TECHNOLOGY FOR ROBOTIC MECHANICAL ASSEMBLY: FORCE-DIRECTED INSERTIONS

David R. Strip

David R. Strip *is a member of technical staff in the Intelligent Machine Principles Division at Sandia National Laboratories in Albuquerque. His current research interests are focused on automatic planning of force-directed robotic manipulation and basic research in solid representation and reasoning. He has a B.A. in mathematics from the University of Pennsylvania and an M.S. and Ph.D. in operations research from Cornell University. He joined Sandia in 1978.*

Definition of robotic primitives for insertions is an essential step in achieving a truly flexible manufacturing environment. We present techniques based on active compliance, implemented with hybrid force-position control, that are capable of inserting a wide variety of shaped pegs. We analyze the forces encountered during the insertions to determine what classes of shapes these techniques will consistently insert. The analysis also gives guidance in selecting parameters of these techniques for specific shapes. These techniques provide a significant step toward simplifying the programming of a flexible manufacturing environment.

Introduction

Mechanical assembly is an important component of contemporary manufacturing technology. In satisfying our design responsibilities to the U.S. Department of Energy, Sandia develops components that will be built in small quantities—frequently fewer than one thousand units. Because of these low production rates, we are concerned with how the design process can be impacted by the use of robotics in small batch manufacturing. Low production rates magnify the impact of costs of special-purpose fixturing and programming (in comparison to mass production), thereby increasing the importance of truly flexible manufacturing systems. In addition, components designed by Sandia may involve the handling of explosive, toxic, or radioactive parts. Application of robotics can reduce personnel exposure to hazardous environments.

In general, robots have not been successfully integrated into the batch manufacturing environment. The reason behind this is that currently available commercial robots are almost exclusively positionally programmed manipulation devices with little capability for sensory feedback. As a consequence, all tasks to be performed are defined by their positions in space, requiring expensive jigs and fixtures. Little or no capability exists for dealing with out-of-position parts. Some assembly operations have tolerances finer than the precision of available manipu-

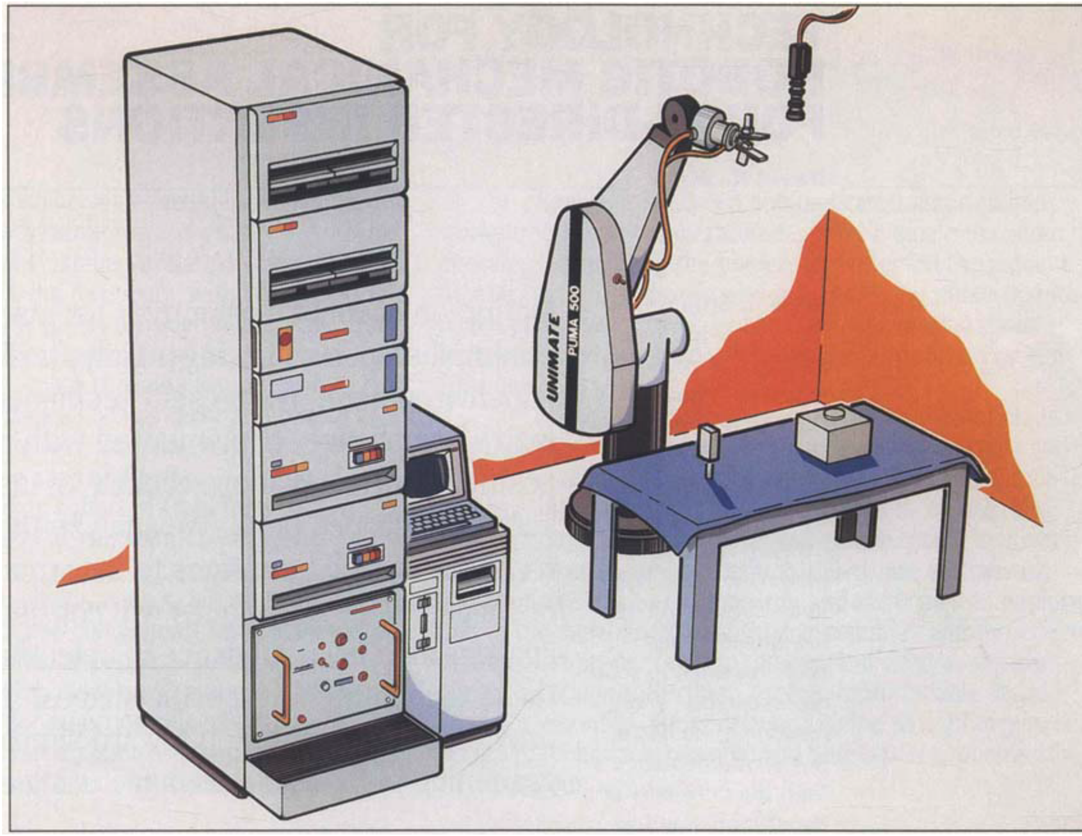

Figure 1. Laboratory setup used in experiments on force-directed manipulation.

lators, and consequently cannot be performed at all with sensorless robots.

A second failing of current commercial robots is the difficulty in programming. The most common method of programming involves guiding the robot through the positions necessary for task performance, recording the positions, and then playing them back to perform the task. This type of programming is tedious and difficult. Programming languages currently available require detailed, step-by-step motion instructions, along with specification of the location and orientation of the manipulator at each stage in the sequence of moves. Programming requires personnel not traditionally found on the factory floor and is very difficult because of the low-level nature of the specification. The work presented in this paper is a step toward providing the primitive operations required to allow task-level programming.

Issues in Mechanical Assembly

Fitting or joining together components is the key idea behind mechanical assembly. These actions generally require control of the full six degrees of freedom available in spatial movement (except in highly structured environments). Because some component positions may be defined by contact with other components, we face the added complexity that motions may be positionally constrained and force-constrained, increasing the dimensionality of our planning space. In contrast, electronics assembly is virtually a two-degree-of-freedom problem; parts lie in a single plane, and orientations are fixed to at most a few directions. Automated machining can often (though not always) be reduced to a lower than six-degree-of-freedom problem by means of a decomposition in which all but one or two degrees of freedom are constrained.

The topics of interest in mechanical assembly can be spread along a spectrum; at one end are the tasks that are performed off-line and represent global planning, and at the opposite end are the real-time tasks that are increasingly manipulative in nature and very local in information use. At the off-line end of the spectrum are issues such as design for assembly, operation sequencing, motion scheduling, and path planning. Moving toward the manipulative end are issues like grasp planning, motion control, error detection, and finally, the topics we will discuss in this paper, fine-motion control and sensory feedback. We will show how these techniques can be combined to attack two important problems in mechanical assembly: development of task-level primitives, and performance of tasks with accuracy requirements exceeding manipulator precision.

Focus: Force-Directed Manipulation

The tolerance and contact constraints common in mechanical assembly combined with a desire to build flexible manufacturing systems preclude the typical robotics approach of using absolute positioning as the basis for robot programs. The approach that we have elected for providing flexible robotic mechanical assembly is *force-directed manipulation*, the use of force and moment measures to provide feedback control to the manipulator. The principal attraction of this approach is the compactness of the data; only six force-moment scalar values are required to describe the forces acting on the end effector. Furthermore, this information is available continuously throughout the insertion procedure.

Contrast this to machine-vision systems, in which the information is contained in hundreds of thousands of pixels containing complex relationships that must be analyzed to understand the scene. For currently available technology, force sensing allows more rapid feedback at lower cost than other sensing techniques. Force sensing provides a direct sensing means for addressing the contact issues that are central to mechanical assembly.

Finally, we note that force sensing seems to be important to the ability of humans to assemble components with clearances much smaller than we can resolve directly with vision or locate via kinesthetic sense (sense of body position). The versatility of force sensing in assembly is indicated by the frequency with which a task is described as being doable "with my eyes closed."

Implementing Force Control. The manipulator used in the experiments described in this paper is a Unimation Puma™ 560 machine with an unmodified controller operating under the VAL-II system.[1] The apparatus is shown in Figure 1. Use of the commercial controller should minimize problems of transferring our results into industrial application, although the ideas presented here are directly applicable to other robot controllers and control systems. A six-axis force-torque sensor is mounted at the end of the Puma and provides the force and moment inputs that are used to determine arm motions. Our implementation provides a force-control loop *outside* of the VAL-II system's position-control loop and is essentially Salisbury's Cartesian stiffness control.[2]

We are not controlling motor torques; force is a function of manipulator stiffness. The resolution with which we can control force is determined by the resolution of the encoders and the stiffness of the manipulator. This approach allows us to perform movements based on either force or velocity (position) specification, which is sometimes referred to as hybrid force-position control. While the implementation used here with the Unimation controller is adequate to demonstrate concepts based on hybrid control, the speed is too slow for commercial application.

Insertions: A Key Assembly Primitive

Studies of assemblies indicate that more than half of the assembly operations are made from a single direction, and that the most common operations can be best described as a form of insertion of a peg in a hole.[3] With proper fixturing, this fact might allow us to reduce the required number of degrees of freedom to four for many tasks. Examples of operations that fit this mold include installing shafts in bearings, mounting gears and pulleys on shafts, and inserting keys into keyways. Not included in this category are screw operations and push-and-twist operations, although learning about peg-in-the-hole operations may tell us much about these as well.

Current robotic technology depends on absolute position control to carry out tasks, restricting our ability to perform insertions to those cases in which the parts tolerance plus location error (determined by fixturing) is sufficiently small and the parts clearance is sufficiently large that the combination is within the positioning resolution of the manipulator. Since the clearance alone may be smaller than the positioning resolution, the task may be impossible using commercially available technology.

In the remainder of this paper we present a strategy that employs force feedback to perform orientable insertions with precision requirements in excess of those achievable by position control. The analysis of contact forces allows us to show that the strategy will work for a large variety of shapes, and as such, constitutes what amounts to a general insertion "reflex." The frequency of peg-in-the-hole insertions in mechanical assembly makes this an important element of a set of robotic primitives for mechanical assembly.

A Strategy for Force-Directed Insertions

Humans can easily perform insertions of parts with extremely small clearances despite the fact that the visual system and kinesthetic sense are not adequate to perform the task based on positional information. In fact, since the task can generally be performed with one's eyes closed, it seems clear that visual information is used only for gross initial positioning. It seems intuitive, then, that the human insertion strategy is based on the reaction forces created by the interference of the involved parts. The strategy we develop in this paper is based on this observation.

A second basic idea underlying this strategy is the restriction of the degrees of freedom of the part being inserted (called a "peg" from here on). By constraining the degrees of freedom, we simplify the problem of interpreting any measured forces, and limit the number of positions that the peg can reach, simplifying analysis. Unlike the strategy of Caine,[4] which inserts rectangular parts by a sequential reduction of degrees of freedom, our strategy creates artificial constraints on degrees of freedom to allow only a single degree of freedom at all times (except for a brief time in the initial phase, which has two degrees of freedom).

The insertion strategy is divided into three phases: initial orientation, final orientation, and insertion. A key concept in this strategy is the notion of a *target point*. The manipulator will apply forces to the peg through a *point of support*, the point about which forces and moments are measured (not the point at which the manipulator is grasping the peg), on the base of the peg and directed toward the target point, a selected point on the peg's edge around the peg base.

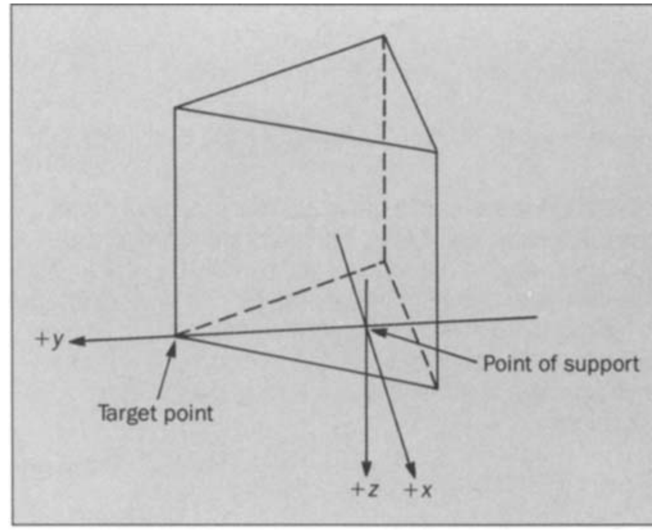

Figure 2. Relation of axes to target and support points.

The selection of the target point is one of the main issues in this paper. Since the hole has a corresponding shape, we can refer to the target point on the hole as well, which is the point on the hole's upper edge that corresponds to the target point on the peg. The target point on the peg is denoted P_t. The point of support is generally near the center of the peg and is denoted P_c. For convenience we adopt a right-handed orthonormal coordinate system with its origin at the point of support P_c. The y axis is defined as the line containing P_c and P_t, the target point, with the positive direction from P_c to P_t. The x axis is in the plane of the base of the peg, and the positive direction of the z axis is in the direction of movement for insertion. This coordinate system is attached to the peg, and moves with the peg. (See Figure 2.)

Initial orientation begins with the peg aligned with the hole in terms of the z axes and rotational orientation about the z axis to the extent that the information is available. Since we are developing this technique with factory environments in mind, this a priori information should be fairly good. Simple fixturing or a crude vision system should allow us to align the z-axis position within 0.25 inches and rotational angular displacement within a degree or two. While these accuracies are sufficient for the approach we present here, they are not good enough to perform the insertion based on positional information alone.

Next the peg is tilted about the x axis so that the target point is the lowest point. The peg is then moved so

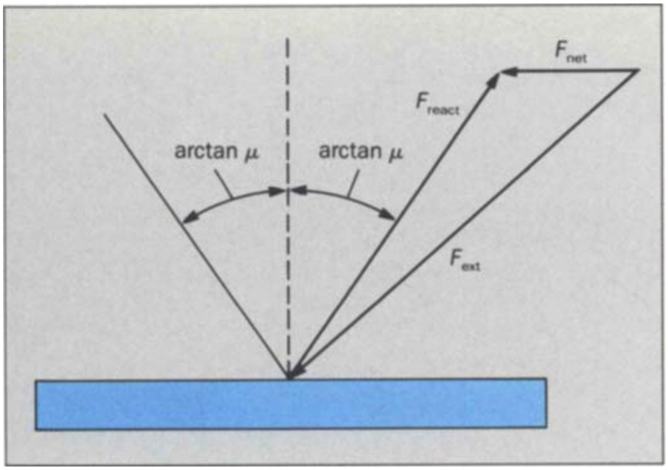

Figure 3. Friction cone.

that the target point is as close as possible to the target point of the hole while ensuring that the peg target point is inside the hole (below the plane of the top of the block containing the hole). A force is applied at the point of support in the direction of the target point. The peg will then move until it is in contact with the edges of the hole. We continue to apply the force, allowing the peg to move compliantly about its z axis and along its x axis. Because of the geometry of the situation, the control algorithm in combination with the physics of the forces will assure that motions are in a direction that reduces position error. Initial orientation ends when a stable equilibrium is achieved.

During final orientation, the peg is uprighted to the a priori estimate of alignment between the z axes of the peg and the hole. During this motion, the peg is allowed to move compliantly about its z axis and along its x axis. Constant force is maintained in the y direction (through the target point) and the z direction (into the hole).

The final stage, insertion, is a translation in the z direction, lowering the peg into the hole. During this motion, the peg is allowed to comply along and about its x and y axes. (This strategy is detailed in Reference 5.)

Analysis of Contacts

Since the strategy we described in the previous section relies on controlling the manipulator on the basis of forces and moments measured by a force-sensing wrist, it is important to understand the forces and reactions involved in inserting a peg. In the orientation stages we can analyze the contacts as a two-dimensional problem since the manipulator constrains the movement to at most two degrees of freedom. We will analyze the one- and two-point cases in this section. In later sections we will apply these results to classes of peg shapes and draw some general conclusions. (Since these classes will require the pegs to be convex, contacts are limited to one or two points.)

In the following discussion we will make frequent reference to the *friction cone.* If we assume that the contact is governed by dry (Coulomb) friction, then for stationary contact, the maximum value for the tangential component of the reaction force is μ, the coefficient of static friction, times the normal force. Thus, if we have a force F applied at angle θ from the normal, the tangential component of the reaction force is min ($F \sin \theta$, $\mu F \cos \theta$). No slipping occurs if $\mu \cos \theta \geqslant \sin \theta$. Thus, the maximum angle from the surface normal at which a force can be applied and be countered by the reaction force is arctan μ.

The cone centered upon the surface normal with half-angle arctan μ is referred to as the *friction cone.* Figure 3 illustrates these concepts. If the force F is applied from outside the friction cone, then there is a net shear force as shown in the figure, resulting in an acceleration. For our experiments we used aluminum parts with a glass-beaded matte finish. For these conditions the coefficient of friction is approximately 0.25 to 0.35. The half angle of the friction cone is approximately 14 to 18°.

Analysis of One-Point Contact. The one-point contact case is shown in Figure 4. We assume that the manipulator is holding the peg with the point of support at P_c and applying a force F_0 toward the target point P_t. If the direction of the applied force lies outside the friction cone at the point of contact, then there must be an acceleration due to the net shear force and moment about P_c as shown in the figure. This holds even if the target point is also the contact point. The force applied at point P_c will cause the point of contact to slip, since the shear force exceeds the friction. In addition, as discussed above, the manipulator is operating under a force program that attempts to maintain a zero z moment and zero x force. The result of this control algorithm is to rotate the peg in the direction of the net moment and move away from the surface contact. The part of the program that attempts to maintain a given y force F_0 will tend to maintain the contact. Thus, depending on the damping in the control system, the contact will either move or contact will be broken. In a later section we will discuss what requirements we must place on the shape of the peg and its position and orientation to ensure that the result of this movement is an improved position and orientation that lies on a path to the desired end state.

Analysis of Two-Point Contact. Figure 5 shows the forces involved in two-point contact. Again we assume that the peg is supported at point P_c, with force F_0 applied toward P_t. If the peg is in static equilibrium, then there are reaction forces F_1 and F_2 such that

$$F_1 \sin \theta_1 - F_2 \sin \theta_2 = 0 \quad (1)$$

$$F_1 \cos \theta_1 + F_2 \cos \theta_2 = F_0 \quad (2)$$

$$l_1 F_1 \sin \theta_1 - l_2 F_2 \sin \theta_2 = 0 \quad (3)$$

If μ is the coefficient of friction at the points of contact, then we must also have

$$\left| \beta_i - \frac{\pi}{2} \right| \leq \arctan \mu \qquad i = 1,2 \quad (4)$$

From eqs. (1) to (3), we see that $l_1 = l_2$. From Figure 5 we see that these equations only have a solution if $\theta_1 + \theta_2 < \pi$. We can interpret this solution graphically as requiring that: F_1 and F_2 must both lie below the line connecting contact points 1 and 2; F_1 and F_2 must both lie within the friction cones at the point of contact; and the extension of vectors F_1 and F_2 must intersect the extension of F_0 at the same point (and therefore Figure 5 is not in static equilibrium).

If these conditions are not met, then there will be a net moment about P_c. The direction of this net moment is determined by where the friction cones at the two contact points intersect the extension of F_0. The net moment will be in the clockwise direction (as drawn) if the friction cone of the contact point on the left intersects the extension of F_0 farther from P_c than does the friction cone of the contact point on the right, where farther is measured in a positive sense for the direction of F_0. As discussed for the one-point contact, the resultant moment will cause the

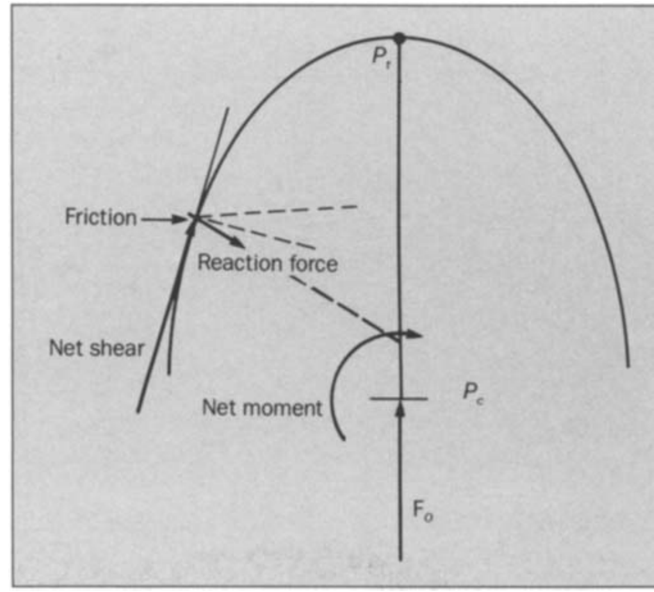

Figure 4. One-point contact.

control algorithm to rotate the peg in a direction to zero out the moment. This will cause the peg to slip at its contact points, reducing the angular error and creating new contact points. In later sections we will discuss how these new contact points are found and under what conditions this procedure leads to an improvement in orientation and position.

Initial Orientation of Convex Polygonal Parts

A large proportion of insertions appear to involve convex parts. This is not surprising if one observes that inserted parts are often shafts that are free to turn, and hence are almost always round. Nonround insertions that are free to turn must have clearances large enough to allow the point farthest from the pivot to move freely. We might think of nonround insertions that do not move as a sort of packing problem. The packing problem may introduce some options about the order of insertion that may affect the task difficulty. In this section we deal with parts that are polygonal in shape and therefore are not free to rotate in their constraining space. The requirement for convexity is not unduly restrictive at this stage since we deal only with local portions of the lower edge of the peg. At the end of this section we discuss the implications of manipulator error for this strategy.

Contact Forces on Polygonal Pegs. In a previous section we analyzed the forces involved in one- and two-point

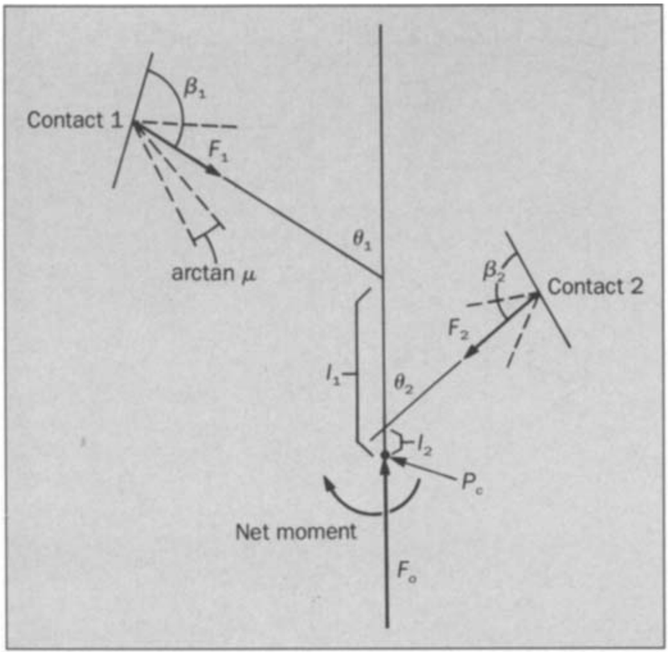

Figure 5. Two-point contact.

contact. If we assume that the target point of a force-directed insertion is one of the vertices of the polygon, then it is clear that in one-point contact between the target point and the hole, we have stability if and only if the direction of the applied force is within the friction cone of the contact. Once outside the friction cone, the peg will rotate and slide along the edge of the hole. The angle of the peg to the hole determines whether the sliding and rotation are an improvement in position. The peg will continue to slide and rotate in one-point contact until it either reaches an orientation at which the angle of contact is within the friction cone (which includes the desired final state), or two-point contact is made, with one of the contact points being the target point.

The analysis of two-point contact showed that static equilibrium exists only if the friction cones of the two contacts contain a point on the line of applied force in their intersection. Since one of the contact points lies on the line of applied force, stability requires that the reaction force from the other contact point be directed through it. While this configuration would satisfy the equation for no net moment, it would not satisfy the condition for no net force, and therefore there can be no static equilibrium.

Thus, in two-point contact the peg will rotate and slide along the edge of the hole in two-point contact until it either achieves one-point contact, the peg becomes geometrically constrained, or the peg achieves its desired orientation (which involves continuous contact but stops because the force is applied within the friction cone of the contact). If the servomechanism controlling the insertion is at least critically damped, then once two-point contact is achieved, it is maintained. Therefore, we can eliminate the case of returning to one-point contact.

Application to Convex Polygonal Pegs. In the previous section we showed how the contact forces cause the control strategy to move the peg. We now examine how we select the target point and determine bounds on the initial position and orientation of the peg relative to the hole. Our goal is to ensure that static equilibrium (the stopping criterion for initial orientation) is reached only when the peg is in the correct orientation. Note that when we constrain the peg to move only in two-point contact (which is the case once we have achieved two-point contact in a system that is at least critically damped), the motion of the peg has only one degree of freedom.

We can plot the possible path of the point of support of the peg as a curve in the plane. Figure 6 shows the path of the point of support for a square peg with target point at one vertex and three cases for the location of the point of support: at the corner opposite the target point, three-quarters of the distance from the target point to the opposite corner, and one-quarter of the distance. If we draw a vector field along the allowable path, the peg will move (if not bound by friction) in the direction along the path that is less than 90° from the direction of the applied force, since the applied force can only perform work in a positive sense. Figure 6 shows the vector field for the square peg superimposed over the path for the point of support at the corner of the peg opposite to the target point. The shape of the path and the nature of the vector field vary with the target point angle for a simple insertion in which only two sides interact (and we may therefore treat the peg as a triangle).

The variations are more complex for an insertion in which the contact involves several of the vertices, and may have several discontinuities in the path derivative. In all cases, however, there is a critical point along the path at which the vector field points in the direction of desired motion and continues to point in that direction up to the point at which the peg is correctly oriented.

Thus, in a practical sense, if we use the manipulator to position the angular orientation of the peg as close as possible to the desired position, and if the contact places the support point within the part of the path between this critical point and the goal, then the insertion will succeed. If the support point is outside that region, the insertion will fail, and would fail for any other point outside that region. Therefore, we need not hunt for another region, possibly farther away (larger initial positioning error), that would allow us a new starting position that would always succeed.

For asymmetric parts, this implies that it may be preferable in some cases to make the initial position goal some other position than best attempt at alignment. This is true if the path on one side of the goal state has a larger area guaranteeing success than the path on the other side. By selecting an initial goal other than correct orientation we may be able to ensure that the contact falls on that guaranteed portion of the path, even if farther from the goal.

Only the portion of the peg tilted below the plane of the top of the hole can interact with the hole. However, the entire top edge of the hole may interact with the peg, so a more precise model requires that the peg be represented at this stage as a portion of the curve representing the hole. Since the hole is represented by a closed curve, for the convex polygonal case the interaction cannot involve an endpoint of the hole curve, as there is no endpoint. Contacts will always be between an edge of the hole polygon and either a vertex or endpoint of the peg polygon.

Once the orientation error in two-point contact has been reduced past a certain limit (which is a function of shape), the contact always involves the endpoint of the peg curve moving along the edge adjacent to the corresponding point on the hole. If this point is a vertex, the corresponding edge is the first edge on the hole polygon not contained in the peg curve; otherwise it is the edge containing that point. This permits us to determine the angle of the reaction force at this final stage. We can select the support point to lie between the target point and the intersection of the contact friction cone with the line of

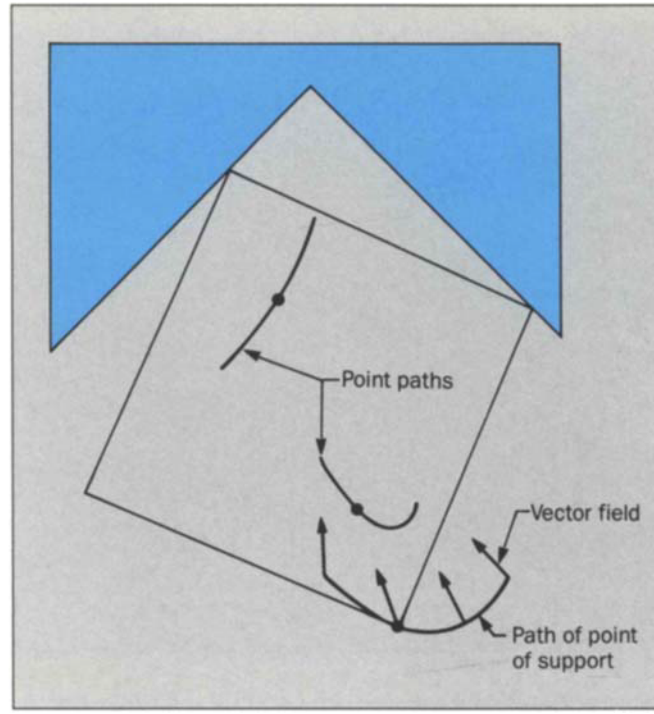

Figure 6. Allowable paths for a square peg.

action of the applied force in order to maximize the reaction moment that drives the control mechanism.

In actual manipulation, it is not possible to ensure that the force can be applied with a line of action through the target point or necessarily along a line bisecting the target point vertex. The direction along the angle bisector has only been used for convenience and has not been critical to the analysis. In fact, by selecting a support point off the bisector such that the line of action of the applied force falls on the side of the angle bisector opposite the second point of contact, we may reduce the possible cases of sticking or jamming. Errors in directing the force through the target point can be similarly compensated for; the line of action should intersect the peg on the side opposite the side connecting the contact points.

Initial Orientation of Smooth Convex Parts

The insertion strategy described above was based on informal observation of human strategies for force-directed insertion. It appears that human strategy selects the point of maximum curvature as the target point. In this section we will show how certain properties of this point make it particularly suitable for the strategy we have developed.

Contacts on Convex Curves. We select as the target point t the point of maximum curvature and restrict f and g, the curves representing the interacting portions of the peg and the hole, to be monotonically decreasing in curvature in both directions from t. We can then show that two-point contacts require that at least one of the contact points be an endpoint of one of the curves.

Let the peg and the hole be represented by two planar convex curves parameterized by arc length, $f(s)$ and $g(s)$. Let $k_f(s)$, $k_g(s)$ be the curvature at some point s for f and g respectively; k_f and k_g are piecewise continuous. Likewise let $\theta_f(s)$, $\theta_g(s)$ be the direction of the tangent at s; θ_f and θ_g are continuous. We assume that f and g are identical up to a rigid transformation of the plane.

Theorem 1. Given $f(s)$, $g(s)$, $0 \leq s \leq 1$ are identical planar convex curves (up to a rigid transformation of the plane), with continuous tangents and piecewise continuous curvatures. Assume that there is a point t such that $k_f(s)$ is monotonically increasing on the interval $0 \leq s \leq t$ and monotonically decreasing on the interval $t \leq s \leq 1$. Given $|\theta_f(1) - \theta_f(0)| \leq \pi$, then f, g cannot be tangent at two distinct points.

Application to Convex Pegs. From the theorem, we know that, for smooth curved convex objects, if we select a section with monotonically increasing curvature up to the target point in both directions, then the contact with the hole will involve the endpoint of either the hole or the peg. As for the polygonal case, the hole cannot have an endpoint; it is a closed curve. Therefore a two-point contact must include the endpoint of the peg. It appears (although we have not yet proven) that the nonendpoint contact moves toward the noncontact endpoint as the angular orientation error is reduced and two-point contact is maintained. Jamming occurs when the friction cones of the two-contact points contain an intersection along the line of action of the applied force. Jamming can be avoided if the interacting portions of the peg and the hole can be constrained to prevent this overlap. For a symmetric peg this may occur when the length of the arc on one side of the target point is larger than on the other.

Different approaches may satisfy this criterion for other shapes. By placing the support point between the intersections of the contact-point friction cones with the line of action of the applied force, we may maximize the reaction moment and reduce problems relating to signal-to-noise ratios in the sensor. In practice, the peg and the hole have a small clearance and are not truly identical curves. In this case the two-point contact will not be continuous. As the contact moves away from the target point, it will eventually reach a point at which any further rotation will break contact. When the contact is broken, the system moves to one-point contact until a second contact is made. This second contact will be near the target point and on the same side of the target point as the contact endpoint. This will result in a net moment and will cause rotation until one-point contact is achieved.

Because the strategy described in "A Strategy for Force-Directed Insertions" is an active one, there is only one stable configuration (other than jamming, which we hope to avoid for most cases). Stability occurs when the line of force intersects the peg boundary at the point of tangency to the hole and is directed within the friction cone at the point of contact on the hole. Any other configuration, even with the force within the friction cone at the contact, is not stable since there is an unresolved moment that will cause the manipulator to rotate. Since the target point is the point of maximum curvature, it can be tangent at any point, not only at the corresponding target point on the hole.

When two-point contact is broken, the state of contact after the discontinuity is a function of the shape of the hole, the clearance with the peg, and the control mechanism. The peg is highly constrained at this point as to possible contacts. However, there is a region about the target point on the hole in which the target point of the peg may make tangent contact and satisfy our stopping criterion. For practical purposes this error is generally of small concern, as the final orientation will be capable of proceeding.

There are positions other than those near the hole target point for which the peg target point is tangent to the hole. However, these generally involve rather gross error. For the small-batch manufacturing environment motivating our research, the assumption on satisfying initial conditions that rules out this case seems reasonable. This is because we assume a moderately structured environment supplemented with sensors in the control system and a manipulator capable of being positioned within a few hundredths of an inch and within 1° angular position. Special situations, such as the assembly of very small components, may require proportionately greater accuracy.

We stated a conjecture that the point of tangent contact moves away from the target point as the orientation error was reduced. We can show that as a point moves farther from the target point along a curve of monotonically decreasing curvature, the intersection of its normal with the normal through the target point moves farther from the target point as well (Figure 7). When the curve is symmetric about the point of maximum curvature, then the only two-point contact on a frictionless surface with no residual moment must involve the corresponding points on both sides of the target point.

For a surface with friction, the distance from the corresponding symmetric point is a function of the coefficient of friction, and therefore goes to zero with μ. If our conjecture is true regarding movement of contact points, then for symmetric curves with no friction, we can assure that jamming does not occur until the peg is properly aligned or close to alignment for cases with low friction. For a given value of μ there must be a clearance at which we can assure no jamming, since the point at which two-point contact is broken occurs before the required point for jamming is reached.

The analysis has assumed that the force was directed at the target point *t* along the normal to the curve at the target point. If our conjecture about the motion of the contact point is true, small errors (relative to the scale of the part) will have little effect on the accuracy of the orientation except at the last stage. The error at the termination (regardless of error in force direction) will be dominated by the clearance between the peg and the hole. Once the peg has broken two-point contact, termination occurs when either the point on the curve through which the force is directed (true target point) is tangent to the hole or there is a two-point contact with the true target point between the contacts. The range of orientations at

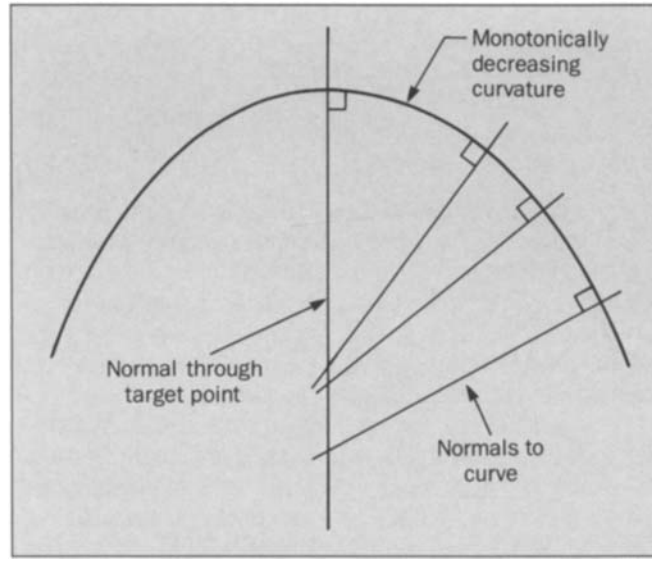

Figure 7. Intersection of normals with normal through target point.

which either of these two conditions may occur is determined by the clearance, not the error in pointing the force.

Final Orientation

Because the target point of the peg has been tilted into the hole, the orientation at the termination of the initial phase may still contain small errors. Figure 8 shows the projection of the base of the peg onto the plane of the top of the hole, as well as the slice through the plane of the top of the hole. Although the orientation through the slice is correct, tilting the peg back to vertical would introduce interference between the peg and the hole.

The motion program for final orientation is to rotate the peg back to vertical while maintaining a constant force in the direction from the support point to the target point ($+ y$), constant force down the peg (roughly, into the hole, $+ z$) and moving compliantly in x translation and z rotation. The analysis of forces and moments in the plane sliced through the peg (perpendicular to its z axis) at the point it contacts the hole is essentially the same as for two-point contact in initial orientation, and therefore there is a net moment about z. (The convexity assures that the contacts are always edge-edge contacts. A nonconvex peg would allow a point-surface contact, which would not produce the required resultant forces and moments and final orientation would fail. In many cases it should be possible to measure moments generated by the surface contact to

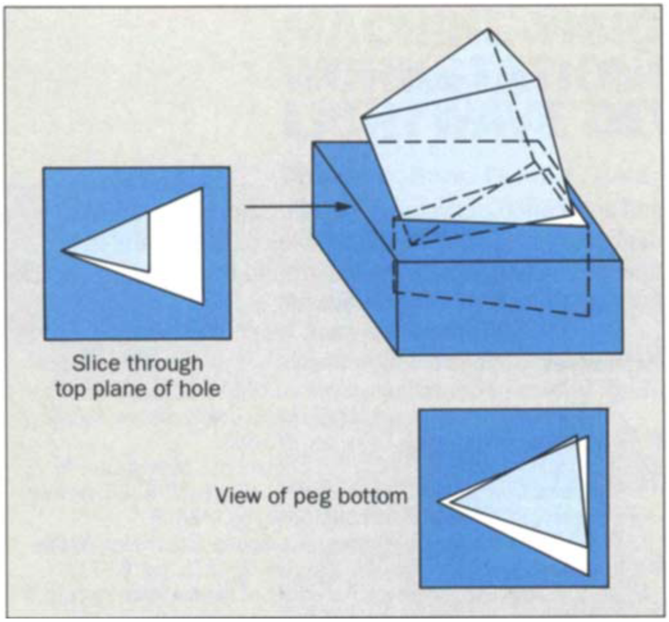

Figure 8. Final orientation beginning position.

determine in which direction to rotate the peg to break the contact.)

If the target point of the peg is inserted too far into a hole of sufficiently tight tolerance, points far from the target point will contact the surface of the hole and prevent continued rotation. The force program of final orientation takes this into account, as the part of the program to maintain constant z force will cause the tip of the peg to lift as the force is applied through the points in contact with the surface of the hole. The force program to maintain constant force in the y direction causes the peg to maintain contact at the target point of the hole. The depth of the target point in the hole is a function of the ratio of the hole size to the peg size. [Depth of the target point is equal to $l_p * \sin(\arccos(l_p/l_h))$, where l_h is the length of the hole and l_p is the length of the peg. Length is the distance in the y direction from the target point to the most distant point on the boundary.] If the hole is oversize by 1 percent, then the distance from the target point to the top of the peg is 14 percent of the length. For a 0.01 percent oversize hole, the target point depression is 1 percent.

Insertion

The forces and moments encountered in insertion have been analyzed in considerable detail for round pegs and holes by Whitney[6] and for rectangular pegs and holes by Caine.[4] These two authors provided a quantitative analysis for special cases. The general nature of peg shape that we have adopted prohibits this sort of quantitative analysis. Instead we will draw upon their results to indicate why our approach works.

Successful completion of final orientation leaves the base of the peg inside the hole. The z axis of the peg is aligned with the z axis of the hole to the extent of the a priori knowledge of the hole's orientation and the positioning accuracy of the manipulator. The control program for the insertion seeks to maintain zero forces and moments along and about the x and y axes (where the origin of the coordinate system is on the base of the peg) at the projection of the support point along the z axis. The possible contacts of the peg with the hole involve the base of the peg in either one-point or edge contact with the side of the hole and the top plane of the hole in edge or point contact with the side of the peg. There is also the possibility of face-face or line contact parallel to the z axis between the peg and the hole. However, this would imply no orientation error in the axis perpendicular to that face, and hence no force on the face, since force is being applied only in the z direction. Jamming can only occur if the reaction forces of the two contacts have friction cones intersecting each other along the z axis. As long as the peg is not jammed, the reaction force from the top edge of the hole intersects the z axis farther from the base of the peg than does the reaction force from the contact at the base of the peg. Therefore the net moment is in a direction that corrects that orientation of the peg. As noted in Reference 6, as the peg progresses farther into the hole, the possibility of jamming is reduced. There is a depth that is a function of the ratio of the peg "diameter" to the hole clearance and the coefficient of friction such that, past that depth, the peg cannot jam, unless external moments are applied.

Summary

The analyses demonstrate the robustness of the proposed insertion strategy and explain the success of the experiments reported in Reference 5. In addition to the rectangular, triangular, and elliptical peg insertions reported there, further tests have demonstrated the ability of this approach to insert a 2-inch-square peg into a

2.0015-inch hole. The insertion approach provides a generic capability for performing automatic insertions, a key primitive in the development of automated assembly systems. This approach has similarities to the techniques of Inoue[7] and Simunovic.[8] Both of these, however, are restricted to round pegs. In addition, the strategy of Simunovic relies upon accurate resolution of moments caused by a tilt when the peg is brought into contact with the hole to determine the direction of motion to align the peg and hole centers. This is still beyond the resolution and signal/noise capabilities of commercial devices. A thorough analysis of the forces involved in the insertion of rectangular solids is used in Reference 4 to develop a strategy in which the degrees of freedom are continuously reduced in a manner that increasingly restricts the possible positions of the peg. In contrast, our technique introduces artificial constraints at each stage to limit motion to a single degree of freedom (although early in initial orientation the peg has two degrees of freedom until two-point contact is achieved). This considerably reduces the difficulty of analysis of forces and produces a robust strategy.

Robots and their controllers will need new designs and capabilities to overcome the shortcomings that have prevented them from making a significant penetration into the manufacturing environment. In this paper we have demonstrated how force-controlled manipulation can simplify the programming of insertion operations. In addition, the technique permits insertion of parts with clearances smaller than the precision of the manipulator or associated fixturing, permitting reductions in system cost and increasing system flexibility. Many other tasks will be simplified by integrating sensing into the manipulator control. This is only one of many changes that will be required to bring robots to batch manufacturing. Task-oriented languages and interfaces to design systems are necessary to reduce programming effort. Education of the designers themselves will be essential to ensuring designs that can be economically produced in robotic factories. Mechanical design of the manipulators will need to change as well. Existing robots generally have a weight-to-payload ratio of up to 100 to 1. New control strategies that employ sensing and intelligence can reduce this ratio, and, in the long run, reduce costs as well. More understanding of the control and application of systems with flexible links, redundant degrees of freedom, and adaptable configurations is needed. Research laboratories around the world are attacking many, if not all, of these issues.

References

1. G. P. Starr, "Edge-Following with a PUMA 560 Manipulator using VAL-II," *Proceedings, 1986 IEEE Conference on Robotics and Automation*, March 1986, pp. 379-393.
2. K. J. Salisbury, "Active Stiffness Control of a Manipulator in Cartesian Coordinates," *Proceedings, 19th IEEE Conference on Decision and Control*, December 1980, pp. 95-100.
3. J. L. Nevins and D. E. Whitney, "Computer Controlled Assembly," *Scientific American*, Vol. 238, No. 2, 1978, pp. 62-74.
4. M. E. Caine, "Chamferless Assembly of Rectangular Parts in Two and Three Dimensions," M.S. thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, June 1985.
5. D. R. Strip, "Insertions Using Geometric Analysis and Hybrid Force-Position Control on a PUMA 560 with VAL-II," *Proceedings, AAAI-87*, American Association for Artificial Intelligence, July 1987, pp. 695-698.
6. D. E. Whitney, "Quasi-Static Assembly of Compliantly Supported Rigid Parts," *Journal of Dynamic Systems, Measurement, and Control*, Vol. 104, 1982, pp. 62-77.
7. H. Inoue, "Force Feedback in Precise Assembly Tasks," *Artificial Intelligence: An MIT Perspective*, Vol. 2, P. H. Winston and R. H. Brown, eds., MIT Press, Cambridge, Massachusetts, 1979, pp. 219-241.
8. S. Simunovic, "An Information Approach to Parts Mating," Sc.D. thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusets, 1979. (Also available as CSDL Report No. T-690.)

(Manuscript received October 27, 1987)