

Mechanics, Planning, and Control for Tapping

Wesley H. Huang
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180

Matthew T. Mason
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

May 25, 2000

Abstract

We present analysis and experimental demonstration of manipulation by tapping. Our problem domain is positioning planar parts on a support surface by a sequence of taps; each tap imparts some initial velocities to the object which then slides until it comes to rest due to friction.

We formulate the mechanics of tapping for circular axisymmetric objects, show how to plan a single tap and a sequence of taps to reach a goal configuration, and present feedback control methods to robustly accomplish positioning tasks. With these methods, we have experimentally demonstrated positioning tasks using tapping, including high-precision positioning tasks in which the object is positioned more precisely than the tapping actuator. We show stability and sensitivity analysis in support of these results.

1 Introduction

We consider *impulsive manipulation* to be any form of manipulation consisting of two phases: a strike to an object, imparting some initial velocities; and free motion of that object, subject to forces and constraints in the environment. Impulsive manipulation has a number of unique characteristics that distinguish it from other modes of manipulation: it produces fast and dynamic manipulation; it can be used to manipulate a wide variety of objects because of its nonprehensile (nongrasping) nature; and it can produce large interaction forces. Impulsive manipulation is generally well suited for micropositioning because impact can both break static friction and deliver a small amount of impulse.

Our present work focuses on one particular form of impulsive manipulation: tapping planar objects which then slide on a support surface and come to rest due to frictional forces. The main problem this work addresses is how to position objects in the plane with a sequence of taps.

This problem arises in a number of situations. For example, in manufacturing, fast but coarse positioning could be used in parts transfer between workcells, and precise positioning could be used for alignment, machining, or assembly operations. Analysis of positioning via tapping could be applied to the design and configuration of certain types of parts feeding systems.

Our focus has been on precise positioning of objects in the plane. Our experimental setup uses a general purpose robotic arm to position a tapping device with respect to the object; an overhead camera measures the position and orientation of the object.

We begin by analyzing the mechanics of tapping and by formulating solutions to the motion planning problem for circular axisymmetric objects; we use this class of objects to make our analysis tractable. Although in practice there are errors due to differences from our model assumptions, we rely upon feedback control to robustly accomplish positioning tasks. Our feedback control methods are built upon the motion planning solutions, and we show conditions for the stability of these control methods.

We report the results of a number of experiments to test the validity of our models and to demonstrate our positioning methods. We also show, both experimentally and analytically, how it is possible to position objects more precisely than the tapping device can be positioned.

A detailed overview of the main problems addressed in this paper appears in Section 2.

1.1 Related work

The first instance of impulsive manipulation in the robotics literature was due to Higuchi (1985) who used an electromagnetic coil to produce an impulsive force for linear micropositioning. Yamagata and Higuchi have since developed related devices for several different micropositioning applications, including precision alignment of optical sensors (Yamagata and Higuchi 1995) and positioning platforms for scanning tunneling microscopes (Yamagata *et al.* 1990).

Whereas Yamagata and Higuchi concentrated on mechanisms and applications, our early work in this area (Huang, Krotkov, and Mason 1995) first showed formal results in the mechanics and planning for tapping. Our analysis of the mechanics of a sliding object is based upon the work of Voyerli and Eriksen (1985) who studied the motion of sliding disks and rings. Goyal, Ruina, and Papadapolous (1991a, 1991b) formulated a limit surface representation of the relationship between net frictional force and torque and the the velocities of a planar slider; this work has provided a number of insights in our analysis.

We use a graphical method for analyzing two dimensional impact with friction due to Routh (1905), more recently presented in (Wang and Mason 1992). Although there has been much recent work on impact (e.g. Bowden and Tabor 1964; Brach 1991; Stronge 1990; Chatterjee and Ruina 1998), the simplicity of Poisson's hypothesis in conjunction with Routh's method makes our analysis simpler and has worked reasonably well in practice.

Our more recent work on tapping (Huang 1997) treats some of the problems presented here more generally, in particular impact analysis and motion planning for "nearly axisymmetric objects" (i.e. objects with a noncircular boundary but an axisymmetric support distribution). We have also studied the case of tapping when the object is not permitted to come to rest in between impacts.

This work has been inspired by the ideas behind minimalist robotics, articulated by Canny and Goldberg (1994a, 1994b) and by work in many areas of nonprehensile manipulation, such as pushing (Mason 1985; Peshkin and Sanderson 1988; Alexander and Maddocks 1993; Balorda 1993; Lynch 1996), palmar manipulation (Trinkle *et al.* 1993; Trinkle, Farhat, and Stiller 1995; Erdmann 1995; Zumel 1997), and paddle juggling (Bühler and Koditschek 1990; Rizzi and Koditschek 1992; Schaal and Atkeson 1993).

Experiments with planar sliders have been performed by Zhu *et al.* (1996) in conjunction with what they call "releasing manipulation." In their experiments, an object was accelerated using a manipulator (instead of striking it) and then released, letting it slide on a support surface until it came to rest due to friction. Their results, like our single-tap trials, show considerable variation, presumably due to variations in initial velocities, friction, and pressure distribution.

Another form of impulsive manipulation has been studied by (Partridge and Spong 1999) and (Bishop and Spong 1999) who have built an air hockey playing robot. They address the problem of using impact to control trajectories of a puck on an air table. In this work, both the control over the puck (with a round striker) and the interaction of the puck with its environment (rebounds off the sides of the table) take the form of impacts.

1.2 Assumptions

This paper will address manipulation of circular axisymmetric objects, i.e. circular objects which have a pressure distribution that is a function of radius only. Examples of this class of objects include disks and annuli. We assume the object is a planar lamina of mass M and moment of inertia I , and the support distribution between the object and the support surface is known. For

this paper, we use a uniform support distribution, though our results hold for any axisymmetric support distribution.

We assume that the support surface is homogeneous and that Coulomb friction acts between the object and the support surface with a coefficient of friction μ .

We use Poisson's hypothesis in conjunction with the classical model of two dimensional impact with friction. This implies that the impact is instantaneous and is characterized by: the geometry and relative velocities of the object and the striker, a coefficient of restitution, and a coefficient of friction.

2 Tapping as a discrete dynamical system

Let the position and orientation of a planar part be described by $\vec{x} = [x \ y \ \theta]^T$. Denote the initial configuration of the object by \vec{x}_0 and the goal configuration by \vec{x}_g . The action of a single tap can be described by the equation:

$$\vec{x}_i = f(\vec{x}_{i-1}, \vec{a}_i) \quad (1)$$

where the vector \vec{a}_i represents the parameters of the i^{th} tap and the function f represents the dynamics of impact and sliding. The particular components of \vec{a}_i will depend upon the type of tapping actuator(s) used. For linear tapping actuators, \vec{a}_i would represent the line of action of the device and the striker velocity.

In the following sections, this paper addresses several problems regarding this system:

- *Motion planning*: For a given \vec{x}_0 and \vec{x}_g , compute a sequence of taps $A = \{\vec{a}_1 \dots \vec{a}_k\}$, such that $\vec{x}_k = \vec{x}_g$.

We describe an algorithm to generate such a plan and show that two taps ($k = 2$) suffice to transfer a part from any configuration to any other configuration in the absence of errors.

- *Feedback control*: A feedback control law specifies how to choose an action at each time step. Suppose we think of the motion planner as a function that maps a configuration \vec{x} and the goal configuration \vec{x}_g to a list of actions $A = \{\vec{a}_1 \dots \vec{a}_k\}$. Denote this function as $A = MP(\vec{x}, \vec{x}_g)$.

Open loop control would create a plan $A = MP(\vec{x}_0, \vec{x}_g)$ and execute each tap a_i of this plan in sequence.

There are many different ways feedback control can be implemented. The basic method that we use is to do motion planning at each step and execute the first tap of that plan. Let $MP_1(\vec{x}, \vec{x}_g)$ denote the function that returns the first step of the plan $A = MP(\vec{x}, \vec{x}_g)$. This type of control policy can then be represented by the equation:

$$\vec{a}_i = MP_1(\vec{x}_{i-1}, \vec{x}_g) \quad (2)$$

We describe several different approaches to this type of control law and describe results of our experiments.

- *Stability and sensitivity*: A particular feedback law determines the behavior of the system in response to some initial configuration. For the type of control policy described above, the closed loop system is governed by the equation:

$$\vec{x}_i = f(\vec{x}_{i-1}, MP_1(\vec{x}_{i-1}, \vec{x}_g)) \quad (3)$$

We establish conditions on the error of a tap in order to guarantee that the system will actually reach the goal \vec{x}_g with any desired precision. In addition, our sensitivity analysis shows that errors in tapping device configuration map to smaller errors in final object configuration, providing analytic support for the results of our high-precision positioning experiments.




Figure 1: Notation for the axisymmetric inverse sliding problem.

3 Mechanics & motion planning

In this section, we describe the mechanics of a sliding planar object and how to plan a sequence of taps to transfer a part from a start to a goal configuration.

After introducing the mechanics, we show how to plan a single tap that will transfer the part from \vec{x}_{i-1} to \vec{x}_i :

$$\vec{a}_i = g(\vec{x}_{i-1}, \vec{x}_i) \quad (4)$$

This problem can be decomposed into two subproblems which we refer to as the *inverse sliding problem* and the *impact problem*. The inverse sliding problem is to determine the initial velocities required for the object to slide the desired displacement (translation and rotation):

$$\vec{v}_i = g_1(\vec{x}_{i-1}, \vec{x}_i) \quad (5)$$

The impact problem is to determine how (if possible) to generate those velocities by striking the object:

$$\vec{a}_i = g_2(\vec{x}_{i-1}, \vec{v}_i) \quad (6)$$

After formulating solutions to these problems, we describe single-tap experiments to validate the mechanics. The results of these experiments led us to slightly modify the models of the mechanics.

Since a single tap cannot transfer a part between two arbitrary configurations, we must do motion planning. Our approach is to form a forward projection from the starting configuration:

$$FP(\vec{x}_0) = \{\vec{x} \mid \exists \vec{a} \vec{x} = f(\vec{x}_0, \vec{a})\} \quad (7)$$

and a backprojection from the goal configuration:

$$BP(\vec{x}_g) = \{\vec{x} \mid \vec{x}_g \in FP(\vec{x})\} \quad (8)$$

These two sets always intersect, so two taps are sufficient to transfer a part from any start to any goal configuration. Once a subgoal has been chosen, the solutions to the inverse sliding problem and the impact problem can be used to determine parameters for the tapping device.

3.1 Notation & basic mechanics

We place the origin of a frame at the object center of mass (COM) in the start configuration such that the x axis passes through the object COM in the goal configuration. Axisymmetric objects have the property that the net force due to friction in the \hat{y} direction is zero. Therefore these objects will slide in a straight line, and we need only consider the translation along the x axis. Let the object's configuration be described by x and θ , its velocities by v and ω . See Figure 1 for notation.

The frictional force load along the \hat{x} axis is:

$$F\left(\frac{\omega}{v}\right) = \mu g \lambda \iint \frac{1 - \frac{\omega}{v} \sin \theta}{\sqrt{1 - 2\frac{\omega}{v} r \sin \theta + \left(\frac{\omega}{v}\right)^2 r^2}} r \, dr \, d\theta \quad (9)$$

where λ is the pressure over the support distribution and the double integral is over this support area. The net torque load due to friction is:

$$T\left(\frac{\omega}{v}\right) = \mu g \lambda \iint \frac{\frac{\omega}{v} r^2 - r \sin \theta}{\sqrt{1 - 2\frac{\omega}{v} r \sin \theta + \left(\frac{\omega}{v}\right)^2 r^2}} r dr d\theta \quad (10)$$

Note that the force and torque are functions of the velocity ratio $\frac{\omega}{v}$, not of the velocities themselves.

The equations of motion of the object are:

$$M\dot{v} = -F\left(\frac{\omega}{v}\right) \quad (11)$$

$$I\dot{\omega} = -T\left(\frac{\omega}{v}\right) \quad (12)$$

which are a pair of coupled first order differential equations. The net translation and rotation of the object are:

$$x_f = \int_0^{t_f} v(t) dt \quad (13)$$

$$\theta_f = \int_0^{t_f} \omega(t) dt \quad (14)$$

where t_f is the time at which the object comes to rest and $v(t)$ and $\omega(t)$ are solutions to the equations of motion, subject to the initial conditions:

$$v(0) = v_0 \quad \omega(0) = \omega_0 \quad (15)$$

We consider x_f and θ_f to be functions over the space of initial conditions v_0 and ω_0 .

3.2 Inverse sliding problem

If the goal configuration is a distance d and an orientation α relative to the start configuration, then the inverse sliding problem is to determine v_0 and ω_0 such that:

$$x_f(v_0, \omega_0) = d \quad \theta_f(v_0, \omega_0) = \alpha \quad (16)$$

Because of the symmetry of the x_f and θ_f functions (for circular axisymmetric objects):

$$x_f(v_0, \omega_0) = x_f(v_0, -\omega_0) \quad (17)$$

$$\theta_f(v_0, \omega_0) = -\theta_f(v_0, -\omega_0) \quad (18)$$

we can consider only $\alpha \geq 0$. The desired solution will lie in the first quadrant of the space of initial conditions.

In general, there are no analytic expressions for the net force and torque load due to friction (Equations 9 and 10), so there will be no analytic form for x_f and θ_f . However, we have been able to prove a number of properties of this system. We state the following without proof; the interested reader can consult (Huang 1997) for details.

Property 1 *The coupling between translation and rotation disappears for $v_0 = 0$ or $\omega_0 = 0$, so we can formulate analytic expressions for x_f and θ_f along these axes.*

Property 2 *The functions $x_f(v_0, \omega_0)$ and $\theta_f(v_0, \omega_0)$ are monotonic in the sense that:*

$$v_{0_1} > v_{0_2} \quad \omega_{0_1} = \omega_{0_2} \quad \implies x_{f_1} > x_{f_2} \quad (19)$$

$$v_{0_1} > v_{0_2} \quad \omega_{0_1} = \omega_{0_2} \neq 0 \quad \implies \theta_{f_1} > \theta_{f_2} \quad (20)$$

$$v_{0_1} = v_{0_2} \neq 0 \quad \omega_{0_1} > \omega_{0_2} \quad \implies x_{f_1} > x_{f_2} \quad (21)$$

$$v_{0_1} = v_{0_2} \quad \omega_{0_1} > \omega_{0_2} \quad \implies \theta_{f_1} > \theta_{f_2} \quad (22)$$

where x_{f_i} and θ_{f_i} is the displacement resulting from the initial velocities v_{0_i} and ω_{0_i} .




Figure 2: Illustration of level curves and the iterative method for finding the coordinates of their intersection point.

Property 3 A level curve of θ_f starts on the ω_0 axis and monotonically decreases, asymptotically approaching the v_0 axis. A level curve of x_f starts on the v_0 axis and asymptotically approaches the ω_0 axis; level curves of x_f are also monotonically decreasing.

Property 4 The value of x_f is strictly monotonic increasing along level curves of θ_f .

Properties 2 and 4 can be shown by reasoning about pairs of object trajectories; Property 3 is implied by Properties 1 and 2.

We can solve the inverse sliding problem by finding the level curves of x_f for a distance d and of θ_f for an angle α . Property 3 implies that these two level curves must intersect and must do so within a rectangle defined by the coordinate axes and the intersection of the level curves with the coordinate axes. Property 4 implies that this intersection point is unique. The coordinates of this intersection point are the desired initial velocities.

There is no analytic form for these level curves, however, so a solution must be found numerically. Reasoning about the monotonicity of the level curves leads to an iterative method for finding the intersection point. At each iteration:

- Subdivide all bounding rectangles.
- For each bounding rectangle, if more than one bounding rectangle remains:
 - Numerically integrate the initial conditions $(\bar{v}_0, \bar{\omega}_0)$ corresponding to corners of the rectangle.
 - If the resulting distance or angle is greater than the desired distance or angle, then any bounding rectangle in the region $v_0 > \bar{v}_0$ and $\omega_0 > \bar{\omega}_0$ can be eliminated from consideration.
 - If the resulting distance or angle is less than the desired distance or angle, then any bounding rectangle in the region $v_0 < \bar{v}_0$ and $\omega_0 < \bar{\omega}_0$ can be eliminated from consideration.

In this manner, we can determine the coordinates of the intersection point to any desired accuracy. After the final iteration, the solution is taken to be the centroid of the remaining bounding rectangles. The algorithm can be terminated when the bounds have reached some desired accuracy.

At each iteration, at least one subrectangle is eliminated. In our experiments, after seven iterations there is most often one but sometimes two subrectangles left.

3.3 Impact problem

Once we know what initial velocities are required, we must determine how, if possible, to generate those velocities via impact. We assume that to generate impact, we will essentially throw a free mass (the striker) at the object. The range of velocities that can be produced are limited only by the friction between the striker and the object.

In general, determining how an object should be tapped in order to generate some desired velocities requires searching the boundary for an appropriate contact point. However, for some simple shapes, including circles, we can write analytic expressions for the impact as a function of the contact point.

Calculation of the striker parameters (velocity and impact point) is relatively straightforward, so we will not go into detail here. We have used Routh’s method of analyzing two dimensional impact with friction as presented in (Wang and Mason 1992).

3.4 Single-tap experiments

We performed experiments consisting of a single tap to the object in order to test how well our analysis could predict object motion despite the strong assumptions in modeling this system. The most significant differences are that objects do not have a uniform pressure distribution, and they are three dimensional, not planar.

We used three different combinations of object and support surface: a plexiglas disk on an aluminum surface, an aluminum square on aluminum, and an aluminum disk on formica. Initial experiments were not successful — the objects did not rotate as much as our model predicted. We then performed more carefully controlled experiments and used a high speed measurement system (an OptoTrak system) to measure the position and orientation of the object at 500 Hz. Data from an example trial appear in Figure 3. From these measurements, we calculated initial velocities and velocity profiles of the object, so we could evaluate the impact model and the sliding model separately.

The impact model worked reasonably well, predicting initial translational velocities to within a few percent and initial rotational velocities to within a few percent up to 33%. There were greater differences between the predictions of the sliding model and the observed velocity profiles. Although the predicted and actual translational velocity profiles matched reasonably well, the rotational velocity profiles did not. Judging by the measured and predicted velocity profiles, there was more frictional torque than predicted.

To make our models more accurately reflect reality, we added an additional parameter to the model that scales the net torque due to friction. For the three sets of experiments we performed, the scale factors were 1.0, 2.17, and 3.64 (determined by fitting the experimental data). With this modified model (which still allows us to use the same method for solving the inverse sliding problem) we were able to attain reasonable matches between predicted and actual velocity profiles: to within 5–10% for translations and to within 10–25% for rotations. This modified model was used for the positioning experiments described in Section 4.2.

3.5 Backprojections and motion planning

In order to do motion planning, we determine forward and backward projections from the start and goal configuration, respectively. The forward mechanics, which map a configuration and a tapping action into a new configuration (the function f in Equation 1), can be broken in to two parts related by the object velocity. The first part, $\vec{v}_i = f_1(\vec{x}_{i-1}, \vec{a}_i)$, determines the initial velocities produced by striking the object; the second part, $\vec{x}_i = f_2(\vec{x}_{i-1}, \vec{v}_i)$ determines the configuration of the object after it has come to rest. We first construct the set of all initial velocities that can be produced by tapping the object:

$$V = \{\vec{v} | \exists \vec{a} \vec{v} = f_1(\vec{x}, \vec{a})\} \quad (23)$$




Figure 3: Measured and calculated velocity profiles for a single tap using the unscaled torque.

We refer to the set V as an impact cone since it forms a cone in the space of initial velocities. The set V can then be mapped to a set of configurations to form the forward projection:

$$FP(\vec{x}) = f_2(\vec{x}, V) \quad (24)$$

The back projection $BP(\vec{x})$ is constructed by “inverting” the forward projection (as in Equation 8).

3.5.1 Impact cones

From our study of Routh’s method for analyzing two dimensional impact with friction, we observe that any impulse within a friction cone at the contact point can be generated by some strike to the object (Huang 1997). We take this as our starting point here.

The impulse generated by impact is related to object velocities by:

$$M\vec{v}_0 = \vec{P} \quad (25)$$

$$I\omega_0 = (\vec{r} \times \vec{P})_z \quad (26)$$

where \vec{r} is the vector from the COM to the contact point. Since the constraint on a possible impulse is on its direction, not its magnitude, it makes sense to consider the ratio of velocities:

$$\frac{\omega_0}{|\vec{v}_0|} = \frac{M}{I}(\vec{r} \times \hat{P})_z \quad (27)$$

Note that ω_0 and v_0 can be made as large as desired by scaling \vec{P} , but their ratio is constrained by the possible directions of \hat{P} .

From our premise, \hat{P} must lie within a friction cone in impulse space (i.e. $|P_t| \leq \mu P_n$). Therefore, there is a continuous range of values for $\frac{\omega_0}{|\vec{v}_0|}$ which can be interpreted as a cone at the origin (i.e. range of slopes) in the space of initial velocities. For circular axisymmetric objects, the contact normal passes through the center of mass, so this cone is symmetric about the v_0 axis. This cone represents all possible velocities (v_0, ω_0) that can be generated by impact, and we refer to it as the impact cone.

3.5.2 Forward and back projections

The forward projection is the image of the impact cone under the forward mechanics of a sliding object, i.e. the function mapping velocities and initial configuration to the final configuration. We require the following property which can be shown with straightforward use of the chain rule.

Property 5 *If $v_0(t)$ and $\omega_0(t)$ are solutions to the equations of motion (Equations 11 and 12) then so are $kv_0(\frac{t}{k})$ and $k\omega_0(\frac{t}{k})$. If the former produce a displacement of x_f and θ_f , then the latter produce a displacement of k^2x_f and $k^2\theta_f$.*

The consequence of this property is that a cone at the origin in the space of initial velocities maps to a cone at the origin in the space of displacements.

The edge of the forward projection cone must be determined by performing a forward integration on some initial velocities that lie on the edge of the impact cone. Since the impact cone, for circular axisymmetric objects, is symmetric about the v_0 axis, the forward projection is symmetric about the x_f axis.

Note that the forward projection actually consists of two individual cones placed back to back because the object can be tapped in either direction along the line from the start to the goal.

The backprojection from a configuration is identical to the forward projection because of the symmetry of the forward projection about the x_f and θ_f axes.

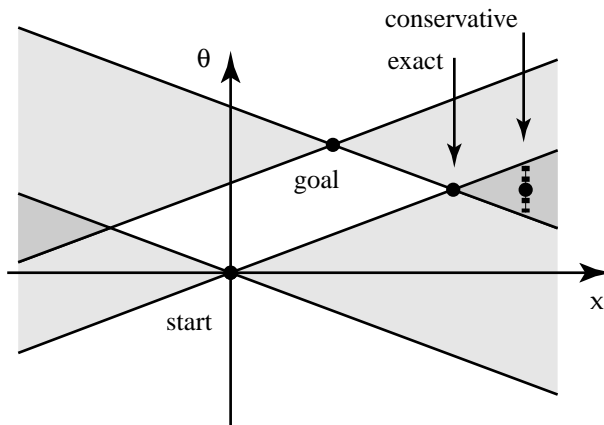


Figure 4: Forward projection from the start and backprojection from the goal. Also shown are the subgoals chosen by the exact and conservative two-tap planning methods.

3.5.3 Basic motion planning

We first construct the backprojection from the goal. If the start configuration is not in this back-projection, then we construct the forward projection from the start configuration. It is easy to see (as illustrated in Figure 4) that the forward projection from the start and the backprojection from the goal will always intersect, so any configuration can be reached with two taps. Furthermore, there will be a two dimensional set of configurations which can be a valid subgoal. A plan can be constructed by picking any subgoal and calculating tapping parameters accordingly.

4 Feedback control

Analysis of the mechanics and formulation of the basic planning method are the basis for feedback control laws that can robustly accomplish positioning tasks. For this problem, the only relevant state is the current configuration of the part. We assume there is no sensing error; errors arise only due to actuation and from inaccurate parameters and modeling. The control problem is to determine the next tap to perform based upon the current configuration of the object.

Our approach is to create a plan from the current configuration (along the line to the goal) using a specific instance of the basic planning method and to execute the first tap of that plan. For the next tap, a new plan is generated from the actual object configuration. In this section, we present several different planning methods and experimental results from corresponding positioning experiments.

In addition, we have found that it is possible to position objects more precisely than the tapping device is positioned, confirming a conjecture we made early in the course of this work. We describe the results of these high-precision positioning experiments.

4.1 Planning methods

The basic planning method does not prescribe which of all possible subgoals should be used for a two-tap plan to reach the goal, and because of errors and tapping device limitations, a two-tap plan is not necessarily the best plan. We have formulated three different planning methods which we describe in this section.

The first two are simple implementations of the basic planning method. For both methods, if the goal can be reached with a single tap, then that tap is used. Otherwise:

- *Exact two-tap planning*: choose the closest possible subgoal.




Figure 5: The first four cones from the goal to the right for the multi-tap planning method.




Figure 6: A trial using the exact two-tap planning method. This figure (and Figures 7–9) shows a positioning trial in configuration space on the left and in cartesian space on the right (note that the axes are not scaled equally). The dotted lines and grey dots are the planned taps; the solid lines and solid dots are the actual taps.

- *Conservative two-tap planning*: choose the closest possible subgoal such that every configuration within an error ellipse about that subgoal is also a valid subgoal.

Because preliminary experiments showed that error in orientation was generally bigger and more important than translation, we use a degenerate error ellipse (a line). The subgoals chosen by these planning methods are illustrated in Figure 4.

The third planning method, which we refer to as *multi-tap planning* plans backwards from the goal. Starting at the goal, a “cone” of configurations is identified that can reach the goal with a sufficiently small tap that the error ellipse is smaller than the positioning accuracy of the task. Successive cones are constructed of configurations which can reach some configuration in the previous cone such that the error ellipse from that tap lies entirely within that cone. Figure 5 illustrates a sequence of these cones extending away from the goal (only considering taps that move the part to the left). For configurations not in some cone, a tap is planned to the closest cone so that the error ellipse from that tap lies entirely within the cone.

4.2 Positioning experiments

We performed experiments using all three planning methods. The experimental setup used an aluminum disk (9.6 cm diameter, 1.3 cm high, 246 grams) sliding on a formica surface ($\mu = 0.19$). An Adept 550 robot was used to position a tapping device with respect to the object, and an overhead camera measured the position and orientation of the object.

The tapping device was similar to a pinball plunger: a steel shaft with a large ball bearing attached to the striking end ($e=0.86$). This shaft slides through linear ball bearings and compresses a spring as it is pulled back. An adjustable backstop controls how much the spring is compressed when the shaft is latched in place.

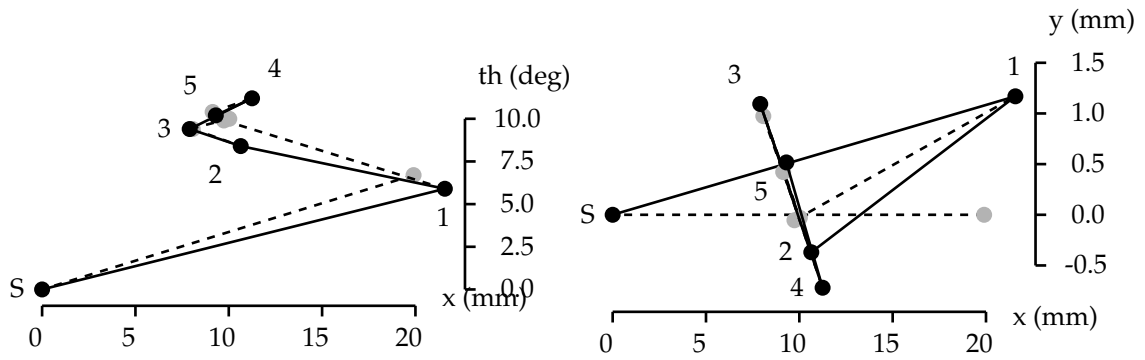


Figure 7: A typical trial using the exact two-tap planning method.




Figure 8: A typical trial using the conservative two-tap planning method.




Figure 9: A typical trial using the multi-tap planning method.

The task for all the positioning experiments was a displacement of 10 mm translation and 10 degrees rotation. This is approximately three times as much rotation as can be achieved in a single 10 mm tap. The desired accuracy was to within 1 mm and 1 degree; this was limited in part by the minimum possible tap of the tapping device and in part by the resolution of the overhead camera.

Overall, the exact two-tap planning method worked the best, taking an average of 4–5 taps. The best run took exactly two taps as illustrated in Figure 6; a more typical run is shown in Figure 7. The conservative two-tap plans were only slightly worse, taking an average of 5–6 taps; a typical run is shown in Figure 8. The multi-tap plans, with one exception, were executed exactly as planned, i.e. they were supposed to take and did take exactly 5 taps to reach the goal. A typical run is shown in Figure 9.

There appears to be a tradeoff among these planning methods in the average number of taps required to reach the goal and the variance in the number of taps. The exact two tap plans have the lowest average number of taps but the highest variance, whereas the multi-tap plans have very little variance but a higher average.

4.3 High-precision positioning

Early in the course of this work, we conjectured that by tapping, an object could be positioned more precisely than the manipulator could position the tapping device with respect to the object. In order to test this idea, we performed a number of experiments using the exact two-tap planning method, but the position and orientation of the tapping device requested by the planning method was rounded to the nearest 15 mm in position and 10 degrees in orientation (with respect to the world coordinates). The experimental setup was otherwise identical to the previous positioning experiments, and the same task (a 10 mm, 10 degrees displacement to an accuracy of 1 mm and 1 degree) was used. In all trials, the task was successfully completed; a typical trial required 5–6 taps — only slightly worse than without coarse positioning of the tapping device. We provide some analytic observations to support this result in the next section.

5 Stability & sensitivity

The positioning experiments described in the previous section experimentally demonstrate the stability of the system, both for different feedback control methods and for when the tapping actuator can only be positioned coarsely. In this section, we first derive conditions for the asymptotic stability of tapping, and then show how the sensitivity of tapping errors permits high precision positioning.

Conditions on the error bounds of a tap are obtained by defining a Lyapunov distance function on the configuration space and requiring that each tap bring the part closer to the goal configuration. We examine the exact two-tap planning strategy here, though this analysis could easily be applied to other planning methods we have used.

The sensitivity analysis examines the Jacobian relating errors in the configuration of the tapping device to errors in the final configuration of the part. We show that the latter are smaller than the former for our experimental setup, which supports our experimental results in high precision positioning.

5.1 Asymptotic stability

Without loss of generality, assume that the goal is the origin of the configuration space (x, y, θ) and that the x axis is oriented so that the current configuration lies on the x - θ plane. We define a Lyapunov distance function based on the backprojection from the goal in this plane.

We define the distance of a configuration \vec{q} from the origin as follows. If the configuration is inside the backprojection, then its distance is $|\vec{q}|$. Otherwise, we express the configuration as $\vec{q} = k_1\vec{c}_1 + k_2\vec{c}_2$, where the vectors \vec{c}_1 and \vec{c}_2 bound the backprojection cones. The distance in this




Figure 10: Concentric “circles” under the Lyapunov distance function defined on the two dimensional configuration space.

case is defined as $|k_1| + |k_2|$. These definitions reflect the “length” of the taps needed with the exact two-tap planning strategy to reach the goal. Figure 10 illustrates these two cases and shows concentric “circles” under this distance function. These curves can be revolved about the θ axis to form the equivalent of spheres under this distance function.

In order to establish conditions for asymptotic stability, we assume that the error for a tap can be bounded by an ellipse whose radii are linearly related to the size of the tap.

$$\begin{bmatrix} r_x \\ r_y \\ r_\theta \end{bmatrix} = \begin{bmatrix} \vec{e}_x^T \\ \vec{e}_y^T \\ \vec{e}_\theta^T \end{bmatrix} \vec{p} \quad (28)$$

where $\vec{p} = [|x_f| \quad |\theta_f|]^T$ is the size of the tap, \vec{r} represents error ellipse radii and the \vec{e} vectors simply relate the tap displacement to an error ellipse radius. Since the error radii must be positive, each component of the \vec{e} vectors must be positive. We can now state the following:

Theorem 1 *Tapping is asymptotically stable under the two-tap planning method if the following hold:*

- for all configurations \vec{q} in the backprojection,

$$\vec{e}_x \cdot \vec{p} < |\vec{q}| \quad (29)$$

$$\vec{e}_y \cdot \vec{p} < |\vec{q}| \quad (30)$$

$$\vec{e}_\theta \cdot \vec{p} < |\vec{q} \cdot \hat{\theta}| \quad (31)$$

- for all other configurations,

$$\vec{e}_x \cdot \hat{c}_2 < \hat{c}_2 \cdot \hat{x} \quad (32)$$

$$\vec{e}_y \cdot \hat{c}_2 < \hat{c}_2 \cdot \hat{x} \quad (33)$$

$$\vec{e}_\theta \cdot \hat{c}_2 < \hat{c}_2 \cdot \hat{\theta} \quad (34)$$

These conditions ensure that the error ellipse from a tap lies completely inside the “sphere” defined by the initial configuration. Thus, the object configuration comes closer to the goal with each tap. There are two cases:

1. When the initial configuration is within the backprojection, the first tap attempts to reach the goal directly. The conditions for any configuration within the error bounds to be closer than the start configuration are Equations 29–31. These conditions stipulate that the error ellipse must be small enough to fit inside the “sphere” of radius $|\vec{q}|$ as illustrated in Figure 11a.
2. When the initial configuration is outside the backprojection, the (first) tap attempts to reach an edge of the cone. The error ellipse must again lie within the “sphere” defined by the start configuration. This gives rise to Equations 32–34 and is illustrated in Figure 11b. (These equations conservatively approximate the “sphere” under this distance function by a cylinder.)




Figure 11: Illustration of the conditions for asymptotic stability for tapping (side and top views). The error ellipse must fit inside the “sphere” defined by the start configuration.

5.2 Sensitivity

The experimental high-precision positioning results confirmed our conjecture that it was possible to position objects more accurately than the tapping device could be positioned. In support of this result, we show that error in positioning the tapping device maps to a smaller error in the final object configuration.

The Jacobian relating errors in tapping device configuration (position and orientation) to error in the final object configuration must be evaluated numerically for a given tap because there is no analytic solution for the inverse sliding problem. For a 10 mm 3 degree tap with the object and materials used in the high-precision positioning experiment described in Section 4.3, the average striker error vector $[x \ y \ \theta]^T = [7.5 \ 7.5 \ 5]^T$ (i.e. 7.5 mm position error and 5 degrees orientation error) is mapped to an object final configuration error of $[1.1 \ 0.9 \ 0.2]^T$. Here, the error is smaller in each dimension.

Another measure to characterize this effect is the norm of the Jacobian. A norm of less than one ensures that the positioning error error will be smaller than the striker configuration error. We have found that there is a broad range of taps for which the norm of the Jacobian is less than one. Figure 12 shows a graph of the Jacobian norm for an aluminum disk sliding on a formica surface.

6 Conclusions

We have analyzed the mechanics of tapping for circular axisymmetric objects and have formulated planning methods and a feedback control strategy for accomplishing positioning tasks using tapping. We have shown conditions for the stability of our control and have experimentally demonstrated the viability of our solutions. In addition, our experiments showed that objects can be positioned with higher precision than the tapping device can be positioned, and we provided some analytic support for this observation.

Our method of solving the inverse mechanics of tapping entails numerically finding initial velocities that cause the object to slide a desired displacement and then computing parameters for a tap to generate those velocities. We performed single-tap experiments to validate our analysis and




Figure 12: Norm of the final configuration Jacobian as a function of displacement. x_f is in mm, θ_f in degrees.

found that our models needed to be modified slightly to account for greater than expected frictional torque. Although the discrepancy is undoubtedly in part due to differences between the assumptions of our model and the physical system, we believe there may be some additional mechanism that results in this increased torque.

Our basic planning method proves that it is possible to reach any configuration in at most two taps. However, because of errors in actuation and modeling, feedback control is required to robustly accomplish positioning tasks. We have developed several different specific planning methods which are used to replan a path to the goal after each tap. We have found them to be fairly robust; the best method takes on average 4–5 taps to accomplish a positioning task. The different planning methods display a tradeoff in the average number of taps required and the variation in the number of taps. We have provided some analysis that results in conditions for the asymptotic stability of tapping.

For most positioning methods, objects can be positioned only as precisely as the manipulator can be positioned. For tapping, this is not the case — we have experimentally demonstrated high-precision positioning in which the object is positioned more precisely than the tapping device is positioned. Through sensitivity analysis, we show that tapping device positioning error result in a smaller error in the final object configuration. It is this latter error, combined with error due to tapping (modeling or parameter errors), that determines the precision to which an object can be positioned.

Although this paper describes results for circular axisymmetric objects, several of these results have been extended to the more general class of nonaxisymmetric objects (Huang 1997). This is also the subject of our ongoing research; some of the key problems are finding some solution to the inverse sliding problem for this class of objects and modifying our planning and feedback control to accommodate such a solution.

Acknowledgements

We would like to thank Jason Powell for building and helping design the tapping devices and the Medical Robotics and Computer Assisted Surgery lab at Carnegie Mellon for use of their OptoTrak system.

This work was supported in part by NASA through the Graduate Student Researchers Program and through grant NAGW 1175 and by the NSF under grants IRI-9318496, IRI-9114208, and IIS-9900322.

References

- Alexander, J. C. and Maddocks, J. H. 1993. Bounds on the friction-dominated motion of a pushed object. *International Journal of Robotics Research*, 12(3):231–248.
- Balorda, Z. 1993. Automatic planning of robot pushing operations. In *IEEE International Conference on Robotics and Automation*, pages 1:732–737, Atlanta, GA.
- Bishop, B. E. and Spong, M. W. 1999. Vision-based control of an air hockey robot. *IEEE Control Systems Magazine*, 19(3).
- Bowden, F. P. and Tabor, D. 1964. *The Friction and Lubrication of Solids: Part II*. Oxford University Press.
- Brach, R. M. 1991. *Mechanical Impact Dynamics: Rigid Body Collisions*. John Wiley & Sons.
- Bühler, M. and Koditschek, D. E. 1990. From stable to chaotic juggling: Theory, simulation, and experiments. In *IEEE International Conference on Robotics and Automation*, pages 1976–1981, Cincinnati, OH.
- Canny, J. and Goldberg, K. 1994a. A RISC approach to robotics. *IEEE Robotics & Automation Magazine*, 1(1):26–28.
- Canny, J. F. and Goldberg, K. Y. 1994b. “RISC” industrial robotics: Recent results and open problems. In *IEEE International Conference on Robotics and Automation*, pages 1951–1958.
- Chatterjee, A. and Ruina, A. L. 1998. A new algebraic rigid body collision model with some useful properties. *Journal of Applied Mechanics*, 65(4):939–951.
- Erdmann, M. A. 1996. An exploration of nonprehensile two-palm manipulation: Planning and execution. In Giralt, G. and Hirzinger, G., editors, *International Symposium on Robotics Research: The Seventh International Symposium*. Springer-Verlag.
- Goyal, S., Ruina, A., and Papadopoulos, J. 1991a. Planar sliding with dry friction. Part 1. Limit surface and moment function. *Wear*, 143:307–330.
- Goyal, S., Ruina, A., and Papadopoulos, J. 1991b. Planar sliding with dry friction. Part 2. Dynamics of motion. *Wear*, 143:331–352.
- Higuchi, T. 1986. Application of electromagnetic impulsive force to precise positioning tools in robot systems. In Faugeras, O. and Giralt, G., editors, *International Symposium on Robotics Research: The Third International Symposium*. MIT Press.
- Huang, W. H. 1997. *Impulsive Manipulation*. PhD thesis, Carnegie Mellon University. Available as Carnegie Mellon Robotics Institute technical report CMU-RI-TR-97-29.
- Huang, W. H., Krotkov, E. P., and Mason, M. T. 1995. Impulsive manipulation. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 120–125.
- Lynch, K. M. and Mason, M. T. 1996. Stable pushing: mechanics, controllability, and planning. *International Journal of Robotics Research*, 15(6):533–556.
- Mason, M. T. and Salisbury, Jr., J. K. 1985. *Robot Hands and the Mechanics of Manipulation*. The MIT Press.
- Partridge, C. B. and Spong, M. W. 1999. Control of planar rigid body sliding with impacts and friction. *International Journal of Robotics Research*, 19(4):336–348.

- Peshkin, M. A. and Sanderson, A. C. 1988. The motion of a pushed, sliding workpiece. *IEEE Journal of Robotics and Automation*, 4(6):569–598.
- Rizzi, A. A. and Koditschek, D. E. 1992. Progress in spatial robot juggling. In *IEEE International Conference on Robotics and Automation*, pages 775–780, Nice, France.
- Routh, E. 1960. *The Elementary Part of a Treatise on the Dynamics of a System of Rigid Bodies*. Dover Publications, eighth edition. Originally published by Macmillan and Company, 1905.
- Schaal, S. and Atkeson, C. G. 1993. Open loop stable control strategies for robot juggling. In *IEEE International Conference on Robotics and Automation*, pages 3:913–918, Atlanta, GA.
- Stronge, W. J. 1990. Rigid body collisions with friction. *Proceedings of the Royal Society, London*, A431:169–181.
- Trinkle, J. C., Farahat, A. O., and Stiller, P. F. 1995. First-order stability cells of active multi-rigid-body systems. *IEEE Transactions on Robotics and Automation*, 11(4):545–557.
- Trinkle, J. C., Ram, R. C., Farahat, A. O., and Stiller, P. F. 1993. Dexterous manipulation planning and execution of an enveloped slippery workpiece. In *IEEE International Conference on Robotics and Automation*, pages 2:442–448.
- Voyenli, K. and Eriksen, E. 1985. On the motion of an ice hockey puck. *American Journal of Physics*, 53(12):1149–1153.
- Wang, Y. and Mason, M. T. 1992. Two-dimensional rigid-body collisions with friction. *ASME Journal of Applied Mechanics*, 59:635–641.
- Yamagata, Y. and Higuchi, T. 1995. A micropositioning device for precision automatic assembly using impact force of piezoelectric elements. In *IEEE International Conference on Robotics and Automation*.
- Yamagata, Y., Higuchi, T., Saeki, H., and Ishimaru, H. 1990. Ultrahigh vacuum precise positioning device utilizing rapid deformations of piezoelectric elements. *Journal of Vacuum Science & Technology A*, 8(6):4098–4100.
- Zhu, C., Aiyama, Y., Chawanya, T., and Arai, T. 1996. Releasing manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 2:911–916, Osaka, Japan.
- Zumel, N. B. 1997. *A Nonprehensile Method for Reliable Parts Orienting*. PhD thesis, Carnegie Mellon University.