

# Uncertainty Reduction Using Dynamics\*

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 an 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 show simulation and experimental results that confirm that our dynamic simulator can be used to find the true pose distribution of an object.*

## 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 not possible to orient a part completely without sensors, but 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 pose distribution. The goal then is to find the pose distribution with minimal entropy, thereby maximally reducing uncertainty.

Suppose a polyhedron is initially in a random configuration and the only force acting on it is gravity. We can then

compute an approximation of the probability distribution function (pdf) of resting configurations. This approximation will not only depend on the geometry and mass distribution of the polyhedron, but also on the physical model (quasistatic vs. dynamic) and the coefficients of friction and restitution. To orient the part, a robot arm with a camera could detect the current orientation, pick up the part and then put it in the desired orientation. This approach can be costly if high throughput is necessary; a robot can typically orient only one part at a time and might have to re-grip to get the part from initial to desired configuration. A more common approach for small parts is to have a particular (moving) surface for a part that can orient many parts at the same time. Examples are SONY's APOS system (Hitakawa, 1988), vibratory bowls and conveyor belts with obstacles that align the parts (Akella et al., 1997; Peshkin and Sanderson, 1988; Wiegley et al., 1996). In the APOS system parts are fed over a vibrating tray with extrusions such that parts will either get stuck in only one orientation, or otherwise are fed over the tray again. Vibratory bowls let parts vibrate to the top of the bowl. In the bowl are obstacles that align the parts in a certain way. Conveyor belts can serve a similar function: parts are put on the belt and are aligned by obstacles (or gates) along the way. The design of APOS trays, vibratory bowls and obstacles on conveyor belts is currently still done by hand by experienced engineers who intuit what could work. Still, it typically takes at least a week to design a good APOS tray or vibratory bowl.

In this article we will discuss briefly the use of dynamic simulation for the design of support surfaces that reduce the uncertainty of a part's resting configuration<sup>1</sup>. As the support surface is changed, the 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:

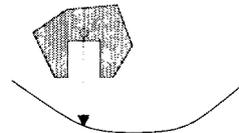


Figure 1: A part is dropped on a surface.

A part is released from a certain height and relative horizontal position with respect to the bowl. The only forces acting on the part are gravity and friction. We assume the bowl doesn't move. We can compute the final rest-

\*This work was supported in part by the NSF under grant IRI-9503648.

<sup>1</sup>For more details see Moll and Erdmann (2000).

ing 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 section 3 we will present some results for this example.

In section 4 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 the pose distribution. In section 5 we will present our simulation and experimental results. Finally, in section 6 we will discuss the results presented in this paper. But first we will give an overview of related work.

## 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). It is part of the automatic assembly system called SMART (Sony Multi-Assembly Robot Technology). One of the strong points of the APOS system is its flexibility: by replacing the tray and fine-tuning the vibrating motion, other parts can be oriented. *How* these trays are designed and *how* to change the motion is not clear. Automating the design of these trays would increase the flexibility even further.

Berkowitz and Canny (1996, 1997) use dynamic simulation to design a sequence of gates for a vibratory bowl. They represent the effects produced by the gates as state transitions in a non-deterministic state automaton. 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.

Goldberg (1993) showed that it is possible to orient polygonal parts with a frictionless parallel-jaw gripper without sensors. Goldberg conjectured and Chen and Ierardi (1995) proved that for every  $n$ -sided polygonal part, a sequence of 'squeezes' can be computed in  $O(n^2)$  time that will orient it up to symmetry. The length of such a sequence is  $O(n)$ . These results might have analogues in three dimensions. 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. Here, however, infinite friction is assumed, whereas Goldberg assumed no friction. In (Rao et al., 1995) an algorithm is described to orient polyhedral parts using so-called pivot grasps. A part is

grasped with two hard finger contacts and is then free to rotate around the axis formed by the contacts. Their algorithm computes an  $m \times m$  matrix of pivot grasps, where  $m$  is the number of stable configurations and each entry corresponds to a transition from one stable configuration to another. In general, there will be some null entries in this matrix. In other words, it is not always possible to go from any stable configuration to any other. A vision system is used to determine a part's location and orientation. In (Gudmundsson and Goldberg, 1997) a similar system is described, where a robot arm picks up parts that pass on a conveyor belt, but here the focus is on tuning the speed of the conveyor belt. Gudmundsson and Goldberg analytically show how to maximize throughput as a function of the speed of the conveyor belt.

Erdmann and Mason (1988) developed a tray-tilting sensorless manipulator that can orient planar parts in the presence of friction. If it isn't possible to bring a part into a unique orientation, the planner tries 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. In particular, for polyhedral parts with  $n$  faces a sequence of 'tilts' of length  $O(n)$  can be found in  $O(n^3)$  time. Zumel (1997) used a variation of the tray tilting idea to orient planar parts. Zumel used two actuated arms connected at a hinge to tilt parts from one arm to the other. The stable poses of the part at different angles were pre-computed. The planner then found a sequence of joint angle pairs for the two arms that would orient the part.

In recent years much work has been done on programmable force fields to orient parts (Böhringer et al., 1997, 1999; Kavraki, 1997). The idea is that some kind of force 'field' (implemented using e.g. MEMS actuator arrays) can be used to push the part in a certain orientation. Kavraki (1997) presented a vector field that induced two stable configurations for most parts. Böhringer et al. used Goldberg's algorithm (1993) to define a sequence of 'squeeze fields' to orient a part. They also gave an example of how programmable vector fields can be used to simultaneously sort different parts and orient them. Luntz et al. (1997) implemented a parcel transport and manipulation system using a distributed actuator array borrowing ideas from Böhringer et al.

### 2.2 Stable Poses

Quasistatic dynamics is often assumed when computing the stable poses of an object. 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. Using potential energy arguments and some simplifying assumptions they were able to get good approximations of the pdf's of some parts. They also introduced a method to get a static solution for the pdf: the probability of coming to rest on a face is simply proportional to the area of the face's projection on a sphere centered at the center of mass. The probability of an unsta-

ble face is added to the probability of the face onto which it rolls. An  $O(n^2)$  algorithm for  $n$ -sided polyhedrons, based on this idea, was implemented by Wiegley et al. (1992). Mirtich et al. (1996) improved this method by approximating some of the dynamic effects. In particular, they computed the area of the intersection of a face with a unit-area circle centered around the center of mass projected on the plane defined by that face. This was then taken as a measure of stability for that face.

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. He also described an algorithm based on Morse theory that computes the maximal capture regions of an object. Note that this work doesn't assume quasistatic dynamics; as long as the part is initially at rest and in contact, and the dynamics of 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

Computing reaction forces for an object in contact with a surface is far from trivial. In fact, Baraff (1993) showed that deciding whether a configuration with dynamic friction is consistent is *NP*-complete (in terms of the number of contact points). Erdmann (1994) introduces the generalized friction cone, which embeds the force constraints that define the Coulomb friction cone into the part's configuration space. The possible motions have a simple geometric interpretation with this representation. Another geometric approach to analyze multiple frictional contacts was proposed by Brost and Mason (1989). Their approach is limited to two dimensions (however, the configuration space is three-dimensional). It represents forces in a dual space as points. A friction cone is then reduced to a line segment in the dual space, and the dual of multiple friction cones is a convex polygon. Trinkle and Zeng (1995) developed a model to predict the quasistatic motion of a planar part in multiple contact. Their analysis yields inequalities that define regions in the space of friction coefficients for which a particular contact mode (i.e., sliding, rolling, separating or a combination thereof) is feasible. Related to this is the work of Wang and Mason (1987). They introduce an impact space, defined as all combinations of orientation and contact motion direction. Within this space one can analytically identify the areas that correspond to the different contact modes.

For rigid body collisions several models have been proposed. Many of these models are either too restrictive (e.g., Routh's model (1897) constrains the collision impulse too much) or allow physically impossible collisions (e.g., Whittaker's model (1944) can predict arbitrarily high increases of system kinetic energy). Recently, Chatterjee and Ruina (1998) proposed a new collision rule, which avoids many of these problems. Chatterjee introduced a new collision parameter (besides the coefficients of friction and restitution): the coefficient of *tangential* restitu-

tion. With this extra parameter a large part of the allowable collision impulse space 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.

Instead of having algebraic laws, one could also try to model object interactions during impact. This approach is, for instance, taken by Bhatt and Koechling (1995a,b), who modeled impacts as a flow problem. While this might lead to more accurate predictions, it is obviously computationally more expensive. Also, in order 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. Most of the complexity in dynamic simulation is due to collision detection. Using a particular quaternion representation for orientation, Canny (1986) reduces the problem of finding the distance between polyhedrons to finding the distance between a point and a number of hyperplanes in 7 dimensions. Lin and Canny (1991) designed a fast algorithm to incrementally find the closest point between two convex polyhedra. In cases where there are a large number of collisions or with contact modes that change frequently one can simulate the dynamics using so-called impulse-based simulation (Mirtich and Canny, 1995). However, there are limits to what 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 impose 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 are analyzed. That is, Krishnasamy discusses conditions for a part to "get trapped" and "stay trapped" in an extrusion in the context of the APOS parts feeder.

In (Lynch et al., 1998) the optimal manipulator shape and motion are 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. Examples of these constraints are: (1) the part cannot break contact, and (2) the part must always be rolling. Although the analysis focuses mainly on the juggling task, it shows that one can simulate and optimize dynamic manipulation tasks using a suitable parametrization of manipulator (or surface) shape and motion.

	Stable Poses						Entropy
	1	2	3	4	5	6	
quasistatic approximation	0.20	0.13	0.16	0.21	0.14	0.16	1.78
dynamic, flat surface, drop height is $h = 0$	0.18	0.16	0.14	0.34	0.05	0.13	1.66
dynamic, 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 rotation is uniformly random.

### 3 Example

Let us continue the example of section 1 to illustrate the general idea before we go into more detail. 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,  $x_0$ .

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 optimal 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 entropy<sup>2</sup>. The table makes it clear that even with a very simple surface we can reduce the uncertainty greatly by taking advantage of the dynamics.

### 4 Analytic Results

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*. (Quasi-capture regions will be defined more precisely later.) Specifically, for a part in contact with a sloped surface, we would like to deter-

<sup>2</sup>This is a local minimum found with simulated annealing and might not be the global minimum.

mine whether it is captured, i.e., whether the part will converge to the nearest stable orientation. 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 orientation. 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*.

So far 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 as a disk with radius  $R$  and uniform mass, but with contact points only at the ends of the rod:

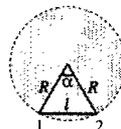


Figure 2: A rod with an offset center of mass.

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 end points of the rod. 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 the contact point. 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. We say the rod is captured if its kinetic energy is less than this amount:

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

Figure 3 illustrates this.

For a surface that is not horizontal the capture condition is not that simple. By bouncing and rolling down the slope,



Figure 3: Capture condition for a rod on a horizontal surface.

the rod can increase its kinetic energy. 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. In what follows the surface is still assumed to be flat but now sloped at an angle  $\phi$ :

What we will show below 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 captured. This bound depends on the current orientation, the current velocity, the slope of the surface and the geometry of the rod. Because of the way we have set up our generalized coordinates, the kinetic energy is  $\frac{1}{2}m\|\mathbf{v}\|^2$ . In other words, the mass is just a constant scalar (as opposed to a mass matrix for the more general case). Without loss of generality we can assume  $m = 2$ . That way the kinetic energy is simply  $\|\mathbf{v}\|^2$ .

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  $v_x$  and  $v_y$  are the translational components of the velocity and  $d_{\theta(t)}$  is a component that depends on the orientation. We can use this equation to derive an upper bound on the maximum increase of the kinetic energy. Let the rod be in contact at  $t = 0$  (so  $d(0) = 0$ ). 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 kinetic energy is  $\Delta\|\mathbf{v}\|^2 = 2g\Delta h$ . To find the maximum  $\Delta\|\mathbf{v}\|^2$  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  $\xi$ . This ignores the rotational component of the velocity, but it can be shown that by assuming the rod can rotate to an 'ideal' orientation this approach will give an upper bound for the true maximal increase of the kinetic energy. In this ideal orientation the rod is parallel to the surface and the center of mass is below the contact points. Let  $\hat{\theta}$  be equal to endpoint 1's relative orientation in the ideal orientation. One can verify that endpoint 1's relative orientation is then equal to  $\frac{\pi}{2} - \frac{\alpha}{2} - \phi$ .

It can also be shown that in order to *roll* to the other side the rod's kinetic energy needs to be larger than

$$-2gR(1 + \sin\theta + (\text{sign}(\cos\theta) - 1)\sin\frac{\alpha}{2}\sin\phi). \quad (3)$$

Here,  $\theta$  is the relative orientation of the contact point with respect to the center of mass. We assume  $0 \leq \phi < \frac{\pi}{2}$ . With these results we can now state our main result; an extension

and relaxation of the notion of capture regions. It is an extension in the sense that it applies to sloped surfaces and it is a relaxation, because capture is only guaranteed under some extra assumptions which will be discussed later.

**Definition 1.** Let a quasi-capture region be defined as a region in phase space such that the rod is in contact with the surface and cannot leave the region with one bounce or roll.

**Theorem 1.** 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:

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

where  $\xi$  is the direction that will result in the largest increase of kinetic energy,  $d_n = R(\cos\frac{\alpha}{2} - \sin(\theta + \phi))$  and  $\epsilon = \cos(\frac{\alpha}{2} + \phi) - \frac{\cos(\alpha/2)}{\cos\phi} + \max(\tan\phi, 2\sin\frac{\alpha}{2}\sin\phi)$ .

*Proof outline.* (For details see (Moll and Erdmann, 2000)). After one bounce the orientation is assumed to be in the ideal orientation, as this will result in the largest increase in kinetic energy. The translation of the center of mass in the direction normal to the surface is equal to  $d_n$ . This can be seen by noting that  $d_n$  is simply the difference between the (signed) distances to the surface in the initial configuration and a configuration with the ideal orientation. The path of the center of mass during the bounce is described by  $\frac{1}{2}gt^2 \cos\phi + v(\sin\xi \cos\phi + \cos\xi \sin\phi)t + d_n$ . We can solve this expression for  $t$  and use the smallest positive solution to find an upper bound on  $\Delta\|\mathbf{v}\|^2$ . Substituting  $\hat{\theta}$  in expression 3 gives  $-2gR(1 + \cos(\frac{\alpha}{2} + \phi))$ . In other words, if the kinetic energy after the bounce is less than  $-2gR(1 + \cos(\frac{\alpha}{2} + \phi))$  and the rod is in the ideal orientation, the rod cannot roll to the other side.

We can combine the two bounds to obtain a sufficient condition to determine whether the rod can rotate to 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. Currently we have a bound of the following form  $f(\hat{\theta}) \leq g(\hat{\theta})$ , where  $f(\cdot)$  computes the kinetic energy after one bounce for a given new orientation and  $g(\cdot)$  computes the energy needed to roll to the other side for a given orientation<sup>3</sup>. We would like to find the smallest possible  $\epsilon$  such that

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

It is not hard to see  $\epsilon$  has to be equal to  $\max_{\bar{\theta}}(g(\hat{\theta}) - g(\bar{\theta}) - f(\hat{\theta}) + f(\bar{\theta})) / (-2gR)$ . It can be shown that this correction

<sup>3</sup>So  $g(\theta)$  equals expression 3.

$\epsilon$  is equal to

$$\cos\left(\frac{\alpha}{2} + \phi\right) - \frac{\cos(\alpha/2)}{\cos\phi} + \max\left(\tan\phi, 2\sin\frac{\alpha}{2}\sin\phi\right)$$

Combining all the bounds we arrive at the desired result.  $\square$

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 for the case  $\phi = 0$ , i.e., for a horizontal table.

For an arbitrary  $d_n$  it is not possible to compute the optimal  $\xi$  analytically. Fortunately, we *can* analytically solve for  $\xi$  if we assume that the bounce consists of pure translation. The resulting  $\xi$  can be used as an approximation. The approximation for the bound for  $\Delta\|\mathbf{v}\|^2$  then simplifies to

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

The relative error in this approximation depends on  $\phi$ ,  $d_n$ ,  $v$  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  $\phi$ , but is still fairly small (on the order of  $10^{-2}$ ).

The theorem above shows 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 capture condition is true the dissipation of kinetic energy is larger than the 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 ranges from 0 to  $\frac{\pi}{2}$  and the orientation ranges from 0 to  $2\pi$ . Note that the orientation of the rod is not the same as the relative orientation of the contact point. However, for each combination of  $\phi$  and  $\theta$  the relative orientation of the contact point can be easily computed. The other relevant parameter values for this plot are:  $R = 1$ ,  $g = -9.81$  and  $\alpha = \frac{\pi}{2}$ . The little bump in the middle corresponds to the rod being captured on the high-energy side. The bigger bumps on the left and right correspond to being captured on the low-energy side.

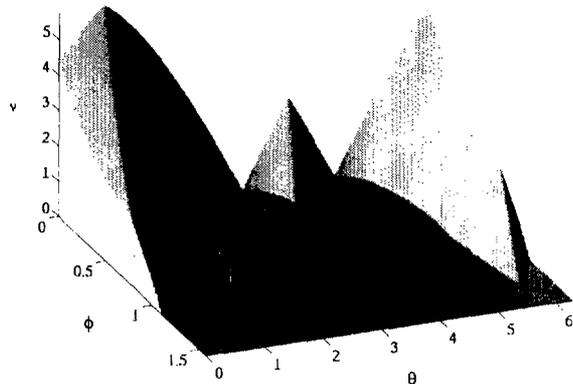


Figure 4: Quasi-capture velocity as a function of the slope of the surface and the orientation of the rod

## 5 Simulation and Experimental Results

We have written a simulator to numerically compute the pose distribution of a polygonal part in a dynamic environment. It uses the analytic results from the previous section to stop simulating the motion of the part once it is captured. The simulator allows us to search for the surface and drop height that reduce the entropy of the pose distribution of a given part maximally.

To verify the simulations we also performed some 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 can vary gravity. At the bottom of the slope is the surface on which the object will be dropped. The angle  $\phi$  of the surface in the plane defined by the air table can, 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 can slide under the surface and will only collide at the pins. Experimentally we determined the pose distribution of the rod for different values for  $g$ ,  $h$  and  $\phi$  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. Note that for a low drop height and a horizontal surface (row 13 in table 2) the pdf is equal to the 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  $\phi$  to appropriate values. In other words we can reduce the uncertainty almost completely.

The differences between the simulation and experimental results can be traced back to several different error sources.

	1	-0.684	0.246	0.467	0.934	0.96
*	2	-0.684	0.186	0.467	0.914	0.93
	3	-0.684	0.122	0.467	0.896	0.94
	4	-0.684	0.0580	0.467	0.850	0.94
	5	-1.53	0.0580	0.467	0.850	0.93
*	6	-1.53	0.122	0.467	0.896	0.92
	7	-1.53	0.186	0.467	0.912	0.97
	8	-1.53	0.246	0.467	0.934	0.97
*	9	-2.62	0.246	0.467	0.912	0.94
	10	-2.62	0.186	0.467	0.914	0.93
	11	-2.62	0.122	0.467	0.896	0.93
	12	-2.62	0.0580	0.467	0.852	0.94
	13	-2.62	0.0760	0.000	0.750	0.75
	14	-2.62	0.156	0.000	0.884	0.83
*	15	-2.62	0.220	0.000	0.918	0.85
	16	-2.62	0.284	0.000	0.874	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  $\phi$ . The drop height is measured from the center of the disk to the surface.

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.

## 6 Discussion

We have shown a sufficient condition on the position and velocity of the simplest possible 'interesting' shape (i.e., a rod with an offset center of mass) that guarantees convergence to the nearest stable pose 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 stable pose. We have coined the term quasi-capture regions for these regions, since they are very similar to Kriegman's notion of capture regions.

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 5 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.

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 (Moll and Erdmann, 2000) 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. We are also planning to do more experiments to verify our current and future analytic results.

## References

- Akella, S., Huang, W. H., Lynch, K. M., and Mason, M. T. (1997). Sensorless parts orienting with a one-joint manipulator. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*.
- 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 *Proc. 1996 IEEE Intl. Conf. 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 *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 2377–2382, Albuquerque, New Mexico.
- 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. (1997). Algorithms for sensorless manipulation using a vibrating surface. *Algorithmica*. Accepted for publication.
- 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. *Intl. J. 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.
- Brost, R. C. and Mason, M. T. (1989). Graphical analysis of planar rigid-body dynamics with multiple frictional contacts. In *Fifth International Symposium on Robotics Research*, pages

- 293–300.
- Bühler, M. and Koditschek, D. E. (1990). From stable to chaotic juggling: Theory, simulation, and experiments. In *Proc. 1990 IEEE Intl. Conf. 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.
- Canny, J. (1986). Collision detection for moving polyhedra. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 8(2):200–209.
- Chatterjee, A. and Ruina, A. L. (1998). A new algebraic rigid body collision law based on impulse space considerations. *ASME Journal of Applied Mechanics*. Accepted for publication.
- Chen, Y.-B. and Ierardi, D. (1995). The complexity of oblivious plans for orienting and distinguishing polygonal parts. *Algorithmica*, 14.
- Christiansen, A. D., Edwards, A. D., and Coello Coello, C. A. (1996). Automated design of part feeders using a genetic algorithm. In *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*, volume 1, pages 846–851.
- Erdmann, M. A. (1994). On a representation of friction in configuration space. *Intl. J. of Robotics Research*, 13(3):240–271.
- Erdmann, M. A. and Mason, M. T. (1988). An exploration of sensorless manipulation. *IEEE J. 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. 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. *J. 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. *J. of Dynamic Systems, Measurement, and Control*, 120:94–102.
- Gudmundsson, D. and Goldberg, K. (1997). Tuning robotic part feeder parameters to maximize throughput. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 2440–2445, Albuquerque, New Mexico.
- Hitakawa, H. (1988). Advanced parts orientation system has wide application. *Assembly Automation*, 8(3):147–150.
- Kavraki, L. E. (1997). Part orientation with programmable vector fields: Two stable equilibria for most parts. In *Proc. 1997 IEEE Intl. Conf. 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. *Intl. J. 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.
- Lin, M. C. and Canny, J. F. (1991). A fast algorithm for incremental distance calculation. In *Proc. 1991 IEEE Intl. Conf. on Robotics and Automation*, pages 1008–1014, Sacramento, CA.
- Luntz, J. E., Messner, W., and Choset, H. (1997). Parcel manipulation and dynamics with a distributed actuator array: The virtual vehicle. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, Albuquerque, New Mexico.
- Lynch, K. M., Shiroma, N., Arai, H., and Tanie, K. (1998). The roles of shape and motion in dynamic manipulation: The butterfly example. In *Proc. 1998 IEEE Intl. Conf. on Robotics and Automation*.
- Marigo, A., Chitour, Y., and Bicchi, A. (1997). Manipulation of polyhedral parts by rolling. In *Proc. 1997 IEEE Intl. Conf. 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*.
- Mirtich, B., Zhuang, Y., Goldberg, K., Craig, J., Zanutta, R., Carlisle, B., and Canny, J. (1996). Estimating pose statistics for robotic part feeders. In *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*.
- Moll, M. and Erdmann, M. A. (2000). Manipulation of pose distributions. Technical Report CMU-CS-00-111, Dept. of Computer Science, Carnegie Mellon University. <http://www.cs.cmu.edu/~mmoll/publications>.
- Peshkin, M. A. and Sanderson, A. C. (1988). Planning robotic manipulation strategies for sliding objects. *IEEE Trans. on Robotics and Automation*, 4(5).
- Rao, A., Kriegman, D., and Goldberg, K. (1995). Complete algorithms for reorienting polyhedral parts using a pivoting gripper. In *Proc. 1995 IEEE Intl. Conf. on Robotics and Automation*.
- Routh, E. J. (1897). *Dynamics of a System of Rigid Bodies*. MacMillan and Co., London, sixth edition.
- Trinkle, J. C. and Zeng, D. C. (1995). Prediction of the quasi-static planar motion of a contacted rigid body. *IEEE Trans. on Robotics and Automation*, 11(2):229–246.
- Wang, Y. and Mason, M. T. (1987). Modelling impact dynamics for robotic operations. In *Proc. 1987 IEEE Intl. Conf. on Robotics and Automation*, pages 678–685.
- 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 *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*.
- 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.