

Manipulation of Pose Distributions

Mark Moll

mmoll@cs.cmu.edu

Michael A. Erdmann

me@cs.cmu.edu

Department of Computer Science, Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

For assembly tasks parts often have to be oriented before they can be put in an assembly. The results presented in this paper are a component of the automated design of parts orienting devices. The focus is on orienting parts with minimal sensing and manipulation. We present a new approach to parts orienting through the manipulation of pose distributions. Through dynamic simulation we can determine the pose distribution for an object being dropped from an arbitrary height on an arbitrary surface. By varying the drop height and the shape of the support surface we can find the initial conditions that will result in a pose distribution with minimal entropy. We are trying to uniquely orient a part with high probability just by varying the initial conditions. We will derive a condition on the pose and velocity of a simple planar object in contact with a sloped surface that will allow us to quickly determine the final resting configuration of the object. This condition can then be used to quickly compute the pose distribution. We also present simulation and experimental results that show how dynamic simulation can be used to find optimal shapes and drop heights for a given part.

1 Introduction

In our research we are trying to develop strategies to orient three-dimensional parts with minimal sensing and manipulation. That is, we would like to bring a part from an unknown position and orientation to a known orientation (but possibly unknown position) with minimal means. In general, it is very difficult to orient a part completely without sensors. Doing so requires accounting for dynamic effects, either by building special shapes that exploit the dynamic effects (Hitakawa, 1988) or by assuming quasistatic motions (Berretty, 2000). In this paper we investigate the design of shapes that can be used to probabilistically orient three-dimensional parts. It is sufficient if a particular orienting strategy can bring a part into one particular orientation with high probability. The sensing is then reduced to a binary decision: a sensor only has to detect whether the part is in the right orientation or not. If not, the part is fed back to the parts orienting device. Assuming the orienting strategy succeeds with high probability, on average it takes just a few tries to orient the part. An alternative view of this type of manipulation is to consider it as manipulation of the

	Stable Poses						
	1	2	3	4	5	6	Entropy
quasistatic approximation	0.20	0.13	0.16	0.21	0.14	0.16	1.78
full dynamics, flat surface, drop height is $h = 0$	0.18	0.16	0.14	0.34	0.05	0.13	1.66
full dynamics, bowl shape is $y = 0.24x^2$, $h = 0.28$, initial hor. pos. $x_0 = -0.41$	0.24	0.03	0.03	0.50	0.08	0.15	1.35

Table 1: Probability distribution function of stable poses for two surfaces. The initial velocity is zero and the initial orientation is uniformly random. (The probabilities in the third row do not add up to 1 due to round-off errors.)

pose distribution. The goal then is to find the pose distribution with minimal entropy, thereby maximally reducing uncertainty.

1.1 Example

In this paper we will discuss the use of dynamic simulation for the design of support surfaces that reduce the uncertainty of a part's resting configuration. As the support surface is changed, the probability distribution function (pdf) of resting configurations will change as well. The pdf will also vary with the initial drop position above the surface. The following figure and paragraph illustrate the basic idea:

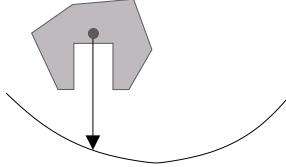


Figure 1: A part with an initially unknown orientation is dropped on a surface.

A part with an initially unknown orientation is released from a certain height and relative horizontal position with respect to a bowl-shaped surface. The only forces acting on the part are gravity and friction. We assume the bowl does not move. We can compute the final resting configuration for all possible initial orientations. This will give us the pdf of stable poses. The goal is to find the drop height, relative position and bowl shape that will maximally reduce uncertainty. In this paper we assume for simplicity that the initially unknown orientation is uniformly random, but our approach also works for different prior distributions.

Table 1 shows three different pose distributions. Each stable pose corresponds to a set of contact points (marked by the black dots in the table). For an arbitrarily curved support surface the stable poses do not necessarily correspond to edges of the convex hull of the part. We therefore define a stable pose as a set of contact points. This means that any two poses with the same set of contact points are considered to be the same as far as the pose distribution is concerned. In our example the support surface is a parabola $y = ax^2$ with parameter a . Other parameters are the drop height, h , and the initial horizontal position of the drop location, x_0 . We limit the surface to

parabolas for illustrative purposes only; in general we would use a larger class of possible shapes (see section 4.1).

The first row in the table shows the pdf assuming quasistatic dynamics. In this case the surface is flat and the part is released in contact with the surface. The second row shows how the pdf changes if we model the dynamics. The initial conditions are the same as for the quasistatic case, yet the pdf is significantly different. The third row shows the pdf for the optimized values for a , h and x_0 .

The objective function over which we optimize is the entropy of the pose distribution. If p_1, \dots, p_n are the probabilities of the n stable poses, then the entropy is $-\sum_{i=1}^n p_i \log p_i$. This function has two properties that make it a good objective function: it reaches its global minimum whenever one of the p_i is 1, and its maximum for a uniform distribution. By searching the parameter space we can find the a , h and x_0 that minimize the entropy. In the third row of the table the pose distribution is shown with minimal entropy¹. The table makes it clear that even with a very simple surface we can reduce the uncertainty greatly by taking advantage of the dynamics.

1.2 Outline

In section 3 we will explain the notion of capture regions and introduce an extension and relaxation of this notion in the form of so-called quasi-capture regions. These quasi-capture regions allow for fast computation of pose distributions. In section 4 we will present our simulation and experimental results. Finally, in section 5 we will discuss the results presented in this paper. But first we will give an overview of related work in the next section.

2 Related Work

2.1 Parts Feeding and Orienting

One of the most comprehensive works on the design of parts feeding and assembly design is (Boothroyd et al., 1982), which describes vibratory bowls as well as non-vibratory parts feeders in detail. The APOS parts feeding system is described by Hitakawa (1988). In this system parts are fed over a vibrating tray. The tray has concavities which are designed in such a way that parts get stuck in one unique orientation. The design of these trays is mostly a process of ad-hoc trial and error. The goal of our research is to facilitate the automated design of these trays. Berretty et al. (1999) present an algorithm for designing a particular class of gates in vibratory bowls. Berkowitz and Canny (1996, 1997) use dynamic simulation to design a sequence of gates for a vibratory bowl. The dynamics are simulated with Mirtich's impulse-based dynamic simulator, *Impulse* (Mirtich and Canny, 1995). Christiansen et al. (1996) use genetic algorithms to design a near-optimal sequence of gates for a given part. Optimality is defined in terms of throughput. Here, the behavior of each gate is assumed to be known. So, in a sense (Christiansen et al., 1996) is complementary to (Berkowitz and Canny, 1997): the latter focuses on modeling the behavior of gates, the former finds an optimal sequence of gates given their behavior. Akella et al. (2000)

¹This is a local minimum found with simulated annealing and might not be the global minimum.

introduced a technique for orienting planar parts on a conveyor belt with a one degree-of-freedom (DOF) manipulator. Lynch (1999) extended this idea to 3D parts on a conveyor belt with a two DOF manipulator. Wiegley et al. (1996) presented a complete algorithm for designing passive fences to orient parts. Here, the initial orientation is unknown. Berretty (2000) showed that under certain assumptions this approach can be extended to polyhedral parts.

Goldberg (1993) showed that it is possible to orient polygonal parts with a frictionless parallel-jaw gripper without sensors. Marigo et al. (1997) showed how to orient and position a polyhedral part by rolling it between the two hands of a parallel-jaw gripper. Grossman and Blasgen (1975) developed a manipulator with a tactile sensor to orient parts in a tray. Erdmann and Mason (1988) developed a tray-tilting *sensorless* manipulator that can orient planar parts in the presence of friction. If it is not possible to bring a part into a unique orientation, the planner would try to minimize the number of final orientations. In (Erdmann et al., 1993) it is shown how (with some simplifying assumptions) three-dimensional parts can be oriented using a tray-tilting manipulator. Zumel (1997) used a variation of the tray tilting idea to orient planar parts with a pair of moveable palms.

In recent years a lot of work has been done on programmable force fields to orient parts (Böhringer et al., 2000a, 1999; Kavraki, 1997; Reznik et al., 1999) The idea is that an abstract force field (implemented using e.g. MEMS actuator arrays) can be used to push the part into a certain orientation. Böhringer et al. used Goldberg's algorithm (1993) to define a sequence of 'squeeze fields' to orient a part. They also gave an example how programmable vector fields can be used to simultaneously sort different parts and orient them. Kavraki (1997) presented a vector field that induced two stable configurations for most parts. In 2000, Böhringer et al. proved a long-standing conjecture that the vector field proposed in (Böhringer et al., 1996) is a universal feeder/orienter device, i.e., it induces a *unique* stable configuration for most parts. Recently, Sudsang and Kavraki (2001) introduced another vector field that has that property.

2.2 Stable Poses

To compute the stable poses of an object quasistatic dynamics is often assumed. Furthermore, usually it is assumed that the part is in contact with a flat surface and is initially at rest. Boothroyd et al. (1972) were among the first to analyze this problem. An $O(n^2)$ algorithm for n -sided polyhedrons, based on Boothroyd et al.'s ideas, was implemented by Wiegley et al. (1992). Goldberg et al. (1999) improved this method by approximating some of the dynamic effects. Kriegman (1997) introduced the notion of a *capture region*: a region in configuration space such that any initial configuration in that region will converge to one final configuration. Note that his work does not assume quasistatic dynamics; as long as the part is initially at rest and in contact, and the dynamics in the system are dissipative, the capture regions will be correct. The capture regions will in general not cover the entire configuration space.

2.3 Collision and Contact Analysis

For rigid body collisions several models have been proposed. Many of these models are either too restrictive (e.g., Routh's model (Routh, 1897) constrains the collision impulse too much) or allow physically impossible collisions (e.g., Whittaker's model (Whittaker, 1944) can predict

arbitrarily high increases of system kinetic energy). Wang and Mason (1992) proposed a collision model that combines Routh’s method with Poisson’s hypothesis. Their model admits tangential impact, i.e., an impact with zero initial approach velocity. Recently, Chatterjee and Ruina (1998) proposed a new collision rule, which is less restrictive than Routh’s model, but does not allow an energy increase. Chatterjee introduced a new collision parameter (besides the coefficients of friction and restitution): the coefficient of *tangential* restitution. With this extra parameter a large part of allowable collision impulses can be accounted for, and at the same time this collision rule restricts the predicted collision impulse to the allowable part of impulse space. This is the collision rule we will use (see (Moll and Erdmann, 2000) for details).

Instead of having algebraic laws, one could also try to model object interactions during impact. This approach is taken, for instance, by Bhatt and Koechling (1995a,b), who modeled impacts as a flow problem. While this might lead to more accurate predictions, it is computationally more expensive. Also, to get a good approximation of the pdf of resting configurations, this level of accuracy might not be required. On the other hand, it is also possible to combine the effects of multiple collisions that happen almost instantaneously. Goyal et al. (1998a,b) studied these “clattering” motions and derived the equations of motion.

Given a collision model and the equations of motion, one can simulate the motion of a part. In cases where there are a large number of collisions or with frequently changing contact modes one can simulate the dynamics using so-called impulse-based simulation (Mirtich and Canny, 1995). However, there are limits to the systems one can simulate. Under certain conditions the dynamics become chaotic (Bühler and Koditschek, 1990; Feldberg et al., 1990; Kechen, 1990). We are mostly interested in systems that are *not* chaotic, but where the dynamics can not be modeled with a quasistatic approximation.

2.4 Shape Design

The shape of an object and its environment imposes constraints on the possible motions of an object. Caine (1993) presented a method to visualize these motion constraints, which can be useful in the design phase of both part and manipulator. In (Krishnasamy, 1996) the mechanics of entrapment were analyzed. That is, Krishnasamy discussed conditions for a part to “get trapped” and “stay trapped” within an extrusion in the context of the APOS parts feeder. Sanderson (1984) presented a method to characterize the uncertainty in position and orientation of a part in an assembly system. This method takes into account the shape of both part and assembly system. In (Lynch et al., 1998) the optimal manipulator shape and motion were determined for a particular part. The problem here was not to orient the part, but to perform a certain juggler’s skill (the “butterfly”). With a suitable parametrization of the shape and motion of the manipulator, a solution was found for a disk-shaped part that satisfied their motion constraints. Although the analysis focused just on the juggling task, it showed that one can simulate and optimize dynamic manipulation tasks using a suitable parametrization of manipulator (or surface) shape and motion.

3 Analytic Results

3.1 Quasi-Capture Intuition

In our efforts to analyze pose distributions in a dynamic environment, we have been working on a generalization of so-called ‘capture regions’ (Kriegman, 1997) that we have termed *quasi-capture regions*. Specifically, for a part in contact with a sloped surface, we would like to determine whether it is captured, i.e., whether the part will converge to the closest stable pose. For simplicity, let the surface be a tilted plane.

Definition 1 Let a *pose* be defined as a point in configuration space such that the part is contact with the surface.

We assume that friction is sufficiently high so that a part cannot slide for an infinite amount of time. (This imposes a lower bound on the friction coefficient for a given slope.) In general capture depends on the whole surface and everything that happens after the current state, but the friction assumption and our definition of pose allow us to derive a sufficient condition for quasi-capture (in section 3.2) of the part in terms of local state. The closest stable pose can be defined as follows:

Definition 2 We define a *stable pose* to be a pose such that there is force balance when only gravity and contact forces are acting on the part. The *closest stable pose* is the stable pose found by following the gradient of the potential energy function (using e.g. gradient descent) from the current pose.

We can now define quasi-capture regions:

Definition 3 For each maximal connected set of stable poses we define an associated *quasi-capture region* as the largest possible region in configuration phase space such that (a) each configuration in this region has as its closest stable pose one of the specified stable poses, and (b) no configuration in a quasi-capture region has enough (kinetic and potential) energy to leave this region either with a rolling motion or one collision-free motion.

Ideally these quasi-capture regions would induce a partition of configuration phase space, so that for each point in phase space we would immediately know what its final resting configuration is. Of course, this is not the case in general, since with a sufficiently large velocity an object can reach any stable pose. But if we restrict the velocity to be small to begin with, then we are able to quickly determine the pose distribution. It has been our experience that without the use of quasi-capture regions a lot of computation time is spent on the final part/surface interactions (e.g., clattering motions) before the part reaches a stable pose. In other words, *with our analytic results it is possible to avoid computing a potentially large number of collisions*.

In our analysis we have focused on the two dimensional case. To illustrate the notion of capture, we will start with another example. Consider a rod of length l with center of mass at distance R from each vertex. One can visualize this rod as a disk with radius R and uniform mass, but with contact points only at the ends of the rod (see figure 2). Let α be the angle between

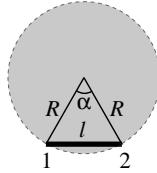


Figure 2: A rod with an off-center center of mass.

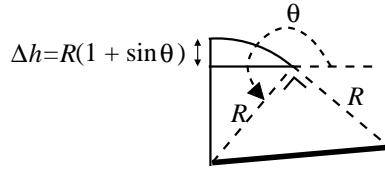


Figure 3: The rod is captured if its center of mass has not enough kinetic energy to roll over the contact points.

the vectors from the center of mass to the contact points.

Note that the endpoints of the rod are numbered. We will refer to these endpoints later. Let the ‘side’ of the rod where the center of mass is above the rod be the *high energy side*, and the other side be the *low energy side*. We can then define that the rod is ‘on’ the high energy side if and only if the center of mass is between and above the endpoints of the rod. The rod is ‘on’ the low energy side in all other cases. Suppose the rod is in contact with a flat, horizontal surface. For the rod to make a transition from one side to the other, it will have to rotate, either by rolling or by bouncing. At some point during the transition the center of mass will pass over one of the endpoints of the rod. Its potential energy at that point will always be greater than or equal to the potential energy at the start of the transition. Hence, to make that transition the rod has to have a minimum amount of kinetic energy. This can be written more formally as

$$\frac{1}{2}m\|\mathbf{v}\|^2 > -mg\Delta h. \quad (1)$$

Figure 3 illustrates the increase in potential energy as the rod rotates over one of the contact points.

For a sloped surface the capture condition is not as simple as for the horizontal surface. By bouncing and rolling down the slope, the rod can increase its kinetic energy. We have derived an upper bound on how far the rod can bounce. This gives an upper bound on the increase in kinetic energy. So the quasi-capture condition can now be stated as: the current kinetic energy plus the maximum gain in kinetic energy has to be less than the energy required to rotate to the other side. To guarantee that the rod is indeed captured, we have to make sure that the maximum gain in kinetic energy is less than the decrease in kinetic energy due to a collision. There are some additional complicating factors. For instance, a change in orientation can increase the kinetic energy, but to rotate to the other side the rod has to rotate back, undoing the gain in kinetic energy.

3.2 Quasi-Capture Velocity

What we will prove is a sufficient condition on the pose and velocity of the rod such that it is quasi-captured. The condition will be of the following form: if the current kinetic energy plus the maximal increase in kinetic energy is less than some bound, the rod is quasi-captured. This bound depends on the current orientation, the current velocity, the slope of the surface and the geometry of the rod.

Let ρ be the radius of gyration of the rod. We will model the dynamics using the generalized coordinates (x, y, q) , where (x, y) is the position of the center of mass and $q = \rho\theta$ represents the orientation of the object. Using these particular generalized coordinates some equations are greatly simplified. For example, the kinetic energy can then be written as

$$KE = \frac{1}{2}m(v_x^2 + v_y^2 + \rho^2\omega^2) = \frac{1}{2}m\|\mathbf{v}\|^2.$$

Without loss of generality we can assume $m = 2$. That way the kinetic energy is simply $\|\mathbf{v}\|^2$. We will write v for $\|\mathbf{v}\|$.

Definition 4 Let the *relative orientation* of a contact point be defined as the angle between the x-axis in the world frame and the vector from the center of mass of the rod to the contact point.

Theorem 5 Let ϕ be the slope of the surface and let θ be the relative orientation of the contact point. The rod with a velocity vector of length v and in contact with the surface is in a quasi-capture region if the following condition holds:

$$\begin{aligned} v^2 + \max_{\xi} \left(\frac{2v \cos \xi \sin \phi}{\cos^2 \phi} (v \sin(\xi + \phi) + \sqrt{v^2 \sin^2(\xi + \phi) - 2gd_n \cos \phi}) \right) - 2g\left(\frac{d_n}{\cos \phi} + R\varepsilon\right) \\ \leq -2gR\left(1 + \cos\left(\frac{\alpha}{2} + \phi\right)\right), \end{aligned}$$

where $d_n = R(\cos \frac{\alpha}{2} - \sin(\theta + \phi))$ and $\varepsilon = \cos\left(\frac{\alpha}{2} + \phi\right) - \frac{\cos(\alpha/2)}{\cos \phi} + \max(\tan \phi, 2 \sin \frac{\alpha}{2} \sin \phi)$.

Proof: See appendix A. □

For the rod on a slope there exist only two quasi-capture regions: one where the rod is on the high energy side and one where the rod is on the low energy side. Note that for $\phi = 0$ this bound reduces to $v^2 \leq -2gR(1 + \sin \theta)$. In other words, this bound is as tight as possible when the surface is horizontal. If the rod is in two-point contact, the above condition has to be satisfied for both contact points.

One can compute ξ numerically, but the appendix also gives a good analytic approximation. The theorem above gives a sufficient condition on the velocity and pose of the rod such that it cannot rotate to the other side during one bounce. But suppose there is a sequence of bounces, each of them increasing the kinetic energy. It is possible that the rod satisfies the quasi-capture condition, but is still able to rotate to the other side in more than one bounce. Thus, the theorem by itself is not enough to guarantee that the rod will converge to its closest stable orientation. In the analysis above we have ignored the dissipation of kinetic energy during collisions. If in the case the quasi-capture condition is true the dissipation of kinetic energy is larger than the

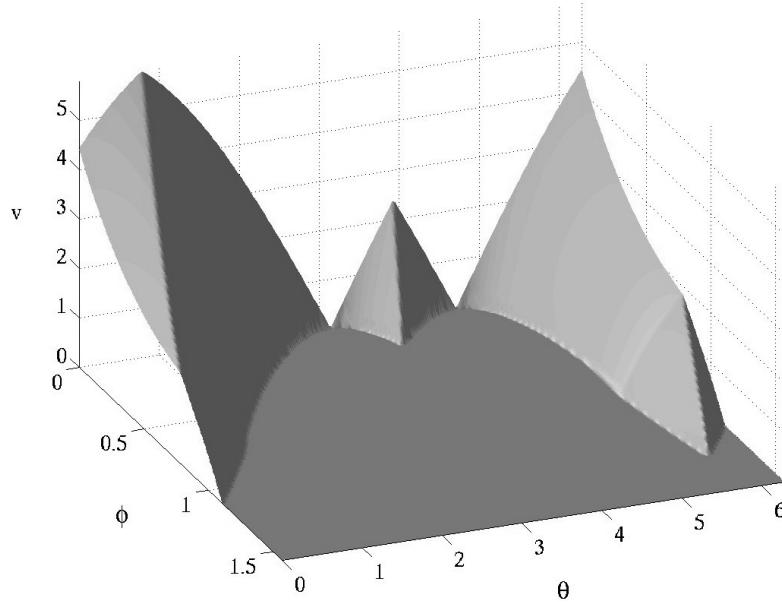


Figure 4: Quasi-capture velocity as a function of the slope of the surface and the orientation of the rod

increase due to the bounce, the rod will indeed be captured after an arbitrary number of bounces. To make sure this is the case the coefficients of restitution can not be too large.

In figure 4 the quasi-capture velocity is plotted as a function of the slope of the surface and the orientation of the rod. The slope ϕ ranges from 0 to $\frac{\pi}{2}$ and the orientation θ ranges from 0 to 2π . At $\theta = 0$ contact point 1 is directly to the right of the center of mass. Note that the orientation of the rod is not the same as the relative orientation of the contact point. However, for each combination of ϕ and θ the relative orientation of the contact point can be easily computed. The other relevant parameter values for this plot are: $R = 1\text{m}$, $g = -9.81\text{m/s}^2$ and $\alpha = \frac{\pi}{2}$. The little bump in the middle corresponds to the rod being quasi-captured on the high-energy side. The bigger bumps on the left and right correspond to being quasi-captured on the low-energy side.

4 Simulation and Experimental Results

4.1 Dynamic Simulation

To numerically compute the pose distribution of parts, we have written two dynamic simulators. One is based on David Baraff's Coriolis simulator (Baraff, 1991, 1993), which can simulate the motions of polyhedral rigid bodies. Coriolis takes care of the physical modelling. Our simulator then computes pose distributions for different (parametrized) support surfaces and different initial conditions.

Our simulator uses simulated annealing to optimize over the surface parameters and drop location with respect to the surface. The objective function is to minimize the entropy of the

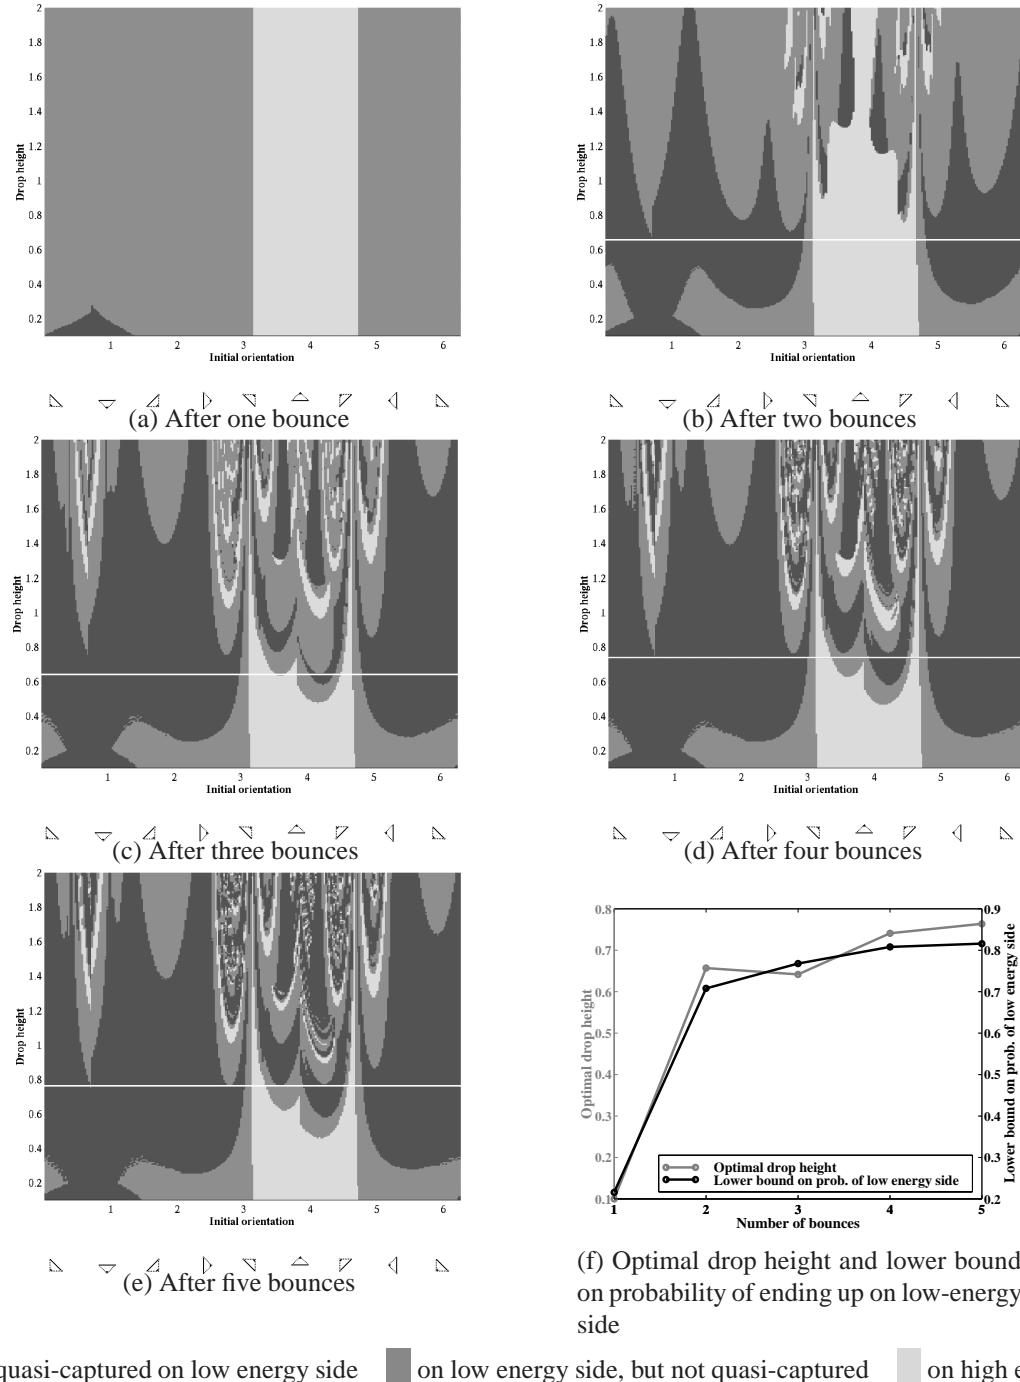
pose distribution. Initially the sampling of orientations of the object is rather coarse, so that the resulting pose distribution is not very accurate. But as the simulator is searching, the simulated annealing algorithm is restarted with an increased sample size and the best current solution as initial guess. This way we can quickly determine the potentially most interesting parameter values and refine them later. Our implementation is based on the one given in (Press et al., 1992, pp. 444–455).

Surfaces are parametrized using wavelets (Strang, 1989; Daubechies, 1993). Wavelet transforms are similar to the fast Fourier transform, but unlike the fast Fourier transform basis functions (sines and cosines) wavelet basis functions are localized in space. This localization gives us greater flexibility in modeling different surfaces compared with the fast Fourier transform or, say, polynomials. There are many classes of wavelet basis functions. We are using the Daubechies wavelet filters (Daubechies, 1993) and in particular the implementation as given in (Press et al., 1992, pp. 591–606). To reduce an arbitrary surface to a small number of coefficients we first discretize the function describing the surface. We then perform a wavelet transform and keep the largest components (in magnitude) in the transform to represent the surface. When we minimize the entropy, we optimize over these components. We can either keep the smaller components of our initial wavelet transform around or set them to zero.

Development of a second simulator was started, because Coriolis had some limitations. In particular, the collision model could not be changed and we wanted to experiment with Chatterjee’s collision model (Chatterjee and Ruina, 1998). The second simulator also allowed us to optimize for our specific dynamics model. In our model there is only one moving object, and the only forces acting on it are gravity and friction. Currently, the simulator only handles two-dimensional objects, but in the future it might be extended to handle three dimensions as well. It uses the analytic results from the previous section to decide whether the part is quasi-captured.

Using the simulator we can compute the quasi-capture regions for the rod. Figures 5(a)–(e) show the quasi-capture regions for the low energy side after one through five bounces. The dark areas correspond to initial orientations and initial velocities that result in the rod being quasi-captured. The zero orientation is defined as the orientation where endpoint 1 is directly to the right of the center of mass. The triangles below the horizontal axes show the pose of the rod corresponding to the orientation at that point of the horizontal axes. It is possible that in some cases the rod is quasi-captured on the high energy side. For simplicity this is not shown in figure 5. The friction and restitution parameters correspond roughly to the ones used in our experiments.

Let the optimal drop height be defined as the drop height that minimizes the entropy. In this particular case this height is equal to the height that maximizes the probability of ending up on the low energy side. Dropping the rod with uniformly random initial orientation from the optimal drop height will reduce uncertainty about its orientation maximally. In figures 5(a)–(e) the drop height that results in the maximum probability of ending up on the low energy side is marked by a horizontal line. After each successive bounce this drop height is likely to be a better approximation of the optimal drop height. In figure 5(f) the approximate optimal drop height and lower bound on the probability of ending up on the low-energy side after $1, 2, \dots, 5$ bounces is shown. One thing to note is that both the optimal drop height and the lower bound on the probability of ending up on low-energy side rapidly converge. This seems to suggest that after only a small number of bounces we could make a reasonable estimate of the optimal drop height



■ quasi-captured on low energy side ■ on low energy side, but not quasi-captured ■ on high energy side

Figure 5: Quasi-capture regions for the low energy side of the rod on a 7° slope. The rod parameters are $\alpha = \pi/2$, $R = 0.05m$, gravity, g , equals $-1.20m/s^2$, the coefficient of friction, μ , is 5 and the restitution parameters are: $e = 0.1$, $e_t = -0.2$.

g (m/s^2)	h (mm)	ϕ	Prob. low energy side	
			Sim.	Exp.
-0.68	58	20°	0.85	0.94
-0.68	122	20°	0.90	0.94
* -0.68	186	20°	0.91	0.93
-0.68	246	20°	0.93	0.96
-1.5	58	20°	0.85	0.93
* -1.5	122	20°	0.90	0.92
-1.5	186	20°	0.91	0.97
-1.5	246	20°	0.93	0.97
-2.6	58	20°	0.85	0.94
-2.6	122	20°	0.90	0.93
-2.6	186	20°	0.91	0.93
* -2.6	246	20°	0.91	0.94
-2.6	76	0°	0.75	0.75
-2.6	156	0°	0.88	0.83
* -2.6	220	0°	0.92	0.85
-2.6	284	0°	0.87	0.89

Table 2: Simulation and experimental results for the rod. Shown are the probabilities of ending up on the low-energy side for different values for g , h and ϕ . The drop height is measured from the center of the disk to the surface.

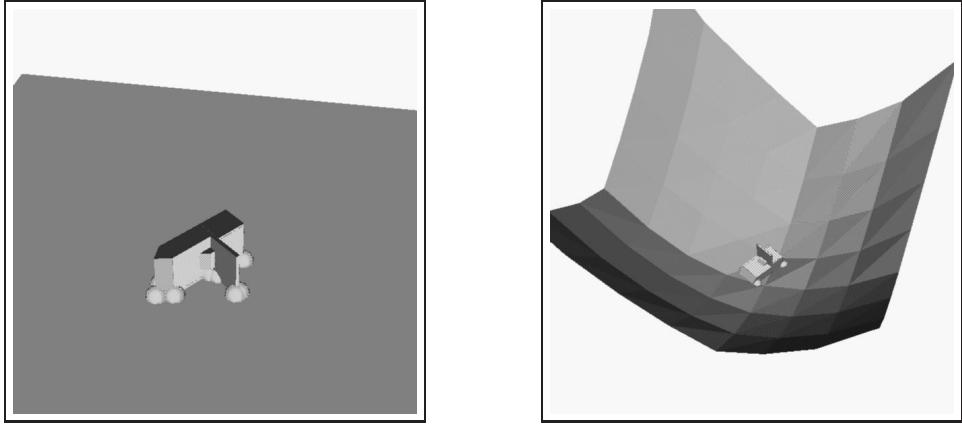
and uncertainty reduction. Further study is needed to find out if this is true in general.

4.2 2D Results

To test the assumptions and simulation results we also performed experiments. Our experimental setup was as follows. We used an air table to effectively create a two-dimensional world. By varying the slope of the air table we could vary gravity. At the bottom of the slope was the surface on which the object would be dropped. The angle ϕ of the surface in the plane defined by the air table could, of course, be varied.

The rod of the previous section has been implemented as a plastic disk with two metal pins sticking out from the top at an equal distance from the center of the disk. When released from the top of the air table the disk could slide under the surface and would only collide at the pins. Experimentally we determined the pose distribution of the rod for different values for g , h and ϕ by determining the final stable pose for 72 equally spaced initial orientations. Our simulation and experimental results of some tests have been summarized in table 2. The rows marked with an asterisk have been used to estimate the moment of inertia of the rod and the coefficients of friction and restitution. The estimated values for these parameters are: $e = 0.404$, $e_t = -0.136$, $\rho = 0.0376$ and $\mu = 4.71$. Note that for a low drop height and a horizontal surface the pdf is equal to a quasistatic approximation, as one would expect. More surprisingly, we see that the probability of ending up on the low-energy side can be changed to approximately 0.95 by setting g , h and ϕ to appropriate values. In other words, we can reduce the uncertainty almost completely.

One can identify several error sources for the differences between the simulation and ex-



(a) Insulator cap on a flat surface

(b) ... and on an optimized bowl

Figure 6: Result of optimizing a surface for the insulator cap.

perimental results. First, there are measurement errors in the experiments: in some cases slight changes in the initial conditions will change the side on which the rod will end up. Second, since the simulations are run with finite precision, it is possible that numerical errors affect the results. Finally, the physical model is not perfect. In particular, the rigid body assumption is just false. The surface on which the rod lands is coated with a thin layer of foam to create a high-damping, rough surface. This is done to prevent the rod from colliding with the sides of the air table. The foam also seems to make the dynamics of the system less chaotic.

4.3 3D Results

We have not generalized our analytic results to three dimensions yet, but we can still use our optimization technique to find a good surface and drop height for a given object. For the dynamic simulation we rely now on Baraff’s Coriolis simulator. Figure 6(a) shows an insulator cap² at rest on flat, horizontal surface. The contact points are marked by the little spheres. In figure 6(b) the bowl resulting from the simulated annealing search process is shown. The shape at the start of the simulated annealing search is a paraboloid: $f(x, y) = (x^2 + y^2)/20$. This shape is reduced to a triangulation of a 8×8 regular grid to obtain a compact wavelet transform. The part is always released on one side of the bowl.

We optimized over the four largest wavelet coefficients of the initial shape and the drop height. The search for the optimal bowl and drop height is visualized in figure 7. The five-dimensional parameter space is projected onto a two-dimensional space using Principal Component Analysis (Jolliffe, 1986). Principal Component Analysis can be thought of as a coordinate transform on a dataset such that the variability in the transformed data is strictly decreasing along successive axes. Hence, by only considering the first two dimensions we are likely to capture the most important features of the original data. Each point corresponds to a bowl shape evaluation, i.e., for each point a pose distribution is computed. The size of each point is proportional to the sample size used to determine the pose distribution. Computing a pose distribution by taking

²This object has been used before as an example in (Goldberg et al., 1999; Kriegman, 1997; Rao et al., 1995).

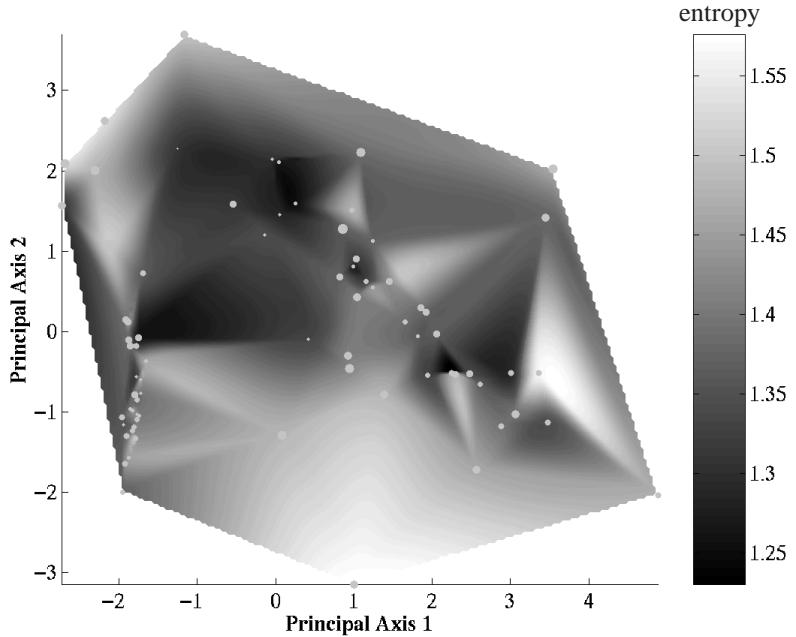


Figure 7: Entropy as a function of the two principal axes of the searched five-dimensional parameter space.

600 samples takes about 40 minutes on a 500 MHz Pentium III. The surface in figure 7 is a cubic interpolation between the points. The dark areas correspond to areas of low entropy. Notice that most of the points are in or near a dark area.

Table 3 compares the simulation results with experimental data from (Goldberg et al., 1999). The format is the same as in table 1, except that the stable poses are now written as vectors. These vectors are the outward pointing normals (w.r.t. the center of mass) of the planes passing through the contact points. That way, a face with many vertices in contact with the surface will always be represented by the same vector, no matter which subset of the vertices is actually in contact. In the experimental setup of (Goldberg et al., 1999) the part was dropped from one conveyor belt onto another. The initial drop height was 12.0 cm. In the experiments the part had an initial horizontal velocity of 5.0 cm/s. The second row corresponds to computing the pose

	Stable Poses						Entropy
	(-1, 0, 0)	(0, -1, 0)	(0, 1, 0)	(.8, 0, .6)	(.7, 0, -.7)	(0, 0, -1)	
experimental, flat (1036 trials)	0.271	0.460	0.197	0.050	0.022	1.58	
dynamic simulation, flat surface	0.355	0.207	0.221	0.185	0.019	0.014	1.48
dynamic simulation, optimal bowl	0.622	0.125	0.154	0.096	0.003	0.000	1.09

Table 3: Probability distribution function of stable poses for two surfaces. The initial velocity is zero and the initial rotation is uniformly random. The experimental data is taken from (Goldberg et al., 1999). There, (0, -1, 0) and (0, 1, 0) are counted as one pose.

distribution when the part is dropped from 12.0 cm (but with initial velocity set to 0). The third row corresponds to a local minimum returned by the simulated annealing algorithm. With the optimal bowl the first pose is significantly more likely to be the resting pose. For the simulation results the initial orientations were drawn from a uniform distribution. The optimal bowl would be different if the initial distribution was equal to the one observed in the experiments, possibly resulting in an even lower entropy.

5 Discussion

We have shown a sufficient condition on the position and velocity of the simplest possible ‘interesting’ shape (i.e., the rod) that guarantees convergence to the closest stable orientation under some assumptions. This condition gives rise to regions in configuration phase-space, where each point within such a region will converge to the same set of stable poses. We have coined the term quasi-capture regions for these regions, since they are very similar to Kriegman’s notion of capture regions. In simulations quasi-capture always seems to imply capture, but further research is needed to derive conditions on the energy dissipation at each impact such that quasi-capture is a sufficient condition for capture.

The quasi-capture regions also apply to general polygonal shapes. However, we can no longer use the symmetry of the rod. So the quasi-capture expressions for general polygonal shapes become more complex. On the other hand, we might be able to orient planar parts by using a setup similar to the one described in section 4 and attaching two pins to the top of the part. Generalizing the quasi-capture regions to three dimensions is non-trivial and is an interesting direction for future research.

The simulation and experimental results show that the simulator is not 100% accurate, but that it is a useful tool for determining the most promising initial conditions for uncertainty reduction. In other words, the optimum predicted by the simulator will probably be near-optimal in the experiments. We can then experimentally search for the true optimum.

Another area where quasi-capture regions may be applied is in computer animation. Before a part comes to rest, there are many interactions between the part and the support surface. It turns out that these interactions are computationally very expensive. With our capture regions we can eliminate the last ‘clattering’ motions of the part, since we can predict what the final pose will be. For applications where fast animation is more important than physical accuracy, a pre-computed motion can be substituted for the actual motion.

With future research we hope to improve the constraints on the quasi-capture velocity by taking into account more information, such as the direction of the velocity vector. If improving the quasi-capture bounds is impossible, it might be possible to get better approximations for pose distributions. As noted in section 4.1 it is possible to get a good estimate of the maximal uncertainty reduction after only a small number of bounces of the rod. So another interesting line of research would be to find out how accurate these approximations are in general.

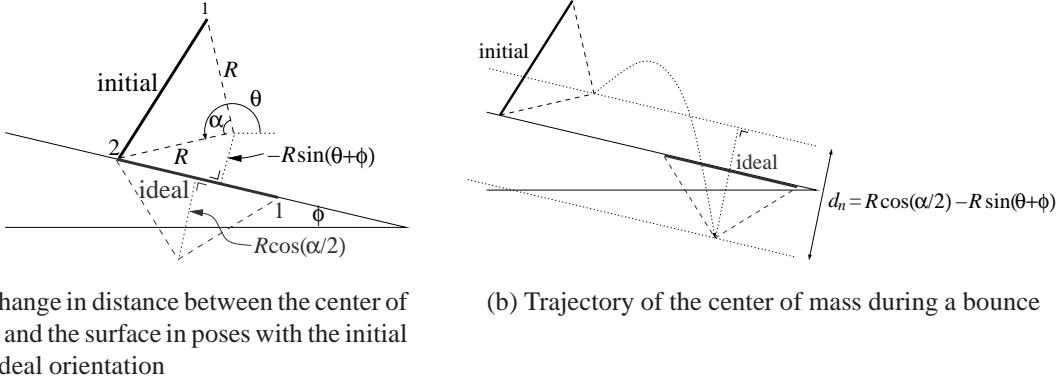


Figure 8: Increase in kinetic energy when rotating to the ideal orientation

Acknowledgments

This work was supported in part by the National Science Foundation under grant IRI-9503648. The authors are grateful to Al Rizzi, Matt Mason, Devin Balkcom and the anonymous reviewers for their helpful comments.

Appendix A: Proof of Theorem 5

Definition 6 Let a *bounce* be defined as the flight path between two impacts.

The closest distance between the rod and the slope during one bounce can be described by

$$d(t) = \frac{1}{2}g(\cos \phi)t^2 + (v_y \cos \phi + v_x \sin \phi)t - d_\theta(t), \quad (2)$$

where g is the gravitational constant (equal to -9.81m/s^2), v_x and v_y are the translational components of the velocity and $d_\theta(t)$ is a component that depends on the orientation. Let the rod be in contact at $t = 0$. Then $d(0) = 0$ (and therefore $d_\theta(0) = 0$). Let \hat{t} be the smallest positive solution to $d(t) = 0$. The change in height is then $\Delta h = \frac{1}{2}g\hat{t}^2 + v_y\hat{t}$, so that the change in v^2 is $\Delta v^2 = 2g\Delta h = g^2\hat{t}^2 + 2v_y g\hat{t}$. To find the maximum Δv^2 for all velocity vectors of length v we can parametrize the translational velocity as $v_x = v \cos \xi$ and $v_y = v \sin \xi$, and maximize over ξ . This ignores the rotational component of the velocity, but the following lemma shows that for a certain value of $d_\theta(\hat{t})$ the resulting solution for Δv^2 is an upper bound for the true maximal increase of v^2 .

Definition 7 Let the *ideal orientation* be defined as the orientation where the rod is parallel to the surface and the center of mass is below the rod.

Lemma 8 We can always increase the rod's kinetic energy after a bounce by allowing it to rotate around the center of mass 'for free' (i.e., without using energy) to the ideal orientation (ignoring penetrations of the surface) and then letting it continue to fall while maintaining this orientation.

However, if the rod is already in the ideal orientation after the bounce, its kinetic energy cannot be increased.

Proof: One can easily verify that rotating around the center of mass to the ideal orientation of the bounce maximizes distance between the rod and the surface. This distance will always be greater than or equal to 0. If we allow the rod to continue to fall until it hits the surface, its kinetic energy will increase. \square

From this lemma it follows that by assuming the rod rotates ‘for free’ to the ideal orientation the increase in kinetic energy due to one bounce is an upper bound on the true increase of kinetic energy. With this lemma computing the next contact point is a lot easier. Let θ be the relative orientation of the contact point at $t = 0$. The case $\theta = 0$ corresponds to the contact point being directly to the right of the center of mass. The signed distance from the center of mass to the surface at $t = 0$ is then $-R \sin(\theta + \phi)$, as shown in figure 8(a). One can easily verify that in the ideal orientation the relative orientation of endpoint 1 is $\frac{\pi}{2} - \frac{\alpha}{2} - \phi$. Let $\hat{\theta}$ be equal to this relative orientation. In the pose where the rod is in contact with the surface and has the ideal orientation the signed distance from the center of mass to the surface is $-R \sin(\hat{\theta} + \phi) = -R \cos \frac{\alpha}{2}$. So during one bounce the total displacement of the center of mass in the direction normal to the surface is at most $R(\cos \frac{\alpha}{2} - \sin(\theta + \phi))$. Let d_n be equal to this distance. To solve for the time of impact we can treat the rod as a point mass centered at the center of mass and replace $d_\theta(\hat{t})$ in equation 2 with $-d_n$. Equation 2 is then simply a paraboloid in t . The distance function now measures the distance between the center of mass and the dotted line parallel to the surface shown in figure 8(b). This approach is not limited to the case where the orientation after a bounce is the ideal orientation. Suppose we knew that the relative orientation of the contact point after a bounce is $\tilde{\theta}$. Then we can solve for the time of impact by substituting $R(\sin(\tilde{\theta} + \phi) - \sin(\theta + \phi))$ for $d_\theta(\hat{t})$ in equation 2.

The following lemma gives a bound on the velocity needed to roll to the other side.

Lemma 9 *If the rod is in rolling contact, then to be able to roll to the other side the following condition has to hold: $v^2 > -2gR(1 + \sin \theta + (\text{sign}(\cos \theta) - 1) \sin \frac{\alpha}{2} \sin \phi)$. We assume $0 \leq \phi < \frac{\pi}{2}$.*

Proof: We can distinguish several cases: endpoint 1 of the rod or endpoint 2 can be in contact with the slope, and the rod can be on the low or high energy side. We will prove the case where endpoint 1 is in contact and the rod is on the high energy side. The proof for the other cases is analogous. The case under consideration is shown in figure 9(a). To roll counterclockwise over to the left side, $v^2 > -2gh_1$. The distance h_1 is simply equal to $R(1 + \sin \theta)$. If the rod rolls clockwise over to two-point contact and continues to roll over endpoint 2, the rod gains kinetic energy because the second contact point is lower than the first contact point. This gain is proportional to h_2 .

One can easily verify that for two-point contact the relative orientations of contact points 1 and 2 are $\frac{3\pi}{2} - \frac{\alpha}{2} - \phi$ and $\frac{3\pi}{2} + \frac{\alpha}{2} - \phi$, respectively. The bound for rolling over endpoint 2 is

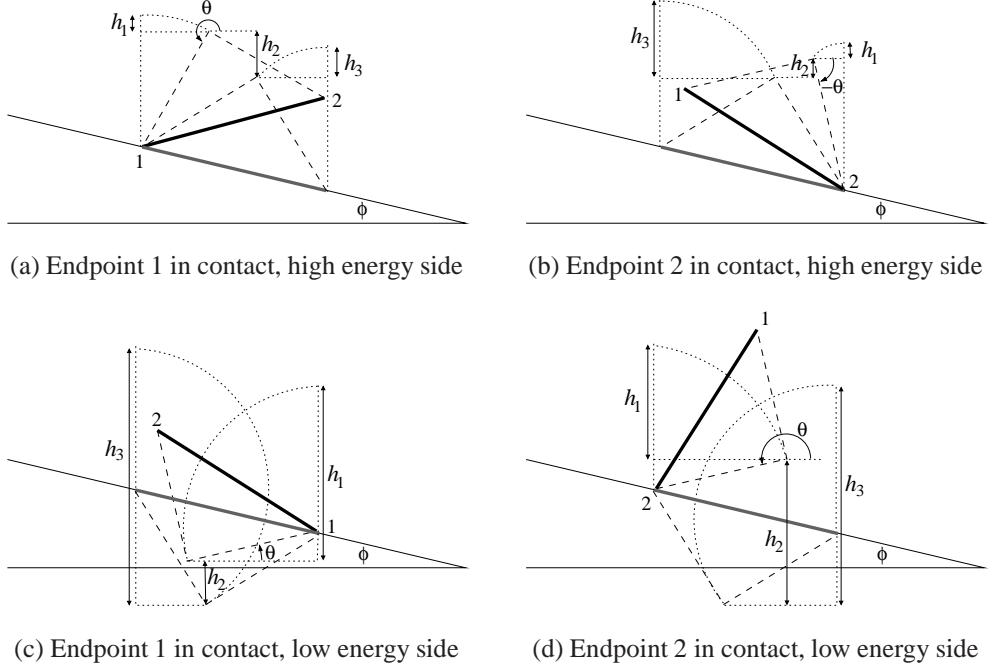


Figure 9: Capture condition for rotation

therefore

$$\begin{aligned}
 v^2 &> -2g(h_3 - h_2) \\
 &= -2gR(1 + \sin\theta - \cos(\frac{\alpha}{2} - \phi) + \cos(\frac{\alpha}{2} + \phi)) \\
 &= -2gR(1 + \sin\theta - 2\sin\frac{\alpha}{2}\sin\phi).
 \end{aligned} \tag{3}$$

By inspection of figure 9 we see that we can use *almost* the same bounds for the other three cases: in the cases corresponding to figures 9(b) and 9(c), the last term of expression 3 will change sign. In those two cases the rod rotates in the opposite direction “up-hill” to make two-point contact. We can combine the two bounds (one for rotating clockwise, and one for rotating counterclockwise) into a bound that covers all four cases:

$$\begin{aligned}
 v^2 &> \min(-2gR(1 + \sin\theta), -2gR(1 + \sin\theta + 2\operatorname{sign}(\cos\theta)\sin\frac{\alpha}{2}\sin\phi)) \\
 &= -2gR(1 + \sin\theta + (\operatorname{sign}(\cos\theta) - 1)\sin\frac{\alpha}{2}\sin\phi).
 \end{aligned} \tag{4}$$

□

Recall that we assume that friction is sufficiently high so that a part cannot slide for an infinite amount of time.

Theorem 5 *Let ϕ be the slope of the surface and let θ be the relative orientation of the contact point. The rod with a velocity vector of length v and in contact with the surface is in a quasi-*

capture region if the following condition holds:

$$\begin{aligned} v^2 + \max_{\xi} \left(\frac{2v \cos \xi \sin \phi}{\cos^2 \phi} (v \sin(\xi + \phi) + \sqrt{v^2 \sin^2(\xi + \phi) - 2gd_n \cos \phi}) \right) - 2g \left(\frac{d_n}{\cos \phi} + R\varepsilon \right) \\ \leq -2gR \left(1 + \cos\left(\frac{\alpha}{2} + \phi\right) \right), \end{aligned}$$

where $d_n = R(\cos \frac{\alpha}{2} - \sin(\theta + \phi))$ and $\varepsilon = \cos\left(\frac{\alpha}{2} + \phi\right) - \frac{\cos(\alpha/2)}{\cos \phi} + \max(\tan \phi, 2 \sin \frac{\alpha}{2} \sin \phi)$.

Proof: The basic strategy of this proof is to derive an upper bound on the increase in kinetic energy acquired by the rod during a bounce and a lower bound on the change in potential energy needed for the rod to rotate from one side to the other after a bounce.

We will first derive an upper bound on the maximum increase in kinetic energy. We maximize over the direction of the translational velocity at the beginning of the bounce and over all possible orientations of the rod at the end of the bounce. Let ξ be defined as the direction of the translational velocity vector at the beginning of the bounce that will result in the largest increase in kinetic energy during the bounce. Let $\tilde{\theta}$ be the relative orientation of the contact point at the end of the bounce.

Since we can choose $\tilde{\theta}$, the rotational velocity of the rod does not contribute to an increase in kinetic energy. We may therefore assume that the rod rotates ‘for free’ to the orientation $\tilde{\theta}$ and that the velocity is pure translation. Assuming that the velocity is pure translation results in a higher upper bound on the increase in kinetic energy than might occur during an actual bounce, and thus is a conservative estimate for quasi-capture.

The signed distance between the center of mass and the surface during a bounce is described by $\frac{1}{2}gt^2 \cos \phi + v(\sin \xi \cos \phi + \cos \xi \sin \phi)t + d_{\tilde{\theta}}$, where $d_{\tilde{\theta}}$ is the distance along the surface normal between the center of mass at the next impact and the center of mass at time t (cf. figure 8(b) and equation 2). The smallest positive solution for t such that this expression is equal to 0 is

$$\begin{aligned} \hat{t} &= \frac{-v(\sin \xi \cos \phi + \cos \xi \sin \phi)}{g \cos \phi} - \frac{\sqrt{v^2(\sin \xi \cos \phi + \cos \xi \sin \phi)^2 - 2gd_{\tilde{\theta}} \cos \phi}}{g \cos \phi} \\ &= \frac{-v \sin(\xi + \phi) - \sqrt{v^2 \sin^2(\xi + \phi) - 2gd_{\tilde{\theta}} \cos \phi}}{g \cos \phi}. \end{aligned} \tag{5}$$

We can therefore bound the maximum change in v^2 by

$$\begin{aligned} \Delta v^2 &= 2g \Delta h = 2g(\frac{1}{2}g\hat{t}^2 + v(\sin \xi)\hat{t}) \\ &= \frac{2v \cos \xi \sin \phi}{\cos^2 \phi} \left(v \sin(\xi + \phi) + \sqrt{v^2 \sin^2(\xi + \phi) - 2gd_{\tilde{\theta}} \cos \phi} \right) - \frac{2gd_{\tilde{\theta}}}{\cos \phi} \end{aligned} \tag{6}$$

$$\leq \frac{2v \cos \xi \sin \phi}{\cos^2 \phi} \left(v \sin(\xi + \phi) + \sqrt{v^2 \sin^2(\xi + \phi) - 2gd_n \cos \phi} \right) - \frac{2gd_n}{\cos \phi}. \tag{7}$$

Let $f(\tilde{\theta})$ be equal to expression 6. Then $f(\tilde{\theta})$ is an upper bound on the increase in kinetic energy if the relative orientation after a bounce is $\tilde{\theta}$. Recall that $\hat{\theta}$ is the ideal orientation, resulting in a maximum increase in kinetic energy. So expression 7 is equal to $f(\hat{\theta})$ and, according to lemma 8, $f(\hat{\theta})$ is an upper bound for all $f(\tilde{\theta})$. Similarly, let $g(\tilde{\theta})$ be the lower bound given by lemma 9 on the kinetic energy needed to roll to the other side after a bounce:

$$g(\tilde{\theta}) = -2gR \left(1 + \sin \tilde{\theta} + (\text{sign}(\cos \tilde{\theta}) - 1) \sin \frac{\alpha}{2} \sin \phi \right).$$

After one bounce the orientation is assumed to be such that rod is parallel to the surface and the center of mass is below the rod, as this will result in the largest increase in kinetic energy according to lemma 8. This means that endpoint 1's relative orientation is equal to $\hat{\theta} = \frac{\pi}{2} - \frac{\alpha}{2} - \phi$. The value of $g(\hat{\theta})$ is $-2gR(1 + \sin \hat{\theta})$. In other words, if the kinetic energy after the bounce is less than $-2gR(1 + \sin \hat{\theta})$ and the rod is in the ideal orientation, the rod cannot roll to the other side. If the rod can not rotate to the other side immediately *after* a bounce, then it certainly does not have enough kinetic energy to rotate to the other side *during* the bounce.

We can combine the two bounds $f(\hat{\theta})$ and $g(\hat{\theta})$ to obtain a sufficient condition to determine whether the rod can rotate to the other side *if its new orientation after one bounce is equal to the ideal orientation*. Unfortunately this condition does not imply a similar condition for the general case where the new orientation is not necessarily equal to the ideal orientation.

Consider the case $v = 0^+$, i.e., v is an infinitesimally small positive number. Substituting this value in equation 7 and expanding the definition of d_n shows that the maximum increase in kinetic energy is then

$$f(\hat{\theta})|_{v=0^+} = -\frac{2gd_n}{\cos \phi} = -\frac{2gR(\sin(\hat{\theta}+\phi)-\sin(\theta+\phi))}{\cos \phi}. \quad (8)$$

Therefore, when $v = 0^+$ and the relative orientation of the contact point after the bounce is equal to the ideal orientation the quasi-capture constraint is

$$-2gR\frac{\sin(\hat{\theta}+\phi)-\sin(\theta+\phi)}{\cos \phi} \leq -2gR(1 + \sin \hat{\theta}). \quad (9)$$

That is, if an upper bound on the kinetic energy after one bounce is less than the energy needed to rotate to the other side, the rod will not be able to rotate to the other side. Now suppose the new orientation is *not* equal to the ideal orientation. Then the increase of kinetic energy will be less, but the energy required to roll to the other side will be less, too. Unfortunately, this bound does not imply that $f(\tilde{\theta}) \leq g(\tilde{\theta})$ for all $\tilde{\theta}$. We would like to determine the smallest possible ε such that

$$f(\hat{\theta}) - 2gR\varepsilon \leq g(\hat{\theta}) \quad \Rightarrow \quad \forall \tilde{\theta}. f(\tilde{\theta}) \leq g(\tilde{\theta}).$$

It is not hard to see ε has to be equal to

$$\max_{\tilde{\theta}}(g(\hat{\theta}) - g(\tilde{\theta}) - f(\hat{\theta}) + f(\tilde{\theta})) / (-2gR).$$

The difference between $f(\hat{\theta})$ and $f(\tilde{\theta})$ is

$$-2gR\frac{\sin(\hat{\theta}+\phi)-\sin(\tilde{\theta}+\phi)}{\cos \phi}.$$

Similarly, the difference between $g(\hat{\theta})$ and $g(\tilde{\theta})$ is

$$-2gR \left(\sin \hat{\theta} - \sin \tilde{\theta} - (\text{sign}(\cos \tilde{\theta}) - 1) \sin \frac{\alpha}{2} \sin \phi \right).$$

The correction ε is therefore

$$\varepsilon = \max_{\tilde{\theta}} \left(\sin \hat{\theta} - \sin \tilde{\theta} - (\text{sign}(\cos \tilde{\theta}) - 1) \sin \frac{\alpha}{2} \sin \phi - \frac{\sin(\hat{\theta}+\phi)-\sin(\tilde{\theta}+\phi)}{\cos \phi} \right).$$

By differentiating the expression inside $\max(\cdot)$ with respect to $\tilde{\theta}$ we find that there is a local maximum at $\tilde{\theta} = 0$. Other local maxima occur when $\tilde{\theta}$ approaches $-\frac{\pi}{2}$ from below or $\frac{\pi}{2}$ from above. The correction ε therefore simplifies to

$$\begin{aligned}\varepsilon &= \max \left(\sin \hat{\theta} - \frac{\sin(\hat{\theta} + \phi) - \sin \phi}{\cos \phi}, \sin \hat{\theta} + 2 \sin \frac{\alpha}{2} \sin \phi - \frac{\sin(\hat{\theta} + \phi)}{\cos \phi} \right) \\ &= \cos\left(\frac{\alpha}{2} + \phi\right) - \frac{\cos(\alpha/2)}{\cos \phi} + \max \left(\tan \phi, 2 \sin \frac{\alpha}{2} \sin \phi \right).\end{aligned}$$

For $v \neq 0^+$ the difference between $f(\hat{\theta})$ and $f(\tilde{\theta})$ is even larger and $g(\tilde{\theta})$ does not depend on v , so the value for ε is an upper bound for all v . Combining all the bounds we arrive at the desired result for a bouncing rod. It is easy to see this also covers the case where the rod is rolling. \square

Note that for $\phi = 0$ this bound reduces to $v^2 \leq -2gR(1 + \sin \theta)$. In other words, this bound is as tight as possible when the surface is horizontal.

For an arbitrary d_n it is not possible to compute the optimal ξ analytically. Fortunately, we can analytically solve for ξ if we assume that the bounce consists of pure translation. The resulting ξ can be used as an approximation. It can be shown that the solution for ξ can be written as

$$\cos \xi = \frac{\cos \phi}{\sqrt{2\sqrt{1-\sin \phi}}} \wedge \sin \xi = \frac{\sqrt{1-\sin \phi}}{\sqrt{2}},$$

assuming $0 \leq \phi < \frac{\pi}{2}$. Substituting these values in equation 5, we find that the approximation for the bound for Δv^2 then simplifies to

$$\Delta v^2 \leq -\frac{2gd_n}{\cos \phi} + \frac{v^2 \sin \phi}{1-\sin \phi} \left(1 + \sqrt{1 - 4d_n g \frac{1-\sin \phi}{v^2 \cos \phi}} \right)$$

The relative error in this approximation depends on ϕ , d_n , v and g and can be computed numerically. Somewhat surprisingly, the relative error appears to be constant in v , d_n and g . The relative error does vary significantly with ϕ , but is still fairly small (on the order of 10^{-2}).

References

- Akella, S., Huang, W. H., Lynch, K. M., and Mason, M. T. (2000). Parts feeding on a conveyor with a one joint robot. *Algorithmica*, 26:313–344.
- Baraff, D. (1991). Coping with friction for non-penetrating rigid body simulation. *Computer Graphics*, 25(4):31–40.
- Baraff, D. (1993). Issues in computing contact forces for non-penetrating rigid bodies. *Algorithmica*, pages 292–352.
- Berkowitz, D. R. and Canny, J. (1996). Designing parts feeders using dynamic simulation. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, pages 1127–1132.
- Berkowitz, D. R. and Canny, J. (1997). A comparison of real and simulated designs for vibratory parts feeding. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 2377–2382, Albuquerque, New Mexico.

- Berretty, R.-P. (2000). *Geometric Design of Part Feeders*. PhD thesis, Utrecht University.
- Berretty, R.-P., Goldberg, K., Cheung, L., Overmars, M. H., Smith, G., and van der Stappen, A. F. (1999). Trap design for vibratory bowl feeders. In *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, pages 2558–2563, Detroit, MI.
- Bhatt, V. and Koechling, J. (1995a). Partitioning the parameter space according to different behavior during 3d impacts. *ASME Journal of Applied Mechanics*, 62(3):740–746.
- Bhatt, V. and Koechling, J. (1995b). Three dimensional frictional rigid body impact. *ASME Journal of Applied Mechanics*, 62(4):893–898.
- Böhringer, K. F., Bhatt, V., Donald, B. R., and Goldberg, K. Y. (2000a). Algorithms for sensorless manipulation using a vibrating surface. *Algorithmica*, 26(3/4):389–429.
- Böhringer, K.-F., Donald, B. R., Kavraki, L. E., and Lamiraux, F. (2000b). Part orientation with one or two stable equilibria using programmable force fields. *IEEE Transactions on Robotics and Automation*, 16(2):157–170.
- Böhringer, K. F., Donald, B. R., and MacDonald, N. C. (1996). Upper and lower bounds for programmable vector fields with applications to MEMS and vibratory plate parts feeders. In Overmars, M. and Laumond, J.-P., editors, *Workshop on the Algorithmic Foundations of Robotics*, pages 255–276. A. K. Peters.
- Böhringer, K. F., Donald, B. R., and MacDonald, N. C. (1999). Programmable vector fields for distributed manipulation, with applications to MEMS actuator arrays and vibratory parts feeders. *International Journal of Robotics Research*, 18(2):168–200.
- Boothroyd, C., Redford, A. H., Poli, C., and Murch, L. E. (1972). Statistical distributions of natural resting aspects of parts for automatic handling. *Manufacturing Engineering Transactions, Society of Manufacturing Automation*, 1:93–105.
- Boothroyd, G., Poli, C., and Murch, L. E. (1982). *Automatic Assembly*. Marcel Dekker, Inc., New York; Basel.
- Bühler, M. and Koditschek, D. E. (1990). From stable to chaotic juggling: Theory, simulation, and experiments. In *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*, pages 1976–1981.
- Caine, M. E. (1993). *The Design of Shape from Motion Constraints*. PhD thesis, MIT Artificial Intelligence Laboratory, Cambridge, MA. Technical Report 1425.
- Chatterjee, A. and Ruina, A. L. (1998). A new algebraic rigid body collision law based on impulse space considerations. *ASME Journal of Applied Mechanics*, 65:939–951.
- Christiansen, A. D., Edwards, A. D., and Coello Coello, C. A. (1996). Automated design of part feeders using a genetic algorithm. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, volume 1, pages 846–851.

- Daubechies, I., editor (1993). *Different Perspectives on Wavelets*, volume 47 of *Proceedings of Symposia in Applied Mathematics*. AMS.
- Erdmann, M. A. and Mason, M. T. (1988). An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4(4):369–379.
- Erdmann, M. A., Mason, M. T., and Vaněček, Jr., G. (1993). Mechanical parts orienting: The case of a polyhedron on a table. *Algorithmica*, 10:226–247.
- Feldberg, R., Szymkat, M., Knudsen, C., and Mosekilde, E. (1990). Iterated-map approach to die tossing. *Physical Review A*, 42(8):4493–4502.
- Goldberg, K., Mirtich, B., Zhuang, Y., Craig, J., Carlisle, B., and Canny, J. (1999). Part pose statistics: Estimators and experiments. *IEEE Transactions on Robotics and Automation*, 15(5).
- Goldberg, K. Y. (1993). Orienting polygonal parts without sensors. *Algorithmica*, 10(3):201–225.
- Goyal, S., Papadopoulos, J. M., and Sullivan, P. A. (1998a). The dynamics of clattering I: Equation of motion and examples. *Journal of Dynamic Systems, Measurement, and Control*, 120:83–93.
- Goyal, S., Papadopoulos, J. M., and Sullivan, P. A. (1998b). The dynamics of clattering II: Global results and shock protection. *Journal of Dynamic Systems, Measurement, and Control*, 120:94–102.
- Grossman, D. D. and Blasgen, M. W. (1975). Orienting mechanical parts by computer-controlled manipulator. *IEEE Transactions on Systems, Man, and Cybernetics*, 5:561–565.
- Hitakawa, H. (1988). Advanced parts orientation system has wide application. *Assembly Automation*, 8(3):147–150.
- Jolliffe, I. T. (1986). *Principal Components Analysis*. Springer-Verlag, New York.
- Kavraki, L. E. (1997). Part orientation with programmable vector fields: Two stable equilibria for most parts. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 2446–2451, Albuquerque, New Mexico.
- Kechen, Z. (1990). Uniform distribution of initial states: The physical basis of probability. *Physical Review A*, 41(4):1893–1900.
- Kriegman, D. J. (1997). Let them fall where they may: Capture regions of curved objects and polyhedra. *International Journal of Robotics Research*, 16(4):448–472.
- Krishnasamy, J. (1996). *Mechanics of Entrapment with Applications to Design of Industrial Part Feeders*. PhD thesis, Dept. of Mechanical Engineering, MIT.
- Lynch, K. M. (1999). Toppling manipulation. In *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, pages 2551–2557, Detroit, MI.

- Lynch, K. M., Shiroma, N., Arai, H., and Tanie, K. (1998). The roles of shape and motion in dynamic manipulation: The butterfly example. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pages 1958–1963.
- Marigo, A., Chitour, Y., and Bicchi, A. (1997). Manipulation of polyhedral parts by rolling. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 2992–2997.
- Mirtich, B. and Canny, J. (1995). Impulse-based simulation of rigid bodies. In *Proc. 1995 Symposium on Interactive 3D Graphics*.
- Moll, M. and Erdmann, M. A. (2000). Manipulation of pose distributions. Technical Report CMU-CS-00-111, Dept. of Computer Science, Carnegie Mellon University.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C: The art of Scientific Computing*. Cambridge University Press, second edition.
- Rao, A., Kriegman, D., and Goldberg, K. (1995). Complete algorithms for reorienting polyhedral parts using a pivoting gripper. In *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*, pages 2242–2248.
- Reznik, D., Moshkovich, E., and Canny, J. (1999). Building a universal part manipulator. In Böhringer, K. and Choset, H., editors, *Distributed Manipulation*. Kluwer.
- Routh, E. J. (1897). *Dynamics of a System of Rigid Bodies*. MacMillan and Co., London, sixth edition.
- Sanderson, A. C. (1984). Parts entropy methods for robotic assembly system design. In *Proceedings of the 1984 IEEE International Conference on Robotics and Automation*, pages 600–608.
- Strang, G. (1989). Wavelets and dilation equations: A brief introduction. *SIAM Review*, 31:613–627.
- Sudsang, A. and Kavraki, L. (2001). A geometric approach to designing a programmable force field with a unique stable equilibrium for parts in the plane. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*.
- Wang, Y. and Mason, M. T. (1992). Two-dimensional rigid-body collisions with friction. *ASME Journal of Applied Mechanics*, 59:635–642.
- Whittaker, E. T. (1944). *A Treatise on the Analytical Dynamics of Particles and Rigid Bodies*. Dover, New York, fourth edition.
- Wiegley, J., Goldberg, K., Peshkin, M., and Brokowski, M. (1996). A complete algorithm for designing passive fences to orient parts. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, pages 1133–1139.
- Wiegley, J., Rao, A., and Goldberg, K. (1992). Computing a statistical distribution of stable poses for a polyhedron. In *30th Annual Allerton Conf. on Communications, Control and Computing*.

Zumel, N. B. (1997). *A Nonprehensile Method for Reliable Parts Orienting*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.