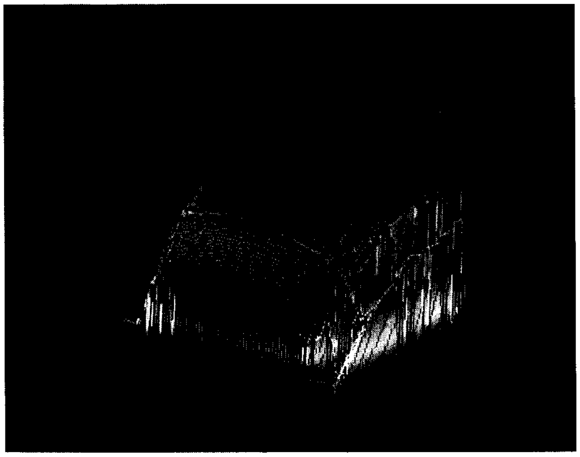

Figure 6. Range image of dented drum.

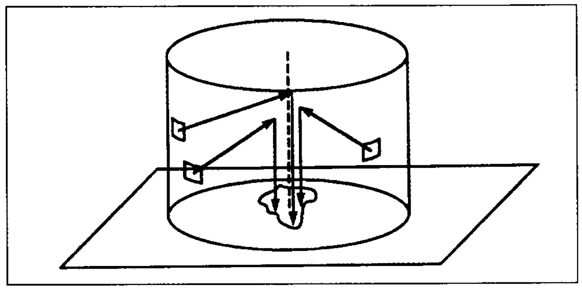

Figure 7. Finding the approximate location of the drum axis, based on measured surface points.

tion from the world model and the laser ranging subsystem. A single video cable carries the composite color signal from the vehicle to the UNIX workstation. Video signals are then decoded into red, green, and blue (RGB) images before digitization on the Androx framegrabber. Vision software is controlled by the system executive and as soon as the current image is digitized, the scan platform and vehicle are released to perform other tasks while the image is processed off-line. We prepared samples of rust, rust streaks, and corroded paint on black, white, yellow, and silver drums to develop and test the vision software.

Image analysis algorithms are developed on a UNIX workstation using the Khoros image processing package, with pattern recognition tools from our internal research on neural networks. Optimized versions of the image processing software will run directly on the framegrabber board in Phase 2 rather than on a workstation. During our image analysis process, a smoothing filter is first applied to the input RGB image to eliminate noise and reduce sharp edges (RGB color edges are not well defined and are not used here). Before feature extraction, the RGB tristimulus response images are orthogonalized (decorrelated) by transforming into the Ohta color space, which is a fast approximation to the Karhunen-Loeve eigenvector transform [7]. The new image bands are "intensity" = $(r + g + b)/3$, "red - blue difference" = $r - b$, and "excess green" = $(2g - r - b)/2$.

A model-based vision approach is used, wherein the drum color is first determined, and then separate modules (with color-dependent parameters) identify rust, streaks, and corrosion. Although rust spots on different color drums appear similar to human vision, color computer vision must proceed with RGB images as the starting point. For example, a small rust spot in the red image from a white drum has a mean intensity of 45 and standard deviation of 4.8, while a similar rust spot in the red image from a black drum has a mean intensity of 92 and standard deviation of 8.5. Because both rust spots were produced in the same way, the spectral difference is a result of different contrast from the surrounding paint. Drum color is determined by a maximum likelihood threshold detector [8] on centered subimages. If the detected color differs from the expected color (available from the world model during normal inspection operations) then the mission executive is notified. Based on the selected color model, adaptive and multispectral histogram thresholding is performed to develop binary image masks.

Streak detection is performed separately from rust and corrosion detection. Although spectral information is important for finding streaks, the dominant feature is shape (narrow, tapered vertical streaks of variable intensity). Thresholding produces a binary image that includes any potential streak structure; morphological shape filters [9] are iteratively applied to identify extended streaks caused by rust or leaks. Color is used to differentiate between rust streaks and streaks in dust or dirt caused by condensation. Streak detection results (a segmented image) are sent to the analysis module for reporting (Figure 9).

Features corresponding to rust and corrosion are extracted from transformed input images. Microfeatures include statistics (mean and standard deviation) and texture (entropy and homogeneity) on small windows (5x5). Macrofeatures include

Table 2. Performance characteristics of structured light range sensor.

Parameter	Value	Description
Depth of Field	90 cm	Maximum Minus Minimum Range
Standoff Distance	70 cm	Minimum Range
Scan Rate	7/sec	Number Of Frames Processed Per Second
Range Resolution	0.5 cm	Smallest Change In Range That Sensor Can Report
Vertical Angular Field Of View	21°	
Vertical Angular Resolution	0.044°	Horizontal Resolution Determined By Scanning; In Phase I System Was 0.25° or 0.5°
Laser Rating	Class II (Low Power, Caution)	20 mW, 680 nm Is Normally Class III; However, Spread Out Into A Stripe Reduces Exit Power To 1mW over 7mm Aperture

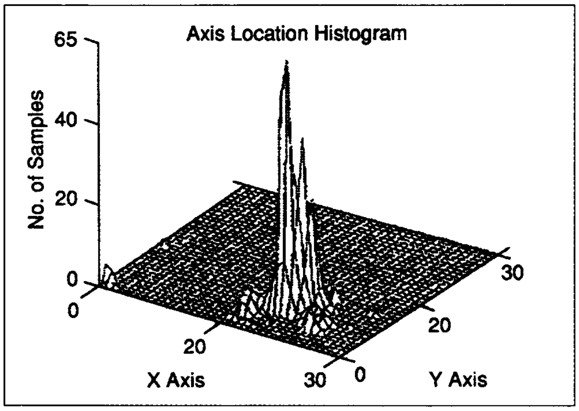

Figure 8. Actual cluster of axis points, for a drum with a 2° tilt; each array element represents 2 cm.

variance and singular value decomposition eigenvalues [8] on larger windows (15x15) from textured images. Feature values are combined into a feature vector (one vector per pixel). Segmented pixels are labeled according to class membership (rust, corrosion, labels, paint) using a supervised clustering algorithm with color-specific parameters (distance measures). It is relatively easy to hand-tune parameters at a computer console to detect rust or corrosion in a particular image. It is more difficult to determine parameter ranges for robust, autonomous inspection outside the laboratory. Hence, two supervised clustering methods for pattern recognition are being evaluated in Phase 2: a conventional k-means algorithm [10] and the Learning Vector Quantization neural network algorithm [11]. The result of feature extraction and classification is a segmented image with regions of interest labeled. The analysis module quantifies detected anomalies (rust, corrosion, missing paint, and streaks) for reporting.

Bar-Code Reading

Bar-code labels are used to identify and track waste storage drums. The bar code is the index into the site database, the local database, and the defect database. The site database contains information about the contents and history of each drum. The local database contains information to be used on the current inspection run including drum locations, drum size, and drum color. The defect database is where inspection results are stored.

For nominal weekly inspections, the bar-code label location on the drum is known from the world model. The vehicle is commanded to move the scanner into location for reading the label (by moving the pan-tilt platform or repositioning the vehicle if necessary). A "laser-on" command from the computer turns on the scanner and an ASCII string is returned on a successful read. The laser automatically turns off after a successful read, or after 1 second, whichever comes first. The bar-code software module is controlled by the system executive.

For new drums or anomalous situations, the label is located using shape analysis of thresholded images from the color camera. This is important if the bar code is not in the expected place on the drum, and therefore, would be missed by pointing the bar-code reader at the expected location. Although bar codes have a well-defined and predictable appearance (i.e., vertical black bars on white paper), drums can vary widely in appearance, thus complicating the problem of locating the bar codes. For example, background color, written markings, and other labels can be present and in a variable configuration.

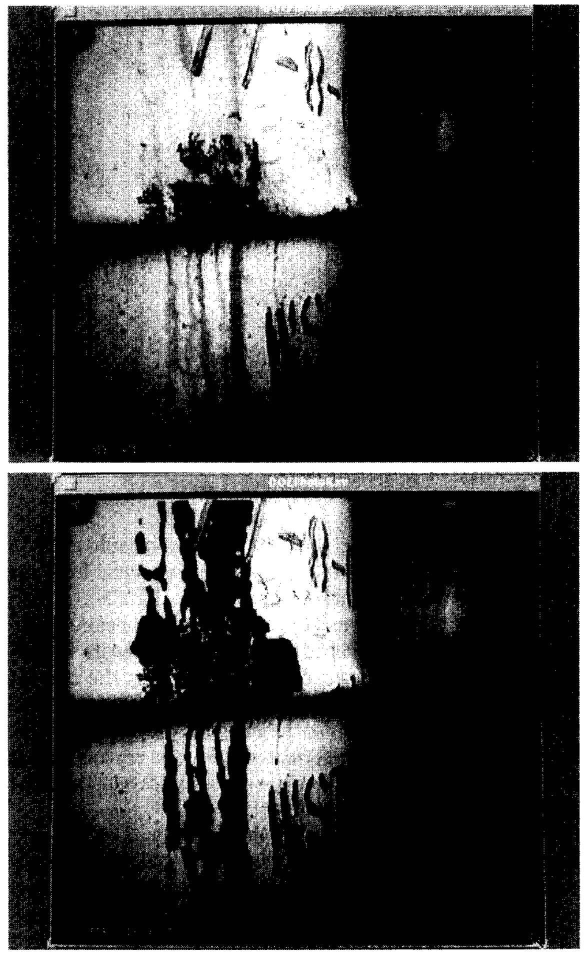

Figure 9. Storage drum with characteristic rust streak and processed image.

The label-locating algorithm consists of a two-stage process. The image is first processed by applying a horizontal gradient filter, which enhances black-and-white vertical edges, and adaptively thresholding the result. The image is then scanned for areas that have a high concentration of these vertical edges under the assumption that this is likely to be a bar code. Results show that the algorithm is reliable as long as the component vertical lines of the bar code can be resolved in the image.

For autonomous label scanning, the Phase 1 IMSS vehicle used a Symbol Technologies Laserscan 6120 visible laser diode scanner. This laser operates at 680 nm with 1.0-mW maximum power and is a Class IIa laser needing no special precautions, other than to avoid staring directly into the light beam. An RS-

232 cable connects the scanner to a UNIX workstation for computer control. The Phase 1 demonstration used bar-code labels patterned after the DOE standard (1-in. tall vertical bars, medium density, Code 3 of 9 symbols, with 10 alpha-numeric characters). Testing shows the scanner consistently reading labels from a 30-cm distance plus 25° in any direction.

OPERATOR INTERFACE

From the control building, the operator will work with an intelligent graphical user interface to initiate and review the inspection process for a number of storage facilities. To initiate the next leg of an inspection, the operator can confirm a preselected mission assignment or override and designate a different mission. The ability to assign the inspection of specific drums or an entire aisle is included. An intelligent planner generates the inspection sequence and allows the operator to preview the plan. Finally, the verified inspection plan is downloaded to the robot, with permission to depart granted after recharging.

Upon return, the robot's collected mission data are offloaded, the operator reviews the actual route of the robot against its assigned route and reviews data collected for any reported defects. A detailed inspection report is generated and printed, notifying the operator of defective drums that must be resolved before the next inspection cycle.

A database of site information is maintained and updated as appropriate after each inspection cycle. At the end of an inspection cycle, a full report of all inspections for each storage facility is compiled and printed. This includes quantitative data on drum defects and a compressed visual image of the defective drum. Other recorded information includes data from radiation and gas detectors, a history of the robot's path, and what portion of the mission may have been modified or aborted because of environmental constraints. Because drum defects are recorded in a database, the system is capable of tracking the condition of drums over a period of time. In addition, identification of incompatible waste storage is enabled.

The operator interface will be discussed in two parts-the controls interface and the database.

Control Interface: The IMSS operator interface allows the storage facility operator to direct and review the inspection process. The Phase 1 operator interface includes operational and developmental capabilities. Operational control functions are essentially task-level commands. Developmental capabilities include direct joystick control of the vehicle, pan/tilt, and individual sensors.

The vehicle control interface is based on a graphical touch screen interface that was developed under internal funding for robotic control applications [12]. The individual screens that support vehicle and sensor control functions are mission executive, facility layout, rover teleoperation, vehicle control, scan platform control, sensor menu, and data display menu. Command buttons that allow the operator to switch between menu screens appear on every screen. The button associated with the current screen is highlighted in color (violet) to identify the active screen. These buttons do not activate functions, they merely switch menu screens. An example of a typical control screen is shown in Figure 10. Note that the display supports multiwindow display of information from sensing subsystems (e.g., 3-D images [raw or processed] and visual images [processed or unprocessed]).

Database: The IMSS database consists of two major data sets-the facility model and the reporting database. The facility model contains information about the facility's fixed components-things that are not expected to change. Examples of facility model data include drum locations, aisle endpoints, obstacle locations and sizes, landmark locations, and basic drum geometry. Drum coordinate frame definitions and a plan view of the vehicle in the mockup storage facility. are also included and easily accessible. These data are used for mission planning and execution monitoring. The facility model is accessed by all other processes other than the vehicle controller. It is only updated under control of the operator interface to maintain integrity.

The reporting database contains information resulting from the inspection. Examples of reporting database data include missing drums, rust patches, and dents. Such defects are indexed by the drum from the facility model on which they were identified. The reporting database is updated by sensor processes and the mission executive. Data are read out to the mission executive and the operator interface.

PROJECT STATUS

We completed the Phase 1 engineering demonstration in June 1993 with a successful demonstration in our waste storage facility mockup. The test setup was composed of three rows of drums with two aisles, a back aisle, and a staging area. The outer rows were stacked two high. This mockup area was chosen because it had characteristics similar to many current DOE storage facilities, including dim lighting, a heavily seamed floor covered with a shiny sealer, and a semi-outdoor environment in terms of temperature, humidity, and dust. The Phase 1 vehicle was not sealed and will not be decontaminated.

Phase 1 Results

Phase 1 culminated in a series of parametric tests to measure sensor and vehicle performance. These data are being used in Phase 2 to build a robust, cost-effective prototype. The level of performance of the system relative to the functional requirements is shown in Table 3. A picture of the vehicle in the storage facility mockup is shown in Figure 11.

A statistical compilation of measured data can be used to determine detection rates. Desired detection rates for the Phase 1 system were greater than 90% detection with less than 20% false positive alarm rate. These are increased to 95%/10% for Phase 2 and 98%/5% for the commercial prototype. The measured detection rates are a function of selected feature size. For instance, 0.25-in. rust spots and 0.25-in. wide streaks were detected at a 98% rate with a false positive rate of 4%. Reducing the feature size to 0.1 in. resulted in a detection rate of 92% and a false positive rate of 14%. Features less than this size were beyond the resolution of the system. Tilt detection was 100% effective down to 2° and dent detection was 100% successful down to 1/2 in.

Phase 2 Progress

We recently completed the detailed design review of the Phase 2 vehicle and are working on system assembly. The new vehi-

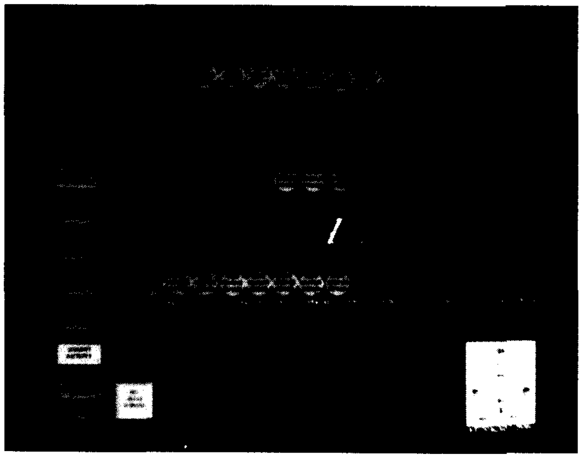

Figure 10. IMSS operator interface.

Figure 11. Phase 1 final demonstration.

cle will be able to inspect drums that are stacked three or four high, to safely navigate 26-in. wide aisles, and inspect 12,000 drums per week. The Phase 2 vehicle will not be tethered. We are anticipating a Phase 2 demonstration in early 1995 at the DOE Hanford Engineering Laboratory.

ACKNOWLEDGMENTS

This work was sponsored by the Department of Energy and the Morgantown Energy Technology Center under Contract DE-AC21-92MC29112. This project is grateful for the coordination and direction supplied by Gary Nelkin, the technical monitor. We appreciate the support of the Hanford Engineering Laboratory, the site of the Phase 2 demonstration, and the Westinghouse Savannah River Company's Waste Operations Research and Development Program. Engineering credit is due to the IMSS team: Tim Charles, Amy Geoffroy, Bob Glade, Brad Haack, Scott Mickey, Mat Ramey, Mark Roberts, John Spofford, and Chris Voth. The bar-code image recognition work was done at the University of Colorado at Denver by Professor William Wolfe, Jim Gunderson, and Matt Walworth.

REFERENCES

[1] J. M. Griesmeyer, M. J McDonald, R. Harriagan, P. L. Butler, and B. Rigdon, "Generic Intelligent System Control (GISC)," SAND92-2159, Sandia National Laboratories, Albuquerque, NM, September 1992.

[2] S. Price, W. Chun, M. Hammond, and A. Hubert, "Wheeled Planetary Rover Testbed," Proceedings of Mobile Robots V, SPIE Vol. 1388, Boston, MA, November 1990.

[3] S. Jonsson, "New AGV with Revolutionary Movement," Proceedings of the 3rd International Conference on Automated Guided Vehicle Systems, Stockholm, Sweden, October 1985.

[4] W. A. Hoff, L. B. Gatrell, and J. R. Spofford, "Machine Vision-Based Teleoperation Aid," Telematics and Informatics, Vol 8, No. 4, 1991, pp 403-423. Also in Proceedings of the 1991 Goddard Conference on Space Applications of Artificial Intelligence, NASA CP-3110, Greenbelt MD, May 1991, pp 199-213.

[5] L. B. Gatrell, W. A. Hoff, and C. W. Sklair, "Robust Image Features: Concentric Contrasting Circles and Their Image Extraction," Proceedings of Cooperative Intelligent Robotics in Space II, SPIE Vol.1612, Boston, MA, November 1991.

[6] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, Numerical Recipes In C, Cambridge University Press, 1988.

[7] Y. Ohta, T. Kanade, and T. Sakai, "Color Information for Region Segmentation," Computer Graphics and Image Processing, Vol 13, 1980, pp 222-241.

Table 3. System performance against requirements.

Drum Defects	**Requirements**	**Capabilities**
Round Or Pointed	Less Than 1 in.	Can Detect Depths Down To 1/2 in.
Superficial Rust	Identify Diameter Size & Increase In Size	Can Detect Diameter Size Down To 1/10 in.; Can Identify Size Increase
Bubble Paint (Corrosion)	Identify Diameter Size & Increase In Size	Can Detect Diameter Size Down To 1/10 in.: Can Identify Size Increase
Rust Streaks	Differentiate Between Water Streaks & Metal Rusting	Can Differentiate Between Water Streaks & Metal Rusting
Tilted Or Bulging Drums	Identify Tilt To Drum Or Bulge In Side	Can Detect Down To 2° Angle Of Tilt
Missing Or Misplaced Bar Codes Or Labels	Identify Missing Or Misplaced Bar Codes Or Labels	Can Identify Missing Bar Codes And Labels: If Misplaced, Identifies Them As Missing
Drum Displacement	Identify Missing Or Misplaced Drums	Can Identify Missing Drums; Can Identify Drum Relocation Down to 1 in.

[8] W. K. Pratt, Digital Image Processing, Second Edition, Wiley-Interscience, 1991.
[9] R. M. Haralick, S. R. Sternberg, and X. Zhuang, "Image Analysis Using Mathematical Morphology," IEEE Transactions of Pattern Analysis and Machine Intelligence, Vol PAMI-9, No. 4, July 1987.
[10] G. Coleman and H. Andrews, "Image Segmentation by Clustering," Proceedings of the IEEE, Vol 67, No. 5, May 1979.
[11] A. Visa, "A Texture Classifier Based on Neural Network Principles," Proceedings of the International Joint Conference on Neural Networks, Vol I, IEEE, 1990.
[12] M. K. Morgenthaler, G. Bruno, J. R. Spofford, R. G. Greunke, and L. B. Gatrell, "A Testbed for Teleautonomous Operation of Multiarmed Robotic Servicers in Space," Proceedings of Cooperative Intelligent Robotics in Space, SPIE Vol 1387, Boston, MA, November 1990.

ABOUT THE AUTHORS

Eric Byler received the B.S. Degree in Ocean Engineering from the Massachusetts Institute of Technology in 1980, and the M.S. in Mechanical Engineering from Stanford University in 1984. Between 1980 and 1984 he developed and integrated CAD design systems for marine applications for Advanced Marine Enterprises, Inc. From 1984 to 1989 he worked for Grumman Space Systems in their Robotics Laboratory in the field of space robotics. Since 1989, Mr. Byler has been as Senior GROUP Engineer at Martin Marietta Astronautics in Denver working on robotics for space applications and environmental remediation applications. He is currently program manager and lead engineer for the Intelligent Mobile Sensing System program.

Wendell Chun received a B.S. in mechanical engineering in 1978 from the University of Hawaii and completed the Carnegie Bosch Program in Engineering Design in 1992 in support of the Mount Erebus robot project to Antarctica. He is currently employed by Martin Marietta Astronautics as a staff engineer in the robotics and automation research group in Denver. For the past eight years, he has co-chaired the SPIE Mobile Robots conference with William J. Wolfe. He was the chief engineer of the Walking Beam project, a seven-legged frame walker that was studied for Mars exploration. Mr. Chun has held lead positions on the NASA Flight Telerobotic Servicer, NASA Satellite Servicing System, ARPA Unmanned Ground Vehicle (UGV) Demo II, U.S. Army Robotics Locating System (RLS) and the DOD Intelligent Task Automation (ITA) programs while at Martin Marietta.

William A. Hoff received the B.S. in physics from the Illinois Institute of Technology, Chicago, in 1978, and the M.S. in physics and the Ph.D. in computer science from the University of Illinois, Urbana-Champaign, in 1981 and 1987 respectively. From 1978 to 1980 he was an engineer at Teletype Corporation, Skokie, Illinois. From 1988 to 1994 he was a staff engineer at Martin Marietta Corporation, where he worked on computer vision systems for space and robotics applications. He is currently an assistant professor in the Engineering Division at the Colorado School of Mines. He is an Associate Editor of Pattern Recognition and a member of the IEEE Computer Society. His research interests include computer vision, robotics, supervisory control, and virtual reality.

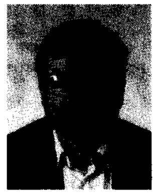

J. Daniel Layne received the B.S. in Mathematics from the University of Southern Colorado in 1975, and the M.S. from Colorado State University in 1977. As a Peace Corps volunteer he then taught for two years at the Catholic University in Guayaquil, Ecuador. He then worked at Ford Aerospace for six years in astrodynamics and software engineering, and at CTA for two years in systems engineering. As a Staff Engineer at Martin Marietta since 1988 he has conducted research and development in artificial neural networks, and developed several image processing applications., including the IMSS color vision system. He has also continued post-graduate studies at the University of Colorado. His research interests include intelligent systems, adaptive control, computer vision, and parallel processing. He has published a dozen technical papers in computation mathematics and neural networks. He is a member of the IEEE Robotics and Automation Society and the International Neural Network Society.