A System for Telerobotic Control of Servicing Tasks in a Nuclear Steam Generator

R. Craig Coulter, Anthony J. Stentz, Paul G. Keller, Gary K. Shaffer,
William L. Whittaker, Barry L. Brummit, William Burky

CMU-RI-TR-90-24

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

December 1990

Table of Contents

Abstract

Many nuclear servicing tasks fall into the tool insertion class. These tasks have principally been accomplished through force and torque feedback paradigms; the peg in the hole problem is the classic example. However, many manipulation problems, such as servicing tasks, require automated movement and/or alignment *before* contact is made. Torque and force feedback are insufficient for providing information to locate the goal, to align the manipulator, and to prevent undesired collisions.

In this work, vision and range sensors are appended to manipulation to enable high accuracy tool insertion, and are used to construct a highly accurate model of the space with which the manipulator interacts. Direct viewing of this model provides a superior man-machine interface; the operator has immediate and intuitive visual confirmation of the manipulator position and configuration.

Once developed, the technical ideas were successfully applied to a current problem of interest: the telerobotic inspection of nuclear steam generators. A complete telerobotic system was developed, built and tested and performed a docking task commonly seen in inspection.

1.0 Introduction

Tool insertion manipulator control has principally been achieved through the use of force and torque feedback control methods. However, many manipulation problems require automated movement and/or alignment *before* contact is made in order to minimize contact or collision forces. Such problems include mating of damageable parts and insertion of damageable tools. Torque and force feedback methods require contact; obstacles can only be detected once a collision has taken place. Torque and force signals can certainly be fed back to the operator, but usefulness is limited because feedback takes place only after contact is made. Alignment might also be accomplished by first making an approximate alignment and then using force-torque feedback; however, this method relies upon a good *a priori* model of the environment and does not offer the robust ability to sense environmental deviations.

In the nuclear field, the servicing of equipment poses significant challenges. The presence of high-level radiation makes human intervention a costly and dangerous undertaking. Teleoperated mechanisms have been used in these environments, preventing human exposure and cutting service costs; however, teleoperation introduces inherent operational difficulties. By removing the human from the work environment, the operator's sensory feedback - most importantly visual - is greatly reduced. Cameras afford only two dimensional views of part of the environment and often insufficient views of the manipulator, making tool insertion manipulation very difficult to achieve. The automation problem is twofold. First, there is the challenge of performing the manipulation task: a system must be developed which is capable of perceiving the state of the environment. Current manipulators are precise enough to perform the physical task, given proper perception and guidance. Second, there is the problem of providing the operator with system information at the appropriate level. The machine performs operations in a remote enclosed environment, thus the operator needs the ability to determine quickly and completely the configuration of the machine and the state of its surroundings. Additional constraints are provided by the need to keep the system simple, robust, and completely compatible with the environment.

2.0 Prior Work

The principal problems that challenged the development of our system are:

- Aligning the manipulator without contact.
- Performing a moderately tight insertion.
- Reconstructing the 3-D environment for the user.
- Sensing appropriate aspects of the environment.

Much work has been done in robotics toward the development of automated manipulation systems and improved man-machine interfaces. The previous work cited here is certainly not all inclusive, but attempts to provide a sample of the technology that influenced our work.

Vision assisted manipulation work performed by Mitchell, Mason and Christiansen [8] focussed on a planning/learning system for manipulation. An arm was given a manipulation task such as moving a block along a wall. A vision system detected the location and orientation of the block in the scene and planning algorithms commanded manipulator moves to achieve the goal.The robot refined its understanding of the manipulation task through its failures. Ikeuchi [5] used a model based vision system to perform bin picking tasks. His work involved interpreting CAD model representations of objects to determine grasping positions. The vision assisted manipulation work demonstrates the usefulness of vision, but was developed for the identification and manipulation of polygonal objects. As will be seen later, insertion manipulation requires precise vision measurement capabilities.

Allen and Bajcsy [1] performed both contact and noncontact sensing of three dimensional objects and proposed the necessity of multisensor integration to achieve reliable reconstruction information. Tactile feedback augmented a stereo vision system to determine object shape. The work demonstrated the advantages of combining multisensor feedback to achieve better reconstruction of physical phenomena.

Probably the best understood method for achieving high accuracy insertion is the force and torque feedback control method. An assembly task is achieved by bringing the peg into contact with the receptor hole, and attempting to zero the forces and torques acting on the peg. The remote center of compliance method [3] is one such example; however, this approach is ineffective because of the inability to locate the insertion goal with sufficient accuracy, necessitating pre-contact alignment.

Sheridan [13] describes the importance of the human/machine interaction in the execution of a robotic task. Sheridan points out that the necessary use of a remotely operated machine introduces some serious problems in providing sufficient sensory information to the operator. He further points out that artificial sensing, intelligence and control can augment a human's remote control capabilities. The key is to allow the robot to perform those tasks for which it is best suited - sensing and precise manipulation, while giving the human the necessary environmental feedback to make good supervisory decisions.

Many groups have proposed methods for reconstructing depth from 2-D images, or providing depth cues from graphics overlays. The Virtual Display Control Interface [6], developed at NASA is one such device. VDCI features a helmet fitted with stereoscopic glasses. Japan's Electrotechnical Laboratory developed the Multi Media Display (MMD), which as part of its capabilities provides additional depth cues by overlaying graphics onto monitor displays, and can provide a 3-D stereo vision effect similar to the VDCI. These extensive capabilities were found to be unnecessary.

Japan's Electrotechnical Laboratory built a telerobotic system [7] which performed manipulation tasks based upon a geometric model, and used the MMD as its user interface. The system demonstrated the ability to remotely assemble and disassemble mechanical components. This work represents the type of system that is necessary for remote inspection. It combines the capabilities of robotic automation with a man-machine interface that gives the user the ability to understand the machine's configuration and surroundings.

Mitsubishi [9] has developed a series of platforms and manipulators which are sold as commercial products. The company offers a variety of automated inspection equipment including rail mounted sensor platforms that inspect for steam and/or radiation leaks, wheeled and tracked mobile inspection and maintenance platforms with manipulation capabilities. These robots have a variety of automated capabilities, some of the mobile platforms have autonomous path following capabilities using vision processing; however, the majority of the systems are teleoperated machines, and none met the need for remote automated inspection.

The state of technology described here demonstrates that there exists no simple solution for low-contact close quarters manipulation that also provides the user with enough information to comfortably assess the state of the machine; however, there is precedence for such a system in the states of component technologies. Each of the above systems demonstrates a technology that is necessary, but no system exists that has successfully integrated these components to perform the task.

3.0 Problem Overview

The development of a method for automated tool insertion was driven by the nuclear industry's desire to automatically inspect nuclear steam generators. A steam generator (See Fig. 1) transfers heat from the radioactive primary loop to the clean secondary loop without mixing the working fluids. Physically, a typical steam generator consists of a 10 foot diameter hemispherical bowl divided in half by a vertical wall, and capped with a 24" thick metal sheet called a tubesheet. The tubesheet is perforated with 3/4" holes spaced 1.25 inches apart in a two dimensional grid. A tall stack of U shaped tubes are pressed into the tubesheet, connecting each hole on one side of the channelhead to a hole on the other side. Water from the reactor is pumped into one side of the bowl, forced through these tubes, collected in the other half of the channelhead, and pumped back to the reactor for reheating. Secondary loop water surrounds the outside of the tubes where it is boiled to pressurized steam to drive the turbine.

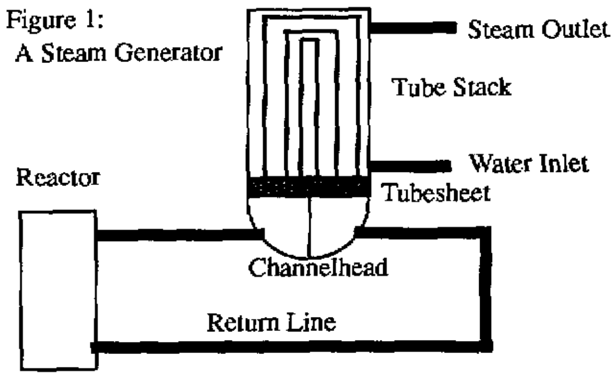

If a crack develops in any of the tubes, radioactive water can leak into the secondary loop and contaminate the turbine. The turbine and generator must be shut down, the leak repaired and decontamination performed -a very costly process, on the order of $500,000 per day of downtime. Because of the extreme cost, preventive maintenance must be performed.

Robotic steam generator servicing (See Fig. 2) with a six degree of freedom robotic arm involves inserting the base through a manway door and docking the base to the tubesheet. After the base has been secured, the remaining arm links are drawn completely into the channelhead. The arm is then in a position which allows its end effector to reach tubesheet hole for inspection and/or repair tasks. The arm is controlled using a joystick and visual feedback from cameras placed on the manipulator and in the channelhead.

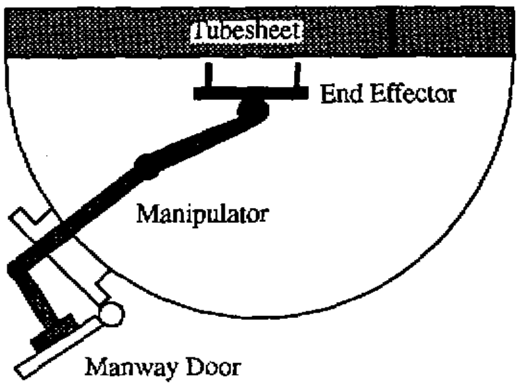

Figure 2: The Channelhead

The desire to automate comes from the tedium of the task. A typical tubesheet has hundreds of holes to be repaired; visually locating the proper set, and performing a teleoperated insertion is slow, cumbersome, and difficult. Automation would not only alleviate operator boredom, but also eliminates operator error.

The docking task and several of the inspection tasks involve variations of inserting a probe or camlock into one or more of the tubes. We chose to focus on the docking task, because it requires the simultaneous insertion of four pegs into four holes. The docking problem is formulated as follows. The arm operates in a world that contains a single plane of holes equally spaced in a 2-D grid (the tubesheet). The operator specifies the docking holes, the robot finds the holes and simultaneously insert four pegs into the four holes. The system must therefore be able to correctly determine the position and orientation of the holes relative to the manipulator, display that information in a meaningful manner to the user and automatically perform the insertion.

4.0 System Overview

The Autonomous Docking System is able to:

- Build an accurate model of a tubesheet.
- Display the arm configuration and environment as a CAD model.
- Calibrate the system camera for high accuracy.
- Form a composite estimate of tubesheet parameters based on numerous readings.
- Reason about the environment based on incomplete sensor readings.
- Generate a set of motion commands to the arm to insert the pegs into the holes.

If the precise location of all of the holes were known and the precise attitude and height of the tubesheet were known, then the arm could be servoed without feedback into a docking position. (pre-supposes an accurate arm.) Because of discrepancies between the manufacturer's specifications and the actual tubesheet, there is uncertainty in the positions of the holes and attitude of the tubesheet. Sensors are employed to determine the position and attitude of the tubesheet and to determine the position of the holes in the tubesheet. To obtain useful estimates from sensors they must be properly calibrated and the readings must be properly filtered. In addition to forming some estimate of the tubesheet parameters, the system takes corrective action when it encounters unexpected readings.

A model describing the position and orientation of the tubesheet, and the location of individual holes within the tubesheet is corrected with sensor updates. The goal is to make the model accurate enough to permit open loop positioning of the arm without collision. For reasons of clarity and safeguarding, it is also desirable to represent the arm and its environment in a manner that permits the operator to quickly extract three dimensional spatial representations. A three dimensional simulation tool that stores accurate model information, can be updated as sensor information becomes available, and can be accessed to provide arm positioning information, meets all of these requirements.

A CAD-based user interface with graphical display demonstrates superior performance to standard camera views. CAD is intuitive; it gives the user the ability to change the viewing angle and to zoom in on interesting features of the model, affording views not available through traditional camera-based systems. Multiple views obtained from panning about in the scene gives the user a pseudo 3D perspective. By using the CAD model as a predictive aid, the user can run the arm model through any number of potentially dangerous maneuvers in simulation and check the trajectory for collisions.

By combining the features of a CAD system with a sensor equipped arm, the robotic system effectively "learns" about its environment. At some point, the law of diminishing returns makes it impractical to continue model updates; the error between model and reality becomes vanishingly small and additional measurements are no longer worth the time that is takes to make them. One must also note that at some point the system is no longer able to return improvements due to its own limitations in calibration and repeatability.

5.0 Simulation Environment

A principal concern of this work is to use the model information to provide the operator with an effective user interface. A safe and successful mission is dependant upon the operator's ability to understand, quickly and completely, the configuration of the arm and its location with respect to its surroundings. The ability of a single video camera to render this 3-dimensional information is extremely limited.

CAD gives the user an intuitive visual model that can be updated as environmental information is obtained. The model affords some views that could not otherwise be obtained. The user can pan and zoom throughout the model, moving the point of view to any location and perspective desired. We chose a CAD system that allows a fully accurate kinematic representation of the mechanism, accurate modelling of the environment, and the ability to check for collisions between the mechanism and the environment. The system provides collision checking by running the proposed motion in kinematic simulation. If no collision occurs, the motion is approved for use, and the arm can be moved. This is a very important capability for manipulation in tight tolerance scenarios.

6.0 Sensing System

In the environment model each hole has five degrees of freedom. As a group, the holes lie at some attitude in a plane which is located some distance from the base frame of the manipulator. Within the plane, each hole has two degrees of freedom. To measure the hole locations with respect to these five degrees of freedom, two different sensors are employed.(Fig. 3) Three piezo-electric sensors return a triad of ranges which describe the attitude and height of the tubesheet with respect to the base frame. A single camera vision system determines the location of the holes in the plane of the tubesheet.

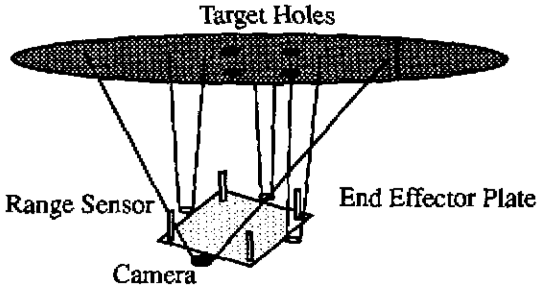

Figure 3: Sensor Configuration

6.1 Range Sensing

To describe the attitude of the tubesheet, the rotations of the sheet about the X and Y axis are calculated from the tubesheet normal (Fig. 4), and transformed to RPY notation. Three piezoelectric range sensors are used to determine the height and plane of the tubesheet. Piezoelectric sensors were found to work well because of their fairly narrow beam (10 deg.). The particular sensor used had a range of 4 to 24 inches, and a measured accuracy of about 0.01 inches.

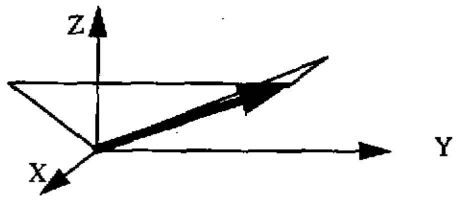

Figure 4: Tubesheet Normal Vector

6.1.1 Error Estimation

Each of the range sensors has an associated inaccuracy which imparts error to the range measurements and consequently to the tubesheet parameters. Euclidean geometry can be used to transform the range readings to tubesheet parameters; the transformation of error is substantially more difficult. Smith-Cheeseman[14] proposed a method for transforming Gaussian uncertainty in geometry. Assume that each reading is a random variable with some mean and variance. The error in the sensor must be transformed to error in the model parameters.The statistical mean and variance are used to describe a current reading and its associated uncertainty; the uncertainty of a set of readings can be described by a variance-covariance matrix.

$$\Lambda_i = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12}^2 & \sigma_{13}^2 \\ \sigma_{21}^2 & \sigma_{22}^2 & \sigma_{23}^2 \\ \sigma_{31}^2 & \sigma_{32}^2 & \sigma_{33}^2 \end{bmatrix}$$

where

$$\sigma^2_{ij}$$

is the variance of a single reading for i=j, and the covariance of reading i with reading j otherwise.

If the error is propagated through a kinematic chain, the transformation between the variance-covariance matrices in the two frames is described by:

$${}^j\Lambda_x = {}^jJ_i \, {}^i\Lambda_x \, {}^jJ_i^T$$

Where J is the manipulator Jacobian, and i, and j denote two general frames of reference. Once the error matrix is transformed to the proper frame, it must be transformed from sensor reading variance to parameter variance. This transformation is achieved using the same method, with a Jacobian describing the sensor to parameter transformation.

Range measurements of the tubesheet are taken from many locations, and at many different orientations during manipulation, requiring a robust method for determining the best possible estimate of the tubesheet parameters, while reducing the total parameter uncertainty. The Kalman filter is a well known technique for merging multiple measurements with unbiased random error. The Kalman filter is appropriately suited to error estimation in this application because a number of discrete readings are taken, a consensus estimate is desired, and a method for dealing with sensor drop-out is desired. The Kalman filter updates the current estimate by performing a weighted average of the current estimate and the new sensor readings; the readings with the lower associated uncertainty are weighted more.

6.2 Vision Sensing

Once range sensing has determined the attitude of the plane of holes in free space, the location of each hole must be determined within that plane. Computer vision techniques are used to find the location of the center of each hole. In this application, vision processing consists of the following sequence of steps:

- Digitize an image.
- Find the edges points in the image.
- Map the edge points from image to scene coordinates.
- Locate holes among the edge points.
- Update the CAD model with the new data.

Sharp discontinuities in light intensity in the digitized image are called edge points. An edge detector locates candidate edge points and assigns a confidence that they compose an edge. A higher value indicates a greater chance that an edge exists at that point. Edge detectors have two important properties: the ability to discriminate between edge points and noise, and the ability to locate the exact position of an edge. The Canny [2] operator was chosen because it optimizes the trade-offs between these two competing metrics. Our implementation demonstrated best performance using a window size of 12 -15 pixels.

A threshold is empirically set to account for lighting conditions (constant in our application) and noise, and to filter out weak edges. The remaining points are run through a calibration formula which maps from image coordinates to robot world coordinates. (See Camera Calibration, Sec 6.3)

The general Hough [4] transform detects a two-dimensional shape in the edge data. The Hough transform is based on a voting scheme. Each edge point "votes" for every circle to which it could belong. Since a range of circle radii is known *apriori*, the point must belong to a circle whose center lies one radius away. The family of all of these circle centers describes a circle (of the known radius) around the given point. Each of these circle centers is given one vote. After iterating over all of the edge points, the centers which receive the most votes are likely to be hole centers. The effects of random error in the hole edge points are reduced through multiple readings. That is, the accuracy of the calculated result is greater than the accuracy of each individual measurement. The Hough transform was chosen over other algorithms because it is robust to noise. Circles with gaps or outlying pixels receive considerably more votes than a random collection of noise points.

After the centers of the image holes have been calculated, the CAD model is updated. A nearest neighbor algorithm matches the holes found in the image to the hole locations predicted by the model. The new locations are sent back to the model for use in future operations.

6.3 Camera Calibration

The camera and its associated electronics introduce lens distortion and electronic timing offsets or uncertainties to the image; high accuracy measurements require the determination of the characteristic camera and error parameters. The use of a camera as a metric device also requires that its location and orientation in free space be determined. Tsai [15] proposed a camera calibration technique which addresses both problems. The parameters which describe the location and orientation of the camera in space are called the *extrinsic parameters*; those which describe the internal camera characteristics are called the *intrinsic parameters*. The six extrinsic parameters are the three translational and three rotational degrees of freedom of the camera. The six intrinsic parameters which serve to characterize the camera model are the focal length, two distortion coefficients, the computer image coordinates for the origin, and the timing uncertainty factor. The focal length is calculated because the focal length supplied by manufacturer's data is not precise enough for accurate measurements. The two distortion coefficients are the coefficients from the first two terms in the Taylor's series expansion which describes radial distortion. Tsai found in his experiments that tangential distortion was not significant enough to warrant computation. The uncertainty scale factor describes the hardware timing uncertainty between the sensing equipment and the image acquisition hardware. A small error (1%) in timing can cause a three to five pixel shift. The strength of this calibration technique lies in the validity of the camera model. The characterization of nonlinear distortion is especially critical to the accurate measurement of object characteristics.

6.3.1 Calibration Algorithm

The mapping of an object from world to computer image coordinates takes place in four steps, each of which has an associated set of calibration parameters. The coordinates of the computer image center are not computed; they are given in manufacturers specifications.

1) Calculate the homogeneous transformation matrices describing frame translation and frame rotation must to transform the image from the object world coordinate system to the camera 3-D coordinate system.

2) Using a pinhole camera model, calculate the camera focal length. This describes the transformation from the 3-D camera coordinate system to the ideal image plane.

3) Calculate the two primary radial distortion coefficients to describe the transformation from the ideal image plane to the distorted image plane.

4) Compute the uncertainty scale factor to describe the transformation from the distorted image plane to computer image coordinates.

Tsai's algorithm allowed calibration to about 30 mils using two planes of twelve points each.

7.0 Planning End Effector Motion

Gross range errors occur when a range sensor is aligned with a hole; i.e. the "footprint" of a sensor overlaps a hole in the tubesheet. The end effector must be moved to the nearest configuration which enables all three range sensors to collect accurate data which will be used to calculate the attitude and height of the tubesheet. This section describes an algorithm for finding such a configuration.

7.1 The Search for a Valid Configuration

The geometry of the sensor layout is compared to the geometry of the tubesheet model to determine whether a sensor footprint overlaps a hole. To make the problem tractable, it is assumed that the plate of the end effector is parallel to the plane of the tubesheet; this reduces the problem from 3D to 2D since only the x, y, and yaw of the end effector are degrees of freedom. We assume hole positions are known accurately enough for a topological search, and that hole positions are known to within half a center to center distance.

The problem is formulated as a search for the nearest end effector configuration (x, y, θ) that causes all three range sensors' footprints to fall "sufficiently" on planar areas of the tubesheet. "Sufficiently" is dependent upon the capabilities of the sensor, and was empirically determined to be about 50%.

The search space is the set of all possible configurations of the end effector. Since the configuration has three components (x, y, θ), the configuration space is three dimensional. A grid structure is imposed on the configuration space in order to tessellate the space into a set of nodes to be searched. To avoid the space efficiency problem incurred by storing every grid cell in memory, a hierarchical data structure (an octree [12]) is used to represent the configuration space. Only the grid cells that are encountered during the search are allocated in memory.

The A* search algorithm [10] is employed to find a sequence of nodes in the search space that connects start node to goal node while minimizing a cost. For this manipulation task, the cost is the length of the end effector path from start to goal; cost minimization yields the shortest path. Since the end effector moves the smallest distance possible, the cycle time of the system is kept to a minimum. The heuristic evaluation function of the A* search is set to zero; the search is reduced to breadth-first search.

The search terminates when a goal node is encountered, necessitating a goal-node test. Recall that the search space (configuration space) is three dimensional; each node (voxel) is a cube that is bounded by minimum and maximum x, y, and θ. A goal-node test is implemented that finds upper and lower bounds on the distances from the range sensor footprints' centers to each hole center in the tubesheet, given that the end effector lies somewhere in the current node being searched. If and only if *all three* of the range sensor footprints' centers are *definitely* outside *all* of the holes, then the current node is a goal node.

7.2 Obstacles

The previous subsection describes the basic algorithm for finding a configuration of the end effector that provides valid range readings. It does not consider obstacles in the environment. This additional functionality is easily implemented.

Obstacles are represented by a list of line segments; the end effector is represented by a polygon. In order to ensure that the end effector does not collide with an obstacle, an additional test must be performed on each search node. This new test (the "admissibility" test) uses the SUP-INF method to determine whether the end effector polygon intersects any of the obstacle line segments, for the end effector located somewhere in the current node. The result of the search is a sequence of admissible configurations that connects the start configuration to a valid goal configuration.

7.3 Fixed vs. Variable Yaw

Since the cost function is dependent strictly upon path length, the current implementation favors paths that only rotate the end effector, with no translation. In some cases, it might be desirable to constrain the end effector to remain at a fixed yaw while traveling from start to goal. Restricting the search to an X-Y slice of configuration space achieves this goal.

7.4 Examples

Figure 5 shows the motion of the square end effector as it moves up and to the right from an initial configuration to a valid configuration, with fixed yaw. The vertices of the triangle are the locations of the range sensors. An L-shaped obstacle is shown in the upper left of the figure. The same scenario is shown in Figure 6 with the fixed yaw constraint removed to allow the end effector to rotate counter-clockwise to a valid configuration. In both figures, the dots in the interior of the triangle show the nodes of the search space that were explored during the search. Note that the reference point of the end effector is at the center of the square.

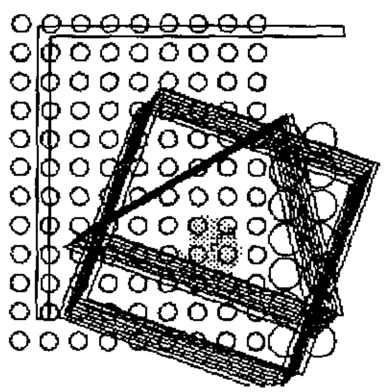

Figure 5.

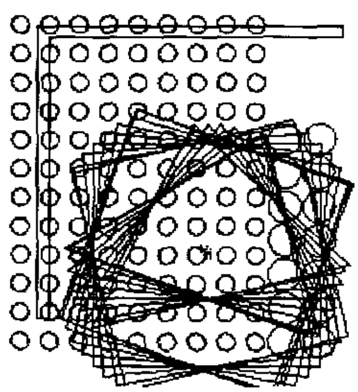

Figure 6.

8.0 Autonomous Docking System

The previous sections describe each of the components necessary to build an autonomous manipulator controller for insertion assembly. It has been stated that by updating a CAD model with filtered sensor data a rendering of the manipulator's environment can be achieved with sufficient accuracy to permit manipulation without feedback. To support that claim, a system was developed that automatically performs the docking task required for nuclear inspection. The docking task is the most challenging of the tool insertion tasks because it requires the simultaneous insertion of four pegs into four holes with about a tenth of an inch tolerance on the diameter.

In our test scenario, a four-pronged plate was mounted on the end effector of a six degree of freedom manipulator. Each of the prongs was lathed to the outside dimensions of a standard camlock found in the nuclear inspection industry: approximately four inches long, the last two inches of which taper by 2 degrees. The largest diameter is 0.70 inches. The tip was chamfered at 45 degrees with a blunt end. Each prong has natural compliance associated with this taper. The receptor holes were each drilled to 0.8125 in.

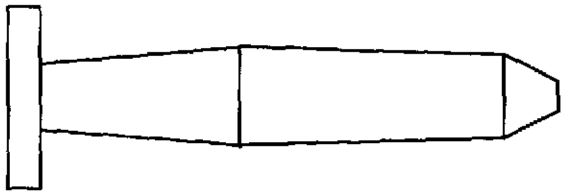

Figure 7: A Camlock.

An American Robot Merlin arm was chosen as the manipulator. The Merlin is a gear driven arm with very low associated compliance, most system compliance is derived from the peg chamfers making the insertion very rigid and unforgiving.

The operator sits at a CAD workstation and views a 3-D CAD model of the Merlin arm and the mock-up environment, with a control window to serve as an interface. The operator can select from the following commands: *calibrate*, *autodock*, *dock* and *move*.

- *Calibrate* performs the camera calibration, leading the operator through the procedure via a number of interactive commands.
- *Autodock* is the fully autonomous docking procedure. The system explores the environment determining the tubesheet parameters, and locates the holes within the plane of the tubesheet. The CAD model is then updated and the peg insertion is performed.
- *Dock* is a peg insertion performed without any sensor updates. A dock is usually performed to demonstrate that once the CAD model is updated with correct information the system need not recheck the environment. It can rely on the model to perform open loop manipulation, using the CAD simulation to check the proposed motion for collision.
- *Move* performs a guarded arm move by running the proposed trajectory in simulation, and then, if no collisions occur, sending a move command to the arm.

The autodock uses a triad of piezoelectric range sensors to measure the three tubesheet parameters, which are filtered to determine a best estimate. If errors associated with range sensing up a hole are detected the system reasons about the readings and moves the end effector to a location where good readings are more probable. Once the tubesheet parameters are calculated the update is sent to the CAD model. Using the updated CAD model, the controller finds the initial estimate of the hole locations and positions the camera under the first hole. An image is digitized and processed, and the updated hole locations are sent to the CAD model for update. The image processing is repeated for each of the remaining holes. When all of the holes are located, the CAD model is queried for the updated locations, and a series of manipulator waypoints are computed. The system runs a trajectory, based on these waypoints, through the CAD simulator; if no collision occurs in simulation, the arm is commanded to follow the trajectory. If a collision does occur the move is aborted and the operator is notified.

9.0 Results

The Autonomous Docking System (ADS) has been demonstrated numerous times, consistently performing docks. The system demonstrates a method for telerobotically performing tool insertion tasks in a remote environment. The ADS also demonstrates the advantages of CAD display as a man-machine interface.

It is important to verify the congruence of our testbed to the actual steam generator environment. The three major concerns are that the tubesheet closely resembles a real tubesheet, that the arm closely simulates the real arm, and that the holes closely simulate real holes of the type found in industry.

Overall we believe that the mock-up environment presented a more difficult problem than that found in the field.The tubesheet plane bowed about 2 degrees; the obvious problem with this considerable bow is that at different positions the range sensors were measuring different phenomena. The system expected to be measuring the attitude of a plane in space, but actually was measuring the local attitudes of a curved surface in space. The arm that we used was an American Robot Merlin. The Merlin has very little compliance compared to the manipulator for which this system is intended. Because of the bow in the tubesheet and the rigidity of the arm, we overdrilled the holes by 1/16th of an inch. Many of the hole edges were frayed and splintered. Even so, the vision routines were able to distinguish them.

Tsai's calibration algorithm allowed us to achieve more than sufficient accuracy for our application needs. We were able to calibrate the camera to an accuracy of about 0.030 inches at 6 inches. Our conclusion is that the calibration's limiting property seems to be the modelling of radial lens distortion. We have since looked into other calibration routines, but have not implemented any.

The man-machine interface was one of the most useful results of this work. The CAD model allowed us to operate the arm remotely while presenting us with a 3-D image of the arm and its environment. One of the best features of the CAD model was the ability to pan and zoom.

10.0 Future Work

Future work in this area includes:

- Expanding the system's servicing capabilities.
- Addition of force and torque feedback methods.

The current system can be easily expanded to allow such capabilities as tubesheet inspection. The end effector would visit each hole and determine whether it is plugged or sleeved, and inspect the hole for cracks. The results of the inspection could then be stored away for future reference.

The addition of force and torque feedback methods can greatly improve the system's ability to mate with tight tolerance holes. The strength of this system is its ability to perform precontact alignment and inspection. Vision feedback has a larger resolution than force and torque methods and once insertion begins is difficult to use. A system using both the vision methods and the feedback methods would use the strength of each to solve a broad range of tough insertion problems.

11.0 Conclusions

This work demonstrates that the current state of robotics technology is capable of supporting autonomous work systems in real environments performing useful tasks. The immediate use for this technology is the autonomous inspection of nuclear steam generators, but the long term applicability is far reaching. The techniques which drive the work are very generally applicable. Any system which performs close order manipulation tasks with minimal contact in an environment not easily viewable by human operators could derive benefit from this work. Tasks such as assembly in space or construction tasks in remote or hazardous areas are candidates.

12.0 References:

[1] Allen, P. and Bajcsy, R. , "Two Sensors are Better than One: Examples of Integration of Vision and Touch", U. Penn MS-CIS-85-29. September 1985

[2] Canny, John. "A Computational Approach to Edge Detection", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-8 No. 6. November 1986.

[3] Drake, S. "Using Compliance in Lieu of Sensory Feedback for Automatic Assembly", Ph.D. Thesis, Mech Eng. Dept. M.I.T., September 1977.

[4] Hough, P.V.C. "Methods and Means for Recognizing Complex Patterns" U.S. Patent 3,069,654. 1962.

[5] Ikeuchi, K., "Generating an Interpretation Tree from a CAD Model to Represent Object Configurations for Bin-Picking Tasks" CMU CS Dept Technical Report, CMU-CS-86-144, 1986

[6] Kim, W. and Stark, L. "Cooperative Control of Visual Displays for Telemanipulation" unknown source.

[7] Matsui, T. and Tsukamoto, M. "An Integrated Robot Teleoperation Method Using Multi-Media Display", Electrotechnical Laboratory. Umezono, Tsukuba-city, Ibaraki, 305 Japan

[8] Mitchell, T., Mason, M., and Christiansen, A. "Toward a Learning Robot", CMU CS Department Technical Report CMU-CS-89-106. Jan. 1989

[9] Mitsubishi Heavy Industries, LTD. Company Brochures. 1990.

[10] Nilsson, N. "Problem Solving Methods in Artificial Intelligence",McGraw-Hill. 1971.

[11] Okano, H. Miyazawa, T. and Shimada, H. "Intelligent Robots for Nuclear Plants", Toshiba Corporation. 9th International Conference on SMIRT., August 1987.

[12] Samet, H. and Webber, R. "Hierarchial Data Structures and Algorithms for Computer Graphics, Part I: Fundamentals", IEEE Computer Graphics and Applications. May, 1988.

[13] Sheridan, T.B. "Telerobotics" Automatica Vol 25, No. 4 pp. 487-507. 1989

[14] Smith, R., Self, M., Cheeseman, P., "Estimating Uncertain Spatial Relationships in Robotics", AAAI Workshop on Uncertainty, 1986.

[15] Tsai, R., "A Versatile Camera Calibration Technique for High Accuracy 3D Machine Vision Metrology using Off-the-Shelf TV Cameras and Lenses", IBM Technical Report RC11413. Sept, 1985.

A System for Telerobotic Control of Servicing Tasks in a Nuclear Steam Generator

R. Craig Coulter, Anthony J. Stentz, Paul G. Keller, Gary K. Shaffer,

William L. Whittaker, Barry L. Brumitt, William Burky

CMU-RI-TR-90-24

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
December 1990