

Limiting Cases of Impulsive Manipulation

Wesley H. Huang Matthew T. Mason
CMU-RI-TR-96-24

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213-3890

September 1996

Contents

1	Introduction	1
1.1	Related Work	1
1.2	Assumptions & Notation	2
2	The Limiting Cases in One Dimension	2
2.1	The Limiting Case for Intermittent Tapping	3
2.2	The Limiting Case of Continuous Tapping	4
3	Preliminaries for 2D	4
3.1	The Physics of a Sliding Rotating Object	4
3.2	Trajectory Scaling	5
3.3	Centers of Rotation (CORs)	5
4	The Limiting Case of Intermittent Tapping in 2D	6
4.1	Net COR in the Limit	6
4.2	Range of Possible CORs	7
5	The Limiting Case of Continuous Tapping in 2D	7
5.1	Impact Constraint for Moving Objects	9
5.2	Range of Possible CORs	9
6	Pushing	12
7	Examples	13
8	Discussion	13
8.1	Limiting Case of Intermittent Tapping	13
8.2	Continuous Tapping Limiting Case	14
8.3	Tracking versus Following the Object	14
9	Vibratory Manipulation	15
9.1	How Vibratory Manipulation is like Bouncing a Ball and Robotic Juggling	15
9.2	Analysis of the Bouncing Ball Problem	16
9.2.1	An Example	17
9.2.2	Experimental Results with the Bouncing Ball	19
9.3	Robotic Juggling	19
10	Summary	19
A	Net Force and Torque Load Due to Friction	21
A.1	General Case	21
A.1.1	Pure Translation	21
A.1.2	Pure Rotation	21
A.2	Rotationally Symmetric Objects	21
A.3	Homogeneous Disc	22
A.4	Ring	22

B	Limit Surface and Eigendirections for Rotationally Symmetric Objects	22
B.1	Eigendirection for Pure Translation	23
B.2	Eigendirection for Pure Rotation	24
C	Distance Traveled and Angle Traversed	24
C.1	Pure Translation or Rotation	25
C.2	Level Curves of Distance and Angle	25
C.3	Intersections of Level Curves of Distance and Angle	25
D	Stopping Time	27
E	Net COR as a Function of $\frac{\omega_0}{v_0}$	28
F	Equi-Energy Ellipses	29

List of Figures

1	One dimensional intermittent tapping limiting case	3
2	One dimension continuous tapping limiting case	4
3	COR construction	6
4	Vector field illustration	11
5	Negative impact cone alignment above stable eigendirection	11
6	Negative impact cone alignment below stable eigendirection	11
7	Maximum $\frac{T}{F}$ for pushing a disc	12
8	Reaching points outside the initial impact cone	14
9	Graph of frequency vs. amplitude for vibratory manipulation	17
10	Continuous and Intermittent Vibratory Manipulation	18
11	Points on x_f level curve for comparison	26
12	Angular velocity trajectories cross first	26
13	Translational velocity trajectories cross first	27
14	Trajectory comparison for points of constant $\frac{\omega_0}{v_0}$	27
15	Behavior of x_f on a θ_f level curve	28
16	Equi-energy ellipse and x_f level curves	29

List of Tables

1	Motion constraints for a disc and a ring	13
---	--	----

Abstract

Impulsive manipulation is the use of impulsive forces to manipulate objects. In particular, we are studying how to manipulate a rigid planar slider by tapping. In this paper, we focus on the limiting case — where the number of taps approaches infinity as the energy of each tap approaches zero. This paper develops two definitions of the limiting case: for *intermittent tapping*, where the object comes to rest between taps, and for *continuous tapping*, where the object does not come to rest. For *intermittent tapping*, we find that the motion obeys a simple scaling law, and for *continuous tapping*, we find that there is a limiting case such that the possible motions of the object are identical to those for pushing for rotationally symmetric objects. We also explore analogous forms of *vibratory manipulation*, where we take the striker behavior foremost and examine the resulting object motion.

1 Introduction

Impulsive forces have many uses in manipulation. Striking an object can be used to instantaneously break static friction. A manipulator can apply larger forces by striking than by static methods. By not requiring that a manipulator directly deposit an object at a goal location, impulsive manipulation effectively increases the workspace of a robot. Impulsive manipulation also provides some degree of separation between the manipulator dynamics and object dynamics. One potential application of impulsive manipulation is in the area of micropositioning; the nonprehensile nature of impulsive manipulation and its ability to deal with high static friction forces make it well suited for this type of task.

Our research has focused on manipulating planar objects by tapping. We have thus far focused on rotationally symmetric objects, though certain aspects of this work apply directly to the general case. The idea of micropositioning has led us to examine the limiting case of tapping — letting the number of taps approach infinity as the energy delivered by each tap approaches zero.

We distinguish between two different ways of tapping an object, which we refer to as *intermittent tapping* and *continuous tapping*. In intermittent tapping, the object comes to rest before the next impact occurs. In continuous tapping, we strike the object before it comes to rest, always keeping the object in motion. Though derived from the same dynamics, the limiting cases for these two methods of tapping result in different behaviors of the object. This paper shows that the limiting case for intermittent tapping follows directly from the observation that the dynamics of a sliding object obey a simple scaling law. We find that there is a limiting case of continuous tapping results in the same motion constraints as quasi static pushing for rotationally symmetric objects.

In developing the limiting cases, we focus primarily on the mechanics of the slider. For vibratory manipulation, we take into account the behavior of the striker as well the mechanics of tapping. The striker may execute some fixed periodic trajectory, may be controlled by feedback, or may use some combination of these methods.

We will first define the two limiting cases of tapping in one dimension before proceeding to the two dimensional case. We present the possible motions for pushing and give two specific examples to compare these modes of manipulation. We outline modes of vibratory manipulation and explore the effect of some simple striker behaviors. We conclude with some further discussion about the limiting cases and vibratory manipulation.

1.1 Related Work

There has been little previous work in the area of impulsive manipulation. Higuchi [6] used an electromagnetic coil to deliver an impulsive force in order to do linear micropositioning. More recently, Yamagata and Higuchi [16] have developed a piezoelectric device to deliver impulsive forces to do micropositioning for precision optical components.

The mechanics of a sliding rotating rotationally symmetric object have been studied more extensively. This work includes Voyerli and Eriksen [14] and Goyal *et al.* [5].

We build upon the results of Huang *et al.* [8] which presented an analytical formulation and solution of the mechanics of tapping a rotationally symmetric object. The limiting cases presented in this paper examine these mechanics as the impulse delivered by a tap approaches zero.

Pushing has been studied in the context of robotic manipulation beginning with Mason [10] and continued by a number of researchers. The mechanics of a pushed or sliding

object involve the net force and torque due to frictional forces over the support distribution of the object. Goyal *et al.* ([4] and [5]) have formulated the limit surface, a graphical representation of the relationship between the motion of a sliding object and the net force and torque on that object.

Mirtich and Canny [11] introduced the idea of impulse-based dynamic simulation. Instead of applying different sets of physics to objects based upon their contact state, they attempt to simplify simulation by treating all contact as collisions. Pushing an object, for example, would be modeled as a series of microcollisions between the pusher and the object.

Vibratory manipulation relates to several areas of work, such as robotic juggling (Schaal and Atkeson [12] and Bühler and Koditschek [2]) and the chaotic dynamics of a ball bouncing on a vibrating table (Holmes [7], Bapat *et al.* [1], Tuffiaro and Albano [13], and Wiesensfeld and Tuffiaro [15]). This work will be discussed in greater detail in Section 9.

1.2 Assumptions & Notation

Although we begin each of the limiting cases with development of the general case (any object), the mechanics that we develop in this paper assume a rotationally symmetric object, i.e. the pressure distribution of the object is a function of only the distance from the center of mass (COM). The object has mass M and moment of inertia $I = M\rho^2$, where ρ is the radius of gyration. The object has a translational velocity v and rotational velocity ω . We denote the initial translational and rotational velocities of the object as v_0 and ω_0 respectively. We also assume that a rotationally symmetric object has a circular boundary of radius R .

Objects slide on a homogeneous support surface subject to Coulomb friction with a coefficient of friction of μ . Coulomb friction also acts between the striker/pusher and the object with a coefficient of friction of μ_r .

We use the classical model of impact — an impact results in an instantaneous change in momentum which is dependent upon the relative velocity of the two objects and the coefficient of restitution. The choice between Newton’s and Poisson’s hypothesis does not affect our results. For simplicity, we assume that the striker has a much greater mass than the object.

2 The Limiting Cases in One Dimension

We begin by defining the limiting cases of tapping in one dimension before proceeding to the two dimensional case. The two dimensional limiting cases are conceptually the same as those for one dimension.

Suppose our task is to move an object a distance d . We could use a single tap to give the object enough velocity so that it will slide and come to rest after traveling a distance d . Alternatively, we could use several taps to cover the distance. We could let the object come to rest after each tap, or we could deliver the next tap before the object stops.

This gives rise to a number of different ways we might define a limiting case of tapping. Regardless of the definition we use, we will take the limiting case to be the limiting behavior as the energy delivered by a single tap approaches zero and as the number of taps approaches infinity.




Figure 1: One dimensional intermittent tapping. We want to move the object a given distance using one or more taps, letting the object come to rest after each tap. As the number of taps increases, the average velocity approaches zero and the total time approaches infinity.

2.1 The Limiting Case for Intermittent Tapping

Suppose we are able to move the object a distance d with a single tap. Consider the tap which produces only $\frac{1}{n}$ the initial velocity. The deceleration due to friction will be the same, so the distance covered by this tap will be $\frac{d}{n}$. It will require n such taps (letting the object come to rest after each one) to move the object a distance d . This is what we refer to as intermittent tapping. See Figure 1 for an illustration.

The initial velocity required to move the object a distance $\frac{d}{n}$ is

$$v_0 = \sqrt{\frac{2\mu g d}{n}} \quad (1)$$

and the time required for the object to come to rest is

$$t_{tap} = \sqrt{\frac{2d}{\mu g n}} \quad (2)$$

Note that the total amount of time for n taps is

$$t = \sqrt{\frac{2dn}{\mu g}} \quad (3)$$

and that the average velocity is

$$\bar{v} = \sqrt{\frac{\mu g d}{2n}} \quad (4)$$

Under this definition of the limiting case, as the number of taps n approaches infinity and the initial velocity (and thus the energy) of each tap approaches zero, the total time required to move the object a distance d approaches infinity, and the average velocity approaches zero.




Figure 2: Velocity profiles for possible definitions of the limiting case of continuous tapping for one dimension. In all cases, we tap the object before it has come to rest. In Figure (a), we maintain a constant average velocity, whereas in Figures (b) and (c), the average velocity approaches a limiting value from above and from below respectively.

2.2 The Limiting Case of Continuous Tapping

Although the limiting case for intermittent tapping has useful properties, we would also like a limiting case which converges to a nonzero average velocity. In this sense, the limiting case approaches a quasistatic equilibrium since a constant velocity gives the appearance that the frictional forces are balanced by the “driving forces”.

In order to converge to a nonzero average velocity, we must tap the object before it has come to rest. There are a number of ways in which we might do this. We could maintain the same average velocity as the energy decreased (Figure 2a), or the average velocity might approach a limiting value from above (Figure 2b) or below (Figure 2c).

This approach to the limiting case for continuous tapping departs somewhat from our stated task of moving an object a distance d , which implicitly assumes that the object starts and finishes at rest. Although it is possible to address the starting and stopping issues in this definition, we are much more interested in characterizing the steady state behaviors we can achieve in the limit.

3 Preliminaries for 2D

3.1 The Physics of a Sliding Rotating Object

In two dimensions, an object can translate and rotate. A unique property of rotationally symmetric objects is that they will slide in a straight line (see Appendix A.2). This allows us to treat translational velocity and net force as scalars, rather than vectors.

As discussed in Huang *et al.* [8] (and also in Appendix A.2), the net force and torque on a rotationally symmetric object are functions of only the ratio $\frac{\omega}{v}$. The equations of motion are:

$$-F\left(\frac{\omega}{v}\right) = m\dot{v} \quad (5)$$

$$-T\left(\frac{\omega}{v}\right) = I\dot{\omega} \quad (6)$$

where F and T are the force and torque loads on the object. The initial velocities are v_0 and ω_0 .

During the motion, the object will translate a distance x_f and rotate an angle θ_f

$$x_f = \int_0^{t_f} v dt \quad (7)$$

$$\theta_f = \int_0^{t_f} \omega dt \quad (8)$$

where t_f is the time for the object to come to rest.

3.2 Trajectory Scaling

In many situations, we find it useful to think of this system as a vector field over the space of velocities:

$$(v, \omega) \mapsto \left(\frac{-F(\frac{\omega}{v})}{M}, \frac{-T(\frac{\omega}{v})}{I} \right) \quad (9)$$

Note that this field is dependent only upon the value of $\frac{\omega}{v}$ (alternatively upon the “angle” of the velocities), and not upon the velocity magnitudes. Consequently, the evolution of this system is not dependent upon the amount of energy of the initial conditions but simply the ratio $\frac{\omega}{v}$. Thus, paths in state space can be scaled.

Alternatively, we can show through straightforward use of the chain rule, that if $v(t)$ and $\omega(t)$ are solutions to the equations of motion, then so are $kv(\frac{t}{k})$ and $k\omega(\frac{t}{k})$. If the first pair of trajectories cover a distance x_f and an angle θ_f , then the second cover a distance k^2x_f and an angle $k^2\theta_f$.

3.3 Centers of Rotation (CORs)

We can describe any single motion of a rigid object in terms of the center of rotation (COR), the point about which the motion of the object appears to be a pure rotation. Given the velocities \vec{v} and ω , the COR is located at $\frac{\vec{z} \times \vec{v}}{\omega}$. We can also examine the net COR by considering the initial and final configurations of an object. Pure translation corresponds to a COR at infinity, pure rotation to a COR at the COM.

The COR implicitly encodes the relative amount of translation and rotation of the object, so we will discuss the motion constraints for the limiting cases and for pushing in terms of the set of CORs that can be achieved. Since we will focus on approximating quasistatic behavior, we are not concerned with the magnitude of the velocities — only with their relative magnitudes. In addition to the location of the COR, we must also know the sign of the rotation in order to fully represent the motion. Since we primarily consider rotationally symmetric objects, the set of positive rotation centers is always the same as the set of negative rotation centers.

We will sometimes compare different motions in terms of the ratio $\frac{\omega}{v}$ which is simply the reciprocal of the distance to the COR. When we say that a motion is “below” another motion we mean that the magnitude of $\frac{\omega}{v}$ for the first motion is less than the value of $\frac{\omega}{v}$ for the other. A value of $\frac{\omega}{v}$ corresponds to a line in the state space. Lines corresponding to lower values of $\frac{\omega}{v}$ lie below lines with higher magnitudes of $\frac{\omega}{v}$. Analogously, we might say that this second motion is “above” the first. The corresponding COR for the first motion is further away from the COM than the COR for the second.

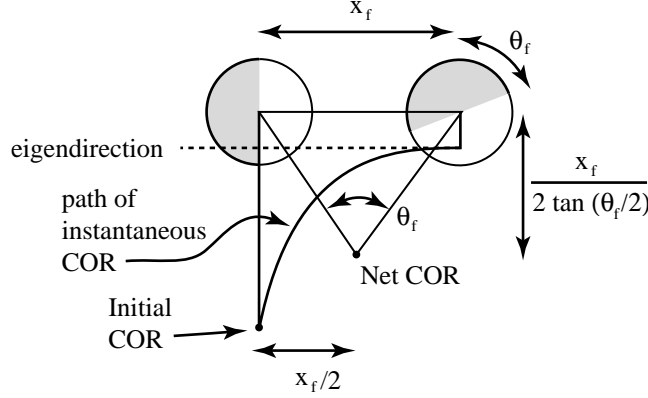


Figure 3: Construction for finding the net COR for a tap where the object moves a distance x_f and turns an angle θ_f . Also illustrated is the path of the instantaneous COR, which asymptotically approaches the COR corresponding to the stable eigendirection.

4 The Limiting Case of Intermittent Tapping in 2D

For intermittent tapping in two dimensions, we consider delivering a sequence of impacts, letting the object come to rest before the next impact. We define the limiting case to be the resulting behavior as the magnitude of the impulse approaches zero. With the observation of Section 3.2 that trajectories can be scaled for rotationally symmetric objects two dimensional tapping, we can directly apply the definition of the limiting case for trajectory scaling we developed in one dimension. Unfortunately, non-rotationally symmetric objects do not obey such a scaling property.

For the remainder of this section, we assume that the object is rotationally symmetric. Since rotationally symmetric objects travel in a straight line, we need only consider one dimension of linear velocity and one dimension of angular velocity.

Suppose we can use a single tap to move an object a distance d and rotate it an angle α . If we tap the object so that the initial velocities are scaled by $\frac{1}{\sqrt{n}}$, then by this scaling property, the object will move a distance $\frac{d}{n}$ and rotate an angle $\frac{\alpha}{n}$.

Like the one dimensional case, the amount of time required for n taps approaches infinity as n approaches infinity. However, this is still a useful result with a number of implications which we discuss in Section 8.1.

Although we know a limiting case exists, we have yet to characterize its motion constraints. We first establish where the net COR is in the limit, and then determine the range of possible CORs.

4.1 Net COR in the Limit

Given the distance traveled and angle turned, we can determine the net COR for the motion. The net COR must lie on the perpendicular bisector of the line connecting the COM at the start and end positions. We also know that the angle formed by the net COR and the initial and final location of the COM must be θ_f (see Figure 3), location of the COR is

$$\frac{\vec{x}_f}{2} + \frac{\hat{z} \times \vec{x}_f}{2 \tan \frac{\theta_f}{2}} \quad (10)$$

As the energy of each tap approaches zero (by scaling the initial conditions), the distance to the COR is

$$\lim_{k \rightarrow 0} \frac{k^2 \vec{x}_f}{2} + \frac{\hat{z} \times k^2 \vec{x}_f}{2 \tan \frac{k^2 \theta_f}{2}} = \lim_{k \rightarrow 0} \frac{\hat{z} \times k^2 \vec{x}_f}{2 \tan \frac{k^2 \theta_f}{2}} \quad (11)$$

Applying l'Hôpital's rule (differentiating both numerator and denominator with respect to k), we get

$$\lim_{k \rightarrow 0} \frac{\hat{z} \times \vec{x}_f}{\theta_f \sec^2 \frac{k^2 \theta_f}{2}} = \frac{\hat{z} \times \vec{x}_f}{\theta_f} \quad (12)$$

Thus, as the energy of a tap approaches zero, the distance to the COR approaches $\frac{x_f}{\theta_f}$ which is a function of $\frac{\omega_0}{v_0}$.

4.2 Range of Possible CORs

Because of the continuity of $x_f(\omega_0, v_0)$ and $\theta_f(\omega_0, v_0)$, the set of limiting CORs will consist of all CORs between the closest and the furthest COR. We can characterize this set by the minimum and maximum values of $\frac{x_f}{\theta_f}$ over the range of possible values of $\frac{\omega_0}{v_0}$.

In Huang [8], we showed that the mechanics of impact limit the range of initial velocities (and thus values of $\frac{\omega_0}{v_0}$ that can be reached). Specifically, we derived the constraint

$$v_0 \geq |\omega_0| \frac{I}{RM} \sqrt{1 + \frac{1}{\mu_r^2}} \quad (13)$$

This equation defines a cone in state space which we refer to as the *impact cone*. This assumes that we can strike the object anywhere on its boundary, so long as it produces a velocity in the proper direction. Each ray in the impact cone corresponds to a different contact point on the object.

It will always be possible to give a rotationally symmetric object with a circular boundary a pure translational velocity, resulting in pure translation of the object. This corresponds to a COR at infinity. In order to characterize the motion constraints for the limiting case of intermittent tapping, we need only find the closest COR, which will correspond to the smallest value of $\frac{x_f}{\theta_f}$.

As shown in Appendix E, the value of $\frac{x_f}{\theta_f}$ as a function of $\frac{\omega_0}{v_0}$ is monotonic decreasing. Therefore, the minimum value of $\frac{x_f}{\theta_f}$ corresponds to the maximum possible value of $\frac{\omega_0}{v_0}$ achievable by an impact, which corresponds to the edge of the impact cone.

The range of COR possible for a given velocity direction can be determined as follows. We pick any initial conditions that lie on the edge of the impact cone (i.e. which satisfy the equality of equation 13). Let $c_{min} = \frac{x_f}{\theta_f}$ for those initial conditions. The set of COR is then

$$\{c \hat{z} \times \hat{v} \mid c \in [c_{min}, \infty)\} \quad (14)$$

5 The Limiting Case of Continuous Tapping in 2D

For continuous tapping, we deliver a sequence of regularly spaced taps to the object such that the object does not come to rest before the next impact. As in the one dimensional case of continuous tapping, we are interested in only the steady state behavior.

We define the limiting case to be the limiting motion as the period of the impacts approaches zero and the magnitude of the impact approaches zero. However, there are many limiting cases, corresponding to the infinitely many ways that the impulse and the impact period can approach zero simultaneously. Many of these do not result in regular motions — if the impulse decreases too quickly relative to the impact period, the object may come to rest before the next tap; if too slowly then object may experience increasing acceleration in the limit.

We are interested in achieving a regular period one behavior, i.e. the velocities of the object (either relative to the world frame or the object frame) must repeatedly follow a path in state space and then be “kicked” back to the original velocities by an impact. Furthermore, this behavior must remain stable as the impulse and impact period approach zero. This will lead to finite non-zero velocities in the limit.

There are other choices to be made when designing this limiting case. We must decide whether to follow or to track the object and what tapping variations to allow in the limit.

By tracking the object as it translates and rotates, we repeatedly tap at the same contact point on the object. By following the object, we deliver the same impulse relative to the COM of the object in the world frame — we tap at a new contact point every time as the object rotates.

In the strictest sense of the limiting case, we allow no variation in the impulse direction (relative to the world frame for following or to the body frame for tracking). We could relax this constraint slightly and allow the direction of the impact to change while keeping the contact point the same (for tracking) or keeping the contact point relative to the COM the same (for following). Finally, we could allow the contact point to change as we approach the limit (though this makes sense primarily for rotationally symmetric objects).

For all possible definitions (e.g. rotationally symmetric or non-rotationally symmetric, tracking or following, and any variation in impulse), so long as they result in some regular periodic motion in state space which remains stable in the limit, we can determine the location of the COR in the limit for a single tap. Relative to the initial position of the COM, the location of the rotation center is:

$$\frac{\vec{x}_f}{2} + \frac{\hat{z} \times \vec{x}_f}{2 \tan \frac{\theta_f}{2}} \quad (15)$$

This is the same relationship developed in Section 4.1 (and illustrated in Figure 3), but for the general case, \vec{x}_f and \vec{v}_0 are not necessarily parallel.

As we decrease the period, \vec{x}_f and θ_f both approach zero:

$$\lim_{t_f \rightarrow 0} \frac{\vec{x}_f}{2} + \frac{\hat{z} \times \vec{x}_f}{2 \tan \frac{\theta_f}{2}} = \lim_{t_f \rightarrow 0} \frac{\hat{z} \times \vec{x}_f}{2 \tan \frac{\theta_f}{2}} \quad (16)$$

Here we use t_f as the time that the object slides before the next impact (which is the period of the impacts). Applying l’Hôpital’s rule, we get

$$\lim_{t_f \rightarrow 0} \frac{\hat{z} \times \frac{d}{dt_f} \vec{x}_f}{\sec^2 \theta_f \frac{d}{dt_f} \theta_f} = \lim_{t_f \rightarrow 0} \frac{\hat{z} \times \frac{d}{dt_f} \int_0^{t_f} \vec{v} dt}{\frac{d}{dt_f} \int_0^{t_f} \omega dt} = \lim_{t_f \rightarrow 0} \frac{\hat{z} \times v(t_f)}{\omega(t_f)} = \frac{\hat{z} \times \vec{v}_0}{\omega_0} \quad (17)$$

Thus, the limiting COR is simply the initial COR.

The remainder of this section is devoted to exploring the least restrictive case for defining a limiting case for continuous tapping. For the remainder of this section, we use a rotationally symmetric object that we follow (instead of track), and we allow the impact point

relative to the COM in the world frame to change in the limit. Since we again consider a rotationally symmetric object, we can use a two dimensional state space (ωv) .

5.1 Impact Constraint for Moving Objects

The approach we will take to developing this limiting case is to determine whether the trajectory starting from some given initial conditions can be stably maintained in the limit. Since we are following the object, we consider velocities relative to the world frame. For regular periodic behavior, we must be able to “kick” the system back to its initial velocities after time t_f , but we cannot change the velocities of an object arbitrarily with a tap.

For the limiting case of intermittent tapping, we used the impact cone (equation 13) to represent the locus of states that we could reach via a single impact. We make the same assumptions here, namely that the object is rotationally symmetric, that we can tap the object at any point on its boundary (so long as the velocity is in the right direction), and that the striker is much more massive than the object. The only difference here is that the object does not start at rest.

For a given impact in the impact cone, we can compute the contact point and striker velocity required in order to impart that impact. Not surprisingly, even if the object has some initial velocity, as long as we recreate the same relative velocities between the striker and the object, we can transfer the same set of impulses to the object. Thus, we can restate the impact cone in terms of the change in velocity

$$\Delta v \geq |\Delta w| \frac{I}{RM} \sqrt{1 + \frac{1}{\mu_r^2}} \quad (18)$$

This implies that we can move the impact cone to any point of the state space; any state within the cone can be reached from the current state through a single impact.

We also make use of the *negative impact cone* — the locus of states from which we can reach the current state with a single impact. The negative impact cone is formed simply by reflecting the impact cone through the current state.

5.2 Range of Possible CORs

The range of CORs for which there is a limiting case is determined by the range of values of $\frac{\omega_0}{v_0}$ for which the path of the system through state space remains inside the negative impact cone (placed at (ω_0, v_0) for some period of time. (Because of the scaling property for rotationally symmetric objects, all initial conditions with the same ratio of $\frac{\omega_0}{v_0}$ will have the same relationship with the negative impact cone.) As long as the path is inside the negative impact cone, we can compute the contact point and relative striker velocity required to send the system back to the initial velocities.

For a given impact period, we choose the contact point and the magnitude of the impulse so that the object returns to its same initial velocities. As we approach the limit, the slope of the line from $(\omega(t_f), v(t_f))$ to (ω_0, v_0) changes, so we must pick a different contact point (relative to the COM of the object) and adjust the impulse accordingly. This is our definition of how the impact period and impulse go to zero for this case.

Determining what paths pass through the negative impact cone requires some knowledge about how the system evolves in state space. The eigendirections of Goyal *et al.* [4] serve to guide our analysis.

For most rotationally symmetric objects, the eigendirections corresponding to pure translation and pure rotation are unstable, and there is a third stable eigendirection somewhere in between. In other words, if the object is given a pure rotational or pure translational velocity, it will continue to rotate or translate (respectively) until it comes to rest. However, for any other initial velocities, the instantaneous COR will approach the COR corresponding to the stable eigendirection as the object comes to rest. All discs and rings fall into this category. This is the general case which we will illustrate.

However, there are objects which do not have a third eigendirection. (See Appendix B for more details.) For these objects, either the eigendirection corresponding to pure translation or that to pure rotation will be stable. The results of this section apply equally to these objects.

There are a number of properties of paths in state space for a sliding rotationally symmetric object. We state them here in terms of the vector field $(\dot{v}, \dot{\omega}) = (-F/m, -T/I)$ over the state space.

- $(\dot{v}, \dot{\omega})$ points towards the origin only at values of $\frac{\omega}{v}$ corresponding to eigendirections.
- at all points below the stable eigendirection, $(\dot{v}, \dot{\omega})$ points above the origin, drawing the system towards the stable eigendirection.
- at all points above the stable eigendirection, $(\dot{v}, \dot{\omega})$ points below the origin, again drawing the system towards the stable eigendirection.
- as $\frac{\omega}{v}$ increases, the orientation of $(\dot{v}, \dot{\omega})$ monotonically increases from π to $\frac{3}{2}\pi$.

See Figure 4 for an illustration.

At $\frac{\omega}{v} = 0$, we have the vector $(-F/M, 0)$ which certainly lies inside the negative impact cone. As we increase $\frac{\omega}{v}$, the vector $(\dot{v}, \dot{\omega})$ will rotate counterclockwise until it aligns with the bottom edge of the cone. At that point,

$$\frac{\dot{v}}{\dot{\omega}} = \frac{M}{I} \frac{R}{\sqrt{1 + \frac{1}{\mu_r^2}}} \quad (19)$$

or

$$\frac{T(\frac{\omega}{v})}{F(\frac{\omega}{v})} = \frac{R}{\sqrt{1 + \frac{1}{\mu_r^2}}} \quad (20)$$

Let $1/c_{min}$ be the value of $\frac{\omega}{v_0}$ which satisfies this equation. In general, there is no closed form solution for $\frac{\omega}{v}$ in equation 20, but it can be solved numerically.

If this alignment occurs above the stable eigendirection (see Figure 5), then the trajectory will curve into the cone as it progresses. Thus, we will always be able to “kick” the system back to its starting point after any amount of time, so the set of possible CORs (for the velocity direction \hat{v}) are

$$\{c\hat{z} \times \hat{v} \mid c \in [c_{min}, \infty]\} \quad (21)$$

If this alignment occurs below the stable eigendirection (see Figure 6), then the trajectory will curve away from the cone as it progresses. For values of $\frac{\omega}{v}$ strictly less than this alignment value, the path of the system will initially lie inside the negative impact cone but may leave after some period of time. Even so, there is still a limiting case for these values. The set of possible CORs (for the velocity direction \hat{v}) is

$$\{c\hat{z} \times \hat{v} \mid c \in (c_{min}, \infty]\} \quad (22)$$

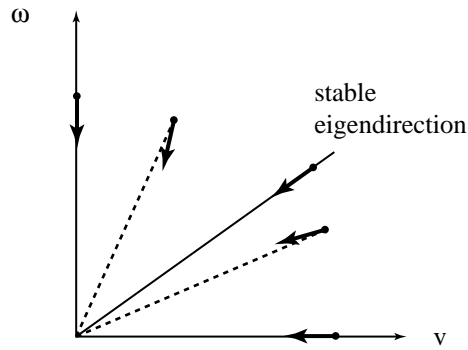


Figure 4: An illustration of a typical vector field for a rotationally symmetric object. The vector field is a function of $\frac{\omega}{v}$, so it is constant along any ray emanating from the origin. The vector points directly towards the origin only at the eigendirection. Typically, there is only one stable eigendirection; all other vectors point towards this eigendirection.

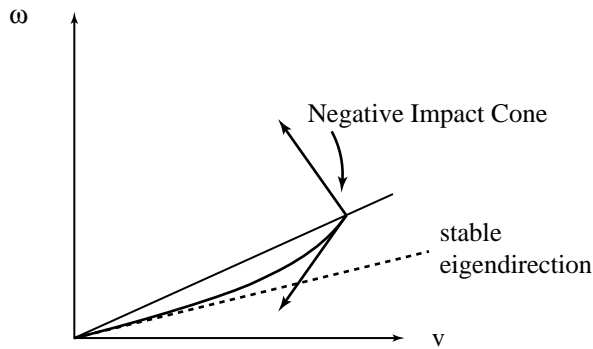


Figure 5: The vector field aligns with the bottom edge of the negative impact cone above the eigendirection, ensuring that the path of the system through state space will remain in the cone as it progresses.

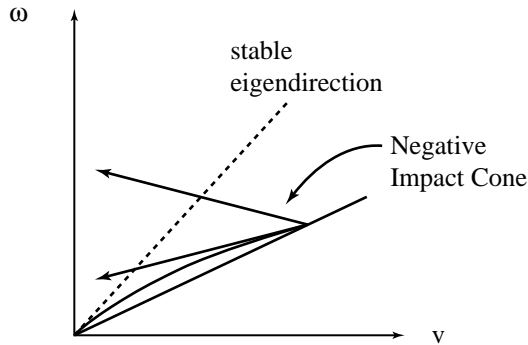


Figure 6: The vector field aligns with the bottom edge of the negative impact cone below the eigendirection. For values of $\frac{\omega}{v_0}$ less than the alignment value, the path of the system will begin inside the negative impact cone but may leave after some period of time.




Figure 7: The maximum value of $\frac{T}{F}$ for pushing a disc is achieved by producing the greatest amount of torque for a given force. Friction cones are drawn here at the contact points which have the maximum lever arm and where the friction cone contains a force in the \hat{x} direction.

6 Pushing

For comparison, we include a discussion of the motions constraints for quasistatic pushing. The pushing force exerted on the object must lie within the friction cone at the contact point. This force and the resulting torque are resisted by frictional forces which are determined by the velocities of the object. Thus, the range of forces and torques that can be exerted on the object indirectly determine the range of velocities and therefore the range of CORs.

Now we consider a rotationally symmetric object under the same assumption as in the mechanics developed for the two limiting cases that we can push anywhere on the object boundary so long as the velocity is in a single direction. We can always exert a force with zero torque on the object; this produces pure translation, corresponding to a COR at infinity. As illustrated in Figure 7, the maximum torque that can be exerted for a given force is determined by the coefficient of friction between the pusher and the object, μ_r .

The net force and torque, $F(\frac{\omega}{v})$ and $T(\frac{\omega}{v})$, are monotonic decreasing and increasing functions of $\frac{\omega}{v}$ respectively. Therefore the maximum torque to force ratio determines the maximum value of $\frac{\omega}{v}$ (the minimum value of $\frac{v}{\omega}$) and therefore the closest possible COR.

The maximum torque that can be generated for a given force is due to a force on the edge of the friction cone. This torque is

$$T = |\vec{r}|F \sin(\tan^{-1} \mu_r) = RF \frac{\mu_r}{\sqrt{1 + \mu_r^2}} \quad (23)$$

we want the value of $1/c_{min} = \frac{\omega}{v}$ that satisfies

$$\frac{T(\frac{\omega}{v})}{F(\frac{\omega}{v})} = \frac{R}{\sqrt{1 + \frac{1}{\mu_r^2}}} \quad (24)$$

The set of rotation centers given the velocity direction \hat{v} is

$$\{c\hat{z} \times \hat{v} \mid c \in [c_{min}, \infty)\} \quad (25)$$

Note that this is the same condition for the minimum COR as in the limiting case of continuous tapping as given in equation 20. With the exception of the single rotation center at $c_{min}\hat{z} \times \hat{v}$, pushing and the limiting case of continuous tapping can maintain the same range of CORs.

	Disc		Ring	
	$\frac{\omega}{v}$	dist. to COR	$\frac{\omega}{v}$	dist. to COR
Eigendirection	30.798	0.032	20.000	0.050
Impact Cone	9.704	0.103	4.851	0.206
Continuous Tapping L.C./Pushing	16.395	0.061	8.951	0.112
Intermittent Tapping L.C.		0.080		0.208
Radius		0.050		0.050

Table 1: Motion constraints for a disc and a ring for both limiting cases (LCs). The distance to the COR given for the impact cone, limiting cases, and pushing is the distance to the closest COR. The distance to the COR corresponding to the stable eigendirection and the radius of the object are given for comparison. quasistatic pushing has the same motion constraints as the limiting case of continuous tapping. These results were numerically generated, and assume that $\mu = \mu_r = 0.25$.

7 Examples

In this section, we present the motion constraints for two different objects, a disc and a ring, for both limiting cases and for pushing. All results were numerically generated. We assumed that both the coefficient of friction between the striker and the object (μ_r) and between the object and the support surface (μ) were 0.25. The results are summarized in Table 1.

This table also includes the size of the impact cone (i.e the maximum value of $\frac{\omega_0}{v_0}$ achievable by impact) which is given by

$$\frac{\omega}{v} \leq \frac{R}{\rho^2 \sqrt{1 + \frac{1}{\mu_r^2}}} \quad (26)$$

and will be a function of the radius of the object. The smaller the radius, the smaller the radius of gyration, and the wider the impact cone will be.

Note that for both examples, the range of CORs possible for the limiting cases is greater than that of the impact cone, i.e. there are states at which we can sustain a limiting case but which we cannot reach directly when starting the object from rest. Also note that the range of possible CORs for both limiting cases in both examples lies below the eigendirection.

8 Discussion

8.1 Limiting Case of Intermittent Tapping

Although the limiting case for intermittent tapping appears to have limited use because the average velocity approaches zero in the limit, it has several useful implications.

First of all, this limiting case would be the simplest and perhaps most practical to implement since the object comes to rest in between taps. This makes it easier to incorporate sensory feedback while executing a series of taps and allows for more exact execution of planned sequence of taps. In contrast, the limiting case of continuous tapping is more dynamic, so it would be difficult to repeatedly strike the object at exactly the same state and “kick” it back precisely to the initial state.

Intermittent tapping also provides some independence between planning a path and executing it because the exact same path that can be executed with a single tap can be

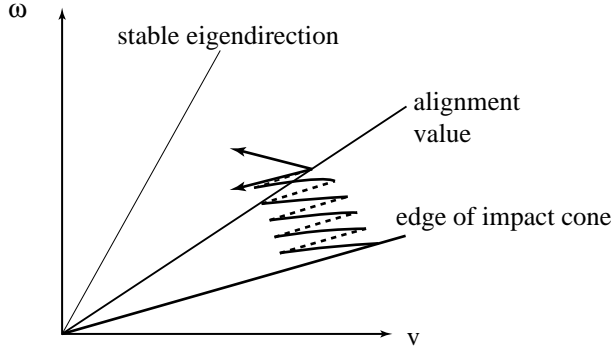


Figure 8: When the initial impact cone lies below the eigendirection, we can reach points in the state space outside the this cone by repeated impacts. Since all points below the eigendirection asymptotically approach the eigendirection, $\frac{\omega}{v}$ increases beyond the impact cone. By repeatedly striking the object, we can maintain the energy of the object and thus reach any point in the state space below the value of $\frac{\omega_0}{v_0}$ corresponding to the minimum COR of the limiting case of continuous tapping. The solid lines show the object sliding; the dotted lines correspond to impacts.

executed by arbitrarily many taps. In fact, since error tends to grow with increased energy, smaller taps yield smaller error, so it may be advantageous to use a sequence of smaller taps, possibly with feedback..

Knowing the range of CORs that are possible for a given object, we could adapt a nonholonomic motion planner to plan complex motions which avoid obstacles, as has been done for pushing by Lynch and Mason [9]. One caveat is that motion about a COR may need to be followed arbitrarily closely in order to avoid obstacles.

8.2 Continuous Tapping Limiting Case

In the examples presented in Section 7, the closest possible COR for the limiting case of continuous tapping was closer than the COR corresponding to the edge of the impact cone, i.e. the limiting case of continuous tapping can be sustained at states which cannot be directly reached from rest.

As illustrated in Figure 8, when the impact cone is below the eigendirection, by repeatedly tapping the object to give it the maximum $\Delta\omega/\Delta v$ possible, we can eventually reach parts of the state space that are outside the initial impact cone.

We can come arbitrarily close to the minimum COR for the limiting case of continuous tapping, but recall that when $\frac{\omega_0}{v_0}$ is below the eigendirection, the trajectory will eventually leave the negative impact cone. If we tap the object again before it leaves the negative impact cone, we will increase the value of $\frac{\omega_0}{v_0}$. However, if it leaves the negative impact cone before it is tapped again, then we will effectively decrease the value of $\frac{\omega_0}{v_0}$.

8.3 Tracking versus Following the Object

In developing the limiting cases, we discussed the choice of tracking versus following the object. For the limiting case of intermittent tapping, this choice does not affect the possible CORs. For the limiting case of continuous tapping, it does.

Regardless of which limiting case is being used, the resulting object behavior is quite different even for the same COR. If tracking the object, the COR is relative to the impulse direction, which rotates with the object, so the object will rotate about a fixed COR. If following the object, the COR is relative to the impulse direction, which translates to follow the object; in this case, the object will appear to be rolling at the same radius as the COR.

9 Vibratory Manipulation

In the limiting cases, we sought to determine what CORs could be sustained in the limit and did so by checking to see whether an impact could produce the required change in velocities. For vibratory manipulation, we will start with the motion of the striker as a given and try to determine the resulting behavior of the object. There are a number of different scenarios we can consider:

- *Feedback control of impulses.* The object is sensed before every impact, and a control law determines an appropriate impulse to drive the object towards some goal configuration. As the object approaches the goal configuration, the impact required becomes smaller and smaller. As the energy of a tap decreases, so does the error associated with that tap; the operation can terminate when the error falls below some threshold.
- *Feedback control of a pusher with superimposed vibration.* The end effector vibrates with some fixed motion as it pushes the slider so that the actual excitation is a high frequency, low amplitude sequence of collisions between the striker and the object. The gross motion of the pusher may be determined by some feedback control law.
- *Open-loop parts transfer.* A sequence of n taps is planned to move an object from start to goal, assuming ideal behavior of the slider and perfect knowledge of all relevant mechanical parameters. The shape of the striker and the details of the striking motion can be designed so that small errors give rise to restoring impulses.
- *Internal hammer platform.* An object is fitted with one or more hammers capable of delivering an impulse to itself (and simultaneously applying a reaction impulse to a reaction mass as in Yamagata and Higuchi [16]). Impacts could be planned with or without feedback.

These behaviors could incorporate either intermittent or continuous tapping. Our interest in vibratory manipulation tends toward the application of the limiting cases of tapping. This has the advantage that as the energy of the impacts decreases, the servo rate for any feedback scheme (active or passive) effectively increases and may thus provide increased stability.

In the remainder of this section, we explore the behavior of an object when the striker motion is a vibration superimposed upon a constant velocity. This is relevant to the behaviors “Feedback control of a pusher with superimposed vibration” and “Open-loop parts transfer” as described above.

9.1 How Vibratory Manipulation is like Bouncing a Ball and Robotic Juggling

We consider in one dimension the problem of a striker moving with some fixed periodic vibration superimposed on a constant velocity. Depending upon the mass and coefficient of

restitution of the object and details of the striker vibration, the object may or may not come to rest before the next impact occurs, giving rise to intermittent and continuous modes of vibratory manipulation.

Under continuous motion, if we view the system relative to a frame moving at constant velocity of the striker, it is equivalent to a ball bouncing on a (fixed) vibrating table. Instead of bouncing a ball up in the air (in a gravity field which produces acceleration g), we strike the object, which is then subject to a frictional force field (which produces an acceleration μg). In this sense, we can consider vibratory manipulation to be planar transverse juggling.

The problem of a bouncing ball on a sinusoidal table has been well studied in the context of analyzing impulsive noise in machinery and in the context of robotic juggling. It is well known that this system has bounded output and can exhibit chaotic behavior.

With a sufficiently low bound, even if chaotic behavior emerges, we can still use this striker motion for vibratory manipulation. However, we may need to be a little more predictable, so one particular point of interest is in the conditions under which the system will exhibit stable periodic behavior.

For application to vibratory manipulation, we can directly apply results from the bouncing ball problem, keeping in mind that the acceleration in our system will be μg instead of g and that we view our system in a frame moving at a constant velocity v . Since we assume continuous motion of the object, its velocity in our moving frame must not dip below $-v$.

In the remainder of this section, we first review results from the bouncing ball problem, show some calculations applying these results to vibratory manipulation, and discuss some experiments in this problem. Finally, we discuss related work in the area of robotic juggling.

9.2 Analysis of the Bouncing Ball Problem

The bouncing ball problem is to determine the behavior of a ball bouncing on a sinusoidally vibrating table. Holmes [7] studied the bouncing ball problem under the assumption that the vibration of the table is small compared to the displacement of the ball. Under this assumption, the velocity of the ball before an impact is the same as its velocity after the previous impact. Bapat *et al.* [1] examined both these approximated mechanics and the exact mechanics. In this paper, we adopt most conventions and notation of Bapat *et al.* but take part of the nondimensionalization from Holmes. We assume that the motion of the table is

$$A \sin \omega t \tag{27}$$

its velocity

$$A\omega \cos \omega t \tag{28}$$

where ω is 2π times the frequency of the table oscillation (not rotational velocity). We denote the velocity of the object just before or after a collision by V .

We first review the results of Holmes [7]. By reasoning about the return map generated by the approximated mechanics, Holmes showed that for a coefficient of restitution less than 1, the trajectory of the ball will be bounded. Given the frequency and amplitude of the oscillation, the velocity of the ball just before an impact (or just after the previous impact) will eventually satisfy

$$|V| < \frac{(1+e)}{(1-e)} A\omega \tag{29}$$




Figure 9: The relationship between the frequency and amplitude for a sinusoidally vibrating table in order to achieve period one, $n = 1$ “bouncing” for vibratory manipulation. Here we have assumed $\epsilon = 0.8$ and $\mu = 0.25$. Increasing the amplitude beyond the range shown here will result in period doubling and will eventually lead to chaotic behavior. Using an amplitude below this range will likely result in the object coming to rest against the striker and being pushed along.

Fixed points of the return map are regions with stable periodic bouncing; we are primarily interested in the simplest mode — what Holmes called period one motions and Bapat *et al.* called $(n,1)$ motions (a collision every n cycles of the table with a collision pattern of length 1). For a given n and ω , stability analysis of the fixed point shows that the amplitude of the table must satisfy

$$\frac{n\pi g}{\omega^2} \frac{(1 - \epsilon)}{(1 + \epsilon)} < A < \frac{g}{\omega^2} \sqrt{n^2 \pi^2 \frac{(1 - \epsilon)^2}{(1 + \epsilon)^2} + 1} \quad (30)$$

The velocity of the ball before and after each impact is given by

$$V = \frac{n\pi g}{\omega} \quad (31)$$

Of course, the value of n is limited by the velocity bound given in equation 29.

Bapat *et al.* [1] found general agreement between the approximated and the exact mechanics (in particular, in the amplitude bounds for stable periodic bouncing) when the coefficient of restitution is 0.8 or greater. At lower values there is some difference.

Bapat *et al.* also studied the regions of attraction for the different stable periodic bouncing modes. Although limited in scope, their results are encouraging; they found that the regions of attraction are larger for the exact mechanics than they are for the approximate mechanics. They also found that the initial velocity of the ball is more influential than the initial phase (when the ball strikes the table) in determining into what mode the system settles.

9.2.1 An Example

In this section, we explore the results of the approximated mechanics as applied to vibratory manipulation. Suppose we take $n = 1$, the coefficient of restitution $\epsilon = 0.8$, and the




Figure 10: Illustrations of an object “bouncing” on a sinusoidally vibrating striker which moves at a constant average velocity v . Here we show the path of the striker and the object in the reference frame moving at velocity v . In Figure (a), the object is in continuous motion. In Figure (b), the object comes to rest before the next impact occurs. In the moving frame, the maximum possible “downward” velocity of the object is $-v$.

coefficient of friction $\mu = 0.25$. Recall that the acceleration field for vibratory manipulation will have magnitude μg (instead of g as in the bouncing ball problem). From equation 30, we find that the nondimensionalized amplitude must satisfy

$$\frac{0.856}{\omega^2} < A < \frac{2.597}{\omega^2} \quad (32)$$

This region is graphed in Figure 9. The corresponding (actual) velocity of the object (before and after each impact) is

$$v_O = \frac{\pi \mu g}{\omega} \quad (33)$$

Note that the velocity of the object v_O (relative to the striker) is only a function of the frequency of the oscillation. For a fixed ω , changing the amplitude of the oscillation (within the range shown in Figure 9) changes the phase at which impact occurs.

The constant velocity component of the striker must be greater than the maximum object velocity ($v > v_O$) in order to ensure that the object does not come to rest before it is struck again. In fact, some tolerance will be required since the mechanics used here are approximate.

Under continuous motion, an object being “pushed” by a sinusoidally vibrating striker moving at a constant average velocity v might appear as shown in Figure 10a (in the frame of reference moving at velocity v).

Under intermittent motion, when the object comes to rest, it will appear to stop “falling” along a parabola and start moving at a constant velocity $-v$ in the moving frame, as illustrated in Figure 10b. Although not nearly as amenable to analysis as continuous motion, we can achieve stable periodic impacts with intermittent motion of the object.

9.2.2 Experimental Results with the Bouncing Ball

Tufillaro and Albano [13] describe an undergraduate physics experiment in chaotic dynamics using the bouncing ball problem. Using fairly simple apparatus (consisting of a speaker, a concave lens, and a ball bearing), they were able to observe period doubling and chaotic motion in the system as the amplitude of the vibration was increased.

Wiesenfeld and Tufillaro [15] describe analysis and experiments in suppressing period doubling in the bouncing ball system by driving the table with two closely spaced frequencies. Such a technique could be useful to stabilize a vibratory manipulation system.

9.3 Robotic Juggling

A number of researchers have done work in the area of robotic juggling which closely relates to the continuous mode of vibratory manipulation. This work centers on studying the stability of such systems and writing controllers to control them.

Schaal and Atkeson [12] have implemented paddle juggling using a “trampoline-like” racket and a ping-pong ball. They observe that negative acceleration of the paddle trajectory at impact seems to provide stability. They calculate the basin of attraction for periodic juggling to be 0.257 the area of the viable space of initial conditions. However, for a number of reasons (including mechanical variability and air resistance) the basin of attraction was significantly larger. In fact, they were unable to avoid getting a periodic (mostly period one) juggling pattern.

Bühler and Koditschek [2] demonstrated a system that would juggle (in the plane) a puck using a bat. The system accomplishes this task through closed loop feedback utilizing a mirror law, i.e. controlling the bat to reflect (nonlinearly) the position of the puck.

Using this method for vibratory manipulation is appealing because a bat would not deliver any rotational velocity to an object with a circular boundary. Thus, two dimensional positioning could be accomplished with the one dimensional mechanics. However, implementation of the mirror law could require sensing the object which would be difficult under these conditions.

10 Summary

In this paper, we have developed the idea of limiting cases for impulsive manipulation, i.e. the mechanics of tapping an object as the number of taps increases and the energy of each tap approaches zero. We began with the observation that the dynamics of a sliding rotationally symmetric object allow trajectories to be scaled, and this led us to the limiting case for intermittent tapping. In this limiting case, the object comes to rest after each tap. As the energy of each tap decreases, however, the average velocity approaches zero. From the intermittent tapping limiting case, we learned that a motion can be executed by any number of taps, providing a degree of separation between planning and execution.

The limiting case of continuous tapping, where we strike the object again before it comes to rest, converges to nonzero average velocities. Essentially, we repeatedly follow a path in state space, letting the system evolve and then “kicking” it back to the starting point. We found that we can define a limiting case of continuous tapping that yields the same motion constraints as quasistatic pushing.

In vibratory manipulation, we considered the behavior of the striker as well as the mechanics of tapping the object. In one dimension, we consider a striker executing a fixed

sinusoid movement (superimposed upon a constant velocity). This system is equivalent to robotic juggling or bouncing a ball on a (fixed) sinusoidally vibrating table — the latter of which can lead to bounded chaotic behavior.

A Net Force and Torque Load Due to Friction

A.1 General Case

We affix a coordinate frame to the object's COM and assume that the object has some translational velocity v along the x axis and some rotational velocity ω measured about the COM. The net velocity at a point \vec{r} is

$$\vec{u} = v\hat{x} + \omega \times \vec{r} \quad (34)$$

Note that the COR corresponding to this motion is located at $(0, \frac{v}{\omega})$.

Under Coulomb friction, the net force and torque load on this object are

$$\vec{F} = \int d\vec{f} = \mu g \int \hat{u} dm \quad (35)$$

$$T = \int \vec{r} \times d\vec{f} = \mu g \int (r_x \hat{u}_y - r_y \hat{u}_x) dm \quad (36)$$

A.1.1 Pure Translation

For pure translation ($\omega = 0$), we have $\hat{u} = \hat{x}$, so

$$\vec{F} = \hat{x} \int \mu g dm = \mu g M \hat{x} \quad (37)$$

$$T = -\mu g \int r_y dm = 0 \quad (38)$$

A.1.2 Pure Rotation

For pure rotation ($v = 0$), we have $\hat{u} = \omega \hat{z} \times \vec{r}$, so

$$\vec{F} = \mu g \omega \int \hat{z} \times \vec{r} dm = 0 \quad (39)$$

$$T = \mu g \int r dm \quad (40)$$

A.2 Rotationally Symmetric Objects

Rotationally symmetric objects are those objects whose pressure distribution is a function of only the radius r . Assuming that the translational velocity is parallel to the x axis, we can express the direction of the velocity of a point in terms of polar coordinates:

$$\hat{u} = \frac{(v - \omega r \sin \theta)\hat{x} + (\omega r \cos \theta)\hat{y}}{\sqrt{v^2 - 2\omega v r \sin \theta + \omega^2 r^2}} = \frac{(1 - \frac{\omega}{v} r \sin \theta)\hat{x} + (\frac{\omega}{v} r \cos \theta)\hat{y}}{\sqrt{1 - 2\frac{\omega}{v} r \sin \theta + (\frac{\omega}{v})^2 r^2}} \quad (41)$$

Note that this is a function of only the ratio of the velocities, $\frac{\omega}{v}$. Thus, the net force and torque loads for a rotationally symmetric object will only be a function of $\frac{\omega}{v}$.

In addition, rotationally symmetric objects have the special property that:

$$F_y = \mu g \int \int_0^{2\pi} \frac{\omega r^2 \cos \theta}{\sqrt{v^2 - 2\omega v r \sin \theta + \omega^2 r^2}} d\theta dr = 0 \quad (42)$$

Thus, rotationally symmetric objects will always travel in the same direction as the velocity. (Note that this is not true in the general case.)

This also means that we can treat the force load as a scalar

$$F = \mu g \int \hat{u}_x dm \quad (43)$$

A.3 Homogeneous Disc

A uniform disc of radius R has a moment of inertia of

$$I = \int r^2 dm = \frac{M}{\pi R^2} \int_0^{2\pi} \int_0^R r^3 dr d\theta = \frac{1}{2} M R^2 \quad (44)$$

Thus, the radius of gyration, ρ is

$$\rho = \frac{R}{\sqrt{2}} \quad (45)$$

The net force and torque load due to friction are:

$$F = \mu g \frac{M}{\pi R^2} \int \frac{v - \omega r \sin \theta}{\sqrt{v^2 - 2\omega r \sin \theta + \omega^2 r^2}} r dr d\theta \quad (46)$$

$$T = \mu g \frac{M}{\pi R^2} \int \frac{\omega r^2 - v r \sin \theta}{\sqrt{v^2 - 2\omega r \sin \theta + \omega^2 r^2}} r dr d\theta \quad (47)$$

For pure rotation, the torque load is

$$T = \mu g \int r dm = \frac{2}{3} \mu g R M \quad (48)$$

A.4 Ring

A ring of radius R with zero width has a moment of inertia $I = M R^2$ and thus a radius of gyration $\rho = R$. The net force and torque load due to friction are

$$F = \mu g \frac{M}{2\pi} \int_0^{2\pi} \frac{v - \omega R \sin \theta}{\sqrt{v^2 - 2\omega v R \sin \theta + \omega^2 R^2}} R d\theta \quad (49)$$

$$T = \mu g \frac{M}{2\pi} \int_0^{2\pi} \frac{\omega R^2 - v R \sin \theta}{\sqrt{v^2 - 2\omega v R \sin \theta + \omega^2 R^2}} R d\theta \quad (50)$$

B Limit Surface and Eigendirections for Rotationally Symmetric Objects

Rotationally symmetric objects always have rotationally symmetric limit surfaces, so instead of considering the entire limit surface, we can simply use its generating curve. This curve can be generated by the functions $F(\frac{\omega}{v})$ and $T(\frac{\omega}{v})$.

Eigendirections are points on the limit surface where the vector (v, ω) is parallel to the vector $(\dot{v}, \dot{\omega}) = (F(\frac{\omega}{v})/M, T(\frac{\omega}{v})/I)$. Goyal *et al.* ([4] and [5]) showed that all objects must come to rest at an eigendirection.

There are always eigendirections corresponding to pure rotation and pure translation, but these are generally unstable, i.e. any deviation from pure translation or pure translation will cause the system to move away from these eigendirections.

There is generally a third stable eigendirection corresponding to a COR at a finite nonzero distance from the object. All discs and rings have this third stable eigendirection.

However, this is not always the case; there may be no third eigendirection, in which case the eigendirection corresponding to either pure translation or rotation would be stable. Goyal *et al.* [5] show that by changing the radius of gyration of an disc or ring without changing its support distribution, the stable eigendirection can be arbitrarily assigned.

In this section, we explore the conditions under which the eigendirections corresponding to pure translation and pure rotation are stable. If both these eigendirections are unstable, then there must be a third stable eigendirection. We follow the mathematics of appendices A.9, A.12, and A.13 of Goyal [3]. We adopt some of their notation for this section.

The stability of an eigendirection can be determined by comparing the meridional radius of curvature in load space, ρ_c with the magnitude of the generalized force vector, $|P|$. If $\rho_c < |P|$, then the eigendirection is unstable.

B.1 Eigendirection for Pure Translation

In the case of pure translation, we have

$$|P| = 1 \quad (51)$$

$$\rho_c = \lim_{s \rightarrow 0} \frac{\sqrt{(\frac{\delta M}{\delta s})^2 + (\frac{\delta F}{\delta s})^2}}{\rho} \quad (52)$$

where $s = \frac{\omega}{v}$, ρ is the radius of gyration, and M and F are the normalized torque and force loads on the object

$$M(s) = \frac{1}{\mu g M \rho} T\left(\frac{\omega}{v}\right) \quad (53)$$

$$F(s) = \frac{1}{\mu g M} F\left(\frac{\omega}{v}\right) \quad (54)$$

Taking the limits,

$$\lim_{s \rightarrow 0} \frac{1}{\mu g M} \frac{\delta F}{\delta s} = 0 \quad (55)$$

$$\lim_{s \rightarrow 0} \frac{1}{\mu g M \rho} \frac{\delta T}{\delta s} = \frac{1}{\rho M} \int r^2 \cos^2 \theta \, dm \quad (56)$$

So, we want to know whether

$$\begin{aligned} \rho_c &< |P| \\ \frac{1}{\rho^2 M} \int r^2 \cos^2 \theta \, dm &< 1 \end{aligned} \quad (57)$$

Substituting for the radius of gyration, this is equivalent to

$$\frac{1}{M} \int r^2 \cos^2 \theta \, dm < \frac{1}{M} \int r^2 \, dm \quad (58)$$

For (rotationally symmetric) objects whose pressure distribution is the “same” as their mass distribution, this condition will be true, so the eigendirection corresponding to translation will be unstable.

B.2 Eigendirection for Pure Rotation

For pure rotation, we have

$$|P| = \frac{1}{\rho M} \int r \, dm \quad (59)$$

$$\rho_c = \lim_{r_c \rightarrow 0} \frac{\sqrt{\left(\frac{\delta M}{\delta r_c}\right)^2 + \left(\frac{\delta F}{\delta r_c}\right)^2}}{\frac{1}{\rho}} \quad (60)$$

where $r_c = \frac{v}{\omega}$. Taking the limits, we have

$$\lim_{r_c \rightarrow 0} \frac{1}{\mu g M} \frac{\delta F}{\delta r_c} = \frac{1}{M} \int \frac{\cos^2 \theta}{r} \, dm \quad (61)$$

$$\lim_{r_c \rightarrow 0} \frac{1}{\mu g M \rho} \frac{\delta T}{\delta r_c} = 0 \quad (62)$$

The condition for an unstable eigendirection becomes

$$\frac{\rho}{M} \int \frac{\cos^2 \theta}{r} \, dm < \frac{1}{\rho M} \int r \, dm \quad (63)$$

Again substituting for the radius of gyration, ρ , we get

$$\frac{1}{M} \int \frac{\cos^2 \theta}{r} \, dm < \frac{\frac{1}{M} \int r \, dm}{\frac{1}{M} \int r^2 \, dm} \quad (64)$$

It is not clear that this condition is satisfied for all rotationally symmetric objects. Goyal *et al.* [5] show that this condition is satisfied for any homogeneous disk or ring, so the eigendirection corresponding to pure rotation is unstable for these objects.

C Distance Traveled and Angle Traversed

The total distance traveled and angle traversed are given by

$$x_f = \int_0^{t_f} v \, dt \quad (65)$$

$$\theta_f = \int_0^{t_f} \omega \, dt \quad (66)$$

where v and ω are solutions to the equations of motion

$$-F\left(\frac{\omega}{v}\right) = m\dot{v} \quad (67)$$

$$-T\left(\frac{\omega}{v}\right) = I\dot{\omega} \quad (68)$$

subject to the initial velocities v_0 and ω_0 .

We first address the special cases of pure translation or rotation before discussing properties of the level curves of x_f and θ_f over the state space.

C.1 Pure Translation or Rotation

If the object is given a pure translational velocity, there will be no rotation, and we can treat the object as a point mass. Given an initial velocity v_0 , it will slide a distance:

$$x_f = \frac{\frac{1}{2}Mv_0^2}{\mu Mg} = \frac{v_0^2}{2\mu g} \quad (69)$$

and this will take time

$$t_f = \frac{v_0}{\mu g} \quad (70)$$

If the object is given a pure rotational velocity, there will be no translation. With an initial velocity ω_0 , it will rotate an angle:

$$\theta_f = \frac{\frac{1}{2}I\omega_0^2}{\mu g \int r dm} = \frac{M\rho^2\omega_0^2}{\mu g \int r dm} \quad (71)$$

which will take time

$$t_f = \frac{\omega_0\rho^2 M}{\mu g \int r dm} \quad (72)$$

C.2 Level Curves of Distance and Angle

In previous work (Huang *et al.* [8]) we examined the total distance traveled and total angle traversed (x_f and θ_f) as functions of the initial conditions ω_0 and v_0 . These functions describe a surface over the space of initial conditions.

We showed that if you travel parallel to the positive v_0 or ω_0 axis on either function, you climb up the surfaces. In other words, increasing the initial translational velocity v_0 increases both the total distance traveled x_f and the total angle turned θ_f . Likewise, increasing ω_0 increases both x_f and θ_f as well.

The level curves of x_f and θ_f , viewed in the $v\omega$ plane, are monotonic decreasing — the level curves of x_f start at infinity and monotonically decrease until they reach the v_0 axis, and the level curves of θ_f start at some point on the ω_0 axis and monotonically decrease, approaching the v_0 axis at infinity.

C.3 Intersections of Level Curves of Distance and Angle

We can also show that θ_f must increase along an x_f level curve and vice versa.

Consider two different points on a level curve of x_f . (See Figure 11.) Since the level curve is monotonic decreasing, we can assume

$$v_{0_1} > v_{0_2} \quad (73)$$

$$\omega_{0_1} < \omega_{0_2} \quad (74)$$

therefore $\frac{\omega_{0_1}}{v_{0_1}} < \frac{\omega_{0_2}}{v_{0_2}}$. As long as $\frac{\omega_1}{v_1} < \frac{\omega_2}{v_2}$, $F_1 > F_2$ and $T_1 < T_2$. This means that the gap between the velocities will close until the velocity trajectories cross or until the objects come to rest.




Figure 11: We compare the velocity trajectories for two different points on a x_f level curve to show that a level curve of x_f and a level curve of θ_f can intersect only once and to show that θ_f increases along a x_f level curve (and vice versa).




Figure 12: The angular velocity trajectories cross first. Since the translational and rotational velocities for a given trajectory must go to zero at the same time, $v_1(t) > v_2(t)$, and therefore $x_{f_1} < x_{f_2}$.

The angular and translation velocity trajectories cannot cross simultaneously because of uniqueness. If the translational velocities cross first, we will have the condition

$$v_1 \leq v_2 \quad (75)$$

$$\omega_1 \leq \omega_2 \quad (76)$$

and if the rotational velocities cross first, we will have the condition

$$v_1 \geq v_2 \quad (77)$$

$$\omega_1 \geq \omega_2 \quad (78)$$

As shown in Huang *et al.* [8], once we have either condition then that condition will be maintained from that point onward.

If ω_1 and ω_2 cross first (Figure 12), then $v_1 \geq v_2$ for the entire trajectory, so $x_{f_1} > x_{f_2}$ which violates our assumption that the initial conditions were on a level curve of x_f .

Therefore, v_1 and v_2 must cross first (Figure 13), which results in the condition $\omega_1 \leq \omega_2$ for the entire trajectory, so $\theta_{f_1} < \theta_{f_2}$.

Similar reasoning can be applied to level curves of θ_f . We conclude that

- θ_f increases along x_f level curves,
- x_f increases along θ_f level curves, and
- any given level curve of x_f and any given level curve of θ_f cross exactly once.




Figure 13: The translational velocity trajectories cross first, so $\omega_1(t) < \omega_2(t)$, and therefore $\theta_{f1} < \theta_{f2}$.




Figure 14: We compare the velocity trajectories for one point (point 1) in the state space with those for a number of different points (points A-E) from a line of constant $\frac{\omega_0}{v_0}$.

D Stopping Time

We also consider the time which it takes the object to come to rest, t_f , as a function of the initial conditions. In particular, we will look at the level curves of this function, i.e. the locus of initial conditions which take the same amount of time for the object to come to rest.

First, however, we compare the trajectories of several points in the space of initial conditions as shown in Figure 14. Point 1 can be chosen arbitrarily and is the same for all comparisons. We move point 2 along a line of constant $\frac{\omega_0}{v_0}$, starting with point A and ending with point E. Since all choices for point 2 have the same $\frac{\omega_0}{v_0}$, the velocity trajectories are all scaled versions of each other.

Point A has the same initial translational velocity as point 1 and the highest energy of all the points we will compare. As the energy decreases, the initial velocities and the stopping time will be scaled. Point B lies on the same x_f level curve as point 1. Point C lies somewhere between points B and D. Point D lies on the θ_f level curve. Point E has the same initial rotational velocity as point 1. Velocity profiles for these points (based on the same reasoning used in Section C.3) are also illustrated in Figure 14.

At point A, the velocities are greater than those for point 1. As we decrease the energy of point 2, the translational velocity profiles intersect, but we are still above the x_f level curve. Then we reach the x_f level curve at point B, where the net area between the two




Figure 15: We consider a level curve of θ_f to show that $\frac{x_f}{\theta_f}$ is a monotonic decreasing function of $\frac{\omega_0}{v_0}$. Since x_f increases along a θ_f level curve, it decreases as we travel toward the ω_0 axis, resulting in a decreasing value of $\frac{x_f}{\theta_f}$ along the level curve. Since the θ_f level curve can be parameterized as a function of $\frac{\omega_0}{v_0}$ and since $\frac{x_f}{\theta_f}$ is constant along a line of constant $\frac{\omega_0}{v_0}$, we conclude that $\frac{x_f}{\theta_f}$ is a monotonic decreasing function of $\frac{\omega_0}{v_0}$.

translational velocity profiles is zero. Continuing to decrease the energy, we pass below the x_f level curve; the net area between the translational velocity profiles for point 1 and point 2 is now positive.

Finally, we reach point C where the velocity trajectories of point 1 and point 2 end at the same time. As we continue to reduce the energy of point 2, the rotational velocity profiles intersect. At point D, we reach the θ_f level curve where the net area between the rotational velocity profiles is zero. As we decrease the energy of point 2 further, we pass below the θ_f level curve, and the net area between the rotational velocity profiles of point 1 and point 2 is positive. Eventually, we reach point E, where the initial rotational velocity profiles are the same.

From this comparison, we can draw several conclusions

- t_f increases along an x_f level curve,
- t_f increases along a θ_f level curve,
- x_f increases along a t_f level curve,
- θ_f decreases along a t_f level curve, and
- at any point in the state space, the t_f level curve lies above the θ_f level curve and below the x_f level curve; because of the monotonicity of the x_f and θ_f level curves, we can conclude that the t_f level curve is also monotonic decreasing.

In addition, from observations about pure translation and pure rotation (Section C.1), we know that t_f level curves begin on the ω_0 axis and end on the v_0 axis.

E Net COR as a Function of $\frac{\omega_0}{v_0}$

We would like to know, for the intermittent tapping limiting case, the behavior of $\frac{x_f}{\theta_f}$ as a function of $\frac{\omega_0}{v_0}$.




Figure 16: Possible relationship between level curves of x_f and an equi-energy ellipse.

Consider a level curve of θ_f . As shown in Section D, x_f monotonically increases along a θ_f level curve. Therefore, the ratio $\frac{x_f}{\theta_f}$ also monotonically increases. Since a θ_f level curve can be parameterized as a function of $\frac{\omega_0}{v_0}$ and since the ratio $\frac{x_f}{\theta_f}$ remains constant along a line of constant $\frac{\omega_0}{v_0}$ (because of trajectory scaling), we can conclude that $\frac{x_f}{\theta_f}$ monotonically decreases as a function of $\frac{\omega_0}{v_0}$, regardless of the energy level. See Figure 15.

F Equi-Energy Ellipses

Curves of constant energy $E_0 = \frac{1}{2}Mv_0^2 + \frac{1}{2}I\omega_0^2$ in the space of initial conditions form equi-energy ellipses.

We know that x_f and θ_f do not monotonically increase or decrease along an equi-energy ellipse. Consider the x_f level curve at the intersection of an equi-energy ellipse and the v_0 axis. The normal to the equi-energy ellipse at this point is parallel to the ω_0 axis, but we know from monotonicity that x_f increases in this direction. It decreases along the negative v_0 axis, so the level curve must initially enter the ellipse. Since the level curve continues to infinity, it must leave the ellipse at some other point. The same argument can be made for θ_f level curves. Thus, x_f and θ_f level curves can intersect an equi-energy ellipse more than once. See Figure 16 for an illustration of one possible relationship between the equi-energy ellipse and x_f level curves.

References

- [1] C. N. Bapat, S. Sankar, and N. Popplewell. Repeated impacts on a sinusoidally vibrating table reappraised. *Journal of Sound and Vibration*, 108(1):99–115, 1986.
- [2] 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.
- [3] Suresh Goyal. *Planar Sliding of a Rigid Body With Dry Friction: Limit Surfaces and Dynamics of Motion*. PhD thesis, Cornell University, Dept. of Mechanical Engineering, 1989.
- [4] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. Planar sliding with dry friction. Part 1. Limit surface and moment function. *Wear*, 143:307–330, 1991.
- [5] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. Planar sliding with dry friction. Part 2. Dynamics of motion. *Wear*, 143:331–352, 1991.
- [6] T. Higuchi. Application of electromagnetic impulsive force to precise positioning tools in robot systems. In *International Symposium on Robotics Research*, pages 281–285. Cambridge, Mass: MIT Press, 1985.
- [7] P. J. Holmes. The dynamics of repeated impacts with a sinusoidally vibrating table. *Journal of Sound and Vibration*, 84(2):173–189, 1982.
- [8] W. Huang, E. P. Krotkov, and M. T. Mason. Impulsive manipulation. In *IEEE International Conference on Robotics and Automation*, 1995.
- [9] Kevin M. Lynch and Matthew T. Mason. Controllability of pushing. In *IEEE International Conference on Robotics and Automation*, Nagoya, Japan, 1995.
- [10] Matthew T. Mason. Mechanics and planning of manipulator pushing operations. *International Journal of Robotics Research*, 5(3):53–71, Fall 1986.
- [11] Brian Mirtich and John Canny. Impulse-based dynamic simulation. In K. Goldberg, D. Halperin, J-C. Latombe, and R. Wilson, editors, *Algorithmic Foundations of Robotics*, pages 407–418. A. K. Peters, Boston, MA, 1995.
- [12] 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.
- [13] N. B. Tufillaro and A. M. Albano. Chaotic dynamics of a bouncing ball. *American Journal of Physics*, 54(10):939–944, October 1986.
- [14] K. Voyerli and E. Eriksen. On the motion of an ice hockey puck. *American Journal of Physics*, 53(12):1149–53, December 1985.
- [15] Kurt Wiesenfeld and Nicholas B. Tufillaro. Suppression of period doubling in the dynamics of a bouncing ball. *Physica D*, 26(1–3):321–335, May-June 1987.
- [16] 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.