T-SMC/19/4//27245

A Dynamic Approach to High-Precision Parts Mating

Susan N. Gottschlich
Avinash C. Kak

Reprinted from
IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS
Vol. 19, No. 4, July/August 1989

A Dynamic Approach to High-Precision Parts Mating

SUSAN N. GOTTSCHLICH AND AVINASH C. KAK, MEMBER, IEEE

Abstract—**The mating of tightly-fitting parts in the presence of significant sensing, model, and control uncertainties is addressed. Using a three-dimensional (3-D) laser range scanning system, we are now able to locate randomly placed parts and, with the method described, mate them even when the clearance between them is only 0.001 in. Our preliminary experiments in peg-in-hole operations have brought to light factors that other researchers in this area are apparently unaware of, and the possibility that these factors may influence robotic operations has guided our overall approach to the assembly problem. First the various sources of error introduced in the fine motions used to execute assemblies are examined. Then a general and still evolving approach to robotic assembly operations is discussed. Finally, a specific instance of this approach is outlined, and with the help of automatic error detection and recovery, the ability to mate tightly fitting parts with a high degree of reliability is demonstrated.**

I. Introduction

TO ADDRESS the subject of high-precision robotic assembly, two distinct but interrelated issues must be considered: parts mating strategies and contact mechanics. Parts mating strategies, consisting of robotic movements for the execution of assembly tasks, usually involve two types of force-guided motions: guarded and compliant. A guarded motion is used most commonly to establish physical contact between a part and some component of its environment. A compliant motion, on the other hand, is typically used to slide a part along or through one or more constraining surfaces while a force constraint is maintained perpendicular to the direction of motion.

Since the complete motion of a part held by a robot is determined by all of the external forces acting upon it, and since some of these forces are a consequence of the physical contact between the part and other object surfaces, the issue of contact mechanics is highly germane to robotic assembly. Considerations that arise out of contact mechanics have a profound bearing on the ability of a robotic system to recognize sensory errors and to invoke corrective actions automatically.

A. Previous Work

Contact mechanics, which involves the examination of impact dynamics as well as sliding dynamics, has been studied for a long time. Only recently, however, have scientists begun to apply the theories of classical contact mechanics to robotic operations. Wang and Mason [1] have modeled the impact dynamics of robotic operations. Rajan *et al.* [2] have discussed the sliding dynamics of rigid bodies in contact with frictional surfaces. Peshkin and Sanderson [3] have considered the effect of sliding dynamics on robotic manipulations. Hogan [4] has proposed a control strategy that reacts stably to the transition between constrained motion where contact is involved and free motion.

Work on parts mating strategies has been going on for about two decades. Several authors at the Charles Stark Draper Labs [5]–[8] have proposed the utilization of passive compliance in parts mating applications. They used a peg-in-hole operation to model a typical assembly task and advocated the use of chamfers to reduce the accuracy required to position a peg over a hole. This work led to the development of a device called the Remote Center of Compliance (RCC) for successful execution of peg-in-hole type assemblies without wedging or jamming.

Soon after the use of passive compliance became feasible, people began studying the use of active compliance. Lee and Smith [9] considered the use of active force feedback for controlling an insertion process. Their work is formulated in terms of conditional probabilities for various possible contact configurations, the probability functions being used for the derivation of decision rules for the recognition of contact configurations.

The use of guarded and compliant motions for assembly was originally proposed by Inuoe [10]. He developed a simple algorithm for inserting a peg into a hole given possible positional errors in the vertical direction (the direction of insertion) and along one horizontal axis.

Lozano-Perez *et al.* [11] developed a theory, later extended by Erdmann [12], that utilizes active compliance to synthesize fine-motion strategies. This theory considers spheres of worst case errors and develops a set of velocity vectors that a robot can follow to accomplish an assembly task. Theoretically at least, these velocity vectors could permit a typical peg-in-hole operation to succeed even if the initial placement of the peg was off-center with respect to the hole—something that would not be possible with an RCC alone. Donald [13] further extended this work by introducing a formal framework for error detection and recovery. In doing so he was able to propose a strategy for

Manuscript received February 24, 1988; revised September 15, 1988 and December 1, 1988. This work was supported in part by the Engineering Research Center for Intelligent Manufacturing at Purdue University and in part by the Purdue CIDMAC program.

The authors are with the Robot Vision Laboratory, School of Electrical Engineering, Purdue University, West Lafayette, IN 47907.

IEEE Log Number 8927245.

0018-9472/89/0700-0797$01.00 ©1989 IEEE

achieving a goal even if its success could not be guaranteed prior to execution.

B. Overview of Work to be Presented

In this paper we begin by examining the work done by Maples and Becker [14]. They present a classification scheme for robot force control algorithms based on two criteria: 1) whether the control is implemented in joint space or Cartesian space; and, 2) the nature of the signals employed for servoing. The motivation for the second criterion is based on the fact that one may servo a robot with respect to either position, torque, or velocity for implementing force control. Maples and Becker concluded that, of the different algorithms available for force control, the best trade-offs with respect to the suppression of internal disturbances are achieved by those algorithms that use position servoing, as opposed to velocity or torque servoing. Position servoing means that errors between the desired forces and the measured forces are converted into position commands; and internal disturbances means disturbances generated within the robot by phenomena such as gear cogging, friction, etc.

Following Maples and Becker, we have carried out a position servoed implementation of force control for a Cincinnati Milacron T^3-726 robot. However, our motivation is less the rejection of internal disturbances and more the diminution of the ill effects of external disturbances generated by the very act of bringing parts into physical contact, external disturbances being usually caused by impacting and sliding motions.

Ultimately, what we really wish to accomplish is the development of an automatic planner capable of using our force control implementation to carry out complex assemblies. On account of the presence of the aforementioned external disturbances, such a planner must of necessity be capable of error detection and recovery, which brings us to the work recently reported by Donald [13].

Donald has put forward a formal framework for fine motion planning in the presence of three types of uncertainties: sensing, control, and model. By using a configuration space representation of these uncertainties, Donald's procedure identifies error detection and recovery (EDR) regions in the space that allow a planner to recognize error conditions and devise automatic recoveries from the errors. Donald has proposed what he calls guiding principles for error detection and recovery, as follows. 1) A strategy should attain the goal when it is recognizably reachable, and signal failure when it is not. 2) A strategy should permit serendipitous achievement of the goal. 3) No motion guaranteed to terminate recognizably in the goal should ever be prematurely terminated as a failure. 4) No motion should be terminated as a failure while there is any chance that it might serendipitously achieve the goal due to fortuitous sensing and control events.

Practical difficulties arise when one tries to implement the EDR framework proposed by Donald, especially if force control is implemented with position servoing. Donald handles control uncertainties by using a cone of velocities where the actual commanded velocity must lie somewhere within this cone. Unfortunately, velocity cones appear to be too simplistic a representation for describing the control uncertainties introduced by electrical and inertial noise in our wrist-mounted force/torque sensor and the external disturbances created by impacts and sliding motions inherent to position servoed implementations of force control.

The goal of this paper is to discuss parts mating in a completely experimental context. We will first delineate the external disturbances that one has to contend with during guarded and compliant moves, especially when such moves are implemented with position servoing. We then discuss how we deal with the detection of errors and automatic recovery from errors. Our discussion of error detection and recovery is mostly within the context of conducting peg-in-hole experiments, although we believe that the proposed strategies are extensible to more complex assembly tasks. For the peg-in-hole problem, our mating strategies are dynamic, in the sense that all decisions about what move to make next are based upon the current position and force/torque readings. For generalization of our work to more complex assemblies, the decisions would have to be augmented with *a priori* knowledge of object geometries.

With our method, we are now able to insert pegs in holes successfully at a clearance level of only 0.001 in. Our overall strategy consists of three phases: 1) make a guarded move of the peg to the surface containing the hole; 2) make translational moves of the peg until the following terminating condition is satisfied: the projection of the bottom face of the peg onto the top opening of the hole is contained in the latter; and 3) after successful termination of the second phase, carry out the insertion of the peg into the hole. Under tight tolerance conditions, we have found the second phase to be the most challenging; the success rate in this phase is a function of the number of allowed EDR cycles, each cycle consisting of a test of the phase-2 terminating condition, followed, upon failure of the test, by a translational compliant move until the satisfaction of certain force and/or position conditions. Note that any rotational discrepancies in the alignment of the axes are automatically corrected during the third phase when the robot must servo with respect to horizontal components of torque, in addition to, of course, the horizontal components of force. (The word "horizontal" only applies to the case when a peg is inserted vertically into the hole.)

To give the reader an idea of the extent of the dependence of the success rate of our experiments on the number of phase-2 EDR cycles, if we allow eight phase-2 EDR cycles to be invoked, the insertion experiments succeed about 80 percent of the time. In these experiments, structured-light-based three-dimensional (3-D) vision is used for locating the parts initially to an accuracy of about 0.25 in. The system then uses the position/orientation data obtained from the structured-light-generated range map to synthesize automatically a manipulation plan that causes

the parts to be brought to the positions/orientations from where the force guided motions can take over for final assembly. In this paper, we will not discuss the details of the structured-light sensing or the generation of manipulation plans; on the other hand, we will focus on the use of force/torque (F/T) sensing for final assembly starting from the peg being above the hole but not necessarily perfectly aligned with it.

Although the phase-2 EDR strategy is central to the assembly process, the EDR strategies used in the first and third phases are also important. In the first phase, the recognition of the termination of the guarded move can get fouled up by inertial and other effects and may require the force readings for verification to be taken under momentarily stationary conditions. In the third phase, jamming and wedging may cause an apparent termination of the compliant move required for insertion; such error states are recognized by a combination of force, torque, and position readings.

II. Potential Sources of Error in Force-Guided Motion

Most work in fine-motion planning up to this point is based on the assumption that motions can take place in a quasi-static manner, meaning that the motions can transform perfectly the present states of the assembly components to the desired states without interacting with the environment. In the next two subsections, we will focus in a qualitative manner on the dynamics of the end-effector and its interaction with a contacting surface to establish that this assumption is not justified. Our discussion will not be quantitative, unlike that of Asada and Ogawa [15], who have analyzed the dynamics of the robot manipulator and the end effector to optimize the orientation of the tool and the configuration of the manipulator arm for experiments such as deburring.

We will discuss the different modes of impact that can occur between mating surfaces, which can lead to unexpected positional changes in the part held by the end-effector as well as unexpected forces.[1] We will show that it is unlikely that during a guarded motion a force constraint can be realized precisely. Therefore, it is unreasonable to expect that a guarded motion will terminate precisely when the desired forces are experienced.

For reasons to be shown, it is also unreasonable to expect a compliant motion to be capable of moving a part while maintaining the force constraint precisely. If the force constraint is not maintained precisely, contact between the mating surfaces may cease temporarily, leading to a premature terminating condition. The opposite condition, which is equally likely to occur, happens when the force normal to the direction of motion becomes too high due to an unexpected collision with a small surface irregularity. When forces normal to the direction of motion become too high, the part held by the robot may stick to the surface along which it is being slid. Since the forces applied by the end-effector are usually at one extremity of the part, the sticking can cause unexpected changes in the position of the part, further exacerbating the high-force condition in some cases. In a dynamic situation, if these forces occur at a time instant that is not coincident with one of the sampling times, the terminating condition may not be recognized.

A. Guarded Motion

Although guarded motion can be considered as a special case of compliant motion, it is more efficient to implement separately those types of guarded moves that the robot is expected to engage in most frequently, the resulting advantage being a speedup in execution. For example, in our assembly experiments, prior to the compliant moves needed for the mating of the parts, we frequently want the robot to move along one of the three world coordinate axes until a force discontinuity occurs along that axis, or a specified position is researched. The necessary guarded moves are executed by moving the robot in small increments in the desired direction until a force or position terminating condition is satisfied.

There are two major effects that cause erroneous terminating conditions in a guarded move: 1) the electrical noise and mechanical vibration in the F/T sensor output, which increases during the motion of the arm due to the increased drive current supplied to the motors; and 2) inertial effects, which depend on the placement of the F/T sensor on the robot arm.

Our F/T sensor is mounted (as most are) between the robot tool plate and the gripper. During a guarded move, while the end-effector is not contacting any surface, the force exerted on the sensor is

$$F_s = m(g + a) + A_p^s F_{\text{part}} \tag{1}$$

where m is the combined mass of the gripper and the part in the gripper, a is the acceleration of the mass due to robot motion, g is the acceleration of the mass due to gravity, and A_p^s transforms the force F_{part} experience by the part to the force F_s experienced by the F/T sensor. Although for simple motions, such as straight up and down, the force on the sensor due to gravity can be accounted for, the acceleration of the mass at the exact moment the sensor is interrogated will in most cases be unknown. Since the inertial force due to this acceleration can be quite large, when a force-terminating condition is satisfied, it may not be that a constraining force is being actually exerted on the part but rather a result of the acceleration of the part and the gripper. To get around the difficulties caused by inertial effects, it often becomes necessary to take a force/torque reading under static conditions after the termination of the move to verify the successful completion of a goal.

[1] When we use the word position, we will, in general, be referring to both position and orientation. Similarly, when we use the word force, we will be referring to both force and torque. Some authors prefer to use the word "pose" to describe both the position and the orientation, and the word "wrench" to describe both the force and the torque.

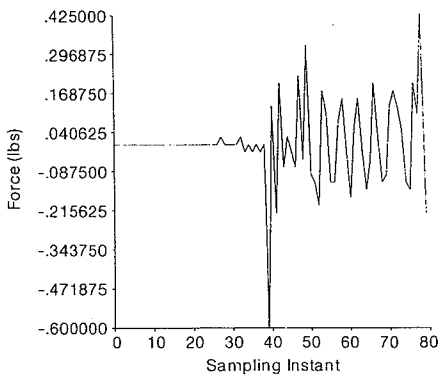

Fig. 1. Readings produced by F/T sensor while end-effector is executing free straight-line motion. Sampling interval for time axis is 20.41 ms.

To demonstrate to the reader the deleterious effects of inertial, vibrational, and electrical noise, we have shown in Fig. 1 a set of readings produced by the force/torque sensor while the end-effector is executing a free straight-line motion; the time interval between consecutive sampling intervals is 20.41 ms. The motion begins in the center of the plot; therefore, a before-and-after comparison of the variability in the force readings can be made. The force component shown in the figure is along the direction of motion. The first large spike, located at the thirty-ninth sample, shows the effects of inertia. Subsequent oscillation probably represents increased electrical noise generated by larger drive currents to the motors while the robot arm is in motion. These oscillations can also be due to the underdamping of the F/T sensor. Note that strain gauges used in F/T sensors are usually mounted on flexible beams, implying that underdamped vibrations could be excited by step inputs.

If the variability in force readings during free motion is as large as that shown in Fig. 1, it is quite possible that the terminating force may be experienced in the middle of a sampling interval. However, since the decision to terminate a motion can only be made at a sampling instant, the actual force when the motion is terminated may be much larger than what is specified by the terminating condition.

When a part contacts another surface it makes an impact. An impact is defined as a collision between two bodies, which occurs in a very short period of time during which the two bodies exert a relatively high force on each other [16]. The line of impact is the common normal of the contacting surfaces of the two bodies during the impact. If the centers of masses of both bodies are along the line of impact, it is referred to as a *central* impact; otherwise, it is an *eccentric* impact. When the velocities of the two bodies are directed along the line of impact, it is a *direct* impact; otherwise, it is an *oblique* impact. For the result of an impact to only be a change in the contact force, it must be a direct central impact.

When an impact is not both direct and central at the same time, the part may slide imperceptibly along the contacting surface. The part may also stick or reverse its sliding direction before the impact ends [17], which can result in the part rotating in an unpredictable direction about the point of greatest compliance; the rotation may bring into play high forces acting on the part. Termination of a guarded move with inadvertently high forces can create difficulties with the initiation of subsequent compliant moves.

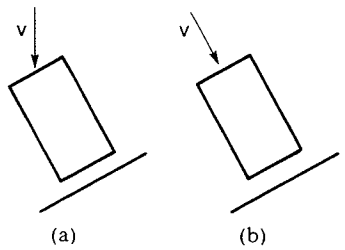

Fig. 2. (a) Motion that will not result in direct impact. (b) Motion that will result in direct impact.

Ill effects associated with guarded motion can be controlled in many ways. The directness of impact can be controlled by specifying only those guarded moves that travel along the line of impact (see Fig. 2). Assuming that direct impacts are attempted, the use of proper gripping fixtures can prevent a part's center of mass from rotating unpredictably, ensuring a central impact. By reducing the incremental distance traveled between force–torque samplings, the terminating force will be much closer to the desired force, and any sliding motions will tend to be greatly reduced. However, this has the potential of reducing significantly the speed of the motion. Finally, if some estimate of the acceleration of the robot can be made, the expected inertial force can be subtracted from the F/T sensor reading.

B. Compliant Motion

As mentioned in the Introduction, compliant motion is a motion made toward a specified goal position while maintaining a given force constraint in a perpendicular direction. A compliant motion terminates when one of a given set of position or force conditions is satisfied. The force control required for compliant motion may be implemented by any of a number of algorithms mentioned in Section I-B, where it was also mentioned that one may use either position, torque or velocity feedback to close the primary control loop. When joint position sensors are used in combination with a force-sensing wrist, then according to Maples and Becker, the predominant force loop closure strategies are 1) those that issue a torque command to adjust forces; 2) those that issue a velocity command to adjust forces; and 3) and those that issue a position command to adjust forces.

The first strategy is typically implemented using a single feedback control loop. The latter two strategies, which have similar control architectures, generally require an inner velocity or position control loop and an outer force control loop. Fig. 3 depicts the difference between a single-loop system and an inner/outer loop system. Fig. 3(b) describes our implementation of force control.

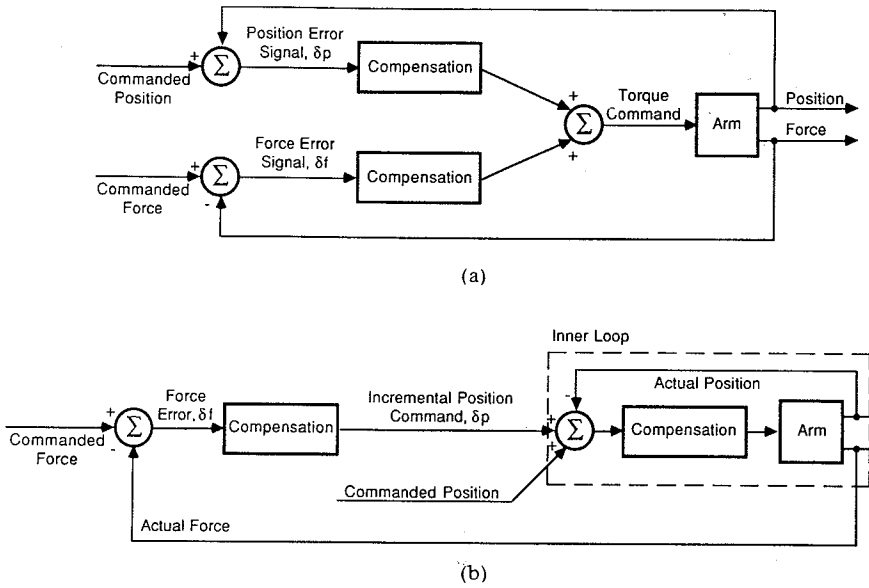

Fig. 3. Control diagram of typical single-loop system that issues torque commands to adjust forces. (b) Control diagram of our inner/outer loop system that issues position commands to adjust forces.

The advantage of the inner/outer loop structure of the type shown in Fig. 3(b) is that, while the inner loop can be used to reject internal disturbances (because it is closed on position sensors which are co-located with the actuators where the disturbances originate), the outer loop can be used to reject external disturbances because it is closed on the F/T sensor, which measures immediately the consequences of physical interactions between the end-effector and the workpiece. Note that another advantage of this type of control structure is that the outer loop can be used for driving a robot with vision-based feedback, as discussed by Khatib [18].

Before discussing the consequences of contact-related external disturbances, we first need to expand upon the conversion of the force error signal in the outer loop to a position command for the inner loop. Using the notation introduced in Fig. 3 where δf is the force error and δp is the corresponding incremental position command, we can relate the two parameters by the equation

$$\delta f = K\,\delta p \tag{2}$$

where K is a stiffness constant. Equation (2) forms the foundation of practically all systems where force errors are converted into position commands. Given its importance, we will first supply the reader with a plausible underlying theoretical rationale for the equation, to be followed by a discussion of the errors to which the resulting compliant motion may become sensitive. Although this first-order approximation is derivable from a spring model, we will use impact-dynamic considerations to arrive at (2), since these considerations, in our opinion, better represent the phenomena underlying the mechanics involved in a compliant motion and, therefore, are better able to lead to insights regarding the dependence of K on the various parameters associated with an impact.

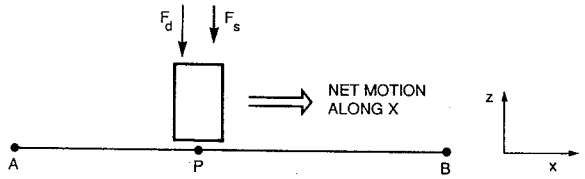

Fig. 4. Compliant motion that slides peg along table in direction $\vec{X}$ while maintaining force in direction $\vec{Z}$.

1) Is the Model Equation $\delta f = K\,\delta p$ Justifiable? Because the specification of force constraints is premised on the occurrence of contact between the part in the gripper and its environment, we can use impact dynamics to examine the nature of the forces at play. The well-known equation that relates the velocity of an object, as it collides with a stationary surface, to the forces exerted on the object is

$$\int_{t_1}^{t_2} F\,dt = m(v_2 - v_1) \tag{3}$$

where, in our case, F is the force exerted on the part held by the gripper, t_1 the time at the start of the impact, t_2 the time at the end of the impact, v_1 the velocity with which the part strikes the contacting surface, v_2 the rebound velocity, and m the combined mass of the robot arm, the gripper, and the part.

Consider the following situation typical of compliant motion. The robot is commanded to move the part in direction $\vec{X}$ while maintaining a force $\vec{F}_d$ in a perpendicular direction, as shown in Fig. 4. Let us say that during this move from A to B, the part is at position P, and let us further say that, as reported by the sensor at this instant, the measured component of the force perpendicular to line AB is F_s. For the sake of discussion, we will also assume

that $F_s < F_d$. (If this is not the case, we will merely have to reverse the sign of δp in the discussion to follow.)

That F_s is less than F_d implies that, at the instant under consideration, although there is contact between the part and the flat plane, the contact is not "forceful" enough since the force constraint is not satisfied. To enforce the force constraint, the robot is commanded to move by an increment δp in the downward direction in Fig. 4 while it continues to move from A to B. The mechanics of what happens next can be studied with the help of (3).

Let us say that the time instant when the end-effector pushes the part further toward the contacting surface is t_1 and that time t_2 is the instant that marks the end of this move. We will assume that the component of the incident velocity owing to the commanded motion δp is v_1. Since, in our case, the robot is holding the part, allowing for only infinitesimal rebound in the form of compliance, we will set v_2 to zero. We may therefore write the following expression for the average force during the impact.

$$F_{\text{av}} = \frac{1}{t_2 - t_1}\int_{t_1}^{t_2} F\,dt = \frac{-mv_1}{\Delta t} \tag{4}$$

where $\Delta t = t_2 - t_1$. Δt is equal to the interval used for sampling the F/T sensor signals. Except for the effects of plasticity, which we will address shortly, the force F_{av} will be the increase in the force experienced by the end-effector at the end of the first impact. Clearly, if $F_s + F_{\text{av}}$ does not equal F_d, the controller will cause the end-effector to be pushed down again and a second impact will follow. If we assume that it takes k impacts to build up to the desired force level, we may write

$$F_{\text{total}} = \sum_k F_{\text{av}} = \frac{-m}{\Delta t}\sum_k v_1. \tag{5}$$

Since velocity integrated over time is net displacement, which we denote by δp, we can write

$$F_{\text{total}} = \frac{-m}{\Delta t^2}\,\delta p. \tag{6}$$

In the absence of any plastic deformations, F_{total} is—although only in an approximate sense—the increase in the force on the part as the end-effector tries to push it by δp into the contacting surface. However, the plastic deformation cannot be ignored, since there is always some slippage between the part and the gripper. Plastic deformations of this type are to be contrasted with the elastic deformations in the robot itself, the F/T sensor, and the contacting surfaces. We will assume the ratio of elastic deformation to total deformation is relatively constant; this constant, which will be denoted by ϵ, is commonly known as the coefficient of restitution.

We will denote by F_Δ the actual increase in the force on the part. F_Δ is given by the expression in (6) multiplied by the coefficient of restitution:

$$F_\Delta = \epsilon F_{\text{total}} = -m\frac{\epsilon}{\Delta t^2}\,\delta p. \tag{7}$$

Since the incremental motion δp was intended to cancel the force error δf, we would want

$$\delta f = F_\Delta = \frac{-m\epsilon}{\Delta t^2}\,\delta p \tag{8}$$

implying

$$\delta f = K\,\delta p \tag{9}$$

where $K = -m\epsilon/\Delta t^2$. Therefore, the overall conclusion of this subsection is that (2) forms a reasonable basis for conducting compliant moves.

In the event that the inner loop is velocity controlled, we could instead use the equation

$$\delta f = K\,\delta v \tag{10}$$

where $K = -m\epsilon/\Delta t$. For the case where force error is directly converted into a motor torque, this derivation has no relevance. Salisbury [19] has found that the expression for the motor torque to be applied to the ith joint in this case is given by

$$T_i = T_{C,i} + G_i\,\delta T_i + K_{V,i}C_{\text{II},i}\,\delta\dot{\theta}_i + V_{0,i}\,\text{sgn}\left(\dot{\theta}_i\right) + C_{I,i} \tag{11}$$

where $T_{C,i}$ is the commanded torque, δT_i the torque error, $\dot{\theta}_i$ the velocity, $\delta\dot{\theta}_i$ the velocity error, G_i the torque compensation function, $K_{V,i}$ the velocity damping term, $C_{\text{II},i}$ the instantaneous inertia, $C_{\text{I},i}$ the gravity loading, and $V_{0,i}$ the friction torque. All of these quantities relate to the ith joint.

2) Sources of Error in Compliant Motion: As is easily seen in the previous derivation, our model for force control is only an approximation. Perhaps the most questionable assumption we made was to use a constant for the coefficient of restitution. Surely, in most cases, the extent to which a deformation is plastic on the one hand and elastic on the other depends upon the magnitude of the deformation already undergone. Therefore, as the deformation continues, the coefficient of restitution would change, making the problem mathematically intractable since plastic deformations are hard to model and predict. The coefficient of restitution must also depend upon the nature of materials used, a fact we and other people ignore in force controlled motion.

Looking at compliant motion from the standpoint of how it is implemented, note that the force is measured at discrete sampling instances. Each incremental force is completely determined by the force measured at the start of the sampling period. Therefore, even if the forces on the part changed drastically during the sample period, the robot would not be able to react to them.

Additionally, factors such as electrical noise and mechanical vibrations affect F/T sensor readings during compliant motion. We will now illustrate the fluctuations in the F/T sensor readings due to the mechanical vibrations induced by the much faster position-controlled inner loop during the successive position updates generated by the relatively slower force-controlled outer loop. As was mentioned before, our force control system is composed of an outer force-servoed loop that issues a differential position command to an inner position-servoed loop for ad-

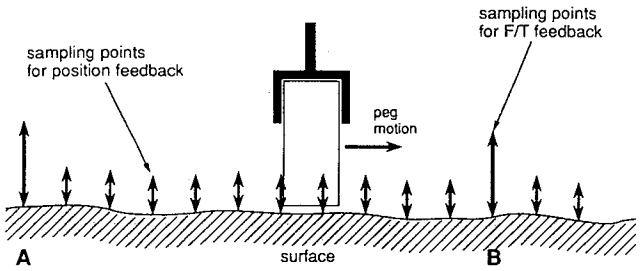

Fig. 5. Motion of peg during force–torque sampling period.

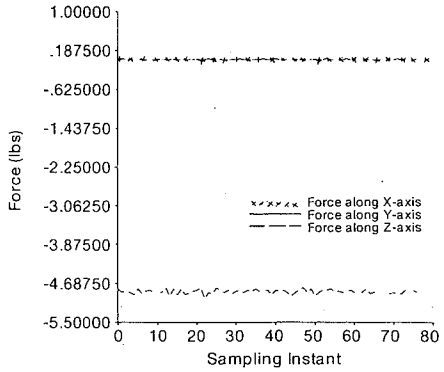

Fig. 6. Effects of electronic noise and noise created by oscillating motion of robot while it is in contact with rigid surface can be seen in this plot. Data were taken with robot stationary and robot end-effector pressing into aluminum block. Time interval between any two consecutive sampling instants is 20.41 ms.

justing the forces. It is important to realize that the servo rate of the inner loop is often much faster than that of the outer loop. In Fig. 5, we show two points, marked A and B, on the path of a part engaged in a compliant move; these two points correspond to the time instants when the position update commands are issued by the outer force-controlled loop; this updating process is portrayed by the large arrows at points A and B. The time that it takes to travel from A to B represents the sampling interval for the outer loop; during this interval, the much faster inner loop is active and issues frequent position updates to the robot, this fact being represented by a series of small arrows in the figure. This updating generated by the inner loop will lead to mechanical vibrations while the part is in transit from A to B. To illustrate the magnitude of such vibrations, we have shown in Fig. 6 the F/T sensor output when the end-effector is stationary at the end of a compliant move. In this situation, the only updates being issued to the robot are from the inner loop, a situation that is very similar to what happens between A and B in Fig. 5. In Fig. 6, the part is only exerting Z-forces on the F/T sensor. Therefore, the mechanical vibration induced by the position updating is most pronounced in the Z-force component, particularly because the part is in contact with a surface below. Fig. 6 demonstrates that even in the absence of any surface irregularities, force readings, which will most likely suffer from the aforementioned mechanical vibrations, will affect the force control during a compliant motion.

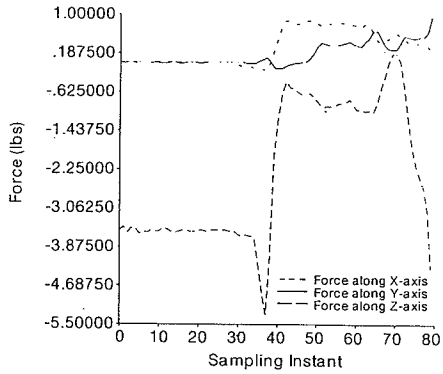

Fig. 7. This plot depicts forces exerted along Cartesian axes as robot is executing compliant motion in Y direction, while attempting to maintain constant force in Z. Time interval between any two consecutive sampling instants is 20.41 ms.

For such reasons, perfect surface following during a compliant move is impossible. Moving a part along a surface imperfectly may cause either the loss of contact with the surface or, conversely, cause the part to exert greater force than desired on the surface. In both cases, a terminating condition may be sensed prematurely. In the latter case, the part may begin to stick to the surface.

Our discussion in Section II-B-1 analyzed forces and velocities along only the force-controlled direction, this assumption being more justified the slower the translational velocity from A to B in Fig. 4. For a completely rigorous analysis, the positional-constraint velocity vector and the force-constraint velocity vector must be considered together. The combined velocities result in impacts being oblique, which may cause the part to slide along an unpredictable direction; in general, there will be a rotation of the end-effector about the most compliant point in the system.

Fig. 7 depicts the forces exerted on the end-effector while the robot makes a compliant move in the Y direction while attempting to maintain a constant force in Z. As can be seen in the plot, some unknown factor, which could be an impact with a surface perturbation or simply a noise spike in the F/T sensor output, causes a force discontinuity midway through the move, and once the actual force strays from the desired force, the robot begins an oscillatory movement along the Z axis as it tries to regain the original force threshold.

Perhaps the best way to deal with all the problems mentioned is to decrease the distance traveled between sampling increments. Also, some filtering could be applied to the F/T sensor readings to reduce the noise. Additionally, the stiffness factor K could be recalculated for each separate task taking into account the materials used.

III. Planning Motions in the Presence of Uncertainties

In any assembly operation, three types of uncertainties will be present to varying degrees; sensing uncertainties, control uncertainties, and model uncertainties. Sensing uncertainties, primarily the result of imperfect sensing, are

uncertainties in the measured position of a part and the measured force acting on the part. Control uncertainties are a consequence of control inaccuracies such as those discussed in Section II. In our case, control uncertainties translate into the commanded destination for the end-effector being different from the position it actually occupies in response to each command. Finally, model uncertainties, arising from imperfect manufacturing processes, are the possible uncertainties due to the finite tolerances allowed between the design dimensions and the actual manufactured dimensions of the features of a part. Model uncertainties manifest themselves as errors in both the shapes and sizes of features and parts.

Given an assembly task with a start state S and goal state G, we would like to attain G from S. To do this, we first introduce the concept of straight line motion goals (SLMG's). There is a sequence of SLMG's, described by the states $S, G_1, \cdots, G_i, G_{i+1}, \cdots, G$, such that a transition between any two subgoal states G_i and G_{i+1} can be accomplished with a single straight line motion in the absence of uncertainties.

Note that while the notion of an SLMG was defined in the absence of uncertainties, its execution in practice would require a nondeterministic plan owing to the presence of uncertainties. Nondeterminism refers to the invocation of one of many possible choices to reach the goal state from the current state.

Our contention is that each SLMG must be planned and executed dynamically, meaning that its execution must depend on the current sensory readings about the state of the robot in relation to its environment. If an error state is reached during the process of achieving an SLMG, the system should be capable of generating dynamically a corrective plan. Therefore, our overall strategy is to use a static planner to decompose an arbitrary assembly task into SLMG's using *a priori* known part geometry (model) information and then depending solely on dynamic execution to carry out each SLMG.

The decomposition of an assembly goal into SLMG's is, of course, not unique, as illustrated in Fig. 8. Of all the possibilities for such a decomposition, some are better than others because they allow us to reduce sensory and model uncertainties. For example, in Fig. 8, of the two possible decompositions, the one drawn in solid lines reduces uncertainties in the Z direction at the end of the first subgoal. Subsequently, the same decomposition attempts to reduce the uncertainties in the X and Y directions at the end of the second SLMG. At this time, there does not exist an automated procedure capable of generating the best decomposition for a given assembly goal; development of such procedures is a part of the ongoing research effort.

As mentioned before, the execution of each SLMG will, in general, require error detection and recovery. It is important to realize that we do not need to know the sources of the errors but merely that the error states be reliably identified. For example, a guarded move intended

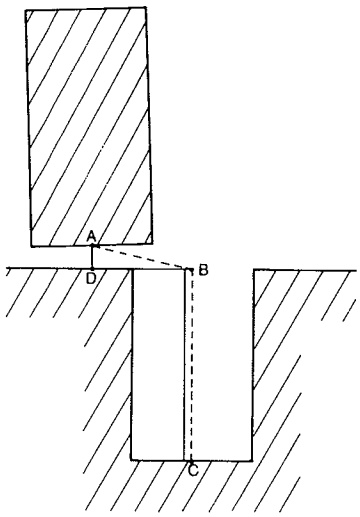

Fig. 8. Path through ADBC here illustrates uncertainty reducing sequence of SLMG's. Another sequence of SLMG's, passing through points A, B, and C would produce no uncertainty reduction, at least in initial phases of plan.

to bring a part into contact with a surface may terminate prematurely due either to inertial forces exerted on the part or a noisy force–torque reading. Regardless, a static force measurement after the apparent termination of the guarded motion will signal the presence or absence of a physical contact. If the static reading indicated zero force, for example, then clearly the part is positioned somewhere above the point of contact, which is all that would be needed to identify the error state.

At this time, the framework espoused in this paper for dealing with uncertainties is more in the nature of developing strategies for the detection of various possible errors, errors that have been observed to date through experimentation. We have not shown in any formal manner that our EDR strategies would be able to deal with the consequences of all possible uncertainties. We believe that we need to gain further experience at the experimental level before launching into a formal development of the subject.

IV. A Dynamic Approach to the Peg-In-Hole Problem

In this section we will focus on the peg-in-hole problem by outlining a method for inserting an unbeveled peg into an unchamfered hole. The less general but common use of a beveled peg or chamfered hole may require that this method be modified slightly, but in general the problem would become easier. We assume that we have a vision system (or some other device) capable of locating the peg accurately enough for the robot to pick it up. In addition, we assume we can locate the hole with sufficient accuracy such that the relative $X-Y$ position of the peg, with respect to the hole, has uncertainty less than the radius of the hole when the peg is held by the robot. We also know some upper bound on the uncertainty of the Z position o'

the peg with respect to the hole.[2] The problem can be solved by carrying out a plan to transition from this initial state through the intermediate SLMG states depicted in Fig. 8 by the solid line, and, finally, to the goal state.

A. Align the Peg with the Hole Along the Axis of Insertion

Two intermediate SLMG's are needed to accomplish the assembly: one to achieve a recognizable Z position; a second to attain the required $X-Y$ position. The first SLMG starts with the peg in the gripper, held some distance above the hole, with the same $X-Y$ position as the hole save sensory and model uncertainties. This SLMG ends either when the peg is in contact with the surrounding surface of the hole, or when the peg has begun to enter the hole. The strategy used to achieve this state is to make a guarded move along the Z axis toward the hole, terminating when either a contact force is sensed, or the position in Z is such that the peg has successfully entered the hole.

Two possible error states can result from this guarded move. The move may end before either terminating condition is satisfied if an invalid contact force is sensed. Reading the F/T sensor after the motion has halted will reveal this situation, and, upon its detection, a plan to continue the guarded move is generated. On the other hand, the robot may continue too far after the force terminating condition is sensed, causing high forces and torques to act on the peg. Again, this state is easily recognized by examining the static force–torque measurements, and, if such a recognition is indeed made, a plan to lift the peg and retry the guarded move is synthesized. This is essentially the phase-1 EDR strategy alluded to in the Introduction.

If, after this SLMG is successfully completed, the peg is not in the hole, we need to invoke the second SLMG to bring it there. The start state of this SLMG has the peg contacting the surface containing the hole and positioned partially over the hole. By assumption, the start state obeys the constraint that the initial $X-Y$ uncertainties in the positioning of the peg relative to that of the hole are bounded by the radius of the hole. At the termination of this SLMG, the peg is positioned over the hole in such a manner that the bottom face of the peg, when projected onto the surface containing the hole, is completely within the hole itself. To carry out this SLMG, we must first ascertain the direction of the error vector of the peg, the error vector, denoted by $\vec{e}$, being defined as the vector from the center of the bottom face of the peg to the center of the hole's opening.

To estimate the direction of the error vector, we measure the torques about the X and Y axes of the peg. As will be derived in the next subsection, in terms of the torques the angle that the error vector makes with the X axis is given by

$$\theta = \tan^{-1}\left(\frac{\tau_{\vec{y}}}{\tau_{\vec{x}}}\right) - 90^\circ \tag{12}$$

where $\tau_{\vec{y}}$ is the torque about the Y axis and $\tau_{\vec{x}}$ is the torque about the X axis. Given the angle θ, the strategy is to slide the peg along the error vector until a discontinuity in the contact force is sensed, or until the move has traveled a distance along the X or Y axis that is greater than a user-supplied threshold corresponding to the maximum possible discrepancy in the X and Y directions between the positions of the peg and the hole.

In the event that the torque vector is not sufficiently accurate the motion is terminated by the satisfaction of the position termination condition. Also a noisy force reading could cause the force termination condition to be satisfied before the peg is entirely over the hole. In either case, θ must be recalculated and the SLMG reexecuted. However, before recalculating θ, the system must check that the previous move did not cause high forces or torques to be exerted on the peg, obscuring the recalculation of θ. These high forces or torques are the result of the peg sticking during the compliant motion causing the peg to rotate slightly. If this is the case, the peg is raised by a small distance, followed by guarded move to the surface containing the hole. Of course, the SLMG is reattempted only if this last guarded move did not serendipitously place the peg into the entrance of the hole.

Another possible error state will arise if the peg loses contact with the surface surrounding the hole as it is being slid. This state may be hard to recognize, but in general, the absence of a contact force and a greater relative Z coordinate will result. Regardless, nothing would be lost by assuming the occurrence of this state whenever the absence of a contact force is sensed. Recovery from such a state consists of making a guarded move along the Z axis toward the hole, and terminating the move when either a contact force is sensed, or the relative Z position is such that the peg can only be in the hole. Once again, this guarded move may have brought the peg into the subgoal state, but if it does not, the strategy is reexecuted. It is possible that repeated attempts to bring the peg over the hole may eventually move the peg completely away from the hole; however, such a divergence can be easily detected and corrected for.

In the following, we will analyze the moments that are created when the peg is partially over the hole and present a derivation of (12). We will also discuss the factors that affect the accuracy of the estimation of the direction of the error vector.

1) Torque Exerted on a Point in the Peg: Let us consider the physics of a peg positioned partially over a hole. We will assume that we have a robot arm with a wrist mounted F/T sensor and that the gripper is attached to the F/T sensor. To simplify matters, we will model the part in the gripper as a continuation of the gripper, such that the gripper and the part are one solid object.

[2] Here the Z axis is vertical, and the X and Y axes are perpendicular and right-handed. We are assuming that the hole is in the XY plane and the axis of the peg is parallel to the Z axis; although, of course, both these assumptions do not have to be satisfied strictly.

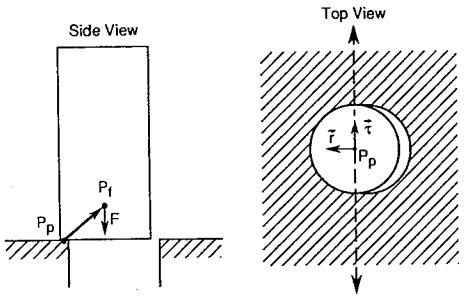

Fig. 9. Configuration of peg positioned over hole.

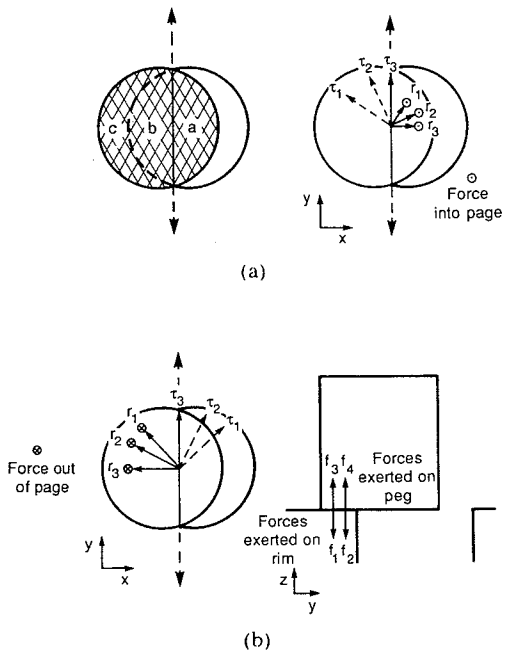

Fig. 10. (a) Definition of regions a, b, and c. (b) Torques resulting from forces in region a. (c) Torques resulting from forces in region b. (d) Forces acting in region c.

Assume there is a force $\vec{F}$ acting on the gripper assembly at point P_f, and let P_p be the fixed point on the peg that is located at the center of the chord joining the two contact points between peg and the rim of the hole (Fig. 9). We now define $\vec{r}$ to be the vector directed from P_p to P_f. The torque acting on the point P_f with respect to the pivot point P_p is given by

$$\tau = \vec{r} \times \vec{F}. \tag{13}$$

2) Torque Exerted on the Peg: We will now consider the peg to be a collection of points, like the one at P_f. We will assume the force exerted on the peg is uniformly distributed throughout the peg; this assumption is an approximation of Hertz's theory. To calculate the total torque on the peg, we need to integrate (13) over all the points in the peg. To do this, we first define the three regions shown in Fig. 10(a). Region a contains the points to the right of the pivot axis, region b contains the points to the left but not over the surface outside the hole, and finally, region c contains the rest of the points, these being the points directly over the surface outside the hole. Without loss of generality, we will only consider the points on the plane of

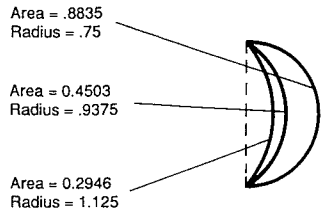

Fig. 11. Numbers on left side are areas and radii of different circular sectors having same chord length.

the bottom of the peg. The direction of the resultant torque vector—note we are not interested in the magnitude—is unaltered by this restriction.

As shown in Fig. 10(b), for all the points in region a, the $\vec{r}$ vectors will point to the right, the force vectors $\vec{F}$ caused by the forces exerted by the robot will all be into the page, and the resulting torque vectors will all have a Y component greater than or equal to zero. If they are added vectorially, the X components will cancel, and there will be a net torque vector pointing in the positive Y direction.

As shown in Fig. 10(c), for all the points in region b, the $\vec{r}$ vectors will point to the left, the $\vec{F}$ vectors will all be into the page, and the resulting torque vectors will all have a Y component less than or equal to zero. Summing the torque vectors yields a net torque vector pointing in the negative Y direction.

For the points in region c, as shown in Fig. 10(d), there is a force pointing downwards and an equal and opposite force pointing upwards. Since no net force is acting on any of these points, no torque is associated with them.

Note that the area of region a is always greater than the area of region b. Although this can be proved analytically, it is easiest to understand it graphically. Given two circular sectors (from two disks of different radii), both with the same chord length, the sector with the smaller radius of the arc will have a larger area. To check the validity of this statement, the reader may visualize the limiting case when the radius of one disk becomes arbitrarily large in relation to the radius of the other disk, making the circular arc of the larger disk nearly parallel to the given chord line of fixed length. In Fig. 11, the numbers on the left-side are the computed areas of sectors corresponding to different radii but with the same chord length.

In our case, since the radius of the arc for region a will always be less than the corresponding radius for region b, it follows that the area of the former region must be greater than the area of the latter. This implies that the torque vector pointing in the positive Y direction, as generated by region a, will always be greater than the torque vector in the negative Y direction generated by the points in region b. Therefore, there will always be a net torque in the positive Y direction.

Therefore, the net $\vec{r}$ vector, corresponding to the net torque experienced by the peg, will point in the direction of the error vector $\vec{e}$. Since the cross product $\vec{r} \times \vec{F}$ will always be perpendicular to $\vec{r}$ and therefore perpendicular to $\vec{e}$ and since the torque measurement will correspond to $\vec{r} \times \vec{F}$, the error vector will be perpendicular to the mea-

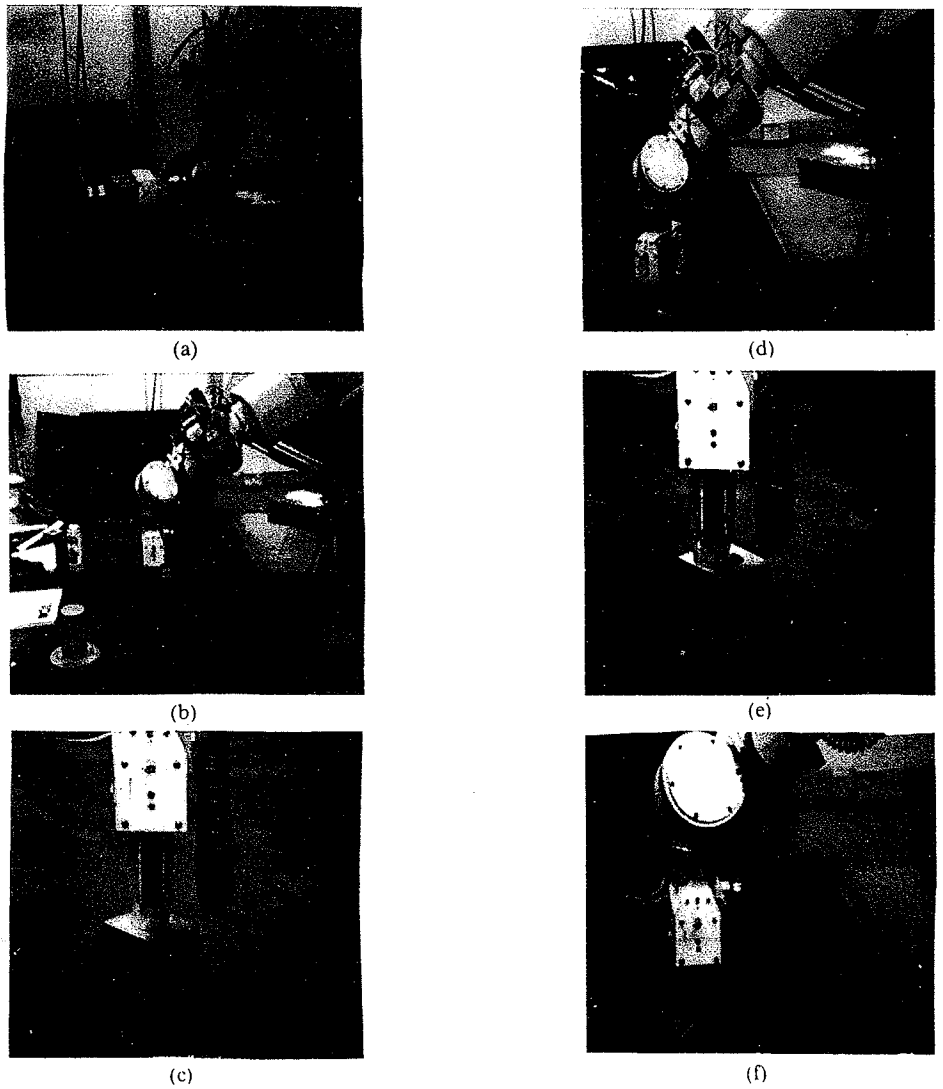

Fig. 12. (a) Peg and block with hole are lying on work platform and are being scanned with structured light scanner held in robot gripper for determination of positions and orientations of two objects. (b) Manipulation of peg. (c) Reorientation of peg prior to insertion. (d) Robot brings peg over to hole with guarded move. (e) Compliant move is being executed for alignment of axes of peg and hole. (f) Finally, insertion is carried out.

sured torque. We can therefore write

$$\vec{e} \| \tan^{-1}\left(\frac{\tau_y}{\tau_x}\right) - 90°. \tag{18}$$

Before concluding this section, we would like to make an important comment about a practical source of difficulty one has to contend with in the measurement of the direction of the error vector. As the clearance between the peg and the hole diminishes, so will the difference between the areas of the regions a and b; the result being a torque that approaches zero. In practice, the smaller the value of a torque, the more error-prone its measurement becomes, requiring recourse to noise reduction techniques such as averaging.

B. Inserting the Peg into the Hole

The final goal of a peg-in-hole operation can be reached with one additional SLMG. Starting with the peg positioned over the mouth of the hole, the robot needs to push the peg into the hole without jamming or wedging. This is done by making a compliant move that attempts to maintain zero forces and torques on and about the X and Y axes while pushing the peg into the hole, stopping when the peg contacts the bottom. An error state will result if a false contact force is sensed or the peg wedges or jams into the hole. In the first case, a static reading of the force will signal the absence of a contact force and the move is simply repeated. In the second case, high forces or torques will be measured. If the peg is wedged or jammed, it is backed out a small distance and the move is repeated.

V. Experimental Results

The goal of our research is to be able to mate tight-fitting parts whose initial positions are unknown. To accomplish this, we use range maps obtained with a 3-D laser scanner to locate each part to an accuracy of roughly 0.2 in. The manner in which the range maps are processed, described fully in [20], consists essentially of first detecting jump and curvature discontinuities, isolating the topmost surface of each object, and extracting various parameters

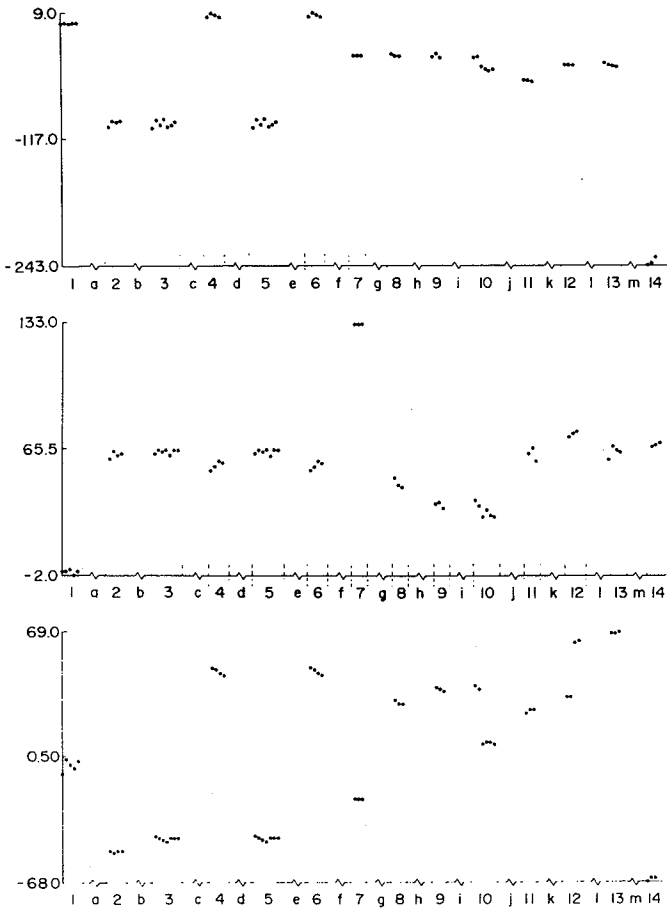

Fig. 13. Top plot depicts readings of force along Z axis taken during one peg-in-hole experiment. Middle plot shows readings of torque about X axis during same experiment. Bottom plot shows corresponding readings of torque about Y axis. Note that at least three readings are taken to ensure that signal is stable; if not, more readings are taken until signal becomes stable. Numbered intervals depict force–torque readings, lettered intervals signify motion or transition between phases. For description of significance of each interval, see text.

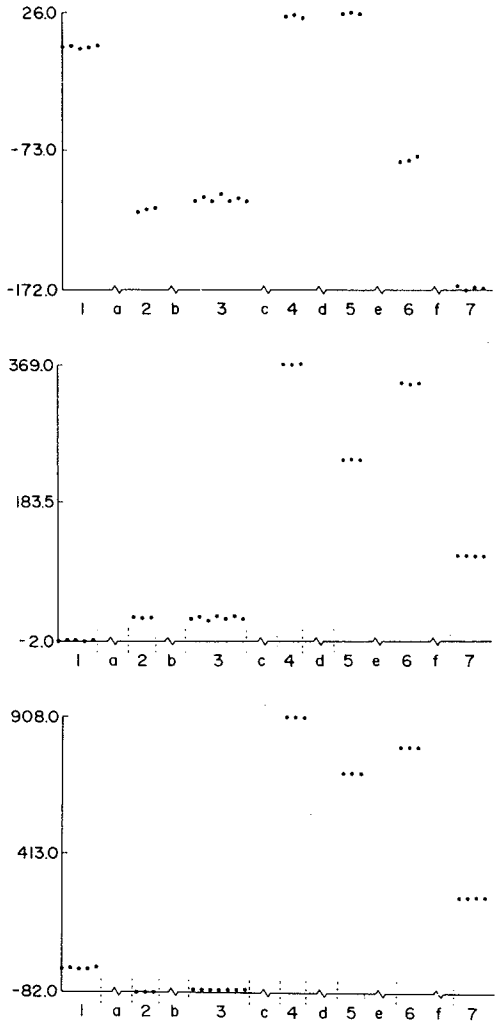

Fig. 14. Top plot depicts readings of force along Z axis taken during second peg-in-hole experiment. Middle plot shows readings of torque about X axis during this experiment. Bottom plot shows corresponding readings of torque about Y axis. As in Fig. 13, numbered intervals depict force–torque readings, lettered intervals signify motion or transition between phases. For description of significance of each interval, see text.

of the topmost surface for recognition and for the computation of position and orientation of each object.

The scene assertions generated by the procedures for range data processing then automatically synthesize a manipulation plan first to grasp the peg, manipulate the peg for the purpose of orienting it properly prior to assembly, and finally to move the peg over the hole. From that point on, the rest of the insertion is carried out under the control of force/torque sensing, as described in this paper. Fig. 12 shows a sequence of photographs taken during such an experiment in which the peg was lying on its side about 4 in away from the block with the hole. Figs. 13 and 14 display the readings taken by the force/torque sensor during two such experiments. Although both experiments were executed under similar conditions, different error situations occurred in each, and we will therefore explain both. Each numbered interval corresponds to the force–torque readings taken; three or more readings were taken each time the F/T sensor was queried to ensure the signals were stable.

In Fig. 13, interval 1 shows the readings at the beginning of the experiment. Interval a signifies the execution of the guarded move to the surface containing the hole. The readings taken during the execution of the phase-1 EDR routine are depicted in interval 2. Here, it was determined that the motion successfully established the phase-1 terminating condition. Interval b denotes the transition between phase-1 and phase-2. The readings taken in interval 3 were used to compute the error-vector direction required for the phase-2 compliant motion. During interval c this compliant motion was executed. Interval 4 displays the readings taken by the phase-2 EDR routine. It was decided that the terminating condition was not met, but rather the peg lost contact with the surface during the preceding compliant motion. This situation was rectified by making a guarded motion to reestablish contact in interval d. The readings taken to recalculate θ are exhibited in interval 5. During interval e the phase-2 compliant motion is reattempted. The readings taken at the completion of this move by the phase-2 EDR routine are displayed in interval 6. The phase-2 EDR routine resolved that the terminating condi-

TABLE I
NUMBER OF PHASE-2 EDR CYCLES AUTOMATICALLY INVOKED FOR SUCCESSFUL ASSEMBLY IN THE 20 PEG-IN-HOLE ASSEMBLY EXPERIMENTS[a]

Number of Attempts Automatically Invoked	Clearance = 0.001 in		Clearance = 0.050 in	
	0.15 < Error < 0.25 in	0.25 in < Error	0.15 < Error < 0.25 in	0.25 in < Error
1	4	3	5	4
2	0	0	0	1
3	0	1	0	0
> 8	1	1	0	0

[a]"Error" refers to the initial offset between the axis of the peg and the axis of the hole. To help the reader understand the table entries, an entry of 4 in the second column means that in four of the experiments, only one EDR cycle was required; the initial offset between the peg and the hole axes in these four experiments was between 0.15 and 0.25 in.

tions had been successfully achieved, and the experiment transitions to phase 3 at interval *f*. Initial phase-3 readings are shown in interval 7. The difference between the final phase-2 reading and the initial phase-3 reading was probably caused when the brake motors were engaged during the phase transition, slightly changing the position of the robot end-effector, therefore slightly altering the force being exerted by the robot on the peg. In interval *g* the phase-3 compliant motion to insert the peg took place. Interval 8 shows the readings taken during the execution of the phase-3 EDR routine. Here it was determined that the previous compliant motion terminated prematurely and the phase-3 terminating condition was not met. Similarly, during the intervals *h*–*m* the phase-3 compliant motion was repeatedly reattempted. In intervals 9–13 the readings taken after each attempt are displayed. In every case it was decided that the terminating condition had not been met and the compliant motion should be reattempted. Finally, in interval 14, it was determined that the phase-3 terminating condition was met and the experiment was terminated.

In Fig. 14, once again interval 1 shows readings at the beginning of phase-1. Interval *a* was used for the guarded move made during phase-1. Following the guarded motion, the readings shown in interval 2 were taken during the execution of the phase-1 EDR routine and it was determined that the phase-1 terminating conditions had been successfully reached. Therefore the system continued on to phase-2, as signified by the transition at interval *b*. Initial phase-2 readings were taken, as displayed in interval 3, and the data used for the computation of the error-vector required for the phase-2 compliant motion, which was executed during interval *c*. At the termination of the compliant motion and through the readings recorded during the execution of the phase-2 EDR routine, which are shown in interval 4, it was determined that the motion had successfully established the phase-2 terminating conditions. The transition between phases-2 and 3 occured during interval *d*. The phase-3 compliant motion was executed during interval *e*, after taking the initial readings shown in interval 5. In this experiment the peg began to jam, which was determined by the high force and torque readings taken in interval 6. In interval *f* this situation was rectified by executing a compliant move similar to the one in interval *e*, but after readjusting the direction of the insertion. The readings taken in interval 7 were used to determine that the phase-3 terminating condition was successfully reached and the experiment ended.

To demonstrate the success rates achieved with our dynamic planning methods, we have summarized in Table I the results of 20 peg-in-hole assembly experiments. Two sets of experiments are shown in the table; ten conducted at a clearance level of 0.001 in. and the other ten at a clearance level of 0.05 in. As mentioned in the table footnote, "error" is the offset between the peg and the hole at the instant the force guided motions take over the assembly task.

As can be inferred from the table, the success rate is a function of the number of cycles allowed for error detection and recovery. We believe these results are quite impressive considering the limitations of our force control system caused primarily by susceptibility of the F/T sensor to electrical noise, inertial effects, and mechanical vibrations.

VI. CONCLUSION

This paper presented seldom discussed sources of error in force-guided motions. Among these sources are noisy F/T readings, mechanical vibrations, the presence of sliding and sticking frictions, and the possibilities of eccentric oblique impacts. We have pointed out how most of these errors might be reduced or eliminated. Although we have ignored issues dealing with stability, we would like to state that empirically we have found our system to be stable, albeit somewhat underdamped for the metal parts that we use in our assembly experiments. For a description of how the stability of force-controlled systems varies with gain, dynamics and environmental stiffness, the reader is referred to [21]. Whitney discusses methods to ensure stability in different force controlled strategies [22].

In Section III of this paper we introduced the notion of straight line motion goals and showed how SLMG's can be considered as the basic building blocks of a dynamic planning strategy. By dynamically building an assembly plan out of SLMG's, we believe it is possible to carry out robotic assemblies in the presence of sensing and model uncertainties while recognizing and recovering from errors introduced by control uncertainties. Finally, this paper

discussed in detail the error detection and recovery strategies required for peg-in-hole type of assemblies and presented experimental results that illustrate the dependence of the success of assembly tasks of the number of error detection and recovery cycles allowed.

ACKNOWLEDGMENT

We would like to thank Prof. George Lee for his helpful comments on an earlier draft of this paper. In addition, we would like to thank Larry Hollingshead, formerly the Director of the Artificial Intelligence Lab at Cincinnati Milacron, for his steadfast support of the Robot Vision Lab. Thanks are also due to all the other members of the Robot Vision Lab for contributing to an intellectually exciting atmosphere in the lab. In particular, we would like to thank Matt Carroll for his valuable hardware support and Bob Cromwell for letting us use his 3-D range data processing software.

REFERENCES

[1] Y. Wang and M. T. Mason, "Modeling impact dynamics for robotic operations," in *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 2, 1987, pp. 678–685.

[2] V. T. Rajan, R. Burridge, and J. T. Schwartz, "Dynamics of a rigid body in frictional contact with rigid walls," in *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 2, 1987, pp. 671–677.

[3] M. A. Peshkin and A. C. Sanderson, "Planning robotic manipulation strategies for sliding objects," *Proc. IEEE Conf. Robotics and Automation*, vol. 2, 1987, pp. 696–701.

[4] N. Hogan, "Stable execution of contact tasks using impedance control," *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 2, 1987, pp. 1047–1053.

[5] P. C. Watson, "The remote axis admittance—A key to robot assembly," in *Industrial Robotics, vol. 1/Fundamentals, Robotics Int. SME*, 1981, pp. 415–423.

[6] ____, "The remote center of compliance system and its applications to high speed robot assemblies," in *Industrial Robotics, vol. 1/Fundamentals, Robotics Int. SME*, 1981, pp. 404–414.

[7] D. S. Seltzer, "Use of sensory information for improved robot learning," in *Industrial Robotics, vol. 1/Fundamentals, Robotics Int. SME*, 1981, pp. 424–431.

[8] D. E. Whitney, "Quasi-static assembly of compliantly supported rigid parts," in *Robot Motion*. Cambridge, MA: MIT Press, 1981, pp. 439–471.

[9] C. S. G. Lee and R. H. Smith, "Force feedback control in insertion process using pattern analysis techniques," in *Proc. American Control Conference*, 1984, pp. 39–44.

[10] H. Inuoe, "Force feedback in precise assembly tasks," in *Artificial Intell.: An MIT Perspective*, vol. 2. Cambridge, MA: MIT Press, 1981, pp. 219–241.

[11] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *Int. J. Robotics Res.*, vol. 3, no. 1, pp. 3–24, 1981.

[12] M. Erdmann, "Using backprojections for fine motion planning with uncertainty," *Int. J. Robotics Res.*, vol. 5, no. 1, pp. 19–45, 1981.

[13] B. R. Donald, "Robot motion planning with uncertainty in the geometric models of the robot and environment: A formal framework for error detection and recovery," *Proc. IEEE Conf. Robotics and Automation*, vol. 3, 1987, pp. 1588–1593.

[14] J. A. Maples and J. J. Becker, "Experiments in force control of robotic manipulators," *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 2, 1986, pp. 695–702.

[15] H. Asada and K. Ogawa, "On the dynamic analysis of a manipulator and its end-effector interacting with the environment," *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 2, 1987, pp. 751–756.

[16] E. J. Routh, *Dynamics of a System of Rigid Bodies*. New York: Dover Publications, 1955.

[17] F. P. Beer and E. R. Johnson, Jr., *Vector Mechanics for Engineers*. New York: McGraw-Hill, 1984.

[18] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *Int. J. Robotics Res.*, vol. 5, no. 1, pp. 90–98, 1986.

[19] J. K. Salisbury, "Active stiffness control of a manipulator in Cartesian coordinates," in *Proc. IEEE Conf. Decision and Control*, 1980, pp. 95–100.

[20] R. L. Cromwell, "Low and intermediate-level processing of range maps," Purdue Univ., School of Elec. Eng., Tech. Rep. 87-41, 1987.

[21] C. H. An and J. M. Hollerbach, "Dynamic stability issues in force control of manipulators," *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 2, 1987, pp. 890–896.

[22] D. E. Whitney, "Historical perspective and state of the art in robot force control," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1985, pp. 262–268.

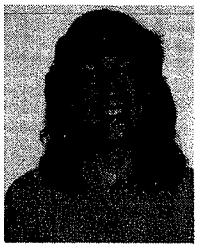

Susan N. Gottschlich received the B.S. degree in computer and electrical engineering in 1986 and the M.S. degree in electrical engineering in 1987, both from Purdue University, West Lafayette, IN, where she is currently working toward the Ph.D. degree in electrical engineering. She is working on the development of a fine motion planner and executor for high-precision assembly tasks.

Avinash C. Kak (M'71) is a Professor of electrical engineering at Purdue University, West Lafayette, IN. His research interests are in reasoning architectures for solving spatial problems, sensor-based robotics, and computer vision. He has coauthored the books *Digital Picture Processing* (Academic Press) and *Principles of Computerized Tomographic Imaging* (IEEE Press).