

Mechanics, Planning, and Control for Tapping

Wesley H. Huang, *Carnegie Mellon University, Pittsburgh, PA, USA*

Matthew T. Mason, *Carnegie Mellon University, Pittsburgh, PA, USA*

We have been studying manipulation using impulsive forces; in particular, we have focused on the manipulation of planar sliders by tapping. These objects then slide on a support surface, slowing down and coming to rest due to frictional forces. In this paper, we present analytical and experimental results in the mechanics, planning, and control of this system.

1 Introduction

We consider *impulsive manipulation* to be any form of manipulation consisting of two phases: a strike to the object imparting some initial velocities and free motion of that object subject to forces and constraints in the environment. In this work, we have focused on manipulation of planar sliders by tapping. After being struck, the object slides on the support surface, slowing down and coming to rest due to frictional forces.

Using impulsive forces for manipulation has several unique characteristics as compared to other modes of manipulation. Tapping is well suited for micropositioning — the impact instantaneously breaks static friction, but the amount of impulse delivered determines how far the object moves. Tapping results in quick dynamic motion, requires a relatively simple actuator, and can manipulate a wide variety of objects due to its nonprehensile (nongrasping) nature.

There are many ways tapping could be instantiated. It could be used for fast and relatively low precision parts transfer between workcells. Our analysis might be applied to the design and configuration of certain types of vibratory parts feeding systems. A “positioning cell” could be built using a number of tapping actuators placed around its perimeter; an object placed in the cell could be precisely positioned by firing an appropriate sequence of actuators. The setup we have used for our experiments features 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.

1.1 Overview

We first consider how to plan a single tap to achieve a desired displacement (translation and rotation). Because the manipulator interacts with the object only at the instant of impact, we can consider separately the mechanics of a sliding object and the mechanics of striking that object to produce some initial velocities. The dynamics of a sliding rotating object give rise to a strongly coupled system for which there is no general analytic solution, so we have developed a numerical technique to determine the required initial velocities for a desired object translation and rotation. Determining how to generate those velocities by impact has been a relatively straightforward application of the classical impact models. In general, we must search the boundary for an appropriate contact point; simple shapes may have analytic solutions for the striker parameters required to generate the initial velocities.

We performed a series of single-tap experiments to validate the models used in our analysis. To our surprise, we found that objects did not rotate as much as our model predicted. Examination of the rotational velocity profiles revealed that there was more frictional torque than predicted, so we added an additional parameter to our model to scale the torque due to friction. Although this discrepancy is in part due to differences between our idealized model assumptions and the actual experimental setup, we believe there may be some additional mechanism giving rise to this increased frictional torque. We are still able to use our original solution to the mechanics of a sliding object with this modified model.

A single tap cannot arbitrarily place an object in the plane; multiple taps may be required. We first charac-

terize the set of states that can be reached with a single tap and then formulate a basic planning method which shows that any configuration can be reached with at most two taps. Given the large variations in single-taps, however, a two-tap plan may not be desirable. We extend this basic planning method to several specific planning methods which essentially serve as feedback control laws. Each planning method has an implicit model of the error from a single tap and makes some trade off between the expected number of taps required to position an object and the variation in the number of taps. The performance of these planning methods was demonstrated in several positioning experiments.

The final result of this paper is an experimental proof of a conjecture we made early in the course of this work: that tapping can be used to position an object more precisely than the manipulator can position the tapping device. Through a basic sensitivity analysis, we show that there are a wide range of taps for which the error in object configuration is smaller than the error in tapping device configuration.

1.2 Related work

The first appearance of impulsive manipulation in the robotics literature was due to Higuchi [10] who used an electromagnetic coil to deliver an impulsive force for linear micropositioning. Since then, Yamagata and Higuchi (e.g. [22]) have developed related devices for several different micropositioning applications.

The present work is the first to show formal results in the mechanics, planning, and control of tapping. Our analysis of the mechanics of a sliding object is built upon the work of Voyerli and Eriksen [20] who studied the motion of sliding disks and rings and the work of Goyal *et al.* ([8] and [9]) who formulated a limit surface representation of the relationship between net frictional force and torque and the the velocities of a planar slider.

Our work makes use of impact analysis due to Routh, more recently presented in Wang and Mason [21]. This work presents a graphical method of using Newton’s and Poisson’s laws of restitution for two dimensional impact with friction. Although there has been much recent work on impact (e.g. Bowden and Tabor [1], Brach [2], Stronge [17], and Chatterjee and Ruina [6]), the simplicity of Poisson’s hypothesis makes our analysis simpler; it seems to work reasonably well in practice.

This work has been inspired by the ideas behind minimalist robotics, articulated by Canny and Goldberg ([4] and [5]) and by work in many areas of non-prehensile manipulation, such as pushing (Mason [14] and many others), palmar manipulation (Trinkle *et al.* ([19] and [18]), Erdmann [7], and Zumel [24]), and paddle juggling (Bühler and Koditschek [3], Rizzi and Koditschek [15], and Schaal and Atkeson [16]).

Experiments with planar sliders have been performed by Zhu *et al.* [23] in conjunction with what they call “releasing manipulation.” In their experiments, objects were accelerated using a manipulator (instead of striking them) and then released, allowing them to slide on a support surface and come to rest due to friction. Their results, like ours, show considerable variation, presumably resulting from variations in friction.

The results presented in this paper appear in greater detail in Huang [11].

1.3 Assumptions

We assume that the support surface is homogeneous. 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 assume a uniform support distribution, though our methods can work with more general support distributions. 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 Mechanics

In this section, we show how to plan a single-tap. This can be decomposed into two subproblems: 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). The impact problem is to determine how (if possible) to generate those velocities by striking the object.




Figure 1: Notation for the axisymmetric inverse sliding problem. The goal configuration is at a distance x_f along the \hat{x} axis and a rotation of θ_f . The object's position and orientation during the sliding motion are x and θ , its velocities v and ω .

2.1 Inverse sliding problem

We address the inverse sliding problem separately for the classes of axisymmetric and nonaxisymmetric objects. Axisymmetric objects are those which have a pressure distribution that is a function of radius only; these objects have the property that they will always slide in a straight line. This is not the case for nonaxisymmetric objects, which we will address at the end of this subsection.

Since axisymmetric objects slide in a straight line, we only need consider the translation along the line from the start to the goal instead of translation in two dimensions. The problem is then:

Determine the initial velocities v_0 and ω_0 such that the object comes to rest after sliding a desired distance d and angle α .

See Figure 1 for notation.

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

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

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

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

Note that the force and torque are functions of the velocity ratio $\frac{\omega}{v}$, not the velocities themselves. Furthermore Mason [14] showed that as a function of $\frac{\omega}{v}$, the force is monotonic decreasing, the torque monotonic increasing.

The equations of motion of this system are:

$$M \dot{v} = -F\left(\frac{\omega}{v}\right) \quad (3)$$

$$I \dot{\omega} = -T\left(\frac{\omega}{v}\right) \quad (4)$$

This is a pair of strongly coupled first order differential equations. The net translation and rotation of the object are given by:

$$x_f(v_0, \omega_0) = \int_0^{t_f} v(t) dt \quad (5)$$

$$\theta_f(v_0, \omega_0) = \int_0^{t_f} \omega(t) dt \quad (6)$$

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:

$$v(0) = v_0 \quad \omega(0) = \omega_0 \quad (7)$$

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

In general, there are no analytic expressions for the net force and torque due to friction, so there will be no analytic expression for x_f and θ_f . However, we have been able to prove a number of properties about this system. We state the following without proof; the interested reader can consult Huang [11] or ([12] and [13]) 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:*

$$v_{01} > v_{02} \quad \omega_{01} = \omega_{02} \quad \implies x_{f1} > x_{f2} \quad (8)$$

$$v_{01} > v_{02} \quad \omega_{01} = \omega_{02} \neq 0 \quad \implies \theta_{f1} > \theta_{f2} \quad (9)$$

$$v_{01} = v_{02} \neq 0 \quad \omega_{01} > \omega_{02} \quad \implies x_{f1} > x_{f2} \quad (10)$$

$$v_{01} = v_{02} \quad \omega_{01} > \omega_{02} \quad \implies \theta_{f1} > \theta_{f2} \quad (11)$$

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

Property 3 *The level curves of x_f and θ_f are strictly monotonic decreasing in the space of initial velocities.*

Property 4 *The value of x_f is strictly monotonic increasing along level curves of θ_f and vice versa.*




Figure 2: Illustration of level curves. The level curve of x_f corresponding to a distance d asymptotically approaches the ω_0 axis and must lie to the left of its v_0 intercept. Likewise, the level curve of θ_f corresponding to a rotation α asymptotically approaches the v_0 axis and must lie below its ω_0 intercept.

These properties imply that the level curve of x_f for a distance d and the level curve of θ_f for a rotation α must intersect exactly once and that this intersection point can be bounded. (See Figure 2 for an illustration.) The coordinates of this intersection point correspond to the initial velocities required to achieve the desired displacement.

These properties also lead us to a method for numerically finding the coordinates of this intersection point:

- Subdivide the bounding rectangle.
- Numerically integrate the initial conditions $(\bar{v}_0, \bar{\omega}_0)$ corresponding to a corner of a subrectangle.
- If the distance or angle is greater than the desired distance or angle, then the region of the bounding rectangle where $v_0 > \bar{v}_0$ and $\omega_0 > \bar{\omega}_0$ can be eliminated from consideration.
- If the distance or angle is less than the desired distance or angle, then the region of the bounding rectangle where $v_0 < \bar{v}_0$ and $\omega_0 < \bar{\omega}_0$ can be eliminated from consideration.

In this manner, we can “zero in” on the coordinates of the intersection point.

2.1.1 Nonaxisymmetric objects

Nonaxisymmetric objects will in general have some component of the net frictional force that is normal to

the net translational velocity, so these objects will not slide in a straight line. The inverse sliding problem is then difficult to solve because the monotonicity properties no longer hold and because the problem is higher dimensional.

In the simulations we have done with nonaxisymmetric objects, the movement of the object in the direction normal to the translational velocity has been very small, suggesting one possible solution to this problem — simply ignore this normal motion, essentially treating the object as axisymmetric. Given the large variations in actual experiments, this appears to be a very reasonable strategy. Errors due to this modeling simplification can be dealt with through feedback control.

The alternative is to employ a more general numerical root finding method; however, we have not been able to prove the existence or the uniqueness of a solution for the general nonaxisymmetric case.

2.2 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. We take as the starting point of our analysis the fact that any impulse within a friction cone at the contact point can be generated. (See Huang [11] for details.)

The friction cone at the impact point constrains only the direction of the impulse. The resulting velocities will scale with the magnitude of the impulse, so we need only consider the set of velocity ratios $\frac{\omega_0}{|\bar{v}_0|}$ and the direction of the translational velocity.

For a given contact point, we can produce a range of velocity ratios, each with translational velocity in a different direction. However, the question we face in planning is whether we can produce a desired velocity ratio in a given direction. This will generally lead to searching the boundary to find a suitable contact point.

For simple shapes, we can write analytic forms for the range of possible velocity ratios as a function of translational velocity direction; it may also be possible to write an analytic function mapping the velocity direction and velocity ratio to the contact point.

Once we have determined the contact point, we must then go through the mechanics of impact to determine the exact striker parameters (velocity and angle of incidence) in order to produce the necessary impulse.

Note that the limitation on tapping is due to the impact process — for any desired translation and rotation, we can determine initial velocities for the object to slide the proper displacement, but we are limited in the velocity ratios that can be produced via impact. Although this is in part due to the shape of the object, a more fundamental reason is the limited amount of rotation per translation that can be produced — friction limits the impulse to lie within a cone about the contact normal.

2.3 Single-tap experiments

Our first experiments were single-tap experiments designed to validate the model used in analysis of the mechanics. There were significant deviations between the idealized model for the analysis and the actual setup that was used in the experiments. First and foremost, the objects do not have a uniform pressure distribution; secondly, the objects are three dimensional, not planar.

We tested the model of the mechanics using three different combinations of object, materials, and support surfaces. Initial experiments were not successful — 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 which measured the position and orientation of the object at 500 Hz).

These measurements allowed us to calculate initial velocities and velocity profiles of the object, which in turn allowed us to evaluate the impact model and the sliding model separately.

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 (see Figure 3 for an example), there was more frictional torque than predicted.

We then added an additional parameter to the model to scale the net torque due to friction. For the three sets of experiments we did, the scale factors were 1.0, 2.17, and 3.64. However, with this modified model (which still allows us to use the same method for solving the inverse sliding problem) we were able to attain




Figure 3: Measured and calculated velocity profiles for the aluminum square experiment using the unscaled torque. Although the translational velocity matches well, the rotational velocity profile reveals that there is more frictional torque than predicted.

reasonable matches between predicted and actual velocity profiles.

In the end, we were able to predict object motion on average to within 5–10% for translations and to within 10–25% for rotations. The impact model generally predicts initial translational velocities to within a few percent and initial rotational velocities to within a few percent up to 33%.

One of the major results of our single-tap experiments is simply that real three dimensional objects (even those that aren't quite axisymmetric) can be modeled reasonably well by the axisymmetric mechanics.

3 Planning & Control

We know that a single tap will not be able to achieve any given displacement because of constraints imposed by impact. In addition, we must be able to deal with the large experimental variations that result from a single tap. We have found that for this problem, planning and control are intertwined — a basic planning method serves as the basis for feedback control.

Since the axisymmetric mechanics are more widely applicable than one would suspect (given the assumptions of the underlying model) we consider planning and feedback control using only these mechanics. Furthermore, we primarily consider manipulation of circular objects in the following subsections.

All our planning methods work in the configuration space of the object, specified by a position x and orientation θ .

3.1 Basic planning method

Before we can consider multiple taps, we must first characterize the displacements we can achieve with a single tap. Through straightforward use of the chain rule, we can show that if $v_0(t)$ and $\omega_0(t)$ are solutions to the equations of motion (Equations 3 and 4) then so are $kv_0(\frac{t}{k})$ and $k\omega_0(\frac{t}{k})$. If the former produces a displacement of x_f and θ_f , then the latter produces a displacement of k^2x_f and $k^2\theta_f$. Therefore, all initial velocities with the same velocity ratio $\frac{\omega_0}{|v_0|}$ produce displacements with the same displacement ratio $\frac{x_f}{\theta_f}$.

For objects with smooth boundaries, we can produce a continuous range of velocity ratios in any given direction, which will result in a continuous range of displacement ratios. For circular objects, these ranges will be symmetric about zero and identical for all translational velocity directions. The range of displacement ratios can be viewed as a range of slopes in configuration space, so the set of states that can be reached in a single tap forms a cone symmetric about the x axis in configuration space; we refer to this cone as the reachable displacement cone.




Figure 4: *Planning for axisymmetric objects. The reachable displacement cones for each direction are placed at the start configuration, the backprojections at the goal configuration. (For circular axisymmetric objects, these cones are identical.) The intersection is the set of all points that may be used as an intermediate configuration in getting from the start to the goal in two taps.*

The set of states then that can be reached in a single tap lie within this cone starting at the current state. The set of states that can reach the goal configuration in a single tap can be determined by reflecting this reachable displacement cone at the goal through the goal state. Since the reachable displacement cone is symmetric, this backprojection looks exactly the same.

It is easy to see (as illustrated in Figure 4) that by intersecting the reachable displacement cone from the start configuration with the backprojection of the goal, we can reach any state within two taps. Furthermore, there will be a two dimensional set of states which can be used as a subgoal to get from the start to the goal in two taps.

3.2 Extensions of the basic planning method for feedback control

The basic planning method does not tell us how to choose a subgoal when two taps are required. In addition, because of tapping device limitations and errors in modeling and actuation, a two tap plan is not necessarily the best plan. We have examined three different schemes based on the basic planning method; these schemes lie somewhere between traditional planners and feedback control laws.

3.2.1 Exact two-tap plans

The simplest way to implement the basic planning method is simply to choose the closest feasible subgoal.

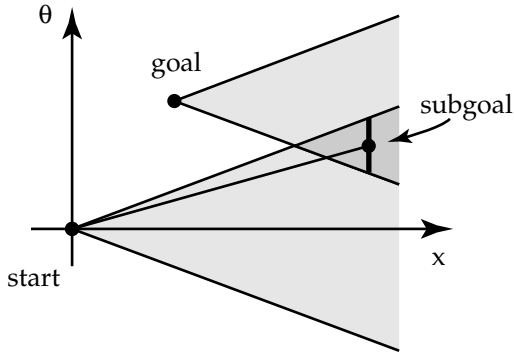


Figure 5: *The conservative two-tap planning method chooses a subgoal such that the (degenerate) error ellipse fits entirely within the space of feasible subgoals.*

We implement this plan as follows: if the goal can be reached in a single tap, then we execute that tap, otherwise, we execute the first tap of an exact two-tap plan to reach the goal from the current configuration.

3.2.2 Conservative two-tap plans

A further refinement we can make is to consider explicitly a model of the error in a single tap. In preliminary experiments, we observed that errors in rotation tended to be much larger (or more important) than errors in translation, so we use the degenerate case of the traditional error ellipse: an error “line” parallel to the θ axis in configuration space.

A potential problem with the exact two-tap plans is that if there is any error in the first tap, the object may not be in a state that can reach the goal configuration in one more tap. So in planning conservative two-tap plans, we select a subgoal so that the entire error ellipse for the first tap fits within the region of allowable subgoals. A conservative two-tap plan is illustrated in Figure 5.

Like the exact two-tap plans, this planning method is executed as follows: if the goal can be reached in one tap, then we execute that tap, otherwise, the first tap of a conservative two-tap plan from the current state to the goal is executed.

3.2.3 Multi-tap plans

One effect of using the conservative two-tap planning strategy is that the excursion in going from the start to the goal increases. Although we plan for the




Figure 6: *The first four cones from the goal to the right for the multi-tap planning method.*

error ellipse from the first tap to lie entirely within a backprojection from the goal, we have not explicitly considered the error ellipse from the second tap. The larger the second tap is, the larger the error ellipse about the goal will be.

Another approach, then, is to consider planning backwards from the goal. We first consider a final tap small enough to have an acceptable error about the goal. There will be a small “cone” of states that can reach the goal with a sufficiently small tap. Then we consider the set of states that can reach this first cone such that the error ellipse from that tap lies entirely within the first cone. In this manner we construct a series of cones in configuration space (illustrated in Figure 6) extending backwards from the goal. We then need to plan a tap that takes us from the start configuration into one of these cones such that the error ellipse from that tap lies completely within the sequence of cones.

This planning method is executed as follows: for every tap, we determine whether the current state is within this sequence of cones. If so, then we execute a tap to take us to the center of the next cone. Otherwise, we execute a tap to take us into the closest cone.

3.3 Positioning experiments

We performed experiments using all three planning methods. The experimental setup used an aluminum disk sliding on a formica surface. 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 task for all the positioning experiments was a displacement of 10 mm translation and 10 degrees rotation. This is about three times as much rotation as can be achieved in a single 10 mm tap. The desired




Figure 7: A nearly ideal trial using the exact two-tap planning method. This Figure (and Figures 8-10) shows a positioning trial in configuration space on the left and in cartesian space on the right. The dotted lines and grey dots are the planned taps; the solid lines and solid dots are the actual taps.

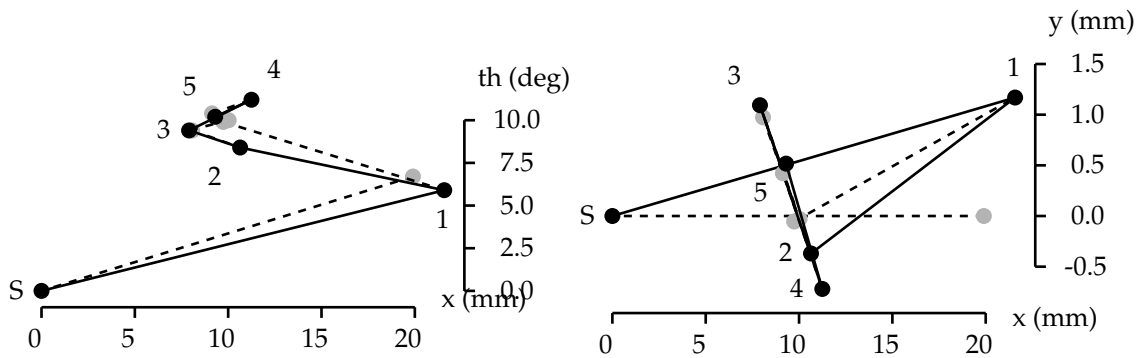


Figure 8: Example of a typical trial using the exact two-tap planning method.




Figure 9: Example of a typical trial using the conservative two-tap planning method.




Figure 10: Example of a typical trial using the multi-tap planning method.

accuracy was to within 1 mm and 1 degree; this was limited in part by the dynamic range 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 7; a typical run is shown in Figure 8. The conservative two-tap plans were only slightly worse, taking an average of 5–6 taps; a typical run is shown in Figure 9. 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 10.

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. The multi-tap plans have very little variance but a higher average. We hope to more formally characterize this tradeoff in future work.

3.4 High-precision positioning

Early in the course of this work, we conjectured that by tapping, an object could be positioned more pre-

cisely 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 in which 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 task was again to displace the object by 10 mm and 10 degrees to an accuracy of 1 mm and 1 degree. In all trials, the task was successfully completed; a typical trial took 5–6 taps.

3.4.1 Sensitivity analysis

In order to characterize this robustness for high-precision positioning, we did a basic sensitivity analysis by examining the Jacobian relating changes in the position and orientation of the tapping device to the resulting change in the position and orientation of the object after a tap. All but one component of the Jacobian can be computed analytically; that component relates the change in initial object velocities to the changes in displacement (recall that there is no analytic form for the displacement functions x_f and θ_f). Since this component of the Jacobian must be evaluated numerically, we must look at variation with respect to a particular tap. For a 10 mm 3 degree tap (using the




Figure 11: Norm of the final configuration Jacobian as a function of displacement. x_f is in mm, θ_f in degrees.

same object and materials as the high-precision positioning experiment), the maximum striker error vector $[x \ y \ \theta]^T = [7.5 \ 7.5 \ 5]^T$ (corresponding to the maximum error due to rounding in this experiment) is mapped to an object error of $[1.1 \ 0.9 \ 0.2]^T$. Here, the error is smaller in each dimension.

Furthermore, we have found that there is a broad range of taps for which the norm of the Jacobian is less than one and thus the error in object configuration is always smaller than that of the tapping device configuration. This is shown graphically in Figure 11.

4 Conclusions

We have shown a framework for manipulation by tapping, one particular form of impulsive manipulation.

Through analysis of the mechanics of a planar slider, we have demonstrated a solution to the single-tap planning problem. We performed single-tap experiments to validate the models used in our analysis and found that there was more frictional torque acting upon the object than our model predicted. An additional parameter in our model scales the frictional torque to bring our model and experimental observations into agreement. Although this deviation from the original model is in part due to differences between the physical experimental apparatus and the idealized model assumptions, primarily uniform support distribution and three dimensional extent, we believe there may be some additional mechanism that results in this increased torque.

We have formulated a basic planning method which shows that it is possible to position an axisymmetric

object anywhere in the plane within two taps. This basic planning method was then adapted to create practical planning methods which essentially serve as feedback control laws; these planning methods specify how to choose intermediate state(s) in order to reach the goal in the presence of variations due to modeling and actuation errors. We demonstrated three different planning methods which display a tradeoff between the average number of taps required to complete a task and the variation in the number of taps.

Finally, we have demonstrated that by tapping it is possible to position an object more precisely than the manipulator positions the tapping device with respect to the object. We have presented a basic stability analysis in support of this result.

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 and IRI-9114208.

References

- [1] F. P. Bowden and D. Tabor. *The Friction and Lubrication of Solids: Part II*. Oxford University Press, 1964.
- [2] R. M. Brach. *Mechanical Impact Dynamics: Rigid Body Collisions*. John Wiley & Sons, 1991.
- [3] M. Bühler and D. E. Koditschek. From stable to chaotic juggling: Theory, simulation, and experiments. In *IEEE International Conference on Robotics and Automation*, pages 1976–1981, Cincinnati, OH, 1990.
- [4] J. Canny and K. Goldberg. A RISC approach to robotics. *IEEE Robotics & Automation Magazine*, 1(1):26–8, March 1994.
- [5] J. F. Canny and K. Y. Goldberg. “RISC” industrial robotics: Recent results and open problems. In *IEEE International Conference on Robotics and Automation*, pages 1951–8, 1994.

- [6] A. Chatterjee and A. L. Ruina. A new algebraic rigid body collision model with some useful properties. *Journal of Applied Mechanics*, July 1997. Submitted.
- [7] M. A. Erdmann. An exploration of nonprehensile two-palm manipulation: Planning and execution. In G. Giralt and G. Hirzinger, editors, *International Symposium on Robotics Research: The Seventh International Symposium*. Springer-Verlag, 1996.
- [8] S. Goyal, A. Ruina, and J. Papadopoulos. Planar sliding with dry friction. Part 1. Limit surface and moment function. *Wear*, 143:307–330, 1991.
- [9] S. Goyal, A. Ruina, and J. Papadopoulos. Planar sliding with dry friction. Part 2. Dynamics of motion. *Wear*, 143:331–352, 1991.
- [10] T. Higuchi. Application of electromagnetic impulsive force to precise positioning tools in robot systems. In O. Faugeras and G. Giralt, editors, *International Symposium on Robotics Research: The Third International Symposium*. MIT Press, 1986.
- [11] W. H. Huang. *Impulsive Manipulation*. PhD thesis, Carnegie Mellon University, Aug. 1997. Available as Carnegie Mellon Robotics Institute technical report CMU-RI-TR-97-29.
- [12] W. H. Huang, E. P. Krotkov, and M. T. Mason. Impulsive manipulation. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 120–5, 1995.
- [13] W. H. Huang and M. T. Mason. Limiting cases of impulsive manipulation. Technical Report CMU-RI-TR-96-24, The Robotics Institute, Carnegie Mellon University, September 1996.
- [14] M. T. Mason and J. K. Salisbury, Jr. *Robot Hands and the Mechanics of Manipulation*. The MIT Press, 1985.
- [15] A. A. Rizzi and D. E. Koditschek. Progress in spatial robot juggling. In *IEEE International Conference on Robotics and Automation*, pages 775–780, Nice, France, 1992.
- [16] S. Schaal and C. G. Atkeson. Open loop stable control strategies for robot juggling. In *IEEE International Conference on Robotics and Automation*, pages 3:913–918, Atlanta, GA, 1993.
- [17] W. J. Stronge. Rigid body collisions with friction. *Proceedings of the Royal Society, London*, A431:169–181, 1990.
- [18] J. C. Trinkle, A. O. Farahat, and P. F. Stiller. First-order stability cells of active multi-rigid-body systems. *IEEE Transactions on Robotics and Automation*, 11(4):545–557, Aug. 1995.
- [19] J. C. Trinkle, R. C. Ram, A. O. Farahat, and P. F. Stiller. Dexterous manipulation planning and execution of an enveloped slippery workpiece. In *IEEE International Conference on Robotics and Automation*, pages 2: 442–448, 1993.
- [20] K. Voyerli and E. Eriksen. On the motion of an ice hockey puck. *American Journal of Physics*, 53(12):1149–53, December 1985.
- [21] Y. Wang and M. T. Mason. Two-dimensional rigid-body collisions with friction. *ASME Journal of Applied Mechanics*, 59:635–641, Sept. 1992.
- [22] Y. Yamagata and T. Higuchi. A micropositioning device for precision automatic assembly using impact force of piezoelectric elements. In *IEEE International Conference on Robotics and Automation*, 1995.
- [23] C. Zhu, Y. Aiyama, T. Chawanya, and T. Arai. Releasing manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, page 1, Osaka, Japan, Nov. 1996.
- [24] N. B. Zumel. *A Nonprehensile Method for Reliable Parts Orienting*. PhD thesis, Carnegie Mellon University, Jan. 1997.