Testing and Control of a Compliant Wrist

Mark R. Cutkosky, John M. Jourdain and Paul K. Wright

CMU-RI-TR-84-4

The Robotics Institute
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

March 1984

Abstract

This report summarizes the evaluation of a controllable, instrumented compliant wrist. The design of the wrist and the concept of a remote center of compliance (RCC) have been described in an earlier report, (CMU-RI-TR-82-9). The wrist is mounted on a large industrial robot used to load precisely machined parts into jigs and fixtures on computer controlled machine tools. The robot and the machine tools form part of an automated cell in which machined parts can be produced with a minimum of human intervention. Loading parts into machine tools is essentially an assembly operation in which the parts are slid into clamps or fixtures and for this reason the RCC techniques that have been developed for assembly robots can also be applied to a machining cell.

The wrist employs spherical springs with an adjustable stiffness that varies between 33 and 450 lb/inch (5.8×10^3 and 7.9×10^4 N/m). This allows the center of compliance to be projected over a range from 1.0 to 6.2 inches (25 to 157 mm) out from the upper platform of the wrist. The wrist has $5\frac{1}{2}$ degrees of freedom, being compliant in each direction except axial extension. Deflections of $\pm$ 0.18 inches (4.6 mm) in the radial plane and 0.20 inches (5.1 mm) in compression are possible. The accuracy of the sensors over this working range is within 0.001 inches (0.025 mm) for translational motions and within 0.001 radians for rotational motions.

Table of Contents

1 Introduction	1
2 The Adjustable Spheres	3
3 Calibration and Sensor Testing	6
4 Robot Control	8
4.0.1 Control of the wrist	8
4.0.2 Control with the wrist	10
5 A Dynamically Controlled Wrist	13
6 Conclusions	14
7 Acknowledgements	15
A- 1 Calibrating a device with several sensors	16
A- 2 Kalman filter equations for contour following	19

List of Figures

Figure 1:	Design Drawing of Compliant Wrist (adapted from [1])	2
Figure 2:	RCC Linkage (from [1])	3
Figure 3:	Hydraulic Schematic of Wrist	4
Figure 4:	Force/deflection curves for different starting pressures	5
Figure 5:	Spring-rate vs deflection for different starting pressures	6
Figure 6:	RCC projection vs spring-rate	7
Figure 7:	Communications scheme for robot sensory control	9
Figure 8:	Robot using the compliant wrist to determine the position, size and orientation of a rectangular hole	11
Figure 9:	Following a contour with a tool held by an RCC wrist	12

1 Introduction

Figure 1 illustrates the instrumented, adjustable compliant wrist discussed in this report. The base of the wrist is bolted directly to the end of a manipulator and the upper plate of the wrist is bolted to the gripper so that the gripper "floats" relative to the arm. Between the base and the upper plate are several elastomeric spheres which function as springs, and a number of cables which hold the base and the upper plate together in tension. Axial spheres give the wrist stiffness in compression and bending while radial spheres give the wrist stiffness against lateral loads. A detailed discussion of the design of the wrist and the principles of adjusting the center of compliance are provided in an earlier report [1]. That report also considers the merits of passive compliance as opposed to active compliance achieved through control.

A *center of compliance* is a point in space for which the force/deflection behavior of a structure becomes decoupled so that a force or torque in any given direction produces a deflection only in the same direction, without any side-effects. The stiffness or compliance matrix of a structure is therefore diagonal at the center of compliance. This is a useful simplification for manipulation problems where contact forces are involved, akin to finding the center of mass and the principal axes of inertia in dynamics. For example, a number of investigators have shown that during parts assembly it is ideal if the center of compliance of the parts is located at or near the initial points of contact [1-4]. Thus, when inserting a peg into a hole it is ideal if the center of compliance is located near the tip of the peg so that contact forces between the peg and the edge of the hole will align the peg correctly for smooth assembly. When the peg and the gripper that holds it are fairly rigid, a wrist unit mounted at the base of the gripper can provide the needed compliance. The purpose of a remote-center-compliance wrist is to project a center of compliance some distance away from itself, near the tip of the part being assembled. The RCC wrist described in this report is instrumented and can be automatically controlled for tasks involving parts of different lengths and weights. Signals from the sensors in the wrist are analyzed and fed back to the robot to improve its performance.

In Figures 1 and 2, the axial spheres and cables are inclined toward the central z axis of the wrist by a slight angle, θ. Like other RCC devices, the wrist can be interpreted as an elastic four bar linkage (Fig 2). The inclination of the cables and spheres allows the center of compliance to be projected a distance l_2 from the upper platform (item 3 in Figure 1). The axial spheres are hollow and their stiffness, k_{vb}, is changed by adjusting the pressure of a fluid within them. This changes the stiffness of the wrist and alters the projected distance, l_2, of the center of compliance. For the two-dimensional linkage shown in Fig 2:

$$l_2 = \frac{a\,k_{vb}\,tan\theta}{k_x + 2k_{vb}\,tan^2\theta}$$

and for the actual device the three-dimensional projection of the compliant center is:

$$l_2 = \frac{3/2\,a\,k_{vb}\,tan\theta}{k_x + 3\,k_{vb}\,tan^2\theta}$$

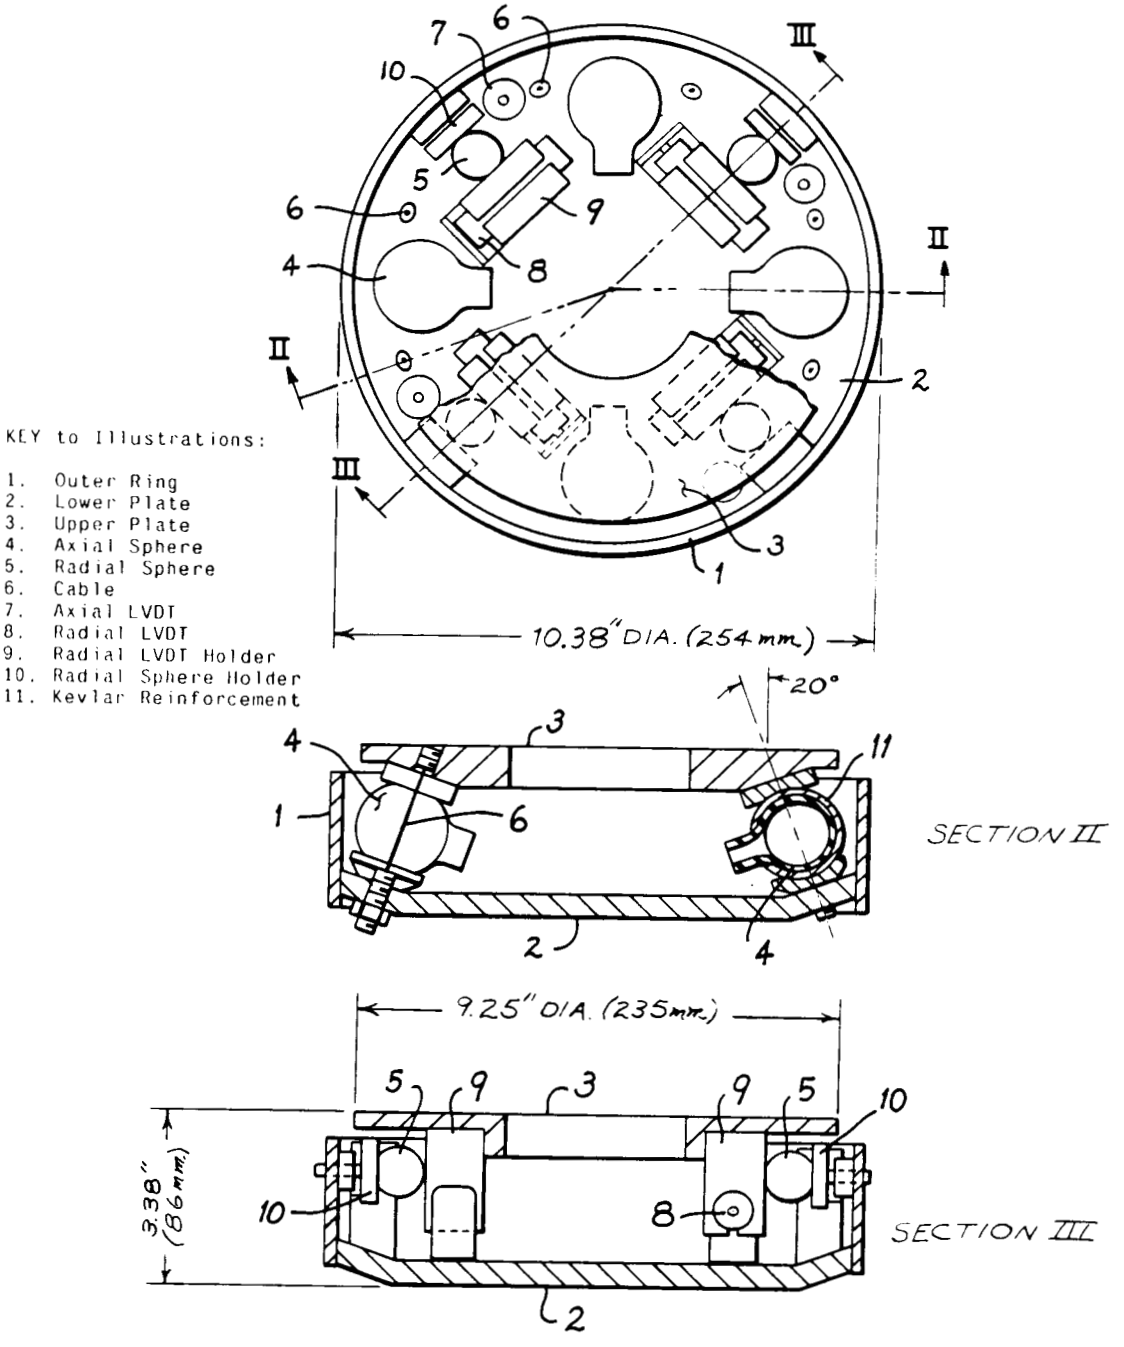

Figure 1: Design Drawing of Compliant Wrist (adapted from [1])

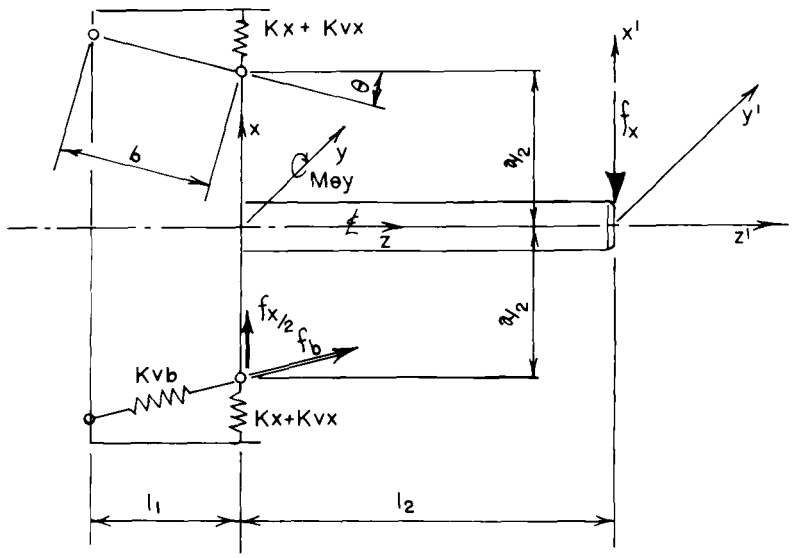

Figure 2: RCC Linkage (from [1])

where k_{vb} is the stiffness of each axial sphere, k_x is the stiffness of each radial sphere, θ is the angle of inclination of the axial spheres and cables and a is the diameter of the wrist.

Plots of the spring stiffnesses for different sphere fluid volumes and of the resulting projected distances for the center of compliance are given in Section 2.

2 The Adjustable Spheres

Prototype axial spheres were made by coating spherical syringe bulbs with rubber cement and wrapping them tightly with 500 denier Kevlar[1] yarn. The Kevlar yarns are extremely strong and flexible but they are vulnerable to abrasion. To protect them, the spheres were given an outer coating of PDS[2] polyurethane.

The axial spheres are connected to solenoid valves using 0.125 inch (3.18 mm) diameter plastic tubing. The tubing is flexible enough to let the spheres roll slightly within the wrist. The solenoid valves are mounted outside the wrist and share a fluid reservoir and a master cylinder as shown in the schematic of Fig 3. The system resembles the hydraulic brake system found in automobiles and, as with the brake system, it is important to bleed all air from the lines. The working fluid is a mixture of water and ethylene glycol (to prevent corrosion). The pressure of each sphere is controlled by opening the valve, adjusting the position of the master cylinder to add or remove fluid, and closing the valve. The individual valves prevent fluid from being pumped from one sphere to another when bending loads are applied to the wrist and allow independent control of the pressure in each sphere. By using individual valves and sharing a single master cylinder we are essentially "multiplexing" adjustment of the spheres. The master cylinder is driven by a stepping motor which, like the valves, is controlled by the dedicated microcomputer attached to the wrist

[1] Kevlar is a trademark of Dupont Inc.

[2] PDS is a trademark of Plasti-Dip International.

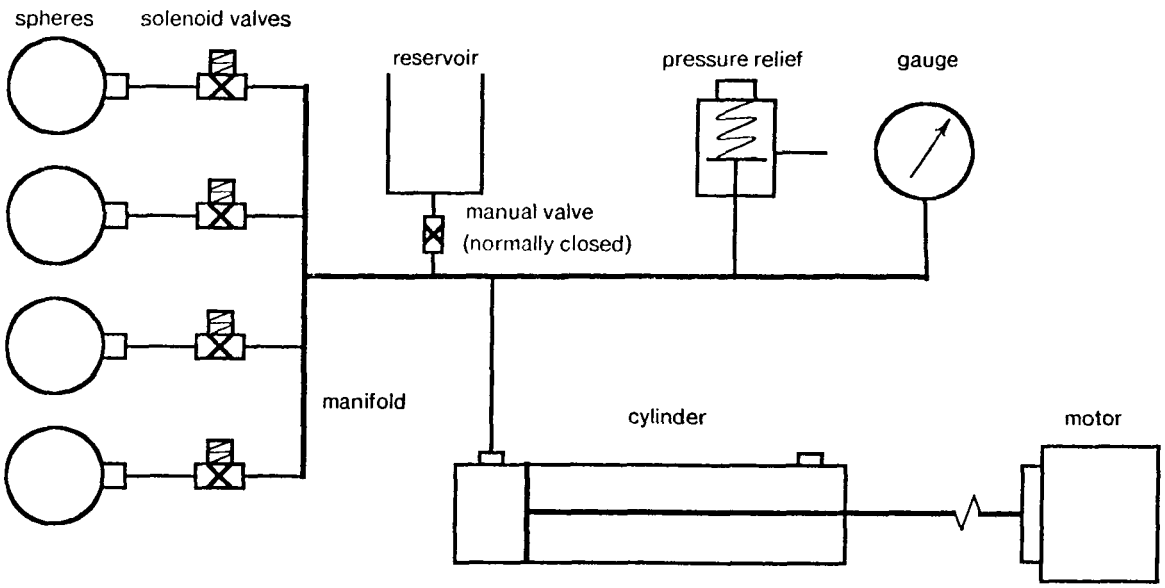

Figure 3: Hydraulic Schematic of Wrist

(Section 4). A faster system, allowing simultaneous adjustment of all four spheres, would be to provide each sphere with its own control cylinder. The advantages of such a system are considered in Section 5.

Figure 4 shows force/deflection plots for the adjustable spheres at three different initial pressures. The amount of fluid added between the second and third tests is the same as the amount added between the first and second, but the difference in slope between the second and third plots is much less noticeable than the difference between the first and second. This illustrates that the stiffening effect of adding extra fluid diminishes as the pressure increases. The derivatives of the force/deflection data in Fig 4 are plotted in Fig 5 as spring-rate *vs* deflection for different initial pressures. As Fig 5 shows, the spring rates for the spheres approach an upper limit of about 450 lb/inch (7.9×10^4 N/meter) beyond which adding more fluid or increasing the load has no additional stiffening effect.

The curves in Figures 4 and 5 represent steady-state data. When the spheres are loaded or unloaded rapidly, viscoelastic effects can change the curves by as much as 10%. Some of the apparent viscoelasticity results from a slight adhesion between the spheres and the upper and lower surfaces touching them. The adhesion can be reduced by dusting the spheres with a dry lubricant or by coating the contacting surfaces with teflon, but care must be taken to keep the spheres from becoming too slippery. If the spheres slip instead of rolling when the wrist is deflected, sliding friction will prevent the wrist from returning to exactly the same equilibrium position each time it is displaced. Hysteresis is undesirable when using sensor data from the wrist to control a robot (Section 4). The wrist was also tested using solid (non-adjustable) spheres. With these spheres the viscoelastic effects were negligible, which suggests that by experimenting with different elastomers and sphere wall-thicknesses one could control the degree of viscoelastic behavior.

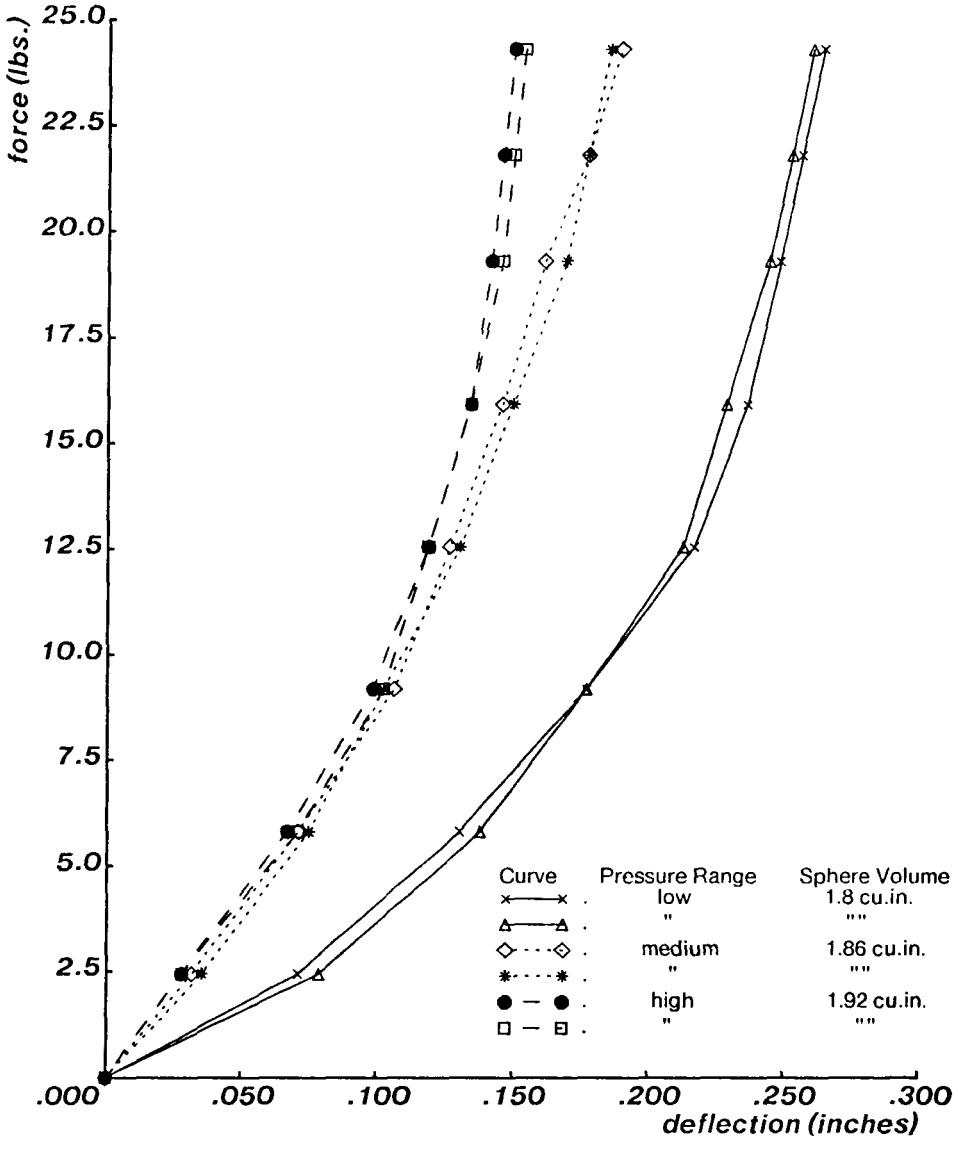

Figure 4: Force/deflection curves for different starting pressures

Figure 6 is a plot of the RCC projection length (l_2 in Fig 2) as a function of the stiffness of the axial spheres. The plot was constructed using data from Fig 5 and the formula for l_2 in Section 1. The range of RCC projection lengths available on the wrist is shown between pairs of dotted lines. As the plot illustrates, the rate of increase of l_2 diminishes noticeably for spring stiffnesses above 600 lb/inch (10.5×10^4N/m).

In fact, if the axial spheres were made infinitely stiff, the RCC projection would asymptotically approach an upper limit of

$$\frac{a}{2\tan\theta} = 9.7 \text{ inches (246 mm).}$$

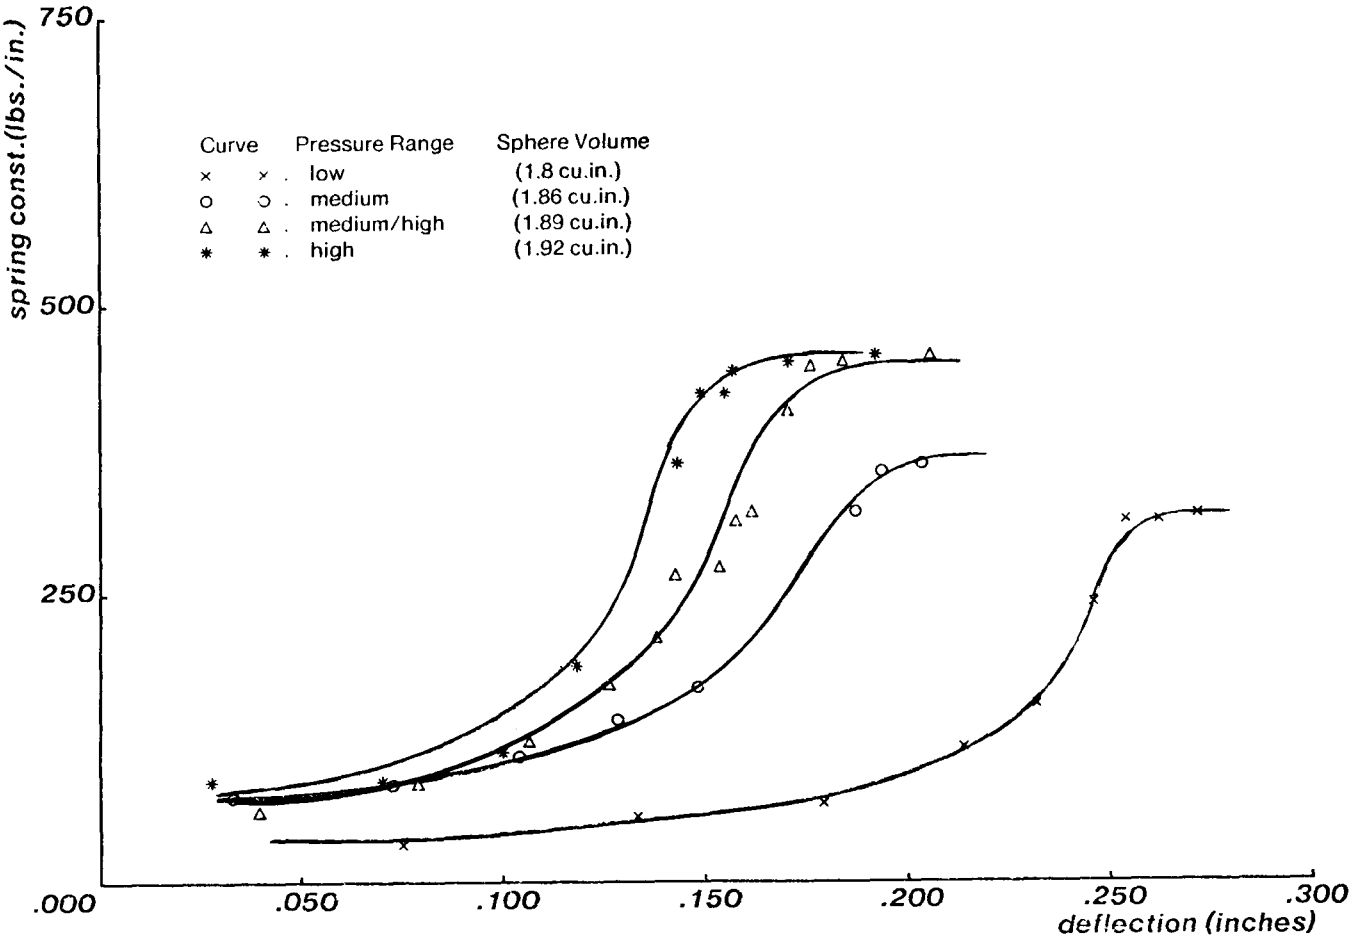

Figure 5: Spring-rate vs deflection for different starting pressures

3 Calibration and Sensor Testing

To calibrate the sensors, the wrist was placed on a three-axis CNC milling machine. A manual rotary table provided a fourth axis. The radial and axial spheres were removed from the wrist during calibration since it was found that the restoring forces they produced in response to even small deflections could seriously degrade the accuracy of the CNC mill. The calibration sequence consisted of jogging the axes of the mill by small amounts, storing the mill coordinates in one array and storing sensor readings from the wrist in a second array. Dial gauges were used to check the accuracy of coordinates read directly from the mill.

The pseudo-inverse of the 6x8 calibration matrix was then determined using a standard least squares technique (Appendix A- 1). The scatter and linearity of the data were checked by multiplying the calibration matrix by the sensor readings and comparing the computed positions with the original recorded positions. The *rms* errors for translational deflections in the x, y and z directions were less than 0.0003 inches (0.008 mm) and the *rms* errors for rotational movements were within 0.0003 radians. These small errors confirm the

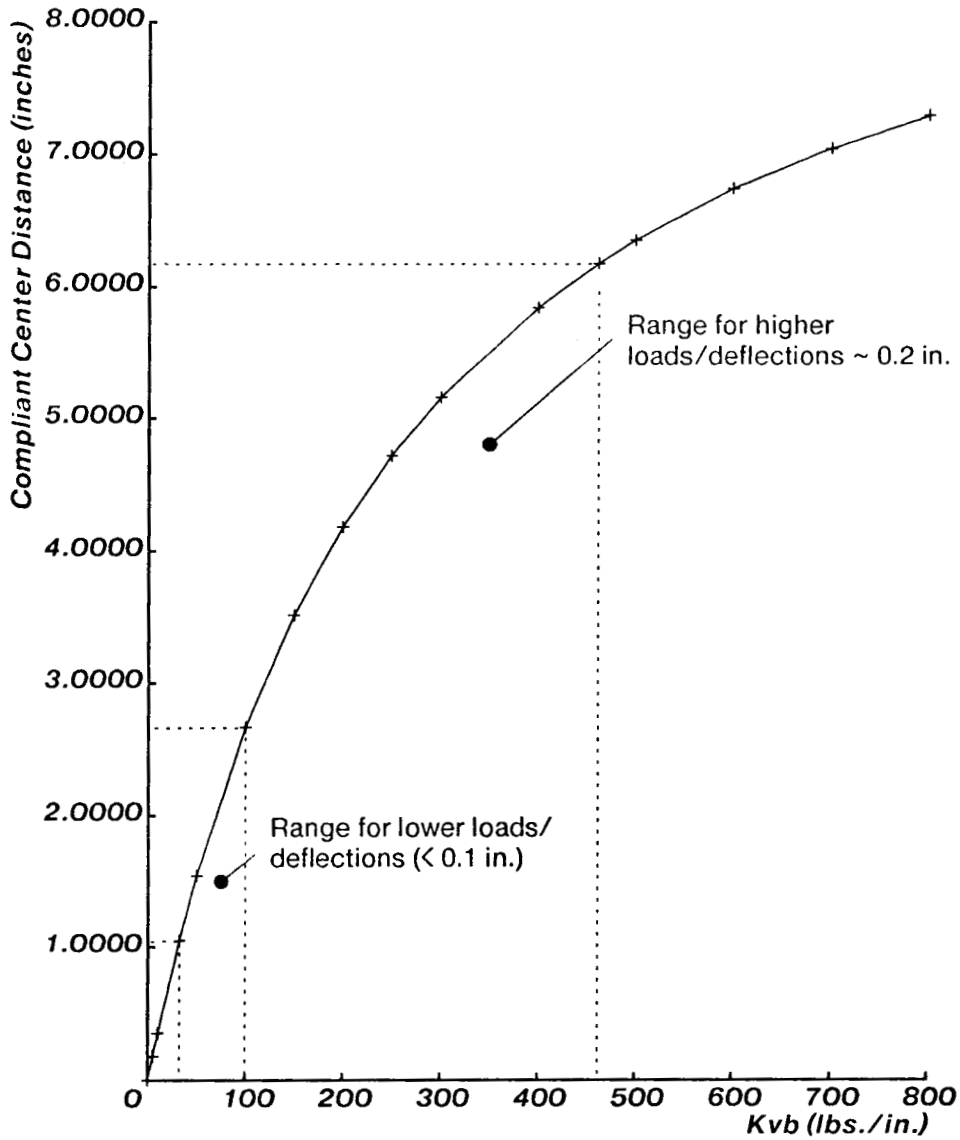

Figure 6: RCC projection vs spring-rate

linearity and repeatability of the LVDT[3] sensors over their relatively small working range of ± 0.18 inch (4.6 mm).

When the wrist is mounted on the robot arm, the accuracy is not quite as good as the above figures would suggest -- primarily because the position and orientation of the wrist with respect to the robot cannot be perfectly established. In addition, the above *rms* errors indicate how well the wrist agrees with the CNC mill and the dial gauges. If the mill and the gauges are slightly miscalibrated then the wrist will be slightly

[3]Linear Variable Differential Transformer

miscalibrated as well. Despite these effects, the working accuracy of the wrist remains within ± 0.001 inch (0.025 mm) for all but large deflections of 0.11 inches (2.8 mm) or greater. For small deflections, where the need for accuracy is highest, the wrist is generally within ± 0.0005 inch (0.013 mm). These accuracies are more than an order of magnitude better than that of the robot and they simplify the control algorithms described in Section 4 because they allow us to treat wrist deflection readings as though they were perfectly accurate compared to position data supplied by the robot controller.

4 Robot Control

Control consists of two functions, control of the robot with the wrist and control of the wrist itself. As Figure 7 shows, these functions are not independent. With the present system, control of the wrist is the simpler problem, and therefore it will be discussed first.

4.0.1 Control of the wrist

As mentioned in Section 2, control of the wrist involves opening and closing the individual valves and moving the master cylinder to add or remove fluid from each sphere. Control of the valves and the stepping-motor is achieved through parallel output ports on the same microcomputer that gathers signals from the wrist sensors discussed in Section 3. The current system, using a stepping motor to adjust the spheres, is open-loop. The system works well since the inertial load is small and the motor velocity is low. Even at the higher pressures there is no danger of the motor missing steps. The correlation between the displacement of the cylinder and the spring rate of the spheres is derived empirically, using the kind of data plotted in Figures 4 and 6. This amounts to a non-linear calibration which, unlike the calibration of the sensors, drifts slightly with time and use as the wrist components settle and trace amounts of fluid leak from the system. For best results, the correlation should be checked periodically and modified as required.

When the robot is about to pick up a new part or perform a new task, the program running on the central computer issues a command to the microcomputer to adjust the initial stiffness of the wrist. The spheres may all be given the same pressure (as is usually the case when the central axis of the wrist is aligned vertically) or they may be given different pressures to allow for static loads. The ability to accommodate static bending loads is especially useful when the central axis of the wrist is aligned horizontally as suggested in Figure 2. Without this feature, most of the compliant range of the wrist would be taken up in gravity loading. As the robot begins to perform a new task, deflection data from the wrist are available to correct the initial settings of the sphere pressures. For example, if the sensors show that the wrist is tilting or bending when contact forces are supposedly acting at the center of compliance this information can be used to increase or decrease the pressure in the spheres, changing the RCC projection distance until the bending or tilting disappears. Deflections recorded by the sensors are also used to correct the initial estimate of the amount of counter-balancing required for static loads.

The discussion above has focused on quasi-static control of the wrist, establishing initial pressures for the spheres and correcting them, if necessary, to accomplish a given task. New possibilities become available

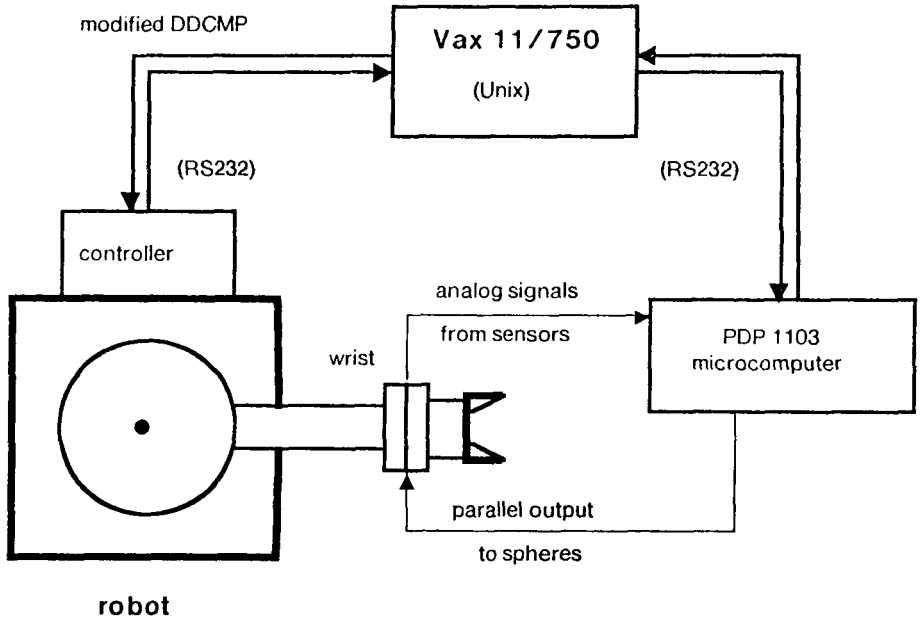

Control Sequence:

robot controller	main program	microcomputer
move robot	send desired robot coordinates	
send coordinates	read achieved robot coordinates	
	get wrist deflections	sample sensors
		send deflections
	adjust wrist if desired	adjust spheres
	determine next move in hand coordinates	
	convert to robot coordinates	
	repeat	

Figure 7: Communications scheme for robot sensory control

when dynamic control of the wrist is considered. The present hydraulic system does not allow rapid readjustment of the spheres but it is possible, for instance, to make the wrist rock slowly from side to side by alternately changing the pressures of spheres on opposite sides of the wrist. This is a useful assembly technique when the amount of angular misalignment is near the limit of what the RCC mechanism can accommodate.

4.0.2 Control with the wrist

Figure 7 shows the computer and communication scheme used to interpret readings from the wrist sensors and modify the robot path. Currently, the most time-consuming part of the control process is the exchange of coordinates between the host computer and the robot controller which, in its present form, only accepts coordinates sent over a serial line at up to 9600 baud. Communication between the central computer and the microcomputer attached to the wrist also uses a serial line, but consists of terse messages that take little time to send. This is possible because the microcomputer does initial processing of the data from the sensors and takes care of the details of adjusting the wrist.

Despite the relatively slow turnaround time, (approximately 1/2 second) the above system is useful for creating and testing new algorithms since the central computer is also the software development computer for the authors' laboratories. Programs running on both the central computer and the microcomputer are written in C and use standard libraries of routines for matrix manipulation and input/output handling. After the control algorithms have been tested, it is possible to combine the programs running on the central computer and on the microcomputer into a single program running on a dedicated microprocessor, but there is little point in doing this until the robot controller is modified to allow rapid communication.

Although low-level routines for filtering data, determining the wrist stiffness and communicating between devices are nearly the same for all robot tasks, it is an unfortunate fact that high-level algorithms are task-specific. A routine that is suitable for determining the position and orientation of a rectangular box must be modified for following an irregular contour and will not work at all for putting a peg in a hole. However, some basic concepts are common to these three tasks which are currently being investigated. The concepts are discussed below, using the examples of determining the orientation of a workpiece and of following an irregular contour. Similar experiments using non-adjustable RCC wrists are discussed in [5]. Some theoretical discussion of the nature of compliant contour-following may be found in [6].

Determining the position and orientation of a workpiece

Figure 8 shows a robot using the compliant wrist to determine the dimensions, position and orientation of a rectangular hole in a box. A peg is mounted to the compliant wrist, and the center of compliance of the wrist has been adjusted to coincide with the tip of the peg. Consequently, contact forces between the peg and the sides of the hole produce only lateral deflections within the wrist (δx, δy). The rough size and orientation of the rectangle are known. The robot begins by moving toward one of the sides until a bump is sensed, at which point the deflection of the wrist and the coordinates of the robot arm are recorded. This process is repeated several times for each side of the rectangle until its position and orientation are established. If the wrist deflection is greater than 0.06 inches (1.5 mm) for any of the contact points, the coordinates returned by the robot for these points may be inaccurate by as much as 0.08 inches (2.0 mm) because the contact forces will have caused the arm to deflect. A more accurate calculation of the position and orientation of the workpiece can be obtained by repeating the touching process for those points. The robot is moved directly to a position just inside the nearest likely position of the contact point and then is slowly advanced toward the side of the rectangle until a bump is sensed. The robot coordinates and the wrist deflection are then recorded and used in place of the earlier ones.

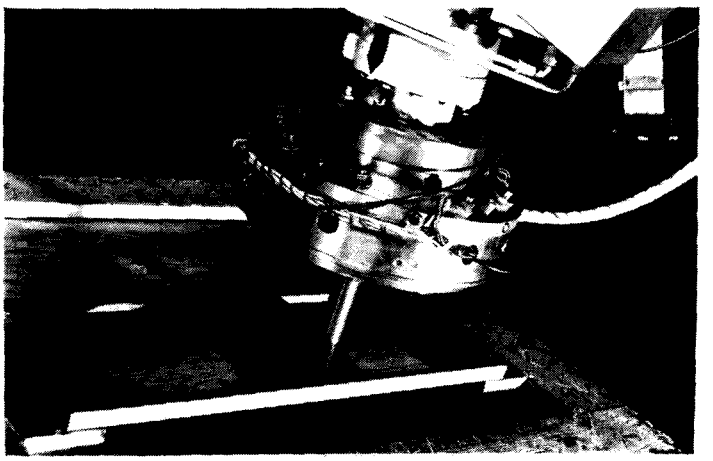

Figure 8: Robot using the compliant wrist to determine the position, size and orientation of a rectangular hole

Similar methods have been developed by colleagues Prinz and Hodges to determine the position and orientation of workpieces prepared for welding. The robot touches plane surfaces of a weldment until enough information has been obtained to uniquely define a transform describing the position and orientation.

For the above tasks, the accuracy of the computed position and orientation of the workpiece can approach the accuracy of the robot which, for constant-force conditions and motions of 10 inches (25 cm) or less, is approximately ±0.03 inches (0.8 mm)

Following an irregular contour

Grinding, polishing and deburring are approximated by a task in which the robot is asked to move a tool along a curved surface while applying a prescribed force against the surface. An instrumented compliant wrist can be used for this task, and if the center of compliance is positioned near the tip of the tool, the problem can be described by lateral deflections in the plane of the wrist and a rotation about the central axis of the wrist (x, y and a in Figure 9).

It is assumed that the position, orientation and curvature of the surface are approximately known and that an initial estimate for the robot path has been computed off-line. As the robot begins to move along the estimated path, differences arise between the desired wrist deflections in the x, y and a directions and the actual values obtained by periodically reading the wrist sensors. The differences result from two kinds of errors:

- Long-term errors produce deflections in the wrist that show a statistically significant trend over several measurements. They may be the result of misalignment of the workpiece, robot miscalibration, or errors in the off-line description of the shape of the part. The effects of misalignment of a workpiece are shown in figure 9. The actual contour is translated by a distance X_s and rotated by an angle A_s with respect to the initial estimate.

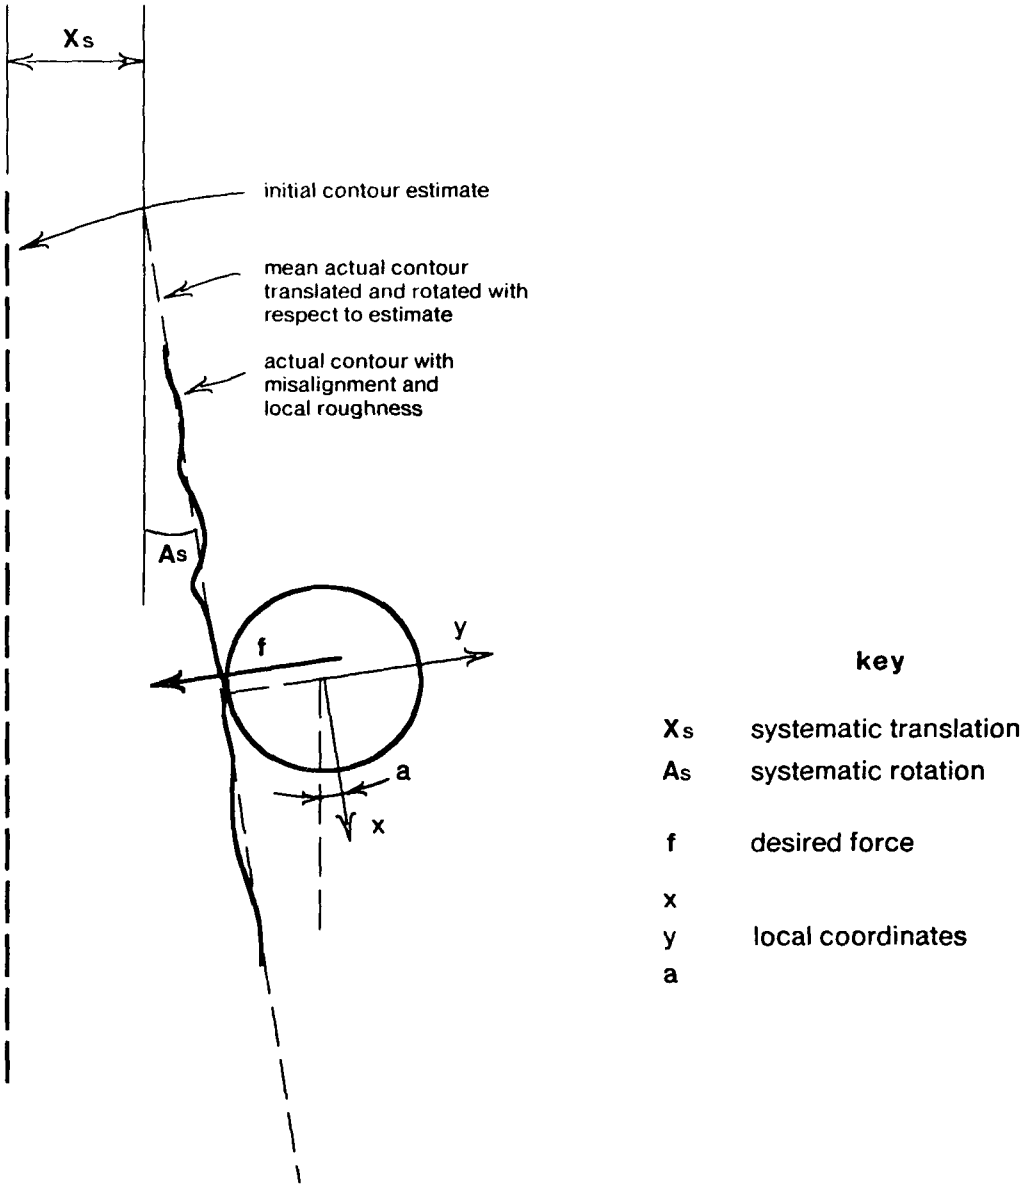

Figure 9: Following a contour with a tool held by an RCC wrist

- Short-term disturbances are also present in the task. As the robot follows the surface, it is subjected to fluctuating forces and torques resulting from local variations in the coefficient of friction and from the actions of the tool as it encounters soft or hard material. In addition to force/torque disturbances there will be position disturbances due to surface roughness and due to the ±0.03 inch (0.8 mm) resolution of the robot. If an instrumented compliant wrist is used, the force/torque disturbances and the position disturbances combine to produce small, fluctuating deflections in the wrist. These essentially random fluctuations are superposed on the long-term errors described above.

The wrist is unable to distinguish between the different sources of deflections, and consequently the errors

and disturbances described above combine to produce wrist deflection data with systematic and random components.

The random fluctuations are unavoidable; they represent the resolution of the robot and the magnitude of disturbances inherent in the task. The wrist must be sufficiently compliant to absorb these disturbances without producing excessive contact forces between the tool and the surface. The task of the control algorithm is to extract long-term trends from the wrist data and to modify the robot path to accordingly.

Extracting long-term trends from sensor data is a filtering problem in which any random fluctuations are thought of as "noise" corrupting a systematic signal. Whitney and Junkel [5] describe similar problems involving noisy sensor data and discuss the advantages of using a Kalman Filter to extract the systematic trends.

A Kalman Filter is being used with the compliant wrist for simple contour-following tasks, using the coordinate system shown in Figure 9. Since the characteristics of the errors do not vary much over the duration of the task, a time-invariant Kalman filter can be used. The random disturbances described above become "process noise" in the Kalman Filter equations (Appendix A- 2). In general, there is also sensor noise but, as shown in Section 3, the sensor errors are small in comparison to the other error terms. This leaves the problem of estimating the process noise. To a first approximation, it is acceptable to assume that the random disturbances in x, y and a are zero-mean and uncorrelated, but determining their magnitudes is more difficult. The variances of these noise quantities will depend upon position of the robot arm (how far it is extended, for example) and on surface properties of the contour which may be expected to vary from piece to piece. For these reasons an adaptive Kalman filter is used in which the process noise covariance matrix is updated using information from the sensors. General methods for updating covariance matrices are given by Mehra [7, 8], but for the contour following problem a simple additive formula is used to compute the variances of the noise in x, y and a. Details are given in Appendix A- 2. In general, all the data points starting from $t = 0$ are used to compute the noise variance until the filter has reached a steady state. However, once the filter has reached steady-state conditions, only the last several sensor readings are used to give an average current value of the noise variance.

5 A Dynamically Controlled Wrist

For more difficult manufacturing tasks, such as positioning a machined part in a fixture while the fixture is tightened, active control of the wrist becomes desirable. In the last section, some techniques for slow control of the wrist are described. A much more rapid response would be possible if each sphere were given a separate servo-controlled pressure source instead of sharing a single cylinder adjusted by a stepping motor. With this modification each sphere becomes a small bladder-type hydraulic cylinder and the wrist becomes a servomechanism with three controllable degrees of freedom. With some minor modifications, such as pressurizing the radial spheres, the wrist could have six controllable degrees of freedom. The sensors already in the wrist would provide the required feedback. Wrist units consisting of active servos have been recently

been discussed in [9, 10], but the present wrist represents a different approach. The mechanical design of the wrist described here is able to passively accomplish a major part of the required behavior. The idea of a servo controlled wrist may be criticized on the grounds that it adds complexity to an already complex manipulator, but there are some important advantages to controlling a compliant wrist instead of controlling the manipulator for fine accommodations:

- The wrist is a mechanism with excellent passive properties, good enough to accomplish many tasks with no active control at all.
- The spheres have no moving seals or pistons and consequently no sliding friction.
- The wrist is well damped.
- The wrist is mechanically simple. The actuators (spheres) are in direct contact with the upper platform. The absence of intermediate linkages, and gears makes the wrist tolerant of overloading and devoid of backlash.
- Most manipulators are "serial-linkage" devices in which the joints and links are connected end-to-end. Consequently, a small positional error in one of the base joints results in large errors at the tip of the arm. The wrist, however, is a device in which the actuators are connected in parallel, so that individual errors are additive and not multiplicative.

Some additional simplifications arise from the fact that the wrist never moves very far or very fast. These are true for all servo wrists, remote-center-compliance or not.

- Large forces and torques can be applied by the wrist without using much power.
- Deflections are measured with a high resolution since the sensors work over a small range.
- The kinematics and dynamics of the wrist are largely linear since small angle assumptions are valid over the working range and since the angular velocities and accelerations are small.

The wrist could be controlled using force or position servos, but the most attractive possibility is to implement active stiffness control following the scheme discussed by Salisbury [11] for manipulators. Since the wrist is essentially a stiffness/compliance device, stiffness servoing is is especially well suited to its characteristics.

6 Conclusions

Initial results of testing and operating the adjustable, instrumented, compliant wrist show that the wrist demonstrates excellent compliant characteristics and provides accurate deflection data for robot control. The ability to adjust the spheres allows the wrist to be tuned for grippers and workpieces of different sizes and weights and allows the center of compliance to be positioned where it will simplify robot assembly and contour- following tasks.

The operating characteristics of the wrist are summarized below:

- The center of compliance can be projected over a range of distances between 1.0 and 6.2 inches (25 to 157 mm) out from the upper platform of the wrist.

- The wrist is capable of compliant deflections in every direction except axial extension. The stiffness of each of the four axial springs can vary between 33 and 450 lb/inch (5.8×10^3 and 7.9×10^4 N/m).

- The wrist provides deflection data accurate to 0.001 inches (0.025 mm) for translations and 0.001 radians for rotations.

- A microcomputer reads the sensors in the wrist, computes deflections and sends them, over a serial line, to a supervisory computer. The supervisory computer interprets the deflections and modifies the path of an industrial robot. The present communications scheme is not fast enough for dynamic control of the robot, but allows the robot to respond to systematic errors or trends.

- The supervisory computer also issues commands to adjust the wrist. The details of adjusting the spheres are carried out by the microcomputer.

7 Acknowledgements

We thank D. T. McKeel and J. R. Dillinger for machining the wrist components and M. Goldstein for advising us on the electronic hardware for the wrist. Professors C. P. Neumann and M. Naguraka provided many valuable comments and suggestions on the draft of this report.

A- 1 Calibrating a device with several sensors

The following procedure may be used to determine the calibration matrix of a deflection-sensing device having j sensors and k degrees of freedom (where $j \geq k$). Since forces and incremental motions transform in exactly the same way [12], the procedure is also valid for force-sensing devices.

For the adjustable compliant wrist discussed in this report, the calibration matrix can be written:

$$\begin{bmatrix} c_{11} & c_{12} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & c_{23} & c_{24} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{35} & c_{36} & c_{37} & c_{38} \\ 0 & 0 & 0 & 0 & c_{45} & c_{46} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & c_{57} & c_{58} \\ c_{61} & c_{62} & c_{63} & c_{64} & 0 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} S_1 \\ S_2 \\ S_3 \\ S_4 \\ S_5 \\ S_6 \\ S_7 \\ S_8 \end{matrix} = \begin{matrix} \Delta_x \\ \Delta_y \\ \Delta_z \\ \Delta_{\theta x} \\ \Delta_{\theta y} \\ \Delta_{\theta z} \end{matrix}$$

where S_i are the readings from each sensor, and Δ_i are the six components of deflection for the compliant platform of the wrist and c_{ij} are non-zero elements of the calibration matrix, [C].

The calibration matrix is similar to a 6x8 calibration matrix predicted by Shimano [13] for the Scheinman Wrist. The Scheinman Wrist uses eight strain gauges symmetrically mounted to a piece of aluminum in the shape of a Maltese Cross. By locating the strain gauges along neutral axes of the arms of the cross it is theoretically possible to eliminate "cross-talk" between measurements of the different components of deflection. Unfortunately, it proved impossible to locate the strain gauges of the Scheinman Wrist perfectly along the neutral axes and therefore the calibration matrix for the Scheinman Wrist contained 48 non-zero elements. By contrast, since it is possible to locate the LVDTs in the wrist design of Fig. 1 with a very high accuracy, and since the quantities they measure are comparatively large, we obtain a sparse calibration matrix. This simple relationship between sensor readings and deflections is consistent with the findings of other researchers who have used LVDTs for deflection measurement [14].

Shimano [13] and Watson [15] describe similar procedures for determining the elements of a calibration matrix. They apply six independent forces of known direction and magnitude to a force sensing wrist and then solve for the pseudo-inverse of the calibration matrix. An analogous procedure for a deflection sensing wrist uses six known deflections. A larger number of independent deflections may be applied to improve the accuracy of the calibration. The coefficients of the pseudo-inverse of the calibration matrix are then determined using a least-squares regression. The equations are as follows:

Nomenclature:

[C] = a k x j matrix of calibration terms relating sensor readings to deflections. (A matrix is justified because the small-motion transformations are linear.)

$\underline{s}$ = column vector of readings from the j sensors

$\underline{si}$ = (n x 1) column vector of n samples from the ith sensor

$\underline{\Delta}$ = column vector of the k force or deflection components

[CI] = j x k pseudo-inverse of matrix [C] above

CI_i = i'th row of [CI]

[D] = an n x k matrix. Each row of [D] is a k element deflection vector, like $\underline{\Delta}$, and represents one of n known deflections or forces applied to the device.

The relation between sensor readings and forces or deflections is:

$$[C]\,\underline{s} = \underline{\Delta}$$

To determine the matrix [C] we first compute each row of the matrix [CI] using a least squares fit:

$$([D]^t\,[D])^{-1}\,[D]^t\,\underline{si} = CI_i^t$$

We must have n independent data vectors where $n \geq k$. If $n = k$, the equations can be solved directly as a set of n simultaneous equations.

When all j rows of [CI] are computed, we take the pseudo-inverse of [CI], again using least squares:

$$([CI]^t\,[CI])^{-1}\,[CI]^t = [C]$$

(if $j = k$ the matrix [CI] can be inverted directly to get [C]).

A slightly more compact solution is possible if we start with a rectangular matrix, [S] such that:

[S] = an n x 8 matrix of sensor data. Each row is a sample and each column

is a set of n data points for a single sensor. Thus $\underline{s}^t$ would be a row of [S] and $\underline{si}$ would be a column.

The relation between sensor readings and forces or deflections now becomes:

$$[C][S]^t = [D]^t.$$

To solve for [C] we apply the same methods as above:

$$([D]^t[D]^{-1})[D]^t[S] = [CI]^t$$

and

$$([CI]^t [CI])^{-1} [CI]^t = [C].$$

A- 2 Kalman filter equations for contour following

As mentioned in Section 4.0.2, a time-invariant Kalman filter has been used with the compliant wrist as an estimator for following a contoured surface. The derivation of the Kalman Filter equations may be found in the text by Gelb *et al* [16] and other texts in optimal estimation. For reference, they are summarized below:

Nomenclature

Φ	The state transition matrix for x,y,a and their derivatives.
s	The state variable estimate vector (x, dx/dt, y, dy/dt, a, da/dt in Figure 9)
w	A component of process noise
Q	The process noise covariance matrix
R	The sensor noise covariance matrix
z	The vector of sensor readings in x,y,a
V	The Kalman filter signal vector matrix
G	The Kalman filter gain matrix
Γu	The control law used for moving the robot in response to deflections (for example, proportional feedback)
I	The identity matrix
k	One time step
v	dy/dt

The Kalman Filter is used as a estimator for the noisy wrist deflection data. We assume that the wrist signals obey the following equation:

$$S_{k+1} = \Phi S_k + \Gamma \mathbf{u}_k$$

where S_{k+1} is the estimate for the next deflection of the wrist and s_k is the current estimate.

For a time-invariant Kalman filter the following three equations are updated, in order, with each each time step.

$$G_k = V_{k-1} H^t (H V_{k-1} H^t + R)^{-1}$$

$$\mathbf{s}_k = \Phi \mathbf{s}_{k-1} + \Gamma \mathbf{u}_{k-1} + G_k(\mathbf{z}_k - H(\Phi \mathbf{s}_{k-1} + \Gamma \mathbf{u}_{k-1}))$$

$$V_{k+1} = \Phi(I - G_k H) V_{k-1} \Phi^t + Q$$

The estimate, $\mathbf{s}_k$, is the output from the filter and $\mathbf{z}_k$ is the input.

The simplest example of a contour-following problem is to use a smooth, round tool to follow an approximately straight line. If the coordinate system is set up as in Figure 9 so that y is perpendicular to the surface, the problem reduces to one with two state variables: y and $v = dy/dt$. Variations in x are unimportant since x is parallel to the surface. There are no significant torques to cause rotational deflections, a, because there is little friction between the smooth tool and the contour. To simplify the example further, we assume that the filter is being used merely to track a sequence of wrist data and that, for the moment, we are not modifying the robot path in response to the deflections. This allows us to look at the filtering part of the problem without having results influenced by the control actions. The state equation is:

$$\begin{matrix} y_{k+1} \\ v_{k+1} \end{matrix} = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} \begin{matrix} y_k \\ v_k \end{matrix} + \begin{matrix} w_y \\ w_v \end{matrix}$$

The noise properties, $[w_y, w_v]$ determine the process noise covariance matrix, Q. For contour-following, the magnitude of the task-related disturbances is typically on the order of 0.05 inches (1.3 mm). The robot resolution is on the order of 0.03 inches (0.8 mm). As discussed in Section 3, the sensor errors are on the order of 0.001 inches (0.025 mm).

Filter adaptation

The filter equations above require initial estimates for Q and R. As discussed in section 4, the process noise, Q, is difficult to estimate. However, for small sensor noise, the sensor readings give an indication of the process noise. In the simplest case, the noise in the measured variables is zero-mean and uncorrelated.[4] The sensor readings can be used to determine a matrix, Λ, containing the noise variances in (x, y, a):

$$\begin{matrix} \sigma^2_x & 0 & 0 \\ 0 & \sigma^2_y & 0 \\ 0 & 0 & \sigma^2_a \end{matrix}$$

where σ^2_x is the variance in x, and so on.

[4]This has been found to be approximately true from measurements with the wrist in contour-following problems.

The variance for each coordinate may be computed using a standard formula [17]:

$$\sigma^2_{k=n} = \frac{1}{n(n-1)} \left(n \sum_{t=1}^{n} z^2 k - (\sum_{t=1}^{n} z_k)^2 \right)$$

where

$$\sum_{t=1}^{n} z_k \quad \text{and} \quad \sum_{t=1}^{n} z^2 k$$

are updated with each step, k, over a time period of n time steps.

In a contour-following problem the last 10 steps are often used to give an current value of the variances. The scalar equation can be written in matrix form to produce a square matrix, Λ, with the dimensions of R:

$z =$ a column vector of measured variables, z.

for $k = 1$, $A_1 = z_1 \cdot z_1{}^t$ and $b_1 = z_1$

then:

$$A_k = \frac{k-1}{k} A_{k-1} + \tfrac{1}{k}(z_k \cdot z_k{}^t)$$

$$b_k = \frac{k-1}{k} b_{k-1} + \tfrac{1}{k} z_k$$

$$\Lambda_k = \frac{k}{k-1} (A_k - b_k \cdot b_k{}^t)$$

References

1. M.R. Cutkosky and P.K. Wright, "Position Sensing Wrists for Industrial Manipulators," Robotics Institute Technical Report CMU-RI-TR-82-9, Carnegie-Mellon University, July 1982.

2. D.E. Whitney *et al*, "Part Mating Theory for Compliant Parts," First Report, The Charles Stark Draper Laboratory Inc., August 1980, NSF Grant No. DAR79-10341

3. H. McCallion, K.V. Alexander and D.T. Pham, "Aids for Automatic Assembly," *1st International Conference on Assembly Automation*, Brighton, UK, March 1980, pp. 87-97.

4. M.S. Ohwovorioli, *An Extension of Screw Theory and its Application to the Automation of Industrial Assemblies*, PhD thesis, Stanford University, April 1980.

5. D.E. Whitney and E.F. Junkel, "Applying Stochastic Control Theory to Robot Sensing, Teaching and Long Term Control," *12th International Symposium on Industrial Robots*, Paris, France, June 1982, pp. 445-457.

6. M.T. Mason, "Compliance and Force Control for Computer Controlled Manipulators," MIT AI Memo TR-515, Massachusetts Institute of Technology, April 1979.

7. R.K. Mehra, "On the Identification of Variances and Adaptive Kalman Filtering," *IEEE Transactions on Automatic Control*, Vol. AC-15, No. 2, April 1970, pp. 175-184.

8. R.K. Mehra, " On-Line Identification of Linear Dynamic Systems with Applications to Kalman Filtering," *IEEE Transactions on Automatic Control*, Vol. AC-16, No. 1, February 1971, pp. 12-21.

9. H. Van Brussel and J. Simons, "The Adaptable or Compliance Concept and its use for Automatic Assembly by Active Force Feedback Accomodations," *9th International Symposium on Industrial Robots*, Washington, DC, 1979, pp. 167-181.

10. H. Van Brussel, H. Thielmans and J. Simons, "Further Developments of the Active Adaptive Compliant Wrist (AACW) For Robot Assembly," *11th International Symposium on Industrial Robots*, Tokyo, Japan, October 1981, pp. 377-384.

11. K. Salisbury, "Active Stiffness Control of a Manipulator in Cartesian Coordinates," *Proc. 19th IEEE Conference on Decision and Control*, Albuquerque, NM, December 1980, pp. 87-97.

12. R.P. Paul, *Robot Manipulators: Mathematics Programming and Control*, The MIT Press, Cambridge, MA, 1981.

13. B.E. Shimano, *The Kinematic Design and Force Control of Computer Controlled Manipulators*, PhD thesis, Stanford University, March 1978.

14. H. Inoue, "Force Feedback in Precise Assembly Tasks," MIT AI Memo 308, Massachusetts Institute of Technology, August 1974.

15. P.C. Watson and S.H. Drake, "Method and Apparatus for Six Degree of Freedom Force Sensing," United States Patent No. 4,094,192, June 1978 .

16. A. Gelb *et al*, *Applied Optimal Estimation*, The M.I.T. Press, Cambridge, MA, 1974.

17. W.H. Beyer, *CRC Standard Mathematical Tables*, CRC Press, Boca Raton, Florida, 1978.