A Nonprehensile Method for Reliable Parts Orienting

Nina Barrameda Zumel

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, 15213

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Support for this research was provided in part through the following grants: NSF Grants IRI-9503648, IRI-9528176, IRI-9157643, IRI-9010686, IRI-9213993, and NASA Fellowship NGT-51272. The views and conclusions expressed in this document are those of the author, and do not necessarily reflect the views of NSF, NASA, or the U.S. Government.

Abstract

Prehension may be defined as "The act of taking hold, seizing or grasping, as with the hand" (Webster's 3rd International Dictionary). *Nonprehensile manipulation*, then, can be defined as the manipulation of objects without grasping them. Manipulation without prehension is a natural way of handling objects for both humans and machines. The ability to manipulate objects which may not be graspable increases the flexibility of a robot interacting with its environment, without adding complexity to the mechanical design. This research analyses the mechanics of nonprehensile contact between a simple, two degree of freedom manipulator and a part. The intent is to develop reliable but sensorless manipulation routines for use in an automated assembly environment. While nonprehensile, sensorless devices are in common use in such environments, existing parts orienting devices, such as bowl feeders or the SONY Automatic Parts Orienting System, must be custom designed for each specific task. To decrease the setup or changeover time for an assembly line, what is needed is a simple but more general device, which can be easily modified or reprogrammed in response to a change in tasks. We present a planning algorithm for sensorless parts orienting in the plane with two one degree of freedom palms. Our method finds feasible paths through the space of equivalent state configurations of the object in the palms, without requiring that the palms maintain stable support of the object over the entire path. We show that such a device can reliably orient parts in the plane. Planning reorientations requires the geometric descriptions of the parts, the part's center of mass, and an upper bound on the coefficient of friction between the part and the palms. The plans produced by our algorithm are robust to uncertainties in the part's initial state and in the coefficient of friction, as well as to small inaccuracies in manipulator calibration.

Acknowledgements

I would like to thank my advisor, Mike Erdmann, for his support, and for all the helpful discussions, some even about research. I would like to thank my parents for having supported me and encouraged me throughout my education. To the rest of my committee: Matt Mason, David Baraff, and Randy Brost, thank you for many useful comments on my dissertation, and on my work. Thanks to the members of the Manipulation Lab, past and present, especially Kevin Lynch, Srinivas Akella, Tammy Abell, Yan-Bin Jia, Wes Huang, Garth Zeglin, and Mark Moll.

My graduate school experience wouldn't have been the same without Tammy, Lin Chase, and Dina Berkowitz — the other three Robobabes. Nor without Garth, Wes, Chris Lee, Michael Witbrock, Tony Brusseau, Scott and Leanne Neal Reilly, John and Mary Mark Ockerbloom, Rich Goodwin, Susan Hinrichs and Alan Carroll.

Finally, thanks to my husband, John Mount, for his patience and his support throughout my years in graduate school.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> As for those things which are moved by other things, this must take place in four ways; for locomotions caused by other things are four in kind (pulling, pushing, carrying, and turning)...
>
> Aristotle
> *Physics*, Book VII [9]

1.1 Why Nonprehensile Manipulation?

Prehension may be defined as "The act of taking hold, seizing or grasping, as with the hand" (Webster's 3rd International Dictionary). *Nonprehensile manipulation*, then, can be defined as the manipulation of objects without grasping them. Manipulation without prehension is a natural way of handling objects for both humans and machines. The ability to manipulate objects which may not be graspable increases the flexibility of a robot interacting with its environment, without adding complexity to the mechanical design.

Nonprehensile manipulation also forms an intermediate point in the range of manipulation operations. At one end of this range we have operations for which complete control is necessary or desired over the complete duration of the operation: either complete control over the object, or complete control over the object's interactions with the environment. Both *power grasps* and *precision grasps* fall into this category. Power grasps are force/form closure grasps. We will follow Nguyen [65] in our definitions of force and form closure.

Definition 1.1 (Force-closure) *A* force-closure *grasp is defined as a grasp such that an arbitrary force and moment can be exerted on the grasped object through the contacts. Equivalently, any motion of the object can be resisted by contact forces.*

Definition 1.2 (Form-closure) Form-closure *grasps are grasps such that the object is totally constrained by the contacts irrespective of the contact force magnitudes.*

In a power grasp, the object's degrees of freedom are completely constrained throughout the manipulation. As a result, the object's motion is insensitive to environmental dynamics, and the problem of controlling the object's motion is equivalent to controlling the manipulator. Precision grasps are different, of course. Take robotic surgery as an example. One would not want the scalpel to be insensitive to its environment. Still, a robotic surgeon must have precise control over the motion of its scalpel, to avoid sensed obstacles, and perhaps some compliance to those parts of the environment it contacts. In other words, the scalpel must be completely controlled throughout the manipulations.

At the other end of the spectrum we have free flight motions of thrown or batted objects: the object's degrees of freedom are only constrained at the beginning of (or at intermittent periods during) its trajectory, and no further control effort is required beyond the initial (or intermittent) forces. The subsequent motion, therefore, is extremely sensitive to environmental dynamics and initial conditions. Control of object motion is then linked to precise control of the manipulator/object interaction over short periods of time. Some examples are tossing or dropping things in a waste paper basket, kicking a door shut, or playing ping-pong.

In between these extremes lies a large class of operations during which complete constraint over the object to be manipulated is either undesirable or impractical, but some control over the object is desired over its entire trajectory, in order to bring the object reliably to a desired final state. Moving objects which are too large to be grasped, but must be pushed or rolled to a desired location, is one example. Other examples include shutting a drawer, or pushing a door shut without slamming it. Yet another example is transferring a portion from a large bag of dried beans into a relatively narrow-mouthed jar. One way to do this is to individually pick and place beans from the bag to the jar until the jar is full. This is an example of the first type of manipulation, form closure grasp manipulation, and it is reliable, but in this particular case, quite tedious and stupid. Another way is to scoop handfuls (or cupfuls) of beans and pour them into the jar. This works. An even better way is to find an appropriate type of funnel, and pour beans from the bag into the funnel until the jar is full. With respect to the beans (rather than the bag), this is a member of the third type of manipulation. There is some control over the trajectory of the beans: the funnel keeps them from spilling all over the floor, and directs them reliably into the jar. But the precise trajectory of each individual bean isn't precisely known or controlled. It is a solution which is fast and simple to execute. Another example is coin sorting. One way to sort coins is to pick each coin up individually, determine the denomination of the coin, and place it into the appropriate stack. A better way is to dump all the coins into a device which can mechanically filter and sort the coins by their size and weight. We claim that nonprehensile manipulation is appropriate for exactly these type of tasks. Nonprehensile manipulation is also appropriate in many situations where, rather than force-closure, it is sufficient to ensure that the contact forces resist a specific external wrench. We will call this *stable support* [1]:

Definition 1.3 (Stable support) *An object is* stably supported *against gravity (or more generally, any known applied wrench) if the contact configuration balances the gravitational force and the object can resist small perturbations in its pose.*

By resisting small perturbations in its pose, we mean that if the object is subjected to a small displacement, it will return to its original pose. A stably supported pose is a stable equilibrium state of the system.

1.2 Why two palm manipulation?

The examples of nonprehensile manipulation given above range from very large ungraspable objects at one end, to many small objects at the other. This work concerns itself with attempting to reliably and quickly orient small objects, although some results may be useful in other domains. Directions of interest include (i) controlling the shape of constraint surfaces of systems in such a way that constraint and external forces naturally attract the system to the desired state, and (ii) exploiting the system dynamics in a similar fashion, much as in the research on underactuated manipulators. *Palmar manipulation* offers a simple domain in which to explore these two directions, although we will mainly concentrate on the first.

Definition 1.4 (Palmar Manipulation) Palmar manipulation *is the manipulation of an object by its interactions with flat, nearly rigid manipulators.*

Palmar manipulation is somewhat analogous to manipulating an object with the palm of a hand, without using the fingers.

Some direct "beans in a funnel" style examples of manipulation are bowl feeders and Automatic Parts Orienting Systems [14], [40], where a large mass of parts in arbitrary orientations are singulated and oriented by their interactions with (in the case of bowl feeders) fences and other obstacles, or (in the case of APOS systems) by an induced vibration and their interaction with pallets of special shapes. As an example, let us imagine such a parts feeder, where parts enter from one end in random orientations, and emerge at the bottom in a single orientation (Figure 1.1). Mani and Wilson [57] designed a programmable parts orienting table where a part in arbitrary and unknown orientation is oriented (without sensors) by the motion of the table, which moves the part into contact with a stationary fence. Peshkin and others have studied the case of parts along a conveyer belt, with stationary fences along the way to do the orientation [68], [20]. One could also design a feeder where a moving part comes into contact with moving fences, to perform the same operation that a string of stationary fences would perform serially. That string of stationary fences would be replaced by some smaller number of moving fences. Akella, *et.al.* [4] have begun work on such a system, the One Joint Over a Conveyer (1 JOC) system, which can arbitrarily orient parts on a moving conveyer belt, using a one degree of freedom arm.

The resulting design can be more compact than a conveyer with stationary fences. It can also be as potentially versatile as the parts orienting table, in that the resulting mechanism can be reprogrammed for other parts, rather than rebuilding a new set of stationary fences.

For the above example task, and for other tasks for which nonprehensile manipulation is appropriate, we would like the final state of the system to be precisely determined, but precise control of intermediate states is less important. We can take advantage of this fact. We will show, using an analysis similar to that described by Trinkle, *et.al.* [88],[90] that the state space of the manipulator and object can be divided into regions which are equivalent, in the sense that a particular state is easily reachable from one of its equivalent states. The manipulation of an object is then a higher level (and coarser) problem of navigating from region to region. When the region containing the goal is achieved, then more careful motions can be used to fine-tune the position of the object. Hence control of the manipulators during the coarse motion need only be precise enough to make the regional transitions correctly, and the object will be propelled along a correct path to the goal.

For the rest of the thesis, we will use a manipulator consisting of two one degree of freedom palms modeled as a cone connected at a central hinge. This is a simple example that still captures some of the basic operations of a general two palm system. For this particular example, we will show that two basic operations into which manipulations can be decomposed fall naturally

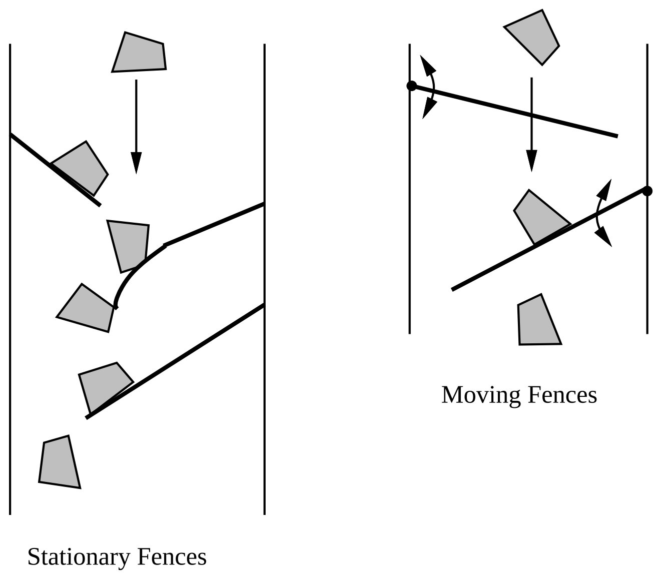

Figure 1.1: Stationary and Moving Parts Feeders

out of the equivalence region analysis described above. We believe that similar decompositions will fall out of the analysis of other types of two palm systems.

When we move beyond the quasistatic domain, inertial forces and collision forces begin to matter. Coriolis and centrifugal forces must also be taken into account, so the object does not fly out of the palms, or out of the desired range of positions during the manipulation. For example, one would like to have the object make contact with the palms in such a way that the object does not bounce away from the expected stable position. Although we will not specifically derive velocity constraints, we will give qualitative examples of the sort of velocity effects most common to this system.

1.3 Related Work

1.3.1 Fine-motion Planning

Our work and much of the sensorless work presented below is based on the framework developed by Lozano-Pérez, *et.al.*[52], for synthesizing compliant, fine-motion strategies for assembly tasks. This framework in turn based its analysis on the *configuration space* [51] approach for analysing geometric motion planning and collision avoidance problems. In configuration space, the pose and motion of an object with N degrees of freedom is represented as the motion of a point in N dimensional space. Although the framework in [52] was originally designed for sensor-based robotic assembly (with uncertainty), it has been commonly used also for sensorless tasks which involve contact between a manipulator and a part, or a part and obstacles. In particular, the idea of *preimages*, that is, regions of a part's configuration and state space for which a commanded velocity (including some bounded imprecision in the execution of the commanded velocity) is guaranteed to reach a goal, has proven to be a useful way of describing sensorless parts orienting strategies.

1.3.2 Prehensile Manipulation and Force-Closure Grasps

The research in prehensile grasping and force-closure grasping is far too extensive to be surveyed here. We cite a few examples which take approaches which we have felt to be in some way similar in flavor to ours.

Trinkle and Paul [89] analyse the interaction of fingers and objects for the purpose of finding squeeze grasps to lift objects, both in the frictionless and frictional cases. They partition the various configurations of contact between the object and fingers into regions where the object will either slide out of the grasp, jam, or be lifted off the support surface when a squeezing force is applied. Kao and Cutkosky [42] study the object-finger interaction for the case of a stable (either stably grasped or stably supported) object and fingers in sliding compliant contact with the object. The emphasis is on determining the relative trajectory and forces between the object and finger during the finger motion. Aiyama, Inaba, and Inoue [2] describe manipulation by *pivoting*, whereby an object is positioned without being fully grasped, but forces are applied in such a way as to "walk" the object into the goal pose. Because the moving contacts are now point contacts, resistance forces due to friction are diminished.

Brock [18] specifically allows slip in a controlled manner to reorient an object in a three fingered dextrous hand. When the fingers are in general position, the three contacts points define a plane, the *grasp plane*. Define the *force focus* to be the point in space through which the three contact forces must pass. This point lies somewhere on the grasp plane, possibly at a point at infinity. If the three contact points are on a line, then, according to Brock's formulation, the force focus is somewhere on that line. For a given object and a given coefficient of friction, one can generate a *constraint state map*, which describes regions in the grasp plane for which the motion of the object has the same constraints and freedoms. One can then cause a desired motion to the object by moving the force focus (varying the forces exerted on the object by each finger) into a region where the constraints on the object will cause the desired motion to happen. The Salisbury Hand, which was used in Brock's experiments [19], had a tactile sensor which directly measured contact location on the fingertips, as well as the normal and tangential forces at the contact. In addition, Brock directly inferred the contact mode (sticking vs. slipping), by looking for high frequency components in the force readings; by comparing the mean values of normal and tangential force, he could also infer whether the presumed slipping was rotational or translational. There was, however, no feedback on the contact force, and no direct information on the object's absolute position in the hand. Yoshikawa, *et.al.* [96] also study the use of controlled slip by a three-fingered hand. They determine a unique force that will cause a finger to slip in a desired direction, while maintaining the stability of the object in the grasp.

Ponce and Faverjon [71] study three finger, force-closure frictional grasps. They develop a conservative method of calculating grasps: *independent contact regions*. In an independent contact region, any triple of contact points such that there is one contact in each region will give force-closure. These contact regions are a generalization of Ngyuen's [65] formulation for two finger frictional force-closure. Ponce and Faverjon's formulation uses a minimum value of the coefficient of friction, μ_{min}, such that the calculated grasps are force-closure for any coefficient of friction $\mu > \mu_{min}$.

Ponce also develops a four-point frictionless grasp planner for the purposes of parts fixturing [72]. This approach is based on the idea of *second order immobility*, due to Rimon and Burdick, [78] rather than on the traditional force-form closure approach. Second order immobility describes grasps which prevent all finite motion of an object, due to curvature effects, despite the fact that differential analysis may not identify the grasp as force or form closure. Fewer frictionless contacts are required for second order immobility than for complete force or form closure.

Brost and Goldberg [23] describe an algorithm for designing planar fixtures from a modular fixture kit. These fixtures are four-point frictionless form-closure, and can be selected based on arbitrary criteria for optimality. Wallack [92] presents two fixturing algorithms, one which completely enumerates all feasible fixtures for a given part and fixture kit, and a more efficient heuristic approximation to the enumeration algorithm.

Rimon and Blake [77] present the *caging problem*: to surround an object with a (frictionless) multifingered hand in such a way that, although the object can still move, it cannot escape the "cage" of the fingers. The configuration of the hand is described as a one-parameter set in such a way that as the parameter is decreased, the fingers approach the part, and ultimately grasp it. This is similar to the work by Brost [21], in that uncertainty in the part's position can be compensated for by an appropriate intial cage, and that uncertainty can be reduced as the cage parameter is decreased until the object is grasped.

Abell and Erdmann [1] use two finger stable support to plan handoffs of an object between sets of frictionless point fingers.

1.3.3 Nonprehensile Manipulation

One of the earliest examples of robotic nonprehensile manipulation is quasistatic pushing in the presence of friction ([58], [21], [53], [69], [20], [56]). In particular, Peshkin [69] works out the minimum distance which a "fence" (or palm) must push an object so that it will be guaranteed to come to rest aligned with the fence at the end of the motion. Goyal [38], [39] shows that the motion of an object on a frictional support surface can be determined if the pressure distribution of the object is known. Peshkin and Sanderson [69] and Lynch [53] analyze this situation when the pressure distribution is not known. Mani and Wilson, Peshkin, and others ([57], [68], [20], [95], [4], [3]) use the mechanics of pushing to design parts orienters or parts filters with a sequence of fences, similar in function to the vibratory bowl feeders and other orienting systems described by Boothroyd, *et.al.* [14].

When pushing a planar lamina, the gravitational and support forces are perpendicular to the plane of interest. Bat juggling [25], [49] may also be considered as palmar manipulation of one (or more) objects, where now gravitational forces also have a component in the plane of interest.

Bicchi, *et.al.* [12], present mobility and manipulatibility analyses for general multi-limb manipulators, including nonprehensile cooperating robots. For a given object/manipulator pose, under the assumption of a force closure grasp, they present, among other things, a way to partition the manipulator joint space into regions of differential motions with different qualitative behavior. Some of the regions identified could include regions where differential motion affects manipulator pose but not object pose, regions where different freedoms of the part remain unconstrained, or regions where there is one-to-one correspondence between manipulator and object velocity. They can also identify allowable and unachievable object motions.

Yun [97] studies two manipulators with open palm end effectors manipulating an object by having one palm push from one end, and another palm push at the other end. The desired effect is to push with both palms hard enough so frictional forces counteract gravitational forces, but not so hard as to damage the object. Coordinated pushing is then used to maneuver the object as desired. This work was extended in Paljug, *et.al.* [66], [67]. Here, two planar palms are used to manipulate large objects in free space. A primary difference between the work of Paljug and Yun and the work described in this dissertation is the emphasis in the former on force closure. Rolling contacts between the palms and the object are permitted, but the contacts are not allowed to slip or break. Hence, direct control over the object state must be maintained at all times. In our present work, however, the object state is not directly controlled at all times, but merely guided towards the desired goal. This type of manipulation, and indeed much of the work on pushing may be considered *passive* manipulation, in contrast to the more *active* manipulation explored in the work of Paljug, *et.al.*, and in much research on grasping.

Examples of passive manipulation include the manipulation techniques of Trinkle, *et.al.* ([90], [36], [85], [86]). Their analysis and planning uses the idea of *contact formations* originally presented by Desai, and incorporated into a planner for dextrous manipulation by Trinkle and Hunter [88]. The work described in this dissertation follows a similar method to Trinkle, *et.al.*. However, the paths through configurations space by [90] and [88] are apparently constrained to always correspond to stable grasps, whereas the method described in our work allows very small unstable motions during the transition from state to state.

Of particular interest in the above set of work is Farahat and Trinkle [36], as well as Trinkle and Zeng [91]. Farahat and Trinkle attempt to find upper bounds on the coefficient of friction for which the plans found by a planner using the frictionless assumption will still work – in other words, all contacts will slide. Trinkle and Zeng treat the coefficient of friction as a variable in the analysis performed by the planner, and determines the ranges of the coefficient of friction for which a particular contact mode will be feasible. These considerations are important, since the coefficient of friction cannot in general be precisely known.

Erdmann [33] is also exploring the use of friction and compliance in nonprehensile palmar manipulation. He studies "slide transfers" of objects from a palm to another planar surface (such as another palm), and the different behaviors of the object for different velocities and "transfer angles" of the palm.

A related concern is that of determining stable poses of assemblies in a gravitational field, when those assemblies are not rigidly attached. In this case, we will say that the assembly is *stably supported* [1] against gravity. The problem, both with and without friction, is addressed by Mattikalli, *et.al.* [61], [62]. They show that for a fixed assembly (to make an analogy with our work: for palms at a fixed relative angle to each other, and a part in a particular equilibrium position in the palms), the orientations of the assembly for which stable support is maintained form a convex region on the surface of a unit sphere [62]. We will take advantage of this result in generating our reorientation plans.

Of particular interest is the effect of impact dynamics in the manipulation domain. Most previous strategies for planning manipulation tasks have relied on an assumption of quasi-static mechanics in the analysis of the physical system. This constrains the plans to situations that are slow moving, and in which contact dynamics can be neglected.

One can imagine situations where one cannot make these assumptions, or when a model of the contact dynamics would be useful. In making contact with an already moving or accelerating object, for instance, the inertial properties of the object affect the motion which results from the applied forces of collision. Knowledge of the magnitudes as well as the direction of forces and velocities becomes important. Juggling and table tennis are two such domains that have been explored in robotics.

Mason and Lynch [59], [60] explore the use of dynamic properties for controlled club throwing. In [54], Lynch analyzes the conditions for desired goal states to be reachable, both for pushing and for dynamic manipulation. In particular, he is interested in the conditions under which a one degree-of-freedom manipulator can take a part to a full-dimensional subspace of goal states. Arai and Khatib [7] studied the exploitation of inertial properties of objects for manipulation of objects in ways other than throwing. They base their approach on work by Arai and Tachi [8] in using dynamic coupling to control passive joints.

Some work has been done to study models of dynamic collision for use in robotic domains. Raibert [73] designed a dynamically stable hopping robot, modelling the bounce as a spring and damper system with perfectly inelastic collision. This system was further analyzed by Koditschek and Buhler [49]. Andersson [6] designed a ping-pong playing robot which used a simple model of point-mass collision to predict the motion of the ball after striking. Wang [93] attempts to characterize the qualitative behavior change in the motion of objects upon collision. In [94], he presents simulations and analysis of systems with intermittent constraints, and uses models of those systems in planning manipulation tasks. In both cases, a rigid body model of impact is used. This model has the advantage of being relatively simple, although other researchers have shown inconsistencies do exist, particularly if one looks at energy dissipation ([82]). Buhler [25] analyses a planar puck juggling system. This system was extended to three dimension by Rizzi [79]. Planar bat juggling for polygonal objects was examined in [98].

1.3.4 Parts Orienting

The work by Brost [21] on the orientation of objects by "squeeze grasps" is analogous to palmar manipulation of one object with two hands, with gravity perpendicular to the plane. He finds sets of actions which reliably orient a part in the presence of uncertainty in the part's location. Goldberg, and later Rao and Goldberg ([37], [74]) found algorithms for determining sequences of squeezes of a parallel jaw gripper which will reliably orient (up to symmetry) frictionless, polygonal or algebraic, planar parts from an arbitrary and unknown initial orientation, without sensors. Mason and Erdmann [34] use gravity to propel parts onto a flat surface, or into a corner formed by two perpendicular flat surfaces, in such a way that the resulting contact forces reliably orient a part. A three-dimensional generalization of this planner was presented in [35]. The computational complexity of a simplified version of this tray tilter, and of other parts feeding systems, was analyzed by Natarajan [64].

Examples of nonprehensile parts orienting systems include bowl feeders and Automatic Parts Orienting Systems [14], [40], where a large mass of parts in arbitrary orientations are singulated and oriented by their interactions with (in the case of bowl feeders) fences and other obstacles, or (in the case of APOS systems) by an induced vibration and their interaction with pallets of special shapes. Krishnasamy, *et.al.* have begun the analysis of APOS pallets, with the goal of automating their design [48]. Peshkin and others have studied the case of parts along a conveyer belt, with stationary fences along the way to orient them [68], [20]. Mani and Wilson [57] designed a programmable parts orienting table where a part in an arbitrary and unknown orientation is oriented (without sensors) by the motion of the table, which moves the part into contact with a stationary fence. Akella, *et.al.* [4], [3] have begun work on the One Joint Over a Conveyer (1 JOC) system, which can, without sensors, orient parts with arbitrary initial state on a moving conveyer belt, using a one degree of freedom arm.

Caine [26] developed a graphical tool which allows a designer of a bowl feeder track, or other parts orienting/fixturing device, to directly manipulate the configuration space obstacles representing the part in the track. Although the inverse transformation from configuration space to real space is not always unique, Caine's system performs an online differential modification of a given real space system in response to the designer's manipulations in configuration space. Joskowicz and Sacks [41] present an algorithm for the qualitative analysis of mechanisms. The qualitative state of a system is described in terms of positions, velocities, and contacts between parts, and by qualititive descriptions such as "engaged" or "disengaged". Their algorithm identifies qualitatively equivalent regions of the system state space, and the transitions between regions. Hence, one can describe the behavior of a system for given initial conditions as a sequence of transitions from one qualitative state to another.

1.3.5 Capture Regions

Parts orienting algorithms, particularly the sensorless sort as described above, take advantage of the fact that stable object poses generally form a discrete set, thus constraining the search over the state space. Boothroyd, *et.al.* [15] empirically derive the distributions of the stable resting aspects of a part, for use in designing bowl feeders. Mirtich, *et.al.* [63] present variations of quasi-static estimation methods for estimating the pose statistics of dropped objects, as well as the use of dynamic simulation to directly simulate the pose statistics.

In his dissertation, Brost [22] describes the notion of a *half-space invariant*. In a system with a half-space invariant, the possible future configurations of a point originally at rest are constrained to lie in some half plane whose boundary goes through the original configuration of the system in configuration space. A variety of physical systems satisfy this condition, including an object falling in a gravitational field. He describes an algorithm to determine preimages of goals for such systems in configuration space.

Kriegman uses potential energy minima to constrain the possible poses of a 3D object resting on a flat surface, this reducing the search space for image recognition algorithms [46]. He later extended this work to the problem of determining *capture regions* of dropped 3D objects. If a part is placed on a plane perpendicular to gravity, in the absence of any initial kinetic energy the part will fall to a unique resting position if its initial pose was in the capture region of this resting position. In [47] he shows how to determine the capture regions of smooth objects. In [75], knowledge of these capture regions is incorporated into a parts feeding system. In our work, we will also use the idea of capture regions to constrain the poses of a part resting in two palms, rather than on a single plane. We will show how moving the palms can move this part from one stable pose into the capture region of another pose.

Böhringer, *et.al.* [13] study the design of programmable vector fields, implemented on arrays of microelectromechanical vibrating actuators, or on macroscopic vibrating plates, to orient or filter parts. They use the vibratory motion to design "potential fields" such that the part is drawn into a local minimum in the appropriate configuration.

1.4 Outline of the Thesis

Chapter 2 will review mathematical tools necessary for the analysis of a nonprehensile manipulation system, and for designing a planner.

Chapter 3 will focus on finding plans for orienting objects reliably to known goal states from known initial states. All reorientation plans can be decomposed into two basic types of manipulator motions. If a certain operation is feasible, all polygonal parts can be oriented: that is, they can be brought from any stable initial state to any stable goal state. Even when this operation is not feasible, we can find conditions under which a polygonal part can be oriented. We show that the plans found are repeatable, and robust to small errors in manipulator calibration.

Chapter 4 will focus on relaxing the requirement that all contacts slide, in order to extend the model to systems with higher coefficients of contact friction. Chapter 5 will present an extension of the basic algorithm to the case of orienting a part from an unknown initial state. The space of all possible manipulator motions can be reduced to a finite set of motions, such that a part can be oriented to a given final state if and only if a plan can be found using this set of motions. Again, these plans are repeatable and robust to calibration errors. Using the considerations of Chapter 4, the plans are also viable for any coefficient of friction between the part and the palms, up to the upper bound estimate of friction used to find the plans.

Chapter 6 will discuss ways of avoiding unpredicted effects due to higher kinetic energy. Chapter 7 will present theoretical results on the algorithmic complexity of the planner, as well as proofs of many of the assertions made in the previous chapters. Finally, Chapter 8 will present conclusions and directions for future work.

Chapter 2

Analysis Tools

> Geometry furnished God with models for the Creation and was implanted in man, together with God's own likeness.
>
> Johannes Kepler
> *Harmonices Mundi*, 1619 [44]

In the previous chapter, we discussed briefly the distinction between what we called active and passive methods of manipulation. *Active methods* are those methods which actively seek to completely control the object at all times throughout the course of a manipulation. *Passive methods*, in contrast, try to exploit mechanical constraints in order to achieve the desired goal, without directly controlling all the degrees of freedom of the object. To maintain force closure is an active concept; to maintain stable support is a passive concept. Sensorless techniques are passive, and sensor based techniques tend to be active. Active methods of manipulation — and, in fact, any closed-loop robotic task, such as mobile robot navigation — require a different style of analysis from passive, open-loop methods, such as those used in sensorless nonprehensile manipulation. Much of the work on pushing and parts orienting may be considered passive manipulation.

With closed-loop methods, sensory information is available, and the instantaneous state of the object is known, at least approximately. Hence techniques which are local or differential in nature are entirely appropriate. Control theoretical methods, or any form of analysis which is based on the (possibly implicit) integration of differential equations fall into this category. Some global properties, such as the asymptotic stability of goal states, or the controllability of the system, are considered in designing the control system. However, the act of manipulation in these cases is essentially that of reacting to the present state of the system, differentially. Furthermore, this style of thinking about the problem naturally leads to the desire to pick a specific trajectory through the state space which the object (not just the manipulator) is obliged to follow exactly on its way to the goal.

Sensorless techniques, in particular passive nonprehensile manipulation, require a different style of analysis. Since moment to moment information about the system state is not available, the analysis techniques are necessarily global and more qualitative in nature. Concepts such as regions of stability or regions of convergence become the only tools available. The evolution of the system state is expressed not in terms of velocity or acceleration, but in terms of directions of motion and lines of force, in a magnitude independent way. The idea of having the object follow an exact trajectory is often less feasible in the passive nonprehensile case.

In principle, one could try to integrate out the differential equations corresponding to all possible scenarios and trajectories. In practice, it is better to try to represent these sets of differential equations as some aggregate geometric entity, which can be reasoned about as a whole. One advantage of viewing the system state space in this inherently global way is that robustness issues automatically come under consideration. Because we have no way of detecting and correcting errors when they occur during a task, we must try to be robust in the face of uncertainty: uncertainty about the system state, errors in manipulator calibration, uncertainty about system geometry and parameters. We have not covered all of these issues explicitly in this thesis. We have covered a few, and have discovered that a few more have come along, "for free." We achieve this robustness because in navigating around in the state of our global geometrical constructs, we can pick trajectories in the interior of "good regions." In this way we avoid singularities and borderline cases. We cannot force the object to follow a predetermined trajectory, but by picking the global constraints correctly (in our case, the direction of the ambient potential field), we can still guide the object to a final desired goal in a reliable way.

In this chapter, we will briefly review some geometric constructs which have proven to be useful for sensorless manipulation. These constructs may also be useful for the analysis of closed loop systems, just as traditional dynamical systems analysis can also be useful in open loop domains. We present them here as having been particularly useful in solving the type of problem presented in this thesis. We will discuss Configuration Space, Friction Cones, Centers of Rotation, and Cone Analysis. We also include a section on Impact Dynamics.

2.1 Configuration Space

Configuration Space as an analysis tool for robotics was first presented by Lozano-Pérez [51]. Using configuration space, the problem of determining the motion and interaction of rigid bodies can be reduced to the problem of determining the motion of a point body in a higher dimensional space. While configuration space analysis is primarily used for motion planning, it can also be used for mechanical analysis, much in the manner of the generalized coordinates of Classical Mechanics.

The configuration space (cspace) of an object has the same dimensionality as the object's degrees of freedom. A planar object, therefore, generally has a three dimensional configuration space, and a three-dimensional object has a six dimensional configuration space. Since this work concerns itself with planar objects, we will discuss three dimensional configuration space, although most of what we say will generalize.

2.1.1 Configuration Space Obstacles

An object moving freely in the plane is equivalent to a point moving about in three dimensional cspace, where the three coordinates of the cspace correspond to the horizontal (x), vertical (y), and rotational (θ) location of the object in the plane. Let the first two components, (x, y), describe the motion of the center of gravity (CG) of the planar object. The third component represents the object's orientation in radians, but in order to make the cspace of the object independent of the measurement units of the planar system, it is normalized by multipying by the object's radius of gyration, ρ. Imagine a planar object, and a planar ring of the same mass, infinitesimally thin, with uniform mass density over the ring. The moment of inertia of the planar object about its CG is I_{obj}, that of the ring is $m\rho^2$, where ρ is the radius of the ring and

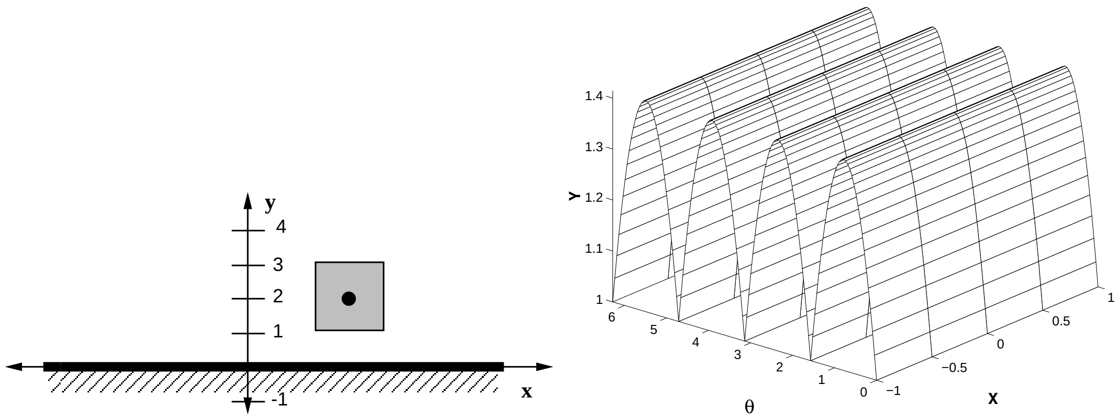

Figure 2.1: left: A square with sides 2 units long in the plane, with a wall at $y = 0$ as an obstacle. right: The corresponding configuration space obstacle.

m is the mass of both the ring and the object. If the object and the ring have the same moment of inertia, then the radius of the ring is the radius of gyration. In other words, $I_{obj} = m\rho^2$. Using these normalized coordinates, the position of an object is described in configuration space by the vector

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ \rho\theta \end{bmatrix}.$$

An object moving in the plane among stationary obstacles is equivalent to a point mass moving through three dimensional cspace among stationary three-dimensional obstacles. The boundary of the cspace obstacles is defined by the positions and orientations of the planar object where it makes contact with the planar obstacles. See [50] for an excellent description of constructing cspace obstacles.

For example, consider a square with sides of length 2 in the plane with an infinite wall along the $y = 0$ axis as an obstacle. The center of gravity of the square is at its exact center. The cspace obstacle would be a curved wall in three space, whose cross section in the $(\rho\theta,\ y)$ plane would be piecewise sinusoidal, described by the equation $y = \sqrt{2}\sin(\theta + \frac{\pi}{4})$, with period $\frac{\pi}{2}$. This cross section would be the same as x varied. See Figure 2.1.

To see how this surface is obtained, suppose that at the orientation $\theta = 0$ the square sits against the wall as shown in Figure 2.2. The vector $\mathbf{v}_0$ from the left vertex of the resting edge to the center of gravity makes an angle $\pi/4$ with the horizontal, and the height of the center of gravity at this state is $y = |\mathbf{v}_0| \sin \pi/4 = \sqrt{2}\, \sin \pi/4$. As θ is increased, with the square maintaining contact with the wall at vertex v_0, the height of the center of gravity will vary as $y = \sqrt{2}\, \sin(\theta + \pi/4)$ until vertex v_3 makes contact with the wall, at $\theta = \pi/2$. By the symmetry of the square, the above curve will repeat itself at each vertex. Since the wall is infinite and flat, it does not matter where along the wall we first make contact, and so the surface of the cspace obstacle will not vary along the x direction. However, moving the wall (tilting it in the plane) would correspond to rotating the obstacle in cspace, since now the y coordinate of the cspace obstacle has an x dependency, as well.

As another example, let us look at a part held in two (infinite length) palms joined at

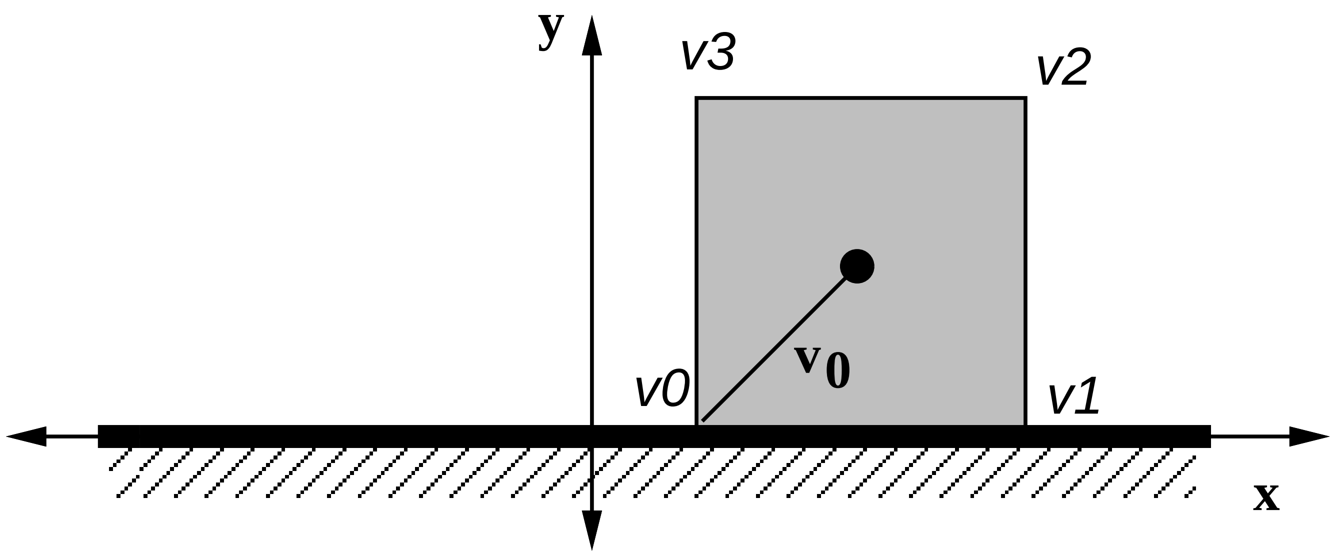

Figure 2.2: Square making contact with wall at orientation $\theta = 0$.

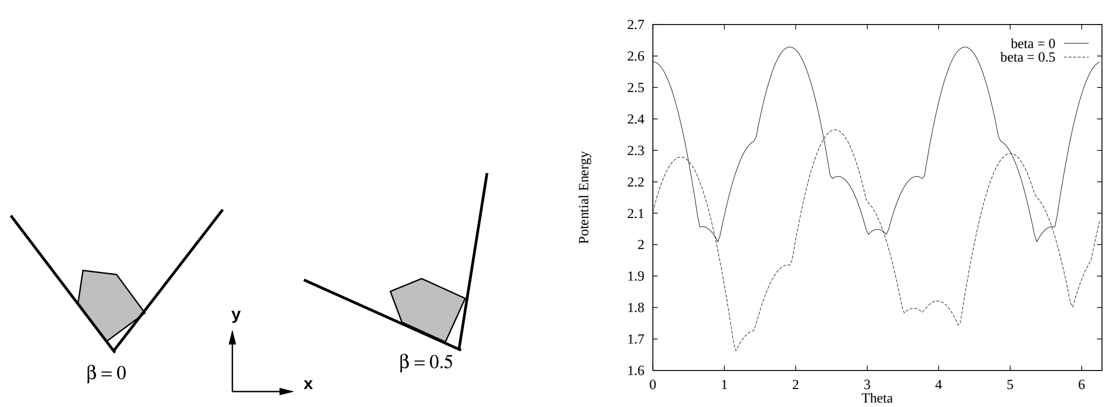

Figure 2.3: left: A part held in two palms, at two different palm orientations (but the same relative angle between the palms). right: The curves in the (θ, y) plane describing the poses of the part in contact with both palms simultaneously.

their intersection. The set of poses for which the object makes contact with both palms[1] can be described as a curve parameterized by θ, the orientation of the object. This curve can be thought of as the intersection of the surfaces of the cspace obstacles which describe each of the palms. If we use the world reference frame of Figure 2.3, where gravity is antiparallel to the y axis, then the y coordinate of this curve is proportional to the potential energy of the object in the palms. In Figure 2.3, we show two such curves for a part in a pair of palms with a fixed angle relative to each other, at two different orientations with respect to gravity. The local minima of the curves correspond to the orientations of the part for which it is in stable equilibrium in the palms.

2.1.2 Force Equations in Configuration Space

Velocity, acceleration, force and torque can also be represented in configuration space. Velocity and acceleration vectors are derived from position vectors by time derivation, as in real space. Force and torque vectors represent the force exerted on the object, and the moment exerted about the object's center of gravity. Again, because force and torque have different units, the torque component of the cspace force vector is normalized by ρ, the object's radius of gyration. This has the advantage that the standard force equation $m\ddot{\mathbf{x}} = \mathbf{F}$ still holds in cspace:

$$m \begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \rho\ddot{\theta} \end{bmatrix} = \begin{bmatrix} F_x \\ F_y \\ \tau/\rho \end{bmatrix}.$$

Frictionless contact forces, which arise when the point mass in cspace contacts a cspace obstacle, are described by the normal to the surface of the cspace obstacle at the contact point, in analogy to real space frictionless contact. The cspace contact force vector describes the forces exerted on the real space object and the moment about its center of gravity which arise from the real space contact forces.

For example, suppose an object makes point contact with an obstacle at a point $\mathbf{r}$ on the surface of the object, as in Figure 2.4. $\mathbf{r}$ is the vector from the object CG to the contact point, in some coordinate frame. Suppose the unit normal to the obstacle surface at the point of contact is $\mathbf{n}$, in the same frame. Then the cspace unit normal contact force vector is given by

$$\mathbf{F_n} = \frac{\rho}{\sqrt{1 + (\mathbf{r} \times_{2D} \mathbf{n})^2}} \begin{bmatrix} n_x \\ n_y \\ (r_x n_y - r_y n_x)/\rho \end{bmatrix} \tag{2.1}$$

where the numerator of the third component is the two dimensional cross product, $\mathbf{r} \times_{2D} \mathbf{n}$, which gives the moment exerted about the CG by $\mathbf{n}$. Any frictionless contact force exerted at this contact will be a nonnegative multiple of the vector $\mathbf{F_n}$.

Tangential forces at a contact, such as those arising from friction, are given similarly. If $\mathbf{t}$ is the unit tangent vector to the surface of the obstacle at the contact point (such that $\mathbf{t} \times_{2D} \mathbf{n} = 1$), then the cspace tangential force is given by

$$\mathbf{F_t} = \begin{bmatrix} t_x \\ t_y \\ \frac{\mathbf{r} \times_{2D} \mathbf{t}}{\rho} \end{bmatrix}. \tag{2.2}$$

Note, however, that $\mathbf{F_t}$ is usually *not* perpendicular to $\mathbf{F_n}$.

[1] For the purposes of this discussion, contact with the vertex will be considered contact with both palms.

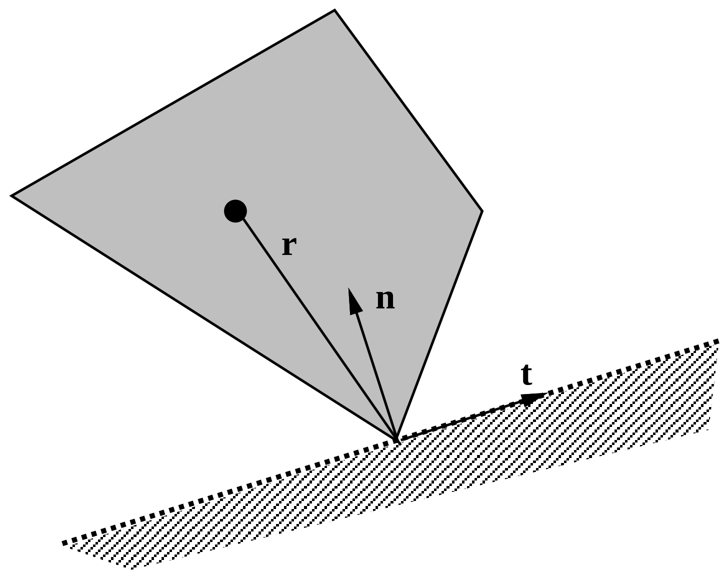

Figure 2.4: Object making point contact with an obstacle

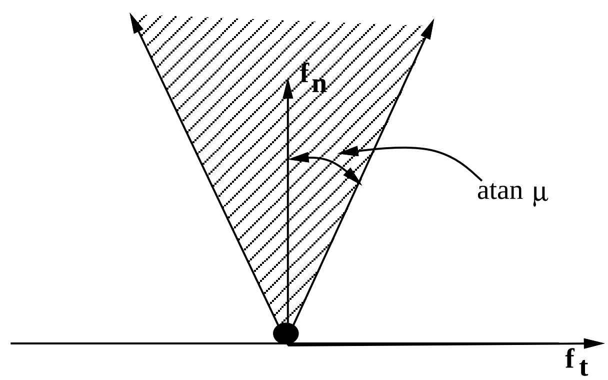

Figure 2.5: Friction cone of point mass on palm

2.2 Friction Cones

2.2.1 Friction Cones in Real Space

We will consider a point mass resting on a flat palm. (See Figure 2.5). If the coefficient of friction between the mass and the palm is μ, then the frictional force exerted on the mass can be described by a *friction cone.* The friction cone is symmetric about the contact normal, with an half opening angle $\arctan \mu$. Let the unit palm normal be $\mathbf{n}$, and let $\mathbf{t}$ be such that the cross product of $\mathbf{t}$ and $\mathbf{n}$ is positive unity. Then the friction cone is described by the relation

$$|f_t| \leq \mu f_n,$$

Where f_t and f_n are the tangential and normal components of the contact force, respectively. If the force $\mathbf{f}$ exerted by the object on the palm lies somewhere inside the friction cone, then friction will completely balance the exerted force, and the mass will remain stationary. Otherwise, friction will resist the exerted force with a contact force $f_n\mathbf{n} + \mu f_n(\operatorname{sgn} f_t)\mathbf{t}$, where $f_n = -\mathbf{f}^T\mathbf{n}$ and $f_t = -\mathbf{f}^T\mathbf{t}$. This force corresponds to the edge of the friction cone which points in the opposite tangential direction as $\mathbf{f}$. For a point mass, it is easy to see that the tilt angle of the palm for which the point mass will begin to slide is given by $|\beta| = \arctan\mu$.

For objects with extent, multiple frictional point contacts can be described by a friction cone at every contact. Whether or not a particular contact slides or rolls (sticks) depends on whether or not the forces experienced at that contact point lie inside the friction cone at that contact. For a line or edge-edge contact between two bodies, under a rigid-body assumption it is sufficient to consider only the friction cones at the endpoints of the contact.

2.2.2 Friction Cones in Configuration Space

Friction cones in configuration space for a single point contact are given analogously to those in real space. They are planar, and their edges are described by the vectors $\mathbf{F_n} \pm \mu\mathbf{F_t}$, where $\mathbf{F_n}$ and $\mathbf{F_t}$ are given by equations 2.1 and 2.2. As noted, however, $\mathbf{F_n}$ and $\mathbf{F_t}$ are not perpendicular, and the cspace friction cone is not symmetric about $\mathbf{F_n}$. In fact, $\mathbf{F_t}$ may actually dip below $\mathbf{F_t}$ in certain cases, causing the frictional forces to actually "pull" the object into the obstacle. See [32] for a more rigorous treatment of cspace friction cones, and [55] for a discussion of the anomalous "pulling" scenario.

In the case of multiple point contacts, we will depart from the procedure used in [32], and treat the configuration space friction cone directly as the *cone combination* of each of the individual cones due to each contact.

Definition 2.1 (Cone combination) *A* cone combination *of a set of vectors is the cone formed by all possible nonnegative combinations of the vectors.*

A cone combination is similar to a convex combination of a set of vectors, except that in a convex combination, the weights on the vectors must sum to unity.

The question of whether or not an object can move, or must remain stationary, becomes the problem of checking whether or not the vector of applied forces lies within the aggregate cspace friction cone. This condition is actually only *potential stability* [62]: there is some nonnegative combination of frictional contact forces which can balance the applied gravitation force. Unlike in the frictionless case, where a part is stable if and only if the contact forces balance the applied forces, a part in a potentially stable frictional grasp or support is not necessarily guaranteed to be stable. For examples of this frictional indeterminancy, see [62]. To determine if a object is in a guaranteed stable configuration in the presence of friction, however, is computationally intractable [10], whereas potential stability, which can be phrased as a question of cone containment, is no harder than linear programming.

2.3 Centers of Rotation and Frictional Contact Problems

Another way of representing the rigid motion of planar objects, which has the advantage of actually being representable in the plane, is as signed centers of rotation. The center of rotation of a rigid motion is that point in the plane which rotates without translation under the action of the rigid motion. For example, if one were to apply a pure torque to an object, the center of rotation of the resulting motion would be the center of gravity of the object, with an associated sign denoting whether the rotation was positive or negative. A pure translation has a center

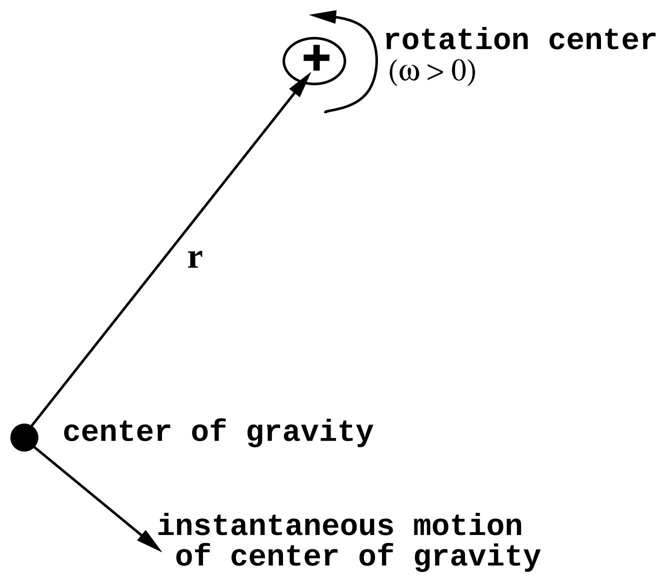

Figure 2.6: Motion of Center of Gravity due to rotation about a rotation center

of rotation at infinity. The method presented here is based on that given by Brost and Mason [24]. Strictly speaking, since we are talking about forces, what we would like is really the center of acceleration: the one point in the plane which is instantaneously unaccelerated due to the action of a particular line of force. In general, the center of rotation and the center of acceleration because the direction of motion and the direction of acceleration are generally not the same. Under the quasistatic assumption, however, the instantaneous velocity of the object is low enough that the motion of the object will be dominated by the applied forces, and the center of rotation and the center of acceleration will coincide.

Given a center and direction of rotation, one can easily determine the resulting motion of an object of interest in the plane. Recall that the cspace force vectors have a third component which is taken with respect to the center of gravity of the object, and which is normalized by the radius of gyration of the object. Therefore, it is quite natural to take the CG of the object as the origin of our reference frame, and to take the radius of gyration as the unit of length. In the following, we will assume that this is so.

The quasistatic equation describing the instantaneous velocity of the object in the plane is

$$\begin{aligned} \mathbf{v} &= -\omega \times \mathbf{r} \\ \begin{bmatrix} v_x \\ v_y \\ 0 \end{bmatrix} &= -\begin{bmatrix} 0 \\ 0 \\ \omega \end{bmatrix} \times \begin{bmatrix} r_x \\ r_y \\ 0 \end{bmatrix} \end{aligned}$$

where $\mathbf{v}$ is the motion of the center of gravity due to some rigid motion, $\mathbf{r}$ is the center of rotation of that rigid motion with respect to the object's center of gravity, and ω is the rotation of the object. $\times$ denotes the cross product. This gives the motion of the object's CG in the plane, and hence the first two components of the configuration space velocity vector. The third component would of course be ω. Under the quasistatic assumption, the instantaneous acceleration of the object is in the same direction as the instantaneous velocity.

This observation gives us a method for determining the feasibility of a motion under quasistatic frictional contact. Suppose we have an object in known contact with an obstacle. We also know the external forces applied to the object. The equation of motion in configuration space is

$$m\ddot{\mathbf{x}} = F_A + \sum_i f_{ci}$$

where F_A is the vector of applied forces, and f_{ci} is the frictional contact force at the ith contact. The problem is that the f_{ci} are generally not determined without knowing $m\ddot{\mathbf{x}}$, but $m\ddot{\mathbf{x}}$ can't be determined without knowing the f_{ci}. As we will see, however, there is a way to determine feasible centers of rotation and their associated rotation senses. Given a center of rotation, $\mathbf{r}$, and a direction of rotation, $\operatorname{sgn}\omega$, the configuration space acceleration which would result is given by the vector

$$m\ddot{\mathbf{x}} \propto \begin{bmatrix} r_y \\ -r_x \\ 1 \end{bmatrix} \operatorname{sgn}\omega$$

as shown in Figure 2.6. The object acceleration is only proportional to this vector, with the constant of proportionality being the magnitude of the object rotation. From a postulated direction for $m\ddot{\mathbf{x}}$, we can determine the directions of the resulting contact forces f_{ci}, as will be shown in the next section. In cspace, these f_{ci} form a cone, $\mathcal{C}_F$, of frictional contact force directions. In order for an acceleration direction to be physically plausible, there must be some nonnegative combination of the force directions F_A and $\mathcal{C}_F$ which sums to form $m\ddot{\mathbf{x}}$. The problem of finding feasible motions of the object subject to frictional contact can now be stated:

$$\text{Determine whether } m\ddot{\mathbf{x}} \text{ is contained in } \mathcal{C}_F \oplus F_A, \tag{2.3}$$

where $\oplus$ denotes the operation of cone combination.

2.3.1 Finding Possible Centers of Rotation

We can use a kinematic analysis due to Reuleaux [76] to determine possible centers of rotation. Consider a point contact against the edge of a body, as shown in Figure 2.7. Any point to the left of the contact can only be a center of rotation with a positive rotation sense, otherwise the contact will penetrate the body. Points to the right of the contact can only be centers of rotation with negative rotation sense. Points on the line through the contact can rotate in either direction. Furthermore, any center of rotation not on this line will cause the object to break contact. Centers of rotation on the line will cause the contact to slide either left or right, depending upon whether or not the center of rotation is above or below the contact, and whether the rotation sense is positive or negative. A center of rotation at the contact point causes the contact point to remain fixed. Note that we are making no statements about the physical plausibility of any of these centers of rotation, merely whether or not they are kinematically plausible, in the sense of not causing the contact to penetrate the object.

If we have more than one contact point, as in Figure 2.8, we can perform the above analysis on each contact separately. Regions where the different contacts give conflicting rotation senses are regions which are not plausible rotation centers. Figure 2.8 shows the regions of plausible centers of rotation, with their rotation directions, as well as the motion of each contact point which would result for the centers of rotation in that region. If the ith character of the motion string is "b", then the ith contact point will break contact. "l" means "slide left", "r" means

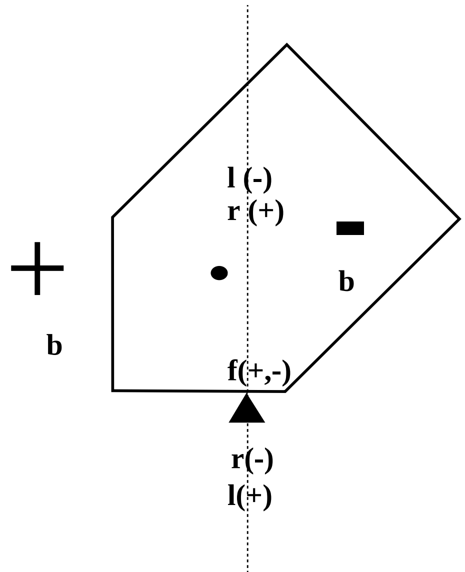

Figure 2.7: Feasible centers of rotation (CORs) for one point contact, along with rotation senses and contact modes. The rotation sense is positive for counterclockwise rotation, negative for clockwise rotation. CORs not on the line of contact cause the contact to break ("b"). CORs on the line of contact can have either rotation sense. They can cause either left ("l") or right ("r") sliding, depending on the rotation sense and their location above or below the contact point. The contact point can have either rotation sense, and always remains fixed ("f").

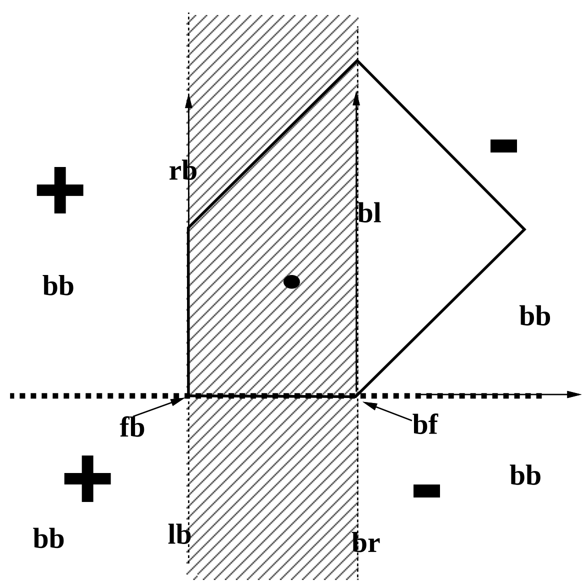

Figure 2.8: Feasible centers of rotation for two point contact

"slide right", and "f" means "stays fixed". For each region of centers of rotation, we can determine the direction of object motion, and the contact forces consistent with that motion. We can determine the contact force cone $\mathcal{C}_F$ from the contact mode as follows. If contact point p_i is in contact mode "l" (left sliding), then that contact force in configuration space is on the right edge of the configuration space friction cone (the edge of the friction cone which opposes the motion). Therefore, we add the vector $\mathbf{F_n}_i + \mu\mathbf{F_t}_i$ to $\mathcal{C}_F$. If p_i is in mode "r" (right sliding) then add the vector $\mathbf{F_n}_i - \mu\mathbf{F_t}_i$. If p_i is in mode "f" (fixed), then the contact force is inside the friction cone. Therefore you add both edges $\mathbf{F_n}_i \pm \mu\mathbf{F_t}_i$ to $\mathcal{C}_F$. If p_i is in mode "b" (break), then the contact has broken, and there is no contact force. Having determined $\mathcal{C}_F$, along with the vector of applied forces and the vector of hypothesized acceleration, we can determine if a particular motion is feasible, as in Equation 2.3.

2.4 Cone Analysis

As stated above, we have expressed the problem of determining the motion of objects subject to contact constraints as the inclusion of certain vectors in convex cones; *i.e.*, as nonnegative combinations of the vertices of those cones. This is a standard problem in linear programming, and in certain special cases, may also be easily solved by linear algebra. For planar problems, Brost and Mason [24] have shown how to solve this problem graphically, with two-dimensional diagrams. They took advantage of the duality between points in velocity space (centers of rotation) and lines of force in wrench space. In this dissertation, the problem will be presented as cones in three dimensional configuration space, since this method will generalize to higher dimensional problems, as well.

Let M be the matrix representing a convex cone with N vertices. Each vertex of the cone is represented by a three dimensional column vector of M. Let $\mathbf{x}_{test}$ be the vector being tested for inclusion. Let $\mathbf{1}$ be the N dimensional vector of all ones, and $\mathbf{c}$ be the N dimensional weighting vector on the columns of M. Then one example of a linear program which may be used to test for inclusion is

$$\text{Minimize } \mathbf{c}^T\mathbf{1}$$

$$\begin{gathered}\text{Subject to:}\\ \mathrm{M}\mathbf{c} = \mathbf{x}_{test}\\ \text{and}\\ c_i \geq 0, i = 1:N.\end{gathered}$$

This linear program will find the minimum nonnegative weights on the vertices of the cone represented by M such that the weighted vector adds up to $\mathbf{x}_{test}$, if such weights exist. Otherwise, the linear program will determine that the problem is infeasible.

The above linear program is a general solution to the cone containment problem; in many of the specific cases which we run into in our system, the test vector $\mathbf{x}_{test}$ is the negative of the gravity vector, and the third component of the test vector is zero. In this case, it is faster in practice to intersect M with the $(x,\ y)$ plane by linear algebra, as in Figure 2.9, and compare the resulting lower dimensional cone with $-\mathbf{x}_{test}$ (where we have now dropped the zero valued third component). Sort this set of vectors in, *e.g.*, counterclockwise order, and check the angles between adjacent vectors. If the maximum of these angles is less than π, then the set of vectors positively spans the plane, and $\mathbf{x}_{test}$ is contained in M. This check can be done in $\mathcal{O}(N \log N)$ time, where N is the number of columns of M.

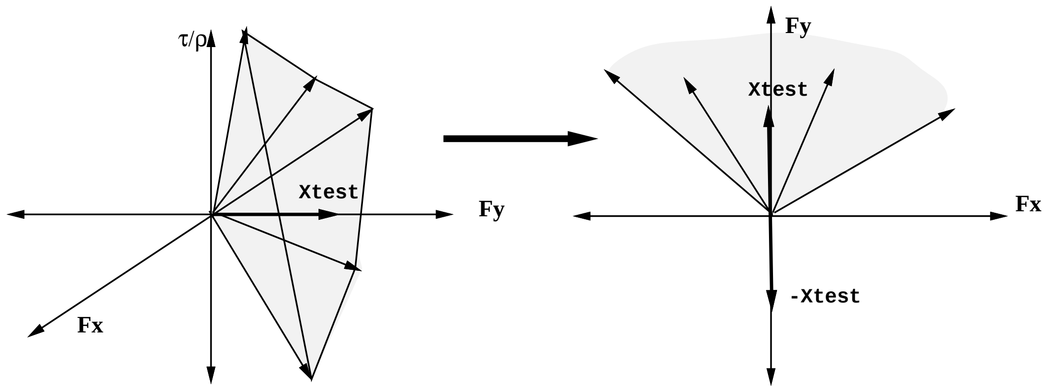

Figure 2.9: If $\mathbf{x}_{test}$ lies in the $(x,\, y)$ plane, we can test inclusion by finding the intersection of M with the $(x,\, y)$ plane

2.5 Impact Dynamics

Following [80], the impulse equations for a body of mass m colliding with point contact against a barrier (where it is assumed that $m_{barrier} >> m$) are given by

$$\begin{aligned} m({v_t}^+ - {v_t}^-) &= P_t \\ m({v_n}^+ - {v_n}^-) &= P_n \\ m\rho^2(\omega^+ - \omega^-) &= P_\omega = P_t l \cos\zeta + P_n l \sin\zeta \end{aligned} \tag{2.4}$$

Here (see Figure 2.10), ρ is the radius of gyration of the object, v_t is the relative velocity of the body tangential to the contact normal, and v_n is the component normal to the contact. The superscript "+" denotes the velocity term immediately after the moment of impact, and "-" denotes the velocity term upon contact, at the moment of impact. P_t and P_n are the components of the impulse in the tangential and normal directions, respectively. ζ is the angle of the line from the center of gravity of the object to the contact point with respect to the contact normal, in the counterclockwise direction. Note that $l \sin\zeta$ is the moment arm of the torque applied by the contact normal force.

For a given cspace velocity $(v_t, v_n, \rho\omega)$, the velocity of the contact point is given by

$$\begin{aligned} v_{tc} &= v_t + l\omega\cos\zeta \\ v_{nc} &= v_n + l\omega\sin\zeta. \end{aligned} \tag{2.5}$$

Combining (2.4) with (2.5) gives us

$$\begin{aligned} {v_{tc}}^+ &= {v_{tc}}^- + \frac{\rho^2 + l^2\cos^2\zeta}{m\rho^2} P_t + \frac{l^2 \sin\zeta\cos\zeta}{m\rho^2} P_n \\ {v_{nc}}^+ &= {v_{nc}}^- + \frac{l^2 \sin\zeta\cos\zeta}{m\rho^2} P_t + \frac{\rho^2 + l^2\sin^2\zeta}{m\rho^2} P_n \end{aligned} \tag{2.6}$$

Newton's model for impact (see [80], [93], or [17]) divides the (normal) impulse into two parts, compression, when the colliding objects are moving into each other, and restitution, when the objects move away from each other. Newton's hypothesis is that the impulse of restitution and

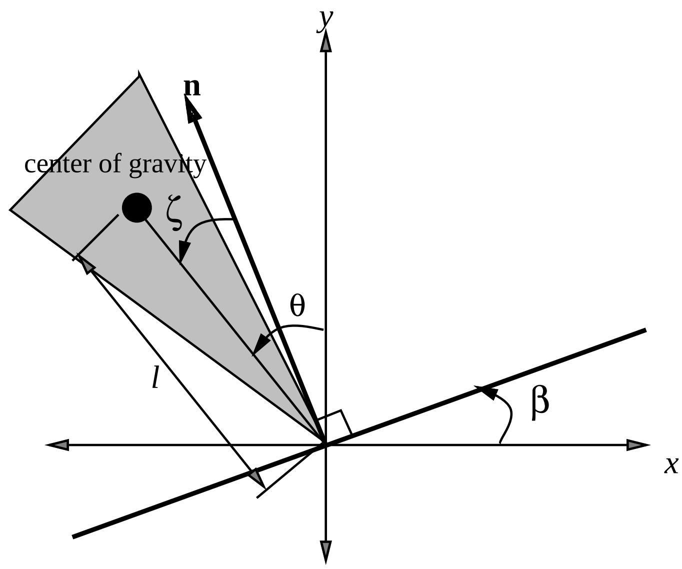

Figure 2.10: System Geometry

the impulse of compression are in the ratio e, called the coefficient of restitution. We define P_{n0} as the impulse at full compression: that is, when $v_{nc} = 0$. If we assume no friction, then P_t must be zero for the entire impact. Hence, we have the expressions for the final total impulse

$$P_t = 0; \quad P_n = (1+e)P_{n0}. \tag{2.7}$$

We can use (2.6) to solve for P_{n0} by letting ${v_{nc}}^+ = 0$. Substituting those results into (2.7), and using the last equation of (2.4) gives us the normal and rotational impulses in terms of the initial normal contact velocity.

$$\begin{aligned} P_n &= -(1+e)\frac{m\rho^2}{\rho^2 + l^2 \sin^2 \zeta}{v_{nc}}^- \\ P_\omega &= -(1+e) l \sin\zeta \frac{m\rho^2}{\rho^2 + l^2 \sin^2 \zeta}{v_{nc}}^-. \end{aligned} \tag{2.8}$$

Substituting this back into equation 2.4 now gives us an expression for the change in velocity due to impact, in terms of the initial velocity. This expression can then be rewritten

$$\Delta \mathbf{v} = -(1+e)(\hat{\mathbf{n}}^T \mathbf{v}^-)\hat{\mathbf{n}} \tag{2.9}$$

by using the following definitions of the configuration space velocities and contact normals:

$$\begin{aligned} \mathbf{v} &= \begin{bmatrix} v_t \\ v_n \\ \rho\omega \end{bmatrix} \\ \Delta \mathbf{v} &= \mathbf{v}^+ - \mathbf{v}^- \\ \hat{\mathbf{n}} &= \frac{\rho}{\sqrt{\rho^2 + l^2 \sin^2 \zeta}} \begin{bmatrix} 0 \\ 1 \\ \frac{l}{\rho}\sin\zeta \end{bmatrix}. \end{aligned} \tag{2.10}$$

Here, ρ is the radius of gyration of the object.

Chapter 3

Manipulation Without Grasping

> A model is by definition that in which nothing has to be changed, that which works perfectly; whereas reality, as we see clearly, does not work and constantly falls to pieces; so we must force it, more or less roughly, to assume the form of the model.
>
> Italo Calvino
> *The model of models* [27]

In this chapter, we make the following assumptions:

- We will restrict ourselves to planar polygonal objects, although our methods should carry over without much difficulty to any planar object whose convex hull has a finite number of stable resting positions. Cylindrical objects can also be modeled as planar.
- Force balance is achieved by the palms stably supporting the object against a known gravitational force. No other external forces are considered, hence complete force/form closure is not necessary.
- We will assume that the motions of the manipulator are slow compared to gravity, so that the kinetic energy imparted to the object by the motion of the palm is dominated by the object's potential energy.
- We will assume that the contacts between the object and the palms are very low friction (*i.e.*, the contacts are all sliding), so that we may approximate the system with a frictionless analysis.
- We will model the two palms as a "cone" manipulator: two palms connected at a central hinge. This is a simple model that still captures many of the basic operations of a general two palm system.

To justify the last assumption, we notice that many (though by no means all) configurations of two palms which are capable of passively supporting an object can be modeled by the behavior of the object resting in a cone (Figure 3.1). This cone is formed by the intersection of the lines along which the palms lie. If we set a frame in the cone such that the y axis is the bisector of the cone, then all possible motions of the cone can be described as a combination of two motions. One motion is holding the y axis fixed and opening or closing the cone opening symmetrically (a *pure squeeze*). The other is holding the cone opening constant and rotating the y axis (a *pure tilt*).

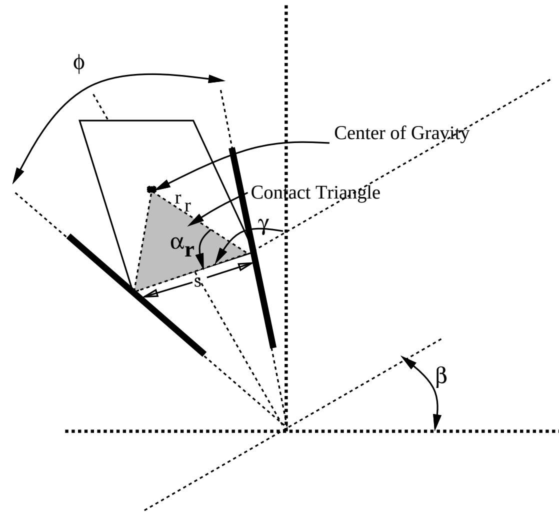

Figure 3.1: Contact triangles and Cone frames

In later chapters, we will show how to relax the assumptions of low friction and low kinetic energy.

3.1 Energy Analysis

The motion of a polygonal object when in contact with the two palms can be determined by constructing the configuration space obstacle formed by the palms. Each palm generates an obstacle in configuration space, and the intersection of the surfaces of these two obstacles form a curve describing the poses of the part when it is making contact with both palms simultaneously. Taking two particular vertices of the polygon to be in contact with the palms, the feasible motions of the polygon can be abstracted by the motion of a triangle in a cone, where one vertex of the triangle is the center of gravity (CG) of the object (Figure 3.1), and the other two vertices are the points of contact with the left and right palms, respectively. Let ϕ be the angle of the cone opening, with the y axis of the cone frame bisecting the cone. The center of gravity location is given by $CG = (CG_x,\ CG_y)$ in the cone frame. If, further, the cone frame is tilted at an angle β counterclockwise with respect to the world frame, then the potential energy of the object is proportional to $CG_x \sin\beta + CG_y \cos\beta$.

We can use the potential energy of the part to identify its *equilibrium* positions.

Definition 3.1 (Stable Equilibrium) *A part is in a* stable equilibrium pose *when the net force on the part is zero, and the part will resist small perturbations of its pose. This corresponds to poses where the potential energy is locally minimized.*

Definition 3.2 (Unstable Equilibrium) *A part is in an* unstable equilibrium pose *when the net force on the object is zero, but the object cannot resist small perturbations of its pose. This corresponds to poses where the potential energy is locally maximized.*

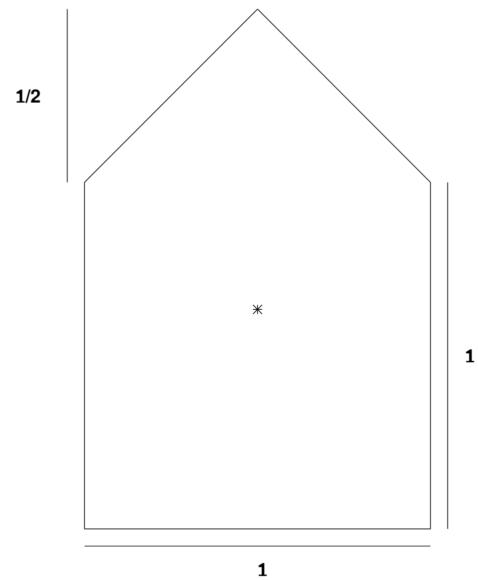

Figure 3.2: Example Object

For a given two point contact configuration (that is, for a given contact triangle), there is an interval I of object orientations for which this contact configuration is maintained as the object rotates. The potential energy of the object as a function of orientation either has a unique maximium in the interior of I, or it is monotonic (The proof is given in Appendix A). This means that if the cone is frictionless, then there is at most one unstable equilibrium orientation of the triangle balanced in the cone, and one or two stable resting orientations, corresponding to the triangle resting on one side or other of the cone. If we look at all the triangles formed by all possible pairs of vertices of the object, we can generate the cspace obstacle formed by the palms, and from the cspace obstacle determine the potential function curve of the object in a given cone tilted at a given orientation. Figure 3.3 shows a scaled example of such a curve, for the example object shown in Figure 3.2, in a cone opened $\pi/3$ radians wide, with $\beta = 0$. In this particular case, the potential surfaces are directly proportional to the y-coordinate of the configuration space constraint surfaces. Since the local minima of this curve represent the orientations where two triangles are simultaneously in contact with the cone, stable orientations of an object correspond to three point contact with the cone. Let x be the horizontal position of the object CG in the world frame, and θ be the orientation of the object in the world frame. For a fixed value of β, the curve defined by the intersection of the palms' configuration space obstacles forms a one dimensional curve in three dimensional space, into which the constraint surface gradients point (Figure 3.3). This curve, in turn, has local minima which attract the system state. Assuming that the kinetic energy of a part is always low compared to to the depth of the potential wells, an object caught in a cone in a particular orientation will settle to a unique resting position determined by the initial position of the object upon contact, and the tilt of the cone with respect to gravity. Ensuring that the kinetic energy is low enough requires that the palms move slowly in comparison with gravity, and that gravity is low compared to some ambient viscosity, or to the coefficient of restitution of the impacts between the palms and the part. This ambient viscosity may even be the friction between the part and the palms, provided that the "all contacts slip" condition is met. Determining the conditions for which

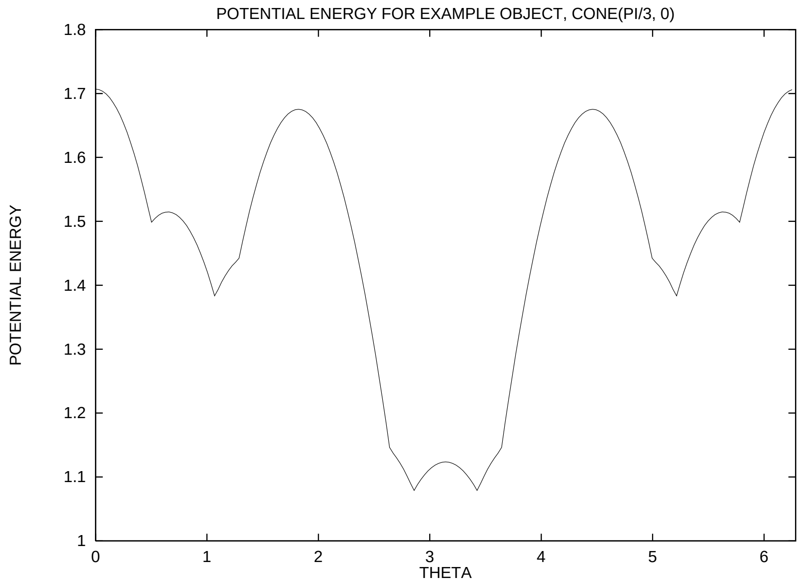

Figure 3.3: For two palms fixed at $\phi = \pi/3, \beta = 0$, the intersection of the configuration space obstacles shown projected into the (θ , y) plane

the coefficient of friction simultaneously meets both the requirements on slipping contacts and energy dissipation will be addressed in Chapter 4. Determining the necessary conditions on the palm velocity and the coefficient of restitution is addressed in Chapter 6.

Once an object is in a stable state, the cone can be tilted back and forth within a certain range of β while maintaining stability. For those regions of β where stable contact is maintained,

$$\frac{\partial\theta}{\partial\beta} = 1 \tag{3.1}$$

where θ here is taken to mean object orientation in world coordinates.

We now look at what happens as ϕ is varied and β is fixed. If the object is already resting stably in edge contact with one palm, then for some range of ϕ, the cone can be widened or narrowed and the object will stay in stable contact with that palm. For the range of ϕ for which the object maintains stability,

$$\frac{\partial\theta}{\partial\phi} = \pm\frac{1}{2}, \tag{3.2}$$

positive if the object is resting on the left palm, or negative if on the right palm[1]. Because we are looking at the low friction case, we are assuming that jamming cannot occur, since the contacts are assumed always to slide.

The change in orientation of an object in response to the tilting and squeezing of the cone can be described by the equation

$$\begin{aligned} \partial\theta &= 1 \cdot \partial\beta - \frac{\varepsilon}{2} \cdot \partial\phi \qquad (3.3) \\ \varepsilon &= -1\text{, resting on left palm} \\ & 1\text{, resting on right palm} \end{aligned}$$

[1]This can be seen by noting that the angle of the left palm in the cone frame is $\frac{\pi}{2} + \frac{\phi}{2}$, and the angle of the right palm is $\frac{\pi}{2} - \frac{\phi}{2}$

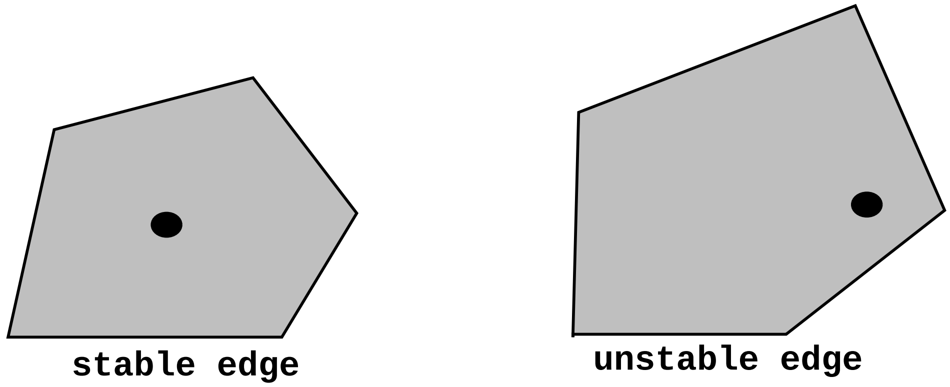

Figure 3.4: The bottom edge of the polygon on the left is stable; the bottom edge of the polygon on the right is not.

over the range of ϕ, β for which the object remains stably supported throughout the cone motion.

3.2 Equivalence Regions

We would like to use the above observations to plan object reorientations automatically. If we look, for the moment, at the object only after it has made contact with the cone, then the cone/object configuration can be characterized as a point in the space (θ, ϕ, β). We will call all stable resting configurations that correspond to a particular side of the object in edge contact with a particular palm *equivalent configurations.* For example, (see Figure 3.5), all stable configurations where rectangle side a rests on the left palm are equivalent. Define *contact formations* (Desai, as cited in [88], [84], [85], [86]) as the set of contact configurations where the same vertices of the object are making contact with the same edges of the cone. Note that the definition of equivalent configurations forms a superset of the sets of stable contact formations corresponding to side a against the left palm, since for some configurations of the cone, the object makes three point contact, for others, two or four point contact.

Definition 3.3 (Equivalence Region) *A set of equivalent configurations will be called an* equivalence region *in (θ, ϕ, β) space. The projection of this set into the (ϕ, β) plane is the* shadow *of the equivalence region.*

Referring to Equation (3.3), we see that an equivalence region in (θ, ϕ, β) space is a planar surface. Figure (3.6) shows the projection of this equivalence region into the (ϕ, β) plane. The equivalence regions of most interest are those which correspond to *stable edges*, since desired goal states will most likely correspond to states in such regions.

Definition 3.4 (Stable Edge) *A* stable edge *of an object is an edge on which the object can rest on a horizontal palm, unsupported by the other palm, without tipping over, in the face of small disturbances.*

In other words, as shown in Figure 3.4, a stable edge is an edge e_s such that, if one were to project the center of gravity onto e_s, this projection would lie in the interior of the line segment e_s. For stable edges, the equivalence regions can be shown to be simply connected (see Chapter

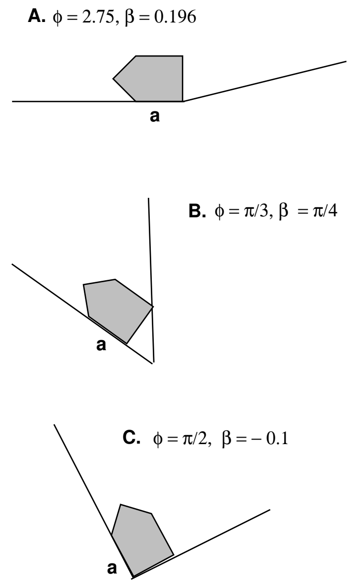

Figure 3.5: Different contact formations which are in the same equivalence region

7): for any two configurations in an equivalence region, there is a stable trajectory from one configuration to another. Hence, if we have the point $(\theta_{des}, \phi_{des}, \beta_{des})$ as our goal state, then an immediate subgoal is to reach the corresponding equivalence region.

We note here that the interior of what we are calling equivalence regions are supersets of what are called in [85] *passive first order stability cells (passive FS cells). First order stability* corresponds to configurations of the object in the manipulator where the potential energy is at an equilibrium point and any feasible infinitesimal perturbation of the object strictly increases the potential energy [84]. An FS cell is *active* if some of the joints of the manipulator must be compliance-controlled in order to maintain stability as the object is manipulated (this is generally true if there are more joints than there are degrees of freedom of the object). An FS cell is *passive* if only position-control of the manipulator is necessary. Specifically, one of our equivalence regions contains a union of one or more passive FS cells, each one corresponding to a different contact formation. An equivalence region corresponds to a region in configuration space for which a completely (first order almost everywhere) stable path exists between any two points in the region.

Recall that for the low friction, low kinetic energy case, every configuration for the object and cone which lies on the intersection of the configuration space obstacles of both palms (except for the unstable equilibrium orientations) is attracted to a unique stable resting configuration. Then for a fixed cone, every stable object orientation θ_s has a neighborhood of orientations which converge to θ_s. Taking the union of all these neighborhoods in (θ, ϕ, β) space gives the preimage $\mathcal{P}_\mathcal{S}$ of an equivalence region, $\mathcal{S}$.

Definition 3.5 (Preimage) *The preimage $\mathcal{P}_\mathcal{S}$ of an equivalence region, $\mathcal{S}$ is the set of all configurations (θ, ϕ, β) which converge to some configuration (θ_s, ϕ, β) such that $(\theta_s, \phi, \beta) \in \mathcal{S}$.*

In other words, $\mathcal{P}_\mathcal{S}$ is a region of state space which is trapped in a potential energy well. The bottom of the energy well is the surface $\mathcal{S}$. For a cone formed by palms of infinite length,

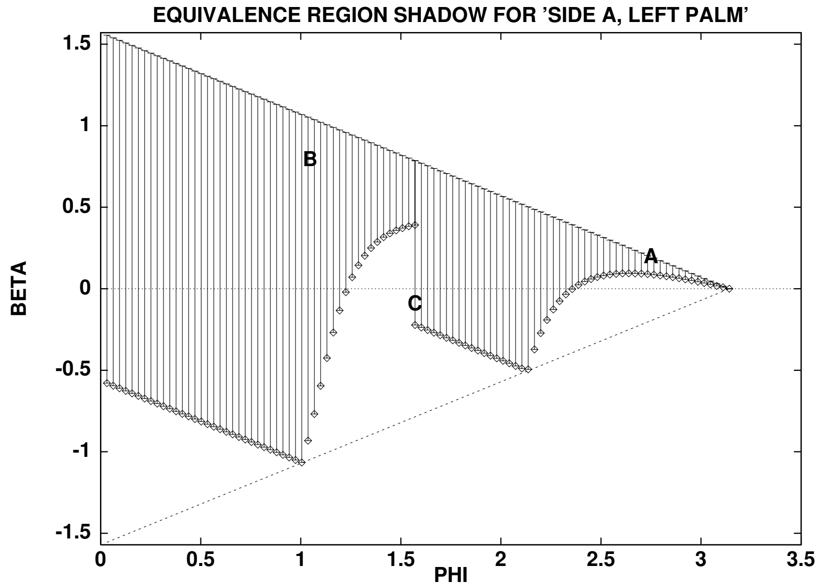

Figure 3.6: Equivalence region shadow for configurations of previous figure. The labeled points show the locations of the corresponding configurations in the previous figure.

the preimages cover the region of state space

$$\begin{aligned} \theta &\in [0, 2\pi), \\ \phi &\in (0, \pi], \\ \beta(\phi) &\in [-\frac{(\pi - \phi)}{2}, \frac{(\pi - \phi)}{2}], \end{aligned} \tag{3.4}$$

For cone configurations (ϕ, β) outside the region above, the mouth of the cone dips below the horizontal, so the object would spill out of the cone. Of course, once the finite length of the palms is taken into account, the preimages will no longer cover this region.

The boundaries of an equivalence region correspond to the edge of stability of a particular contact class, $\mathcal{S}_0$: there will be some directions of movement of the cone which will take the system state out of $\mathcal{S}_0$, causing the edge contact of interest to be lost. If that particular boundary region of $\mathcal{S}_0$ lies in the preimage of another equivalence class, $\mathcal{S}_1$, then the object will fall into the stable contact corresponding to equivalence region $\mathcal{S}_1$. In other words, an object's orientation can be brought from $\mathcal{S}_0$ to $\mathcal{S}_1$ by bringing the object to the appropriate boundary of $\mathcal{S}_0$ and moving the cone in such a way that the cone state moves out of the shadow of $\mathcal{S}_0$, and into the shadow of $\mathcal{S}_1$. This transition is reliable even though the manipulator does not maintain stable support of the object during the transition from $\mathcal{S}_0$ to $\mathcal{S}_1$, as long as the object's kinetic energy is low compared to its depth in the potential energy well.

3.3 Planning in the Frictionless Domain

We present here some theorems which will be relevant to planning reorientations. The proofs of these theorems will be presented in Chapter 7, and we have numbered them here as they are numbered in that chapter.

Let $\mathcal{S}$ be an equivalence region corresponding to a stable edge, and p, q be points in $\mathcal{S}$.

Theorem 1 *$\mathcal{S}$ and its projection into the (ϕ, β) plane are simply connected.*

By a *connected region* we mean a region X such that if $x_1 \in X$ and $x_2 \in X$, there is a curve segment whose endpoints are x_1 and x_2, and which is completely contained in X. By a *simply connected region*, we mean a connected region X without holes. For the rest of this section we will refer to both the equivalence region $\mathcal{S}$ and its projection into the (ϕ, β) plane interchangeably as $\mathcal{S}$, since (as discussed in Chapter 7) the two are homeomorphic, or topologically equivalent.

Theorem 5 *If $p = (p_x, p_y), q = (q_x, q_y) \in \mathcal{S}$, and $p_x = q_x = \phi$, then the line segment $\overline{p\,q}$ is completely contained in $\mathcal{S}$.*

In other words, if p, q are both in $\mathcal{S}$, and both correspond to points with the same cone opening ϕ, then the pure tilt motion between them preserves the stable configuration of the part.

In addition to pure tilts (ϕ constant, β varying) and pure squeezes (β constant, ϕ varying), there is another type of simple movement of the palms related to pure squeezes: a *fixed-θ squeeze*. In a fixed-θ squeeze, the palm making edge contact with the object stays fixed, and the other palm opens and closes, resulting in the cone configuration changing, but the object orientation relative to the world frame remaining fixed as long as the state is stable. In the parameter space, this corresponds to $\frac{\partial \beta}{\partial \phi} = \varepsilon \frac{1}{2}$, where $\varepsilon = -1$ for the left palm held fixed, 1 for the right palm held fixed. A fixed-θ squeeze between two points in an equivalence region (one which corresponds to a stable edge) will be completely contained in that equivalence region, if at both endpoints of the squeeze, the vertex contacts are the same: not only is the same resting edge on the same palm (which is the definition of an equivalence region), but the other palm is making contact with the same vertex at both endpoints of the squeeze.

Theorem 6 *If $p, q \in \mathcal{S}$, p and q are vertex-equivalent, and $\overline{p\,q}$ has slope $\frac{\partial \beta}{\partial \phi} = \varepsilon \frac{1}{2}$, then $\overline{p\,q}$ is completely contained in $\mathcal{S}$.*

So, for any start point and any subgoal point in an equivalence region corresponding to a stable edge e resting on palm k (where k is left or right), one of the following is true:

- There is a pure tilt connecting the two points, or there is a fixed-θ squeeze connecting the two points. Then there is a single operation that will bring you from the start state to the subgoal state.

- Define the *baseline* of an equivalence region to be the curve in the equivalence region corresponding to palm k being horizontal. If there is not a pure tilt or fixed-θ squeeze connecting the two points, then draw the pure tilt line from the start state to the baseline, and the pure tilt line from the goal state to the baseline. These are both guaranteed to be contained in the shadow of the equivalence region by Theorem 5. There will be some line (a fixed-θ line) contained in $\mathcal{S}$ which is parallel to the baseline and which passes through these two pure tilt lines. (At the very least, the baseline will satisfy this requirement). See Figure 3.7. Therefore, the subgoal can be achieved from the start state in at most two tilts and one fixed-θ squeeze.

For an unstable edge, equivalence regions as we have defined them may no longer be connected. However, in each component of the equivalence region, the above results about pure

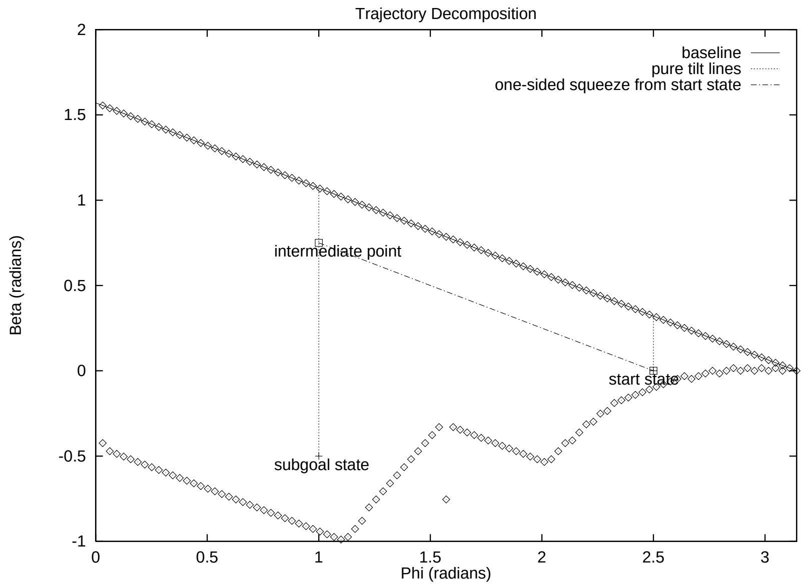

Figure 3.7: Trajectory decomposition

tilts and fixed-θ squeezes still hold.

The results are summarized below.

Theorem 8 *Given any two points in an equivalence region $\mathcal{S}$ corresponding to a stable edge of a polygon P, there exists a path entirely contained in $\mathcal{S}$. which can be decomposed into at most two tilts and one fixed-θ squeeze.*

If the equivalence region does not correspond to a stable edge, but the two points are in the same component of the equivalence region, there still exists a path between the two points composed only of pure tilts and fixed-θ squeezes which is entirely contained in that component.

Therefore, if there exists a path from an arbitrary start state to an arbitrary end state, where the transitions between equivalence regions are given by pure tilt motions, the entire path between the start state and the goal state can be decomposed naturally into pure tilts and fixed-θ squeezes.

Intuitively, pure tilts roughly correspond to pure rotations of the objects, and fixed-θ squeezes roughly correspond to pure radial translation. Just as pure rotations and pure radial translation span the plane when using polar coordinates, pure tilts and pure rotations span the motion space of the object in the palms.

In addition to moving around within a stable equivalence region, we also would like to know what motion will take us to a new equivalence region. Suppose that we orient the edges of the polygon so that edge e_i[2] is a ray whose origin is vertex v_i, and the vertices (and edges) are ordered counterclockwise with respect to the center of gravity.

Lemma 9 *For a cone with infinite length palms, suppose polygon P is resting stably on edge e_0 on the left palm. Let the right palm make contact with vertex v_k. Let the vector from the center of gravity of P to vertex v_k be $\mathbf{v}_k$, and ${\phi^\perp}_k$ be the cone opening such that the right palm*

[2]In the subscripts for the edges or vertices of N-gons, $i \pm k$ is shorthand for $(i \pm k)$modulo N.

is perpendicular to $\mathbf{v}_k$*. Then:*

- *For all cone openings such that* ${\phi^\perp}_{k+1} < \phi < {\phi^\perp}_k$*,* P *will rotate to edge* e_k *on the right palm, upon executing a pure clockwise tilt. If* e_k *is stable, the part will stay in that configuration.*

- *For all cone openings such that* ${\phi^\perp}_k < \phi < {\phi^\perp}_{k-1}$*,* P *will rotate to edge* e_{k-1} *on the right palm, upon executing a pure clockwise tilt. If the* e_{k-1} *is stable, the part will stay in that configuration.*

In Chapter 7, it will also be shown that the arcs from one equivalence region to another are bidirectional.

In order to determine which orientations of a particular part can be brought to which other orientations:

1. First, determine all the equivalence regions, (two for every flat face of the convex hull of the object) and their preimages.

2. Determine the boundary of each equivalence region, and divide each boundary into segments, according to which preimage of another equivalence class that segment is contained in. If we use pure tilt motions to transit out of an equivalence region into the preimage of another equivalence region, then Lemma 9 shows that these boundary segments can be described by intervals of cone openings, ϕ. For example, suppose we start in an equivalence region, $\mathcal{S}_0$, and divide the boundary into cone opening intervals I_i as described in Lemma 9. Then picking any cone opening ϕ in a particular interval I_i and tilting in the appropriate direction will cause a transition to the same new equivalence region, $\mathcal{S}_i$.

3. Construct the graph $\mathcal{G}$ whose nodes are the equivalence regions, with arcs denoting which equivalence regions transit into another. Each arc is labeled with the appropriate set of cone configurations (for the example above, the arcs would be labeled by the appropriate ϕ-interval), and the direction in cone configuration space in which the cone must be moved. Figure 3.8 shows $\mathcal{G}$ for our example object. The arcs in $\mathcal{G}$ were determined by using only pure tilts at the equivalence region boundaries to transit from region to region. Different contact formations in each equivalence region are shown with an arc into the contact formation of another region which it transits to.

The graph $\mathcal{G}$ in Figure 3.8 was generated assuming that all configurations (ϕ, β) are achievable. In reality, not all cone tilts and orientations will be achievable, because the finite length of the palms will prevent the object from being held in some cone configurations, and because of other physical limitations of the device. One possible limitation is the ability or inability to slide the part from one palm to the other without changing the part's orientation relative to the palm on which it rests; we will refer to such a motion as a *sliding transfer*. Sliding transfers may not be possible on a specific device because of the space between the palms, or because of an obstruction (such as the joint) at the vertex of the cone formed by the palms. As shown in Section 3.4, such physical limitations can cause certain arcs of $\mathcal{G}$ to be completely eliminated.

The planning problem has now been segmented into two parts. Given the initial and desired final configurations of the system, the high level problem is how to get from the initial to the final equivalence region. This can be determined by straightforward graph search on $\mathcal{G}$. If a

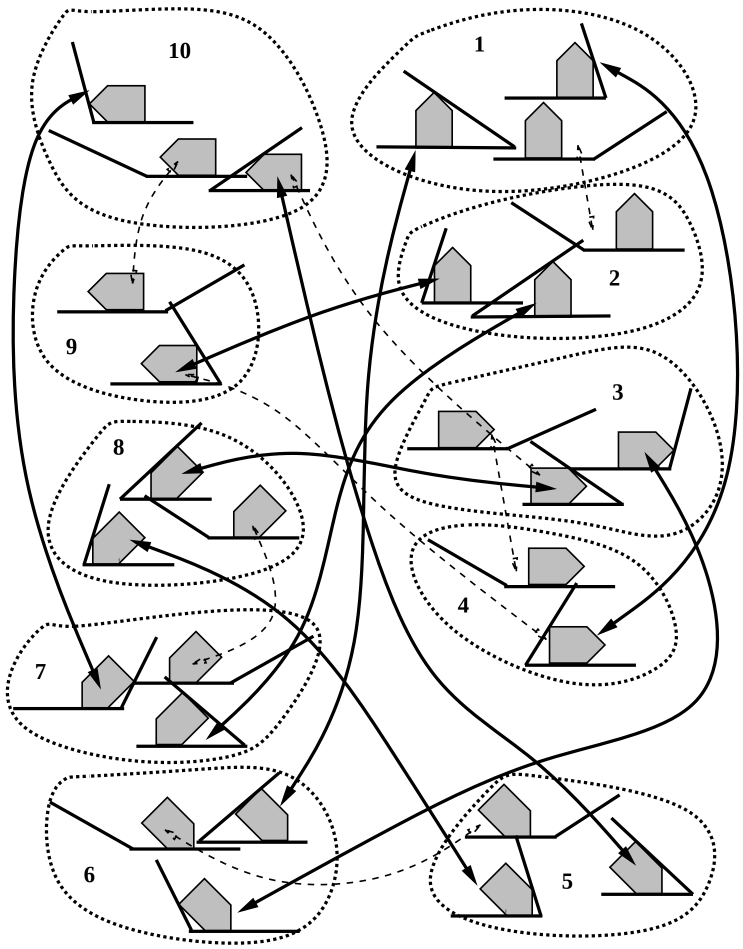

Figure 3.8: Transition graph for example object. Transition motions are made by pure tilts in the appropriate direction. The arcs shown as dashed lines correspond to transitions which were eliminated by the planner, as explained in Section 3.4.

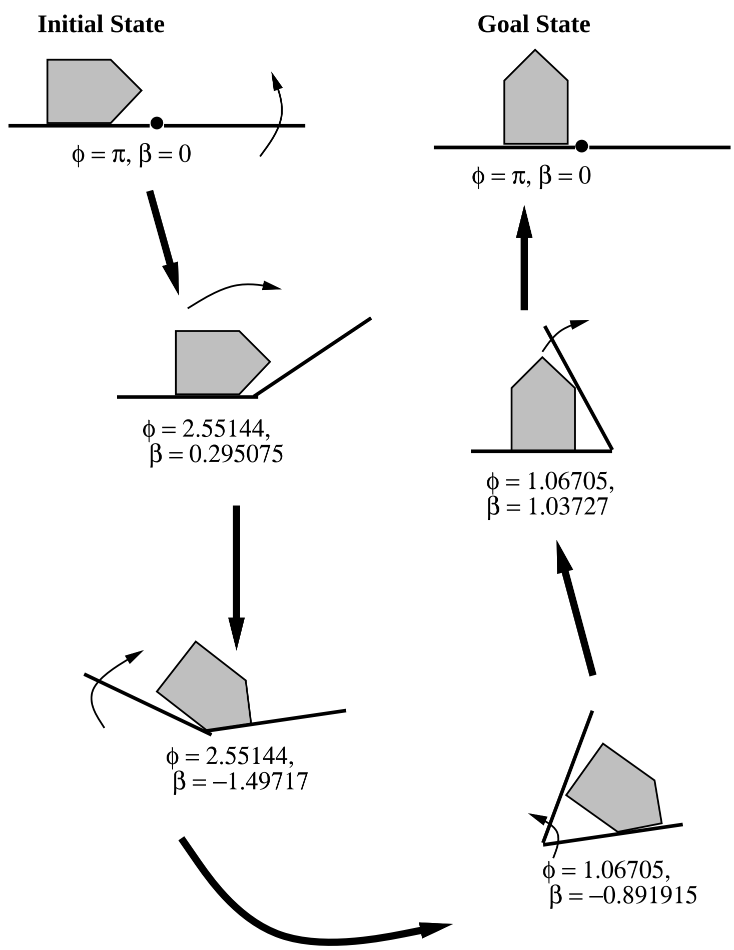

Figure 3.9: Example reorientation: from $\theta = -\pi/2$ on left palm to $\theta = 0$ on left palm

path through the graph exists, the reorientation is in principle possible, and the path determines a series of sets of equivalence regions which the cone trajectory must go through.

Once it has been established that a high level path exists, the lower level trajectory planning problem for each equivalence region (node) is to determine the trajectory which the cone must follow to reorient the part. The motions to transit from one equivalence region to another are given by the arcs of $\mathcal{G}$. To determine trajectories through equivalent regions, we can take advantage of the fact that equivalence regions are piecewise straight-line connected, as described in Theorem 8. Figure 3.9 shows an example reorientation for our example object.

Sufficient conditions for the existence of plans are summarized below. Again, the proofs will be presented in Chapter 7.

Definition 3.6 (Orientable) *We will call a polygon* orientable *if it can be brought from any equivalence region associated with a stable edge, to any other equivalence region associated with a stable edge.*

The *relative angle* α_{ij} between two edges e_i and e_j is the counterclockwise angle formed between e_i and e_j.

Theorem 12 *Let P be a convex polygon with N all stable edges, G the corresponding transition graph. In order for P to be orientable, it is sufficient that either:*

1. *Sliding transfers be possible, or*

2. *For every edge e_i of P, there be at least two edges e_j, e_k such that $\alpha_{ij}, \alpha_{ik} \leq \pi$.*

If either of the above two cases is true, then the length of any reorientation plan is bounded by N. This plan can be found in $\mathcal{O}(N^2)$ operations.

A triangle with all stable edges is an example of a part which satisfies the predicates of Theorem 12, but not Condition 2. Note, however, that Theorem 12 gives sufficient, but not necessary conditions. Hence it may be possible that a part with all stable edges which is orientable, even if neither Condition 1 nor 2 are true.

Theorem 14 *If sliding transfers are possible, any polygon P is orientable. The path length from any initial state to any goal state is bounded by $2M$, where M is the number of stable edges. It can be found in $\mathcal{O}(M^2)$ operations.*

The following is a paraphrase of Theorem 15.

Theorem *If sliding transfers are not possible, in order for P to be orientable it is sufficient that every stable edge e_{α_i}, when resting on the left palm, can reach two other stable edges $e_{\alpha_{i+1}}$ and $e_{\alpha_{i+2}}$ on the right palm by a clockwise tilt. If P has M stable edges, the path from any initial state to any goal is bounded by $2M$ and can be found in at most $\mathcal{O}(M^2)$ operations.*

An analogous result naturally holds in the counterclockwise case, that is, going from the right palm to the left palm, as well. The difference between Theorem 15 and Theorem 12 is that now P need not have all stable edges, and in Theorem 15, it is not the case that both the transitions from a given edge are bidirectional.

3.4 Experimental Results

The preceding algorithm was implemented in C on a Dec-station 5000/20. For the example object, the transition graph $\mathcal{G}$ (Figure 3.8) can be generated in about one minute. Once $\mathcal{G}$ is generated, reorientation plans can be found in one or two seconds. Plans were generated to bring the object from the initial stable orientation on a flat palm, to the goal stable orientation on the goal palm, much as in Figure 3.9. The plans were tested both in simulation and on a plastic cone manipulator (Figure 3.10), mounted on a tilted air table to reduce support friction. Note that for the transition graph Figure 3.8, which was generated assuming an ideal manipulator, every equivalence region can be reached from every other equivalence region. To generate plans which could be executed by the manipulator, motions corresponding to the object sliding from one palm to the other over the central hinge point, and configurations where $\phi < 0.5$ were disallowed, and $\mathcal{G}$ regenerated. This resulted in certain arcs being completely eliminated from the graph, and they are shown as dashed arrows in Figure 3.8.

The planner can also be generalized to return a path from a given initial state to a generalized goal state, for example "resting flat on side a, either palm". This can be done by adding nodes to $\mathcal{G}$ representing this general state, and reachable only from the nodes "side a, left palm" and "side a, right palm". Naturally, if the original graph is already strongly connected, any

Figure 3.10: Plastic cone manipulator used to test plans

generalized state is reachable from any initial state, where the initial state specifies both the resting side and resting palm.

For the example object, the plans simulated tended to either be "robust" to friction between the part and the palms as high as about $\mu = 0.25$, or be extremely sensitive to the frictionless approximation, failing for friction higher than $\mu = 0.02$. Simulation of the plan shown in Figure 3.9 showed that a static coefficient of friction $\mu \leq 0.25$ would permit enough of the contacts to slide for the predictions from the frictionless approximation to be valid, and for the plan to succeed.

In the experiments conducted on the air table, the static coefficient of friction was approximately 0.19, low enough for the "robust" plans to succeed. The palms were actuated by servo motors, run open loop, and only a minimum of effort was spent on calibration. The motor behavior was assumed to be linear. The motor encoder signals were read with the palm at two positions — for the right palm motor the encoder was read with the palm at 0 degrees and 90 degrees; for the left palm motor the encoder was read with the palm at 90 degrees and 180 degrees — and these readings were used to fit a line between palm orientation and encoder values for each motor.

To evaluate the reliability of the example plan, we ran 50 trials, starting the object in its initial orientation, $\phi = -\pi/2$, at different arbitrary locations on the left palm. The varying initial configuration of the part led to variation of the part's trajectory through its configuration space, as expected. Nonetheless, of the 50 attempts, the manipulator failed to correctly reorient the object only 4 times. Despite the variation in the object's trajectory, the object always stayed trapped in the correct region of the state space and hence would be propelled along to the correct final orientation. Each of the failures seemed to be due to a single rough spot on the right palm, which caused a contact to roll rather than slide.

Figure 3.11 shows a plan for our example part which tended to fail due to frictional instability. During the transition from the state ($\theta = 3.92699$, $\phi = 2.67254$, $\beta = -0.234525$) to the state ($-\pi/2$, 2.67254, 0.234525), the part would fall off its resting edge to another edge, as

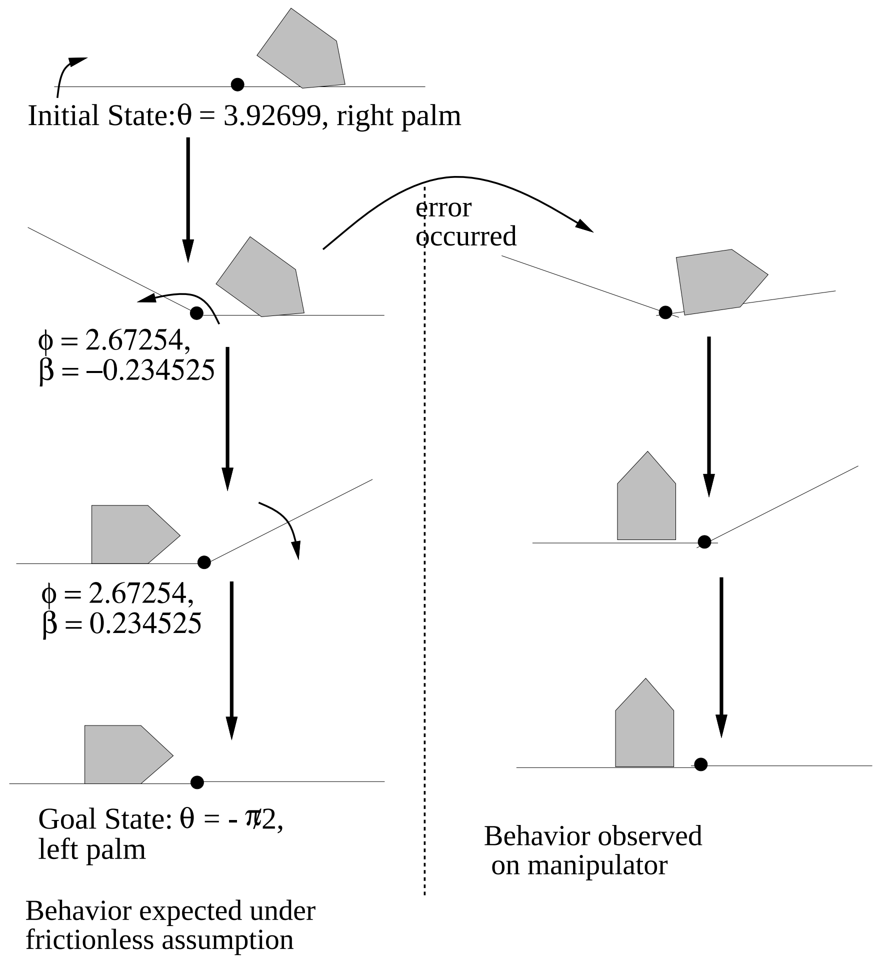

Figure 3.11: A plan which tended to fail.

shown in the right column of the figure, even though the frictionless analysis showed that the pure tilt should have been entirely stable until the part was transferred to the left palm. Figure 3.12 shows another plan which was consistently successful for the same part.

Figures 3.13, 3.14, and 3.15 show successful plans for another part (Part 2), which was basically a triangle with its tips cut off. This part had three stable edges and two unstable edges. The plans for this part tended to be quite stable with respect to friction, and failure modes tended to be caused by the narrow portions of the part getting caught in the gap between the palms. The plan shown in Figure 3.14 is one such example. The plan shown in Figure 3.15 sometimes failed due to jamming: when the palms attempted to execute the squeeze to $\phi = 0.675$, sometimes the part would become jammed in the palms, and the palms would not be able to squeeze all the way down to the desired ϕ. When this occurred, often upon executing the next step in the plan (a pure tilt), the part would transit to the orientation $\theta = 0$ on the left palm, rather than the desired $\theta = 2.05$. This problem was intermittent, and when the palms were recalibrated, actually went away.

Although one of the plans for Part 2 had jamming problems, the plans for Part 2 did not suffer from the frictional instability which some of the plans for Part 1 did. This is because for

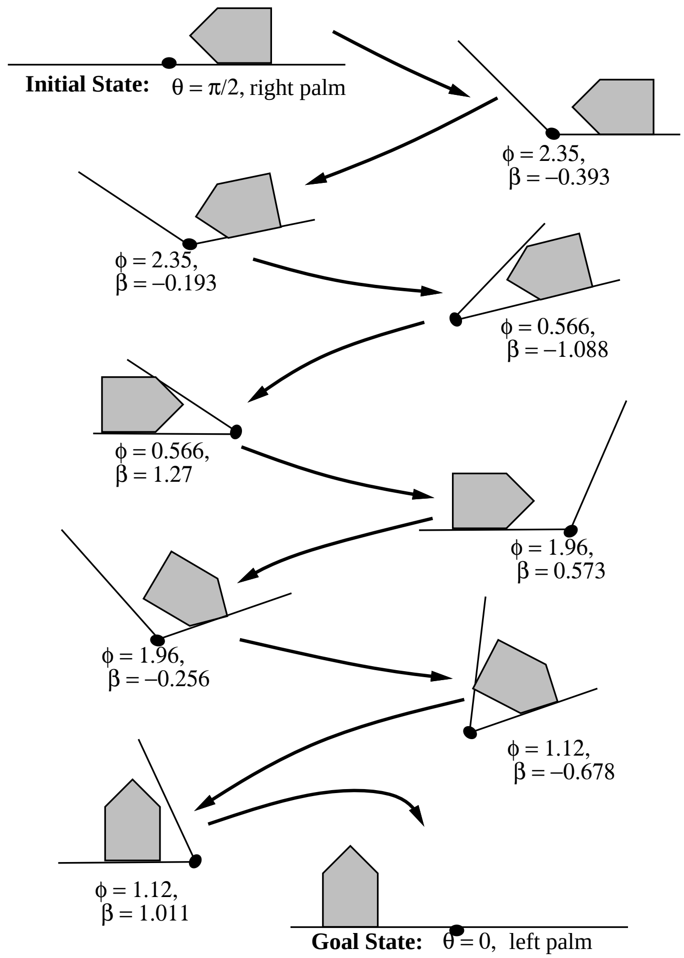

Figure 3.12: A plan which tended to succeed.

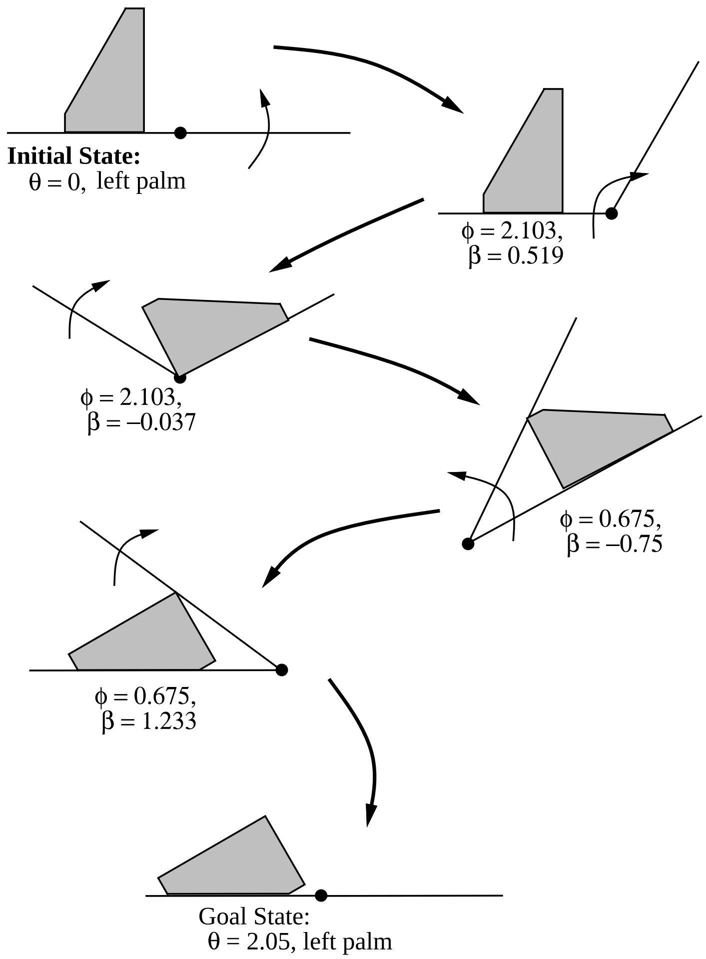

Figure 3.13: A plan which generally succeeded.

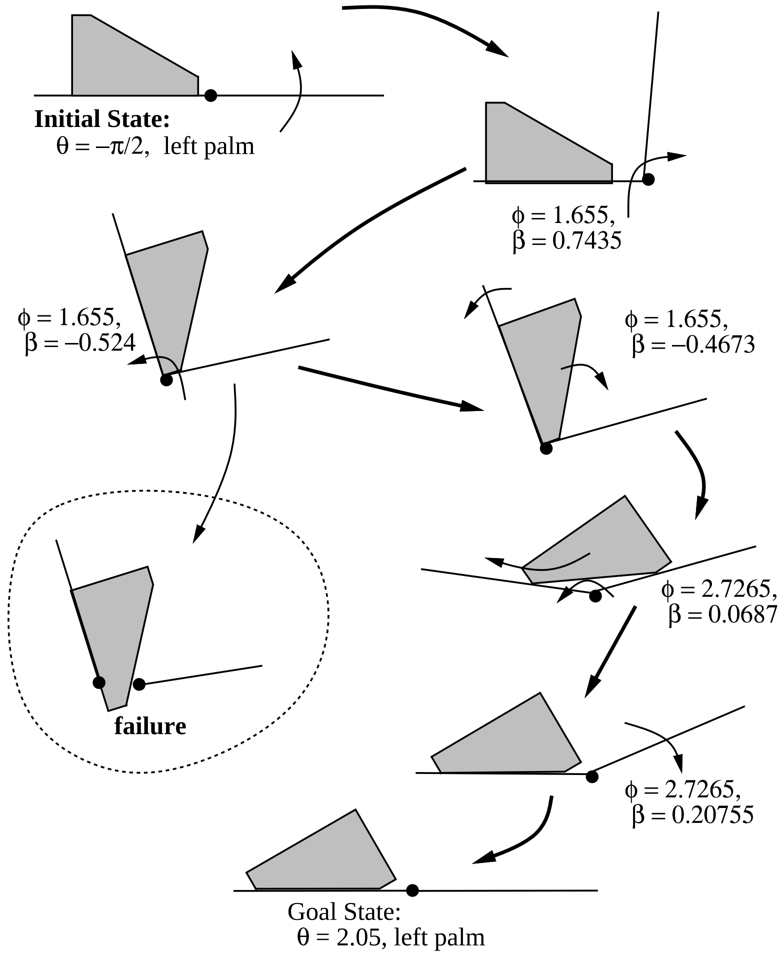

Figure 3.14: This plan would fail because the part would get caught in the gap between the palms.

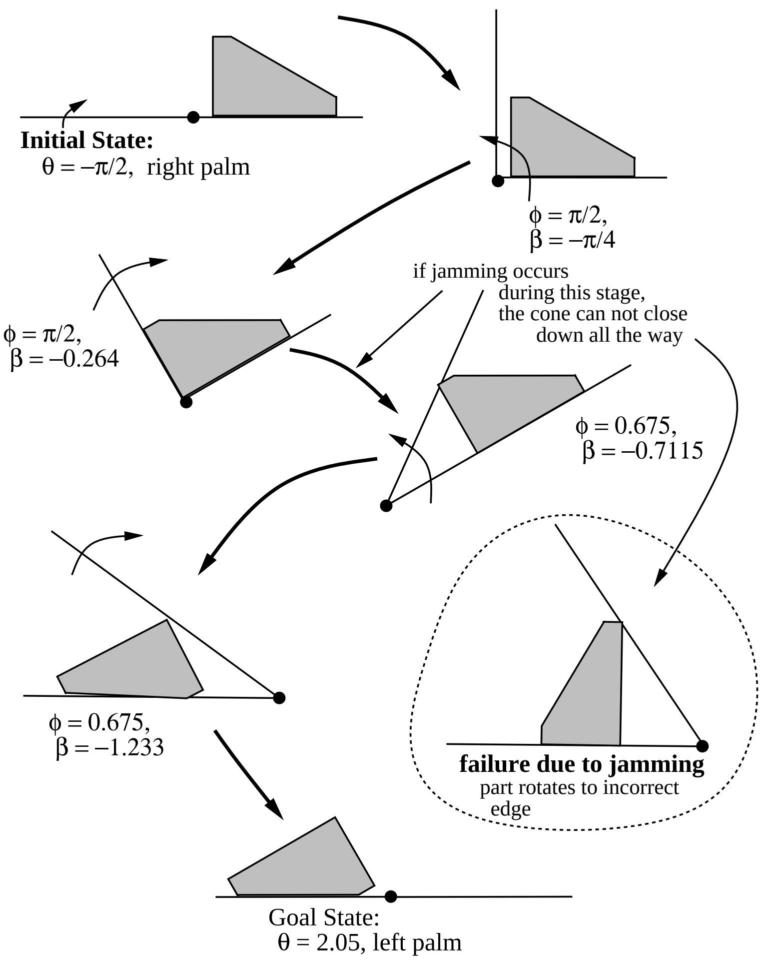

Figure 3.15: A plan which would fail due to jamming.

Run	Left Palm			Right Palm			Success Ratio
	at 90°	at 180°	RMS error	at 0°	at 90°	RMS error	
Nominal	1450	2250	-	1025	1825	-	13/15
1	1450	2225	1.44°	1025	1800	2.71°	14/15
2	1400	2225	5.43°	1025	1775	5.43°	15/16
3	1375	2250	8.14°	1025	1750	8.14°	13/16
4	1375	2225	7.98°	1000	1750	7.98°	9/16
5	1350	2250	10.66°	1050	1725	11.23°	13/16
6	1450	2250	0.00°	1025	1700	13.92°	15/15
7	1325	2175	13.4°	1025	1825	0.00°	4/15
8	1550	2250	10.85°	1025	1925	10.85°	15/15

Table 3.1: Trials of the plan in Figure 3.12 at different manipulator calibrations

all the stable edges, Part 2 is "long and low": that is, with the center of gravity as the origin, the tangential distances to the vertices of a given edge are long compared to the perpendicular distance to the edge itself. This is less true for Part 1; in particular, it is not true for the edges labelled "c" and "d" in Figure 3.2. A quantifiable definition of "long and low" is dependent on the coefficient of friction of the palms, and a analytical condition for frictional instability will be given in Chapter 4.

3.4.1 Robustness to Calibration

A useful but unexpected side effect of Lemma 9 is that orientation plans display robustness to small miscalibrations of the manipulator. To test this, we took the plan shown in Figure 3.12 and ran it with different palm calibrations. The results are summarized in Table 3.1. The motors are assumed to be linear, so each calibration is a line describing the mapping between the motor encoder signal and the palm orientation. The calibration for the left palm was made with the palm at 90° and 180° degrees, and for the right palm at 0° and 90° degrees. The encoder signals at these positions for each test calibration are given in the Table 3.1, for comparison of the calibrations. The *RMS error* of a palm calibration C with respect to a reference calibration is the square root of the mean squared distance, in degrees, between corresponding points on the two lines, taken over the same interval of encoder signals. The errors are taken with respect to a nominal calibration, which is shown in line 1 of Table 3.1. Roughly, RMS error is the average magnitude of the discrepancy in degrees between the commanded position of the palm and the actual position of the palm. The right palm orientation was assumed to vary between 0 and 150 degrees, and the left palm orientation was assumed to vary between 30 and 180 degrees. The success ratio is the number of successful reorientations over the the number of total trials.

The manipulator trajectory for the reorientation plan we used in these trials is generated by interpolating between the control positions shown in Figure 3.12. The interpolations are either pure tilts or fixed-θ squeezes, depending upon whether the quantity ϕ or $\partial\beta/\partial\phi$ is preserved between two given points (in other words, whether the line in (ϕ, β) connecting the two points is of the form $\phi =$ constant or $\varepsilon\frac{\phi}{2} - \beta =$ constant. In Table 3.2 we give the maximum deviation from the nominal for the trial trajectory, taken over the control positions. The maximum deviation is the difference in degrees $\phi_{test} - \phi_{nom}$ that would result from the motor commands for the test and nominal calibrations. Note that in Tables 3.1 and 3.2 we have taken the difference, not from

Run	Maximum ϕ Error (degrees)	Nominal Position (ϕ, β) in radians
Nominal	-	-
1	2.7948	(0.5664 , -1.0878)
2	-8.4011	(0.5664 , -1.0878)
3	-16.7844	(0.5664 , -1.0878)
4	-11.1896	(0.5664, -1.0878)
5	-19.6084	(0.5664, -1.0878)
6	0.3702	(0.5664, 1.2712)
7	-19.6521	(0.5664, -1.0878)
8	22.3793	(0.5664, -1.0878)

Table 3.2: Maximum deviation from nominal calibration for the trajectory in Figure 3.12 at different manipulator calibrations

the absolutely "correct" calibration, but from a nominal calibration that we have determined to be approximately correct. The nominal position (over the control positions) for which the maximum error is achieved is given as (ϕ, β) in radians, for compatibility with Figure 3.12.

The sign and positions of maximum error are consistent with the fact that many of the failures observed were due to the palms squeezing to such a small angle that the part was forced off the palms. Note that Run 8, where the maximum error is large but positive, performed quite well. Other failures were due to centrifugal force causing the part to slide off the palms; this was not observed in the trials of the nominal calibration. This mode of failure is probably also due to the palms' relative angles being too small, forcing the part further out away from the cone vertex and increasing the centrifugal effects. Few of the errors observed (mostly in Run 7) were actually due to the part landing on the wrong side after a tilt.

Chapter 4

Friction

> In other words, if the model does not succeed in transforming reality, reality must succeed in transforming the model.
>
> Italo Calvino
> *The model of models* [27]

Empirically, some edges which are considered stable in the frictionless analysis may be unstable in practice, in the sense that the object may not stay on that edge even when the palms make motions which stay within the corresponding equivalence class. This may be because the object has sufficient energy to escape its energy well, or it may be because friction causes a rolling contact where sliding was expected. In this chapter, we will evaluate the stability of edges with respect to their tendency to roll rather than slide.

Definition 4.1 (Frictional Instability) *Suppose we have a stable edge e_s, and its associated equivalence region, $\mathcal{S}$, when resting on palm k. We will say e_s is* frictionally unstable on palm *k if there are trajectories through (θ, ϕ, β) space which are contained in $\mathcal{S}$, but which fail in practice because of frictional effects: contacts roll which are expected to slide.*

After identifying frictionally unstable edges, the planner can be augmented with this information, in order to take frictional effects into account.

There is a fundamental tradeoff between the completeness of a given model of the world and its complexity. Because in this work we are interested in designing a planner which is fast and relatively simple to use, in terms of the number and type of its input parameters, we must choose computational simplicity over completeness. The extent to which we consider friction in our world model is far from complete. We consider only two point contact (that is, contact with the resting palm, ignoring the other palm) when looking for rolling contact points, and so will be unable to detect jamming or wedging. We attempt to justify this on two grounds. The first, as we have mentioned, is computational speed. From an offline computation point of view, we will only consider the part's resting edge and the resting palm (the two pieces of information which define an equivalence class), rather than consider all possible contact configurations, which gives us $2N$ cases to consider, rather than $\mathcal{O}(N^2)$. From an online computation point of view, we do not wish the overhead of sensors, so we must forgo exact knowledge of state; in particular, whether we are actually in two or three point contact. One can argue that this type of sensory information is binary in nature, and low overhead — a valid argument — but by not having it, we can also dispense with those model parameters which would be required to carry this information around. A lower parameter model of the world is more tractable, and naturally faster to compute. In our case, the frictional model is linear in the number of stable edges, rather than polynomial. We must distinguish between two and three point contact if the difference in these states crucially affects a decision which the planner must make. If in fact the decision can be made independently of this information, and still be a correct decision, then the information is not necessary to the planner.

The second point is that this manipulator technique fundamentally relies on low frictional contact in order to reliably reach potential energy minima without jamming. The potential energy based method of finding stable states is both easy to compute and robust. A minimal energy state in a frictionless world is still a minimal energy state in the presence of friction. An object which is stably supported in the absence of friction will still be stably supported in its presence. A grasp or stable support which is predicted using a particular coefficient of friction, on the other hand, may not be stable if the actual coefficient of friction differs from that of the model. Hence, it can be brittle in the face of model uncertainty. Furthermore, because the location of minimal energy states varies smoothly (linearly) with the palm positions (as long as we remain inside a single equivalence region), the frictionless approach is also robust to small errors in palm calibration. A friction based system may not be. We will return to these points in Chapter 5, when we will compare our system to another sensorless two palm system [33], which uses frictional grasps. For now, we will present the frictional contact model which we incorporated into our planner.

4.1 Analysis techniques for planar systems

A planar system has a three dimensional configuration space: two dimensions for location in the plane, and one for rotation. Similarly, it has a three dimensional force space: two components for translational force, and one for moment about the origin. In order to keep the dimensional units of the configuration space consistent, the orientation axis is given in units of orientation (in radians) times the radius of gyration of the object in question. Similarly, the torque component of the corresponding force space is given in units of (force times moment arm) divided by the radius of gyration. In both cases, this normalization ensures that all three axes of the space are given in the same units: length, for configuration space, and force for force space. Therefore, the radius of gyration of the object of interest is a natural unit of length for performing force analysis. Since torques and rotations will generally be given with respect to the center of gravity (CG) of the object, the CG is the natural origin for our coordinate frame. In all the analysis given in this section, we will therefore use a coordinate system with the center of gravity of the object of interest as the origin, and the object's radius of gyration as the unit of length.

Suppose we are given a diagram such as those in Figures 4.1 and 4.2, and then asked the question, "What happens to this system when a particular force is applied?" In the quasistatic frictionless situation, there is a one-to-one mapping between forces and motion, so the question is easily answered. If there is friction, however, this mapping is no longer one-to-one, and the question becomes somewhat complicated. We would like to come up with a method for answering this question, at least in the qualitative sense. In Chapter 2, we reviewed some techniques for planar frictional quasistatic analysis, based on those presented in [51], [76], [24], and [32]. We will now apply those techniques to the problem at hand.

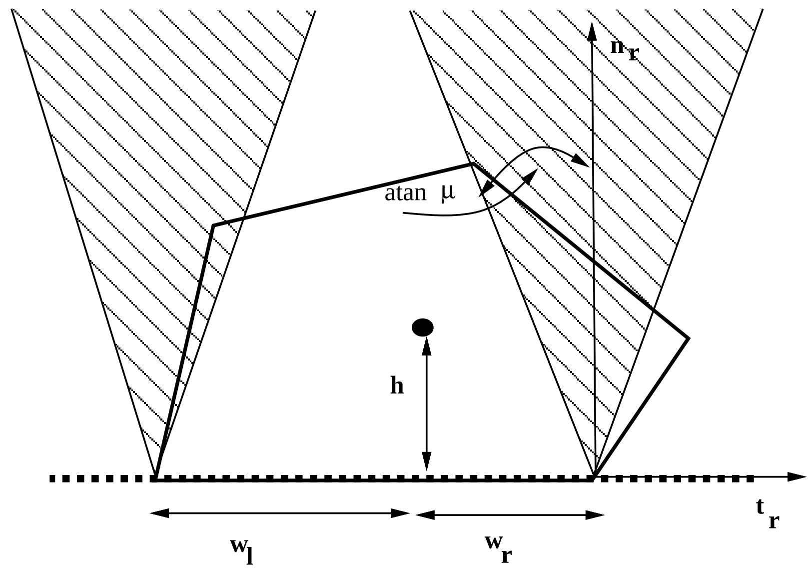

Figure 4.1: Friction cones at contact vertices

4.1.1 Friction Cones

We will consider a edge resting in two-point contact on a single palm. (See Figure 4.1). If the coefficient of friction between the object and the palm is μ, then the frictional force at each contact can be described by the *friction cone* at the contact vertex whose cone edges are

$$\mathbf{n} \pm \mu \mathbf{t}$$

Where $\mathbf{n}$ is the outward unit normal to the palm, and $\mathbf{t}$ is the unit tangent to the palm. For a point mass, it is easy to see that the tilt angle of the palm for which the point mass will begin to slide is given by $|\beta| = \arctan \mu$.

4.1.2 Convex cones

Now consider a body with extent. The contact forces on the body can be adequately represented, under the rigid body assumption, by only considering point contact forces; for edge-edge contact one need only consider the endpoints of the contact, as shown in Figure 4.1. The friction cones can be extended to three dimensional configuration space by representing each edge of each friction cone as a ray in three dimensional cspace

$$\mathbf{f} = \begin{bmatrix} u_x \\ u_y \\ r \times_{2D} u \end{bmatrix},$$

where u is a unit vector in the direction of the friction cone edge, and r is the location of the contact point (the vertex of the friction cone) with respect to the CG, and $\times_{2D}$ represents the two dimensional cross producti: $r_x u_y - r_y u_x$. In configuration space, all forces that can be exerted by the frictional contacts must be a nonnegative combination of these rays; in other words, they must be contained in the convex cone (the *configuration space friction cone*) formed by the rays. If we can represent all forces and accelerations of the object in configuration space, then the dynamic equations become questions of cone containment. To illustrate the point, let $\mathbf{x}$ represent the position of the object in configuration space, F_A be the net external applied force and torque on the object, $F_{contact}$ be the force applied to the part by the manipulator, and let $\mathcal{C}_F$ be the configuration space friction cone. Then the dynamic equation describing the system is

$$m\ddot{\mathbf{x}} = F_A + F_{contact}.$$

We can restate this equation as

$$m\ddot{\mathbf{x}} - F_A \text{ is contained in } \mathcal{C}_F.$$

In the problems we will be solving, we will generally know $\mathcal{C}_F$ and F_A and will be hypothesizing object accelerations (actually, we will be hypothesizing object velocities, but under the quasistatic assumption, velocities and accelerations will be in the same direction. See the discussion in Section 2.3 of Chapter 2). We can use the kinematic analysis due to Reuleaux which is described in Chapter 2 to find all feasible centers of rotation for the object motion. By converting the centers of rotation into cspace acceleration vectors, we can state the problem thus:

$$\text{Determine whether } m\ddot{\mathbf{x}} \text{ is contained in } \mathcal{C}_F \oplus F_A, \tag{4.1}$$

where the left hand side of the relation can be derived from a postulated center of rotation. and $\oplus$ denotes the operation of cone combination. In fact, we do not consider all the vectors which make up the cone $\mathcal{C}_F$, since for a postulated center of rotation, we can also determine the motion of all contact points, and hence which edges of the friction cone are active.

4.2 When does it slide and when does it roll?

Returning to Figure 4.1, we wish to determine the behavior of the object on the palm under the influence of gravity. For what orientations of the palm will the object remain stationary? For orientations of the palm steeper than that, will the object slide down the palm, or roll about a vertex? We can use the formulation from the previous section to determine the stability of an equivalence class with respect to the contact friction. That is, for each equivalence class, we can find an upper bound on the coefficient of friction such that for values of friction lower than this bound, if the part is in this equivalence class, it will behave as if its contacts with the palm which defines the equivalence class were frictionless: the contacts will slide, not roll.

Suppose we are tilting the palm clockwise (so that the tilt angle of the palm, β, is negative). Then the critical contact is the right vertex of the resting edge. Let h be the height of the center of gravity, and w_r be the tangential distance from the center of gravity to the right vertex. Consider the case shown in Figure 4.1, where the friction cone at the right contact does not contain the center of gravity. It is easily shown, using the technique described above, that if the magnitude of the palm's tilt angle is less than $\arctan \mu$, the object will stick, as expected. For steeper tilt angles, the object will slide without rolling. Informally, the object will not rotate because the contact force at the right vertex always points to the right of the center of gravity, exerting a counterclockwise moment about the center of gravity, which will be compensated for by the contact force at the left vertex.

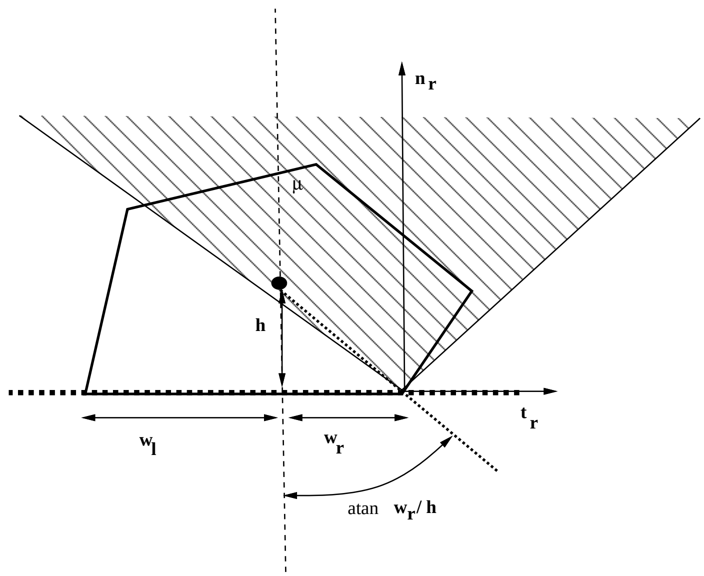

Figure 4.2: Considering the case of a clockwise tilt ($\beta < 0$): Friction cone contains CG

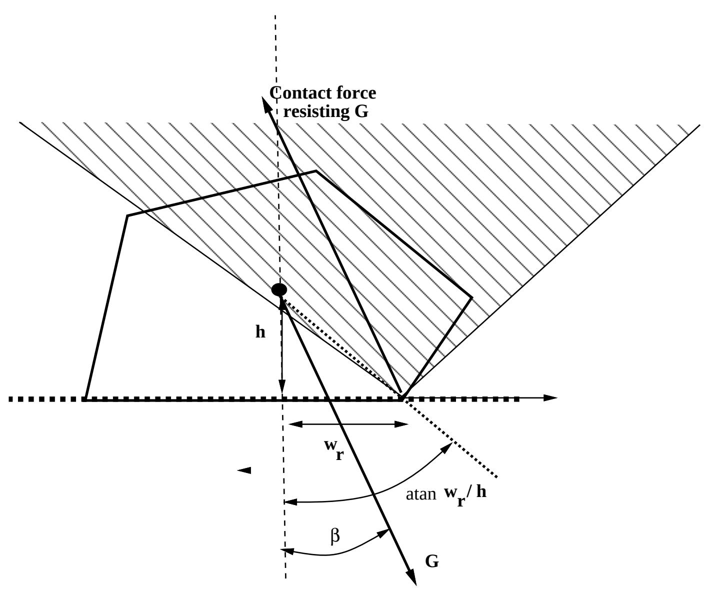

Figure 4.3: Again considering a clockwise tilt ($\beta < 0$): Friction cone contains CG, and $|\beta| < \arctan \frac{w_r}{h}$.

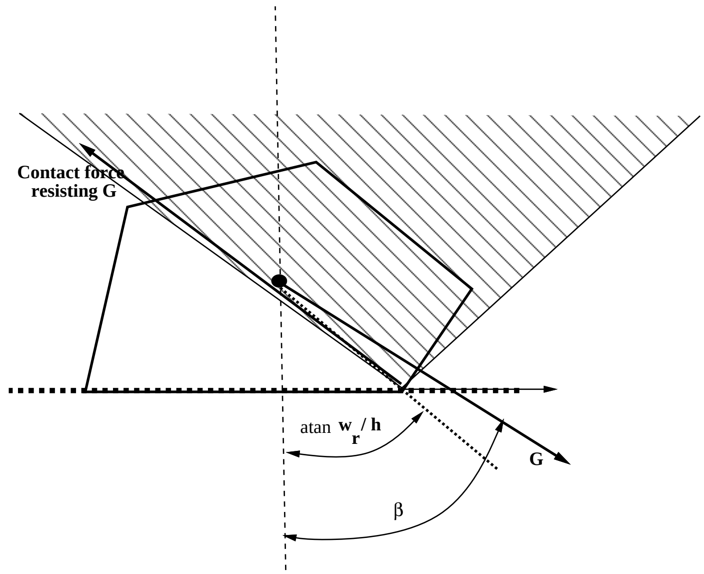

Figure 4.4: For a clockwise tilt: Friction cone contains CG, and $|\beta| > \arctan \frac{w_r}{h}$.

Now consider the case shown in Figure 4.2, where the right friction cone contains the center of gravity. Recall that we are considering a clockwise tilt, so β is negative. In this case, it can be shown that for $|\beta| < \arctan \frac{w_r}{h}$, the object will stick, and for steeper tilt angles, the object will roll clockwise about the right vertex. This is because when $|\beta| < \arctan \frac{w_r}{h}$ (Figure 4.3), the gravity vector points to the left of the right vertex, and so the compensating contact force at the right vertex will point to the right of the center of gravity. This contact force will exert a counterclockwise moment about the center of gravity, which will be compensated for by a contact force at the left vertex. Hence the object will not rotate. Because the applied forces at both contacts will be inside the friction cone, the object will stick, as well. For steeper tilt angles (Figure 4.4), gravity will point to the right of the right vertex, and the compensating force will point to the left of the center of gravity. This will exert a clockwise moment about the center of gravity, which will cause the left contact to break, and the object will rotate clockwise about the right vertex.

The above can be summarized as follows:

Consider an object resting on stable edge contact with a single palm. Let $\mu_{crit} = \frac{w_r}{h}$, and $\alpha_{crit} = \arctan \mu_{crit}$. When the palm is tilted clockwise, one of the following cases occurs:

- If $\mu < \mu_{crit}$, the object will stick if $|\beta| < \arctan \mu$. For steeper tilt angles, the object will slide without rotating.

- If $\mu > \mu_{crit}$, the object will stick if $|\beta| < \alpha_{crit}$, For steeper tilt angles, the object will rotate clockwise about the right vertex.

Similar results hold for the left vertex when the palm is tilted counterclockwise.

Using this result, we now have a stability criterion for a particular equivalence class with respect to the contact friction of our system. If the equivalence class corresponds to resting on the left palm, we consider μ_{crit} for the right vertex of the resting edge. If the object is resting on the right palm, we consider μ_{crit} for the left vertex of the resting edge. In either case, if

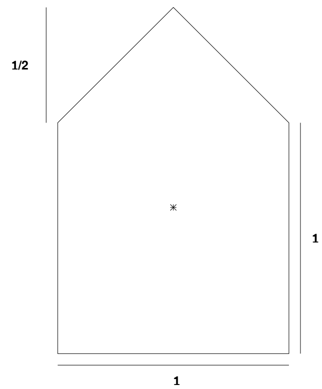

Figure 4.5: Object 1: "House"

μ_{crit} is less than the system contact friction, the resting edge will have a tendency to roll as the palm tilts out of the horizontal, and the equivalence class may be considered to be frictionally unstable. Frictional instability is determined using only the *direction* of the next tilt motion (which we know), not the *magnitude* of the tilt (which we do not know, until after we have found the plan).

Although this stability criterion cannot be incorporated directly into the transition graph of Chapter 3 in a useful way, we can use this criterion to augment the transition graph which will be described in Chapter 5. This larger transition graph takes into account uncertainty in the state of the world, not only the uncertainty associated with friction, but that of the initial state of the system as well. If a certain motion causes the object to fall into an equivalence class which is frictionally unstable, we consider this motion to have two possible outcomes. One is the equivalence class which is determined by the frictionless analysis. The other is the equivalence class which the object may roll to due to friction. For example, if a certain motion has as one of its possible output states the state (e_i, left), and this state is frictionally unstable with respect to the contact friction of the system, then the object in state (e_i, left) will have a tendency to roll clockwise. The next edge clockwise is edge e_{i-1}, and so the state (e_{i-1}, left) should also be considered as a possible output state. This augmentation will not be sufficient if the object has enough energy to roll over past edge e_{i-1} to another subsequent edge, but is valid under the assumption of low kinetic energy.

By considering both the sliding and the rolling possibilities, the resulting transition graphs will be valid for any coefficient of friction less than the coefficient used to actually construct the transition graph. It will not be valid for coefficients greater than this, however, since a greater coefficient of friction may cause previously stable edges to become frictionally unstable.

4.3 Verification

To test our friction model, we took three objects, shown in Figures 4.5, 4.6, 4.7, and tested them on the air table. The first object is the five sided figure we have been using as an example part. The second part has three stable and two unstable sides, and was also used as an example part in Chapter 3. The third part is a rectangle with width to height ratio of 3:2.

For all three of the objects, we tested all stable sides under two conditions: sliding friction coefficient of approximately 0.2, and of approximately 0.9. We then compared the sliding vs.

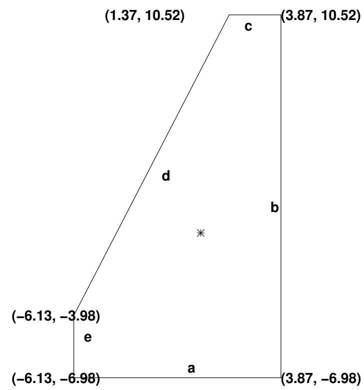

Figure 4.6: Object 2: "Pseudotriangle"

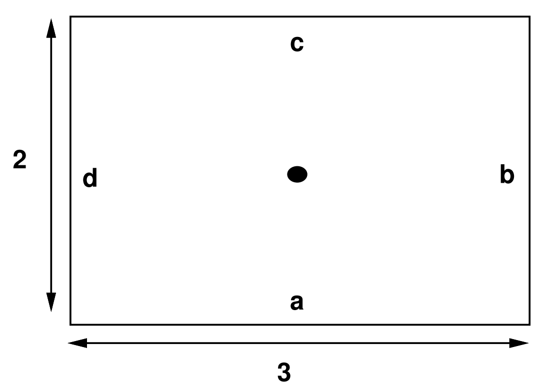

Figure 4.7: Object 3: "Rectangle"

Side	μ_{crit}	$\mu = 0.2$	$\mu = 0.9$
a	0.79	slide	roll
b	1.27	slide	slide
c	0.15	roll	roll
d	1.0	slide	slide
e	0.73	slide	roll

Table 4.1: Sliding trials for “House” on right palm

Side	μ_{crit}		left palm		right palm	
	left palm	right palm	$\mu = 0.2$	$\mu = 0.9$	$\mu = 0.2$	$\mu = 0.9$
a	0.55	0.88	slide	roll	slide	roll
b	2.7	1.8	slide	slide	slide	slide
d	1.75	2.17	slide	slide	slide	slide

Table 4.2: Sliding trials for “Pseudotriangle”

rolling behavior of each side of each part with the calculated values of μ_{crit} for both the left and right palms. The results are summarized in Tables 4.1, 4.2, and 4.3. Since Object 1 is symmetric about the y axis, and Object 3 is symmetric about the x and y axes, we have only recorded the outcome to within symmetries.

Each table shows μ_{crit} for each side, and the behavior of the part (rolling or sliding), when placed on the appropriate palm on that side. The palm was pre-positioned to an angle where sticking was no longer feasible. The outcomes observed were all consistent with the predictions made by the μ_{crit} calculation.

Side	μ_{crit}	$\mu = 0.2$	$\mu = 0.9$
a	1.5	slide	slide
b	0.67	slide	roll

Table 4.3: Sliding trials for “Rectangle” on right palm

Chapter 5

Uncertainty: Planning from an Unknown Initial State

> In all fictional works, each time a man is confronted with several alternatives, he chooses one and eliminates the others; in the fiction of Ts'ui Pên, he chooses—simultaneously—all of them. He *creates* in this way, diverse futures, diverse times which themselves proliferate and fork....In the work of Ts'ui Pên, all possible outcomes occur; each one is the point of departure for other forkings. Sometimes, the paths of this labyrinth converge...
>
> Jorge Luis Borges
> *The Garden of Forking Paths* [16]

Chapter 3 described a method for planning reorientations from arbitrary but known initial configuration, to a desired goal configuration. In many applications, such as parts feeding, the initial state of the object may not be known. If a reorientation plan can be found to reliably bring the object from any initial state to a single known final state, then the method from the previous section can be applied to bring the object from that known state to any desired goal state. In this chapter we will focus on the problem of determining a palm trajectory which will always bring the part to a single final state.

5.1 Frictionless, low energy case

Under the assumptions about the system used in Chapter 3 (all contacts slide, kinetic energy is low), the transition from initial equivalence region to final equivalence region is unique for a given cone opening ϕ. If there are $2N$ equivalence regions, then we can in principle build the 2^{2N} elements of the *power set* of the equivalence regions: that is, the set of all possible combinations of the equivalence regions. For instance, if we have a set of equivalence regions {A, B, C, D}, then the power set of this set of equivalence regions would be { {A}, {B}, {C}, {D}, {A, B}, {A, C}, {A, D}, ... {A, B, C, D}}. Given some set of palm motions, we can build a larger transition graph, $\mathcal{GG}$ where each node is an element of the power set, and each arc corresponds to a palm motion. Each arc then connects a set of initial states to its corresponding set of final states. For instance, suppose for a given palm motion, if the object was initially in state A, the motion will transfer the system state to C. If the object started in state B, the same palm motion will transfer the system to state D. Then the graph $\mathcal{GG}$ would include a node corresponding to the set {A} with an arc to a node corresponding to the set{C}, a node correspondong to the set {B} with an arc to a node corresponding to the set {D}, and a node corresponding to the set {A, B} with an arc to a node corresponding to the set {C, D}. All the arcs in this example correspond to the same palm motions. We can then do breadth-first search over $\mathcal{GG}$, starting from the set of all possible initial states, in the hope of finding a sequence of arcs that will take the system to a node where only one state is possible. We will call such a sequence of arcs a *homing sequence.*

Definition 5.1 (Homing Sequence) *A* homing sequence *is a sequence of manipulator motions such that no matter what the initial state of the part is, the final state of the part is always the same.*

If we wish to find a homing sequence which homes to a particular final state, we can do a breadth-first backchaining search from our desired final state, hoping to find a path backwards to the set of all possible initial states.

In practice, since an arc will generally correspond to bringing an object from resting on one palm to resting on the other palm, we do not have to consider all 2^{2N} elements of the power set. We will generally have to consider the set of all initial states, all combinations of equivalence regions corresponding to resting on the left palm, and all combinations of equivalence regions corresponding to resting on the right palm (excluding the empty set), for a total of $2(2^N - 1) + 1 = 2^{(N+1)} - 1$ sets of combinations of states.

Figure 5.1 shows a homing sequence found for our example object. For each cone opening considered, there were two possible motions: One type of motion considered was to start with the left palm horizontal and the palms fixed with cone opening ϕ, and tilt both palms clockwise, keeping ϕ fixed, until the right palm is horizontal. The other type of motion was to start with the right palm horizontal, and tilt clockwise until the left palm is horizontal. In Figure 5.1, each arc is labeled with the ϕ used, and the direction of the tilt.

In Chapter 7, we will address the questions of knowing whether or not a homing sequence is possible for a given object, and of finding a set of palm motions (or in our example, a set of ϕs) which is sufficient to produce a homing sequence.

5.1.1 Experimental Results

A planner to find homing sequences using the frictionless quasistatic assumption was written. The above plan was one of the sequences found for the example object. However, when the plan was tried on the airtable system, it failed regularly. The problem was the frictional instability of the circled state in the third stage of the plan shown in Figure 5.1. The object would often roll off this edge, into another state not anticipated by the planner. Interestingly, the object was reliably brought to the same final state from all initial states, just not the state anticipated by the planner. This suggests the possibility of learning the dynamics of the interaction of parts with the palms by observation, and constructing homing sequences from the learned transition graph. This approach was in fact used in [28] for another sensorless part orienting system.

5.2 Extensions with Friction

If the assumptions about the system in the previous section are violated, then the object can end up in a state not anticipated by the planner. In terms of the power set graph, this means that an arc out of a particular node may not end up at a node that fully describes the possible states

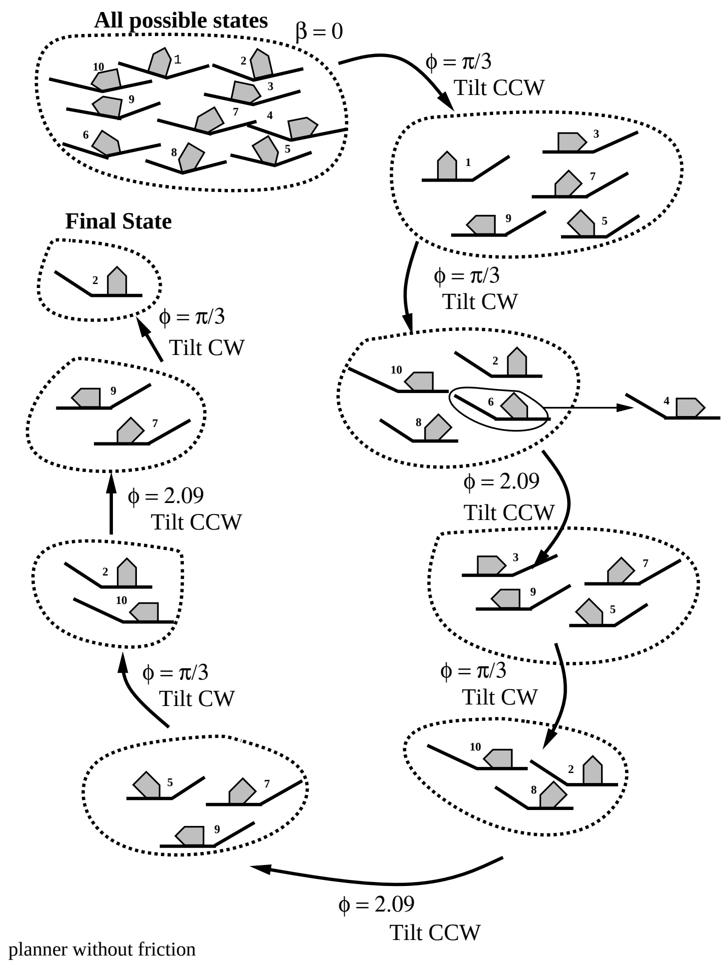

Figure 5.1: Example homing sequence. The circled configuration in the third stage would often collapse to another state, shown, which was not predicted by the planner.

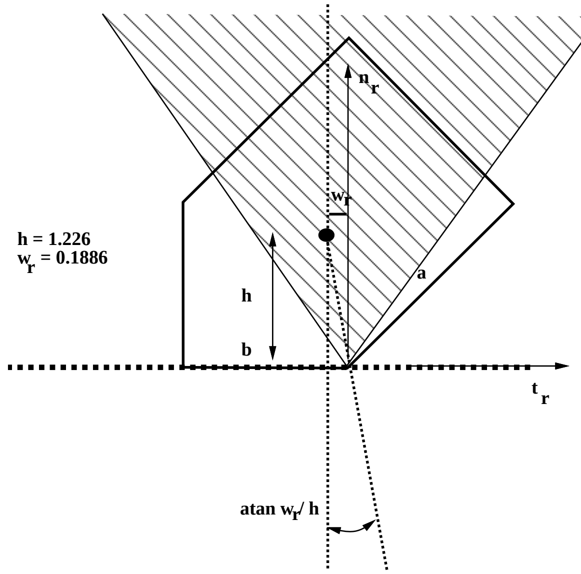

Figure 5.2: Friction cone contains CG ($\mu > 0.154$). Edge b will either remain stationary or rotate about its right vertex for clockwise rotations of the palm.

of the system. One reason for this could be the frictional instability of a particular nominally stable edge. If one of the edge contacts rolls, rather than slides, the palms may follow a path which is entirely contained in the equivalence region according to the frictionless analysis, and yet the object may transit to another, unpredicted state. This was one of the reasons that the plans generated by the algorithm in the previous section would fail in practice.

Fortunately, frictional instabilities can easily be incorporated into the power set approach, with the additional assumption that the object will roll to the next stable edge and stay there, rather than tumble further beyond to yet another stable edge. Chapter 4 describes an analysis for identifying edges which are potentially unstable with respect to the contact friction of the system. By knowing which edges will likely have rolling contacts, and in which direction they will roll, we can identify the additional states which may result for a given palm motion. For example, in the configuration shown in Figure 5.2, $h = 1.226$ and $w_r = 0.1886$, hence $\mu_{crit} = 0.154$. If the system coefficient of friction $\mu_0 > \mu_{crit}$, then we know that when the polygon is resting on the left palm on edge b it will have a tendency to roll over the right vertex, onto edge a. Note that this example is the symmetric counterpart to the problem case in Figure 5.1.

While generating the transition graph for the power set of the object states, we take both (a, left) and (b, left) as possible end states for motions which the frictionless analysis says should result in a transition to state (b, left). This covers both the possibilities that the object will either slide or roll. We can then search this graph for a homing sequence. Because we incorporate the possibility of both sliding and rolling contacts, the transition graph and any resulting homing sequences will be valid for any coefficient of friction from 0 to μ_0. For frictional coefficients higher than μ_0, the graph may no longer be valid, since edges which were assumed to always slide may become frictionally unstable.

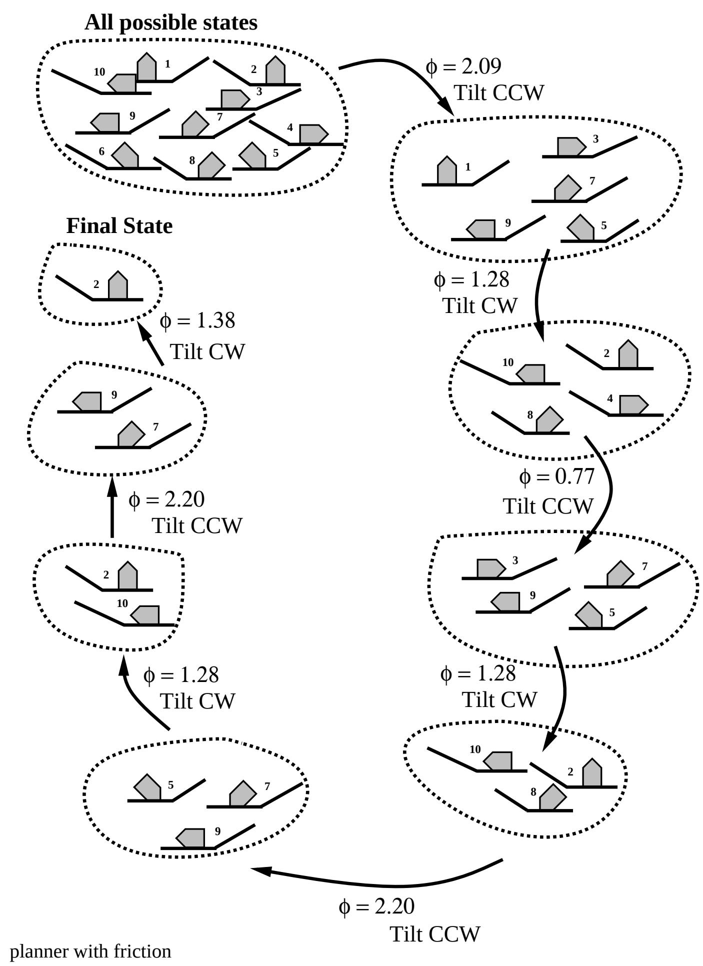

Figure 5.3: Example homing sequence

Goal State	(# successes)/(# trials)
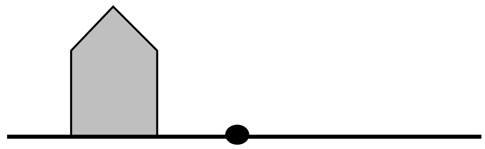	26/30
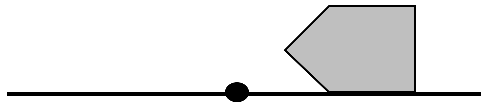	12/15
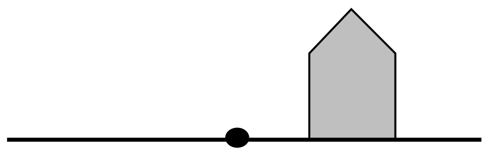	14/25
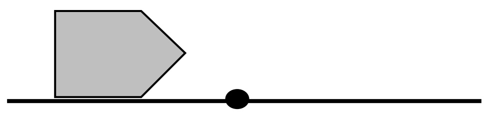	15/15

Table 5.1: Trials of homing sequences

5.2.1 Experimental Results

The planner was extended to take frictional instability under consideration. Figure 5.3 shows a homing sequence found for our example object, using the set of cone openings shown to build the arcs. The coefficient of friction, μ_0, was taken to be about 0.2. The sequence was run with the object started in all ten of the possible starting conditions, and was successfully brought to the goal state from each initial state. Then the object was dropped into the palms, into an arbitrary initial state, and the sequence was executed. Out of 30 such trials, the object failed to reach the desired goal state 4 times.

The planner attempted to find homing sequences to home the example object to all possible final states. However, it could only find sequences for four of the possible ten final states, even though only two of the edges of the object had a frictional instability for a coefficient of friction of 0.2. Each of these sequences was tested, as above. The results are summarized in Table 5.1.

Although orientation plans for the other six resting states of the object were not found, those which were found were quite repeatable, as shown in Table 5.1.

This version of the planner can also be used to find frictionally reliable paths from known initial states as well as from unknown ones. Instead of backchaining all the way back to the set of all states, the backchaining terminates upon reaching a node which contains the desired initial state.

We also evaluated the plan shown in line 1 of Table 5.1 with respect to manipulator calibration errors, using the same set of calibrations as in Section 3.4.1. The results are summarized in Table 5.2.

Table 5.3 gives the maximum ϕ error and its position over the control positions of the trajectory. The same comments apply here as in Chapter 3. Many errors were from tight

Run	Left Palm			Right Palm			Success Ratio
	at 90°	at 180°	RMS error	at 0°	at 90°	RMS error	
Nominal	1450	2250	-	1025	1825	-	15/15
1	1450	2225	1.44°	1025	1800	2.71°	14/15
2	1400	2225	5.43°	1025	1775	5.43°	15/15
3	1375	2250	8.14°	1025	1750	8.14°	6/15
4	1375	2225	7.98°	1000	1750	7.98°	9/16
5	1350	2200	10.66°	1050	1725	11.23°	7/15
6	1450	2250	0.00°	1025	1700	13.92°	12/15
7	1325	2175	13.4°	1025	1875	0.00°	5/15
8	1550	2250	10.85°	1025	1925	10.85°	8/15

Table 5.2: Trials of a homing sequence at different manipulator calibrations

Run	Maximum ϕ Error (degrees)	Nominal Position (ϕ, β) in radians
Nominal	-	-
1	2.7884	(0.7725, 1.1845)
2	-8.4011	(0.7725 , -1.1845)
3	-16.8026	(0.7725, -1.1845)
4	-11.2017	(0.7725, -1.1845)
5	-19.6392	(0.7725, -1.1845)
6	0.3702	(0.7725 , 1.1845)
7	-19.6392	(0.7725, 1.1845)
8	22.4034	(0.7725, -1.1845)

Table 5.3: Maximum deviation from nominal calibration for the plan in line 1 of Table 5.1 at different manipulator calibrations

squeezes and centrifugal force causing the part to fall from the palms. The failures due to the part landing in an unpredicted state were seen mostly in Runs 5, 7 and 8. For Run 8, in fact, all the failures were due to unpredicted states.

5.3 Is this Friction Model Adequate?

As discussed in Chapter 4, the friction model which we use in our planner is far from complete. One notable disadvantage is that larger feasible coefficients of friction translate directly into larger model uncertainty. The drawback to using this model of frictional uncertainty is that it can render certain parts unorientable. For instance, the planner cannot find any homing sequences for our example part if the friction parameter is higher than 0.73. This is the value of μ_{crit} for which the side labelled a in Figure 5.2 becomes frictionally unstable on the right palm. Also, by taking only two-point contact into account, we cannot recognize situations where perhaps the third contact will prevent the part from rolling. More importantly, we cannot recognize jamming.

In principle, orientation plans for more resting states can probably by found by setting a tighter bound on the friction parameter used by the system (specifying an exact coefficient of friction, or a minimum as well as a maximum bound). A tighter bound would eliminate some of the uncertainty in the outcome of certain motions in our model. However, assuming an exact value for the coefficient of friction is a nontrivial assumption — Kao and Cutkosky note that "the coefficient of friction can easily vary by 30% with changes in surface texture or cleanliness" [42] — which we believe will reduce the robustness of the resulting plans. As it is, the planner will produce fewer plans as the frictional upper bound increases, but the plans are guaranteed to be frictionally robust up to that value of friction. Similarly, by only considering the two points of contact that are guaranteed to occur for a part in a given equivalence region, we do not have to worry about the exact position of the part in the palms at any given moment. As we mention in Chapter 4, we have chosen to use friction independent criteria to find stable regions. The friction model we have chosen to use is only detailed enough to identify certain places where our frictionless planner model may break down, but no more.

Let us compare our treatment of friction with the models used in other systems. The two palm manipulation system described by Erdmann in [33] assumes an exact knowledge of the coefficient of friction. Starting with the part in a known pose, and known contact configuration with the palms, Erdmann's system splits the orientation space of the palms into sectors, each of which causes the same qualitative behavior of the part (slide, roll, etc.). Palm orientation is taken with respect to the part frame, with the center of gravity fixed at the origin. This plan is then fed into a simulator to verify the predicted pose of the part in the palms, as well as the palm coordinates in the world frame. From the simulations the trajectory of the palms is determined, assuming perfectly known position of the part at every step.

A significant cause of plan failure is the slipping of a contact which is expected to roll, either because the coefficient of friction is not as the planner expected, or because of inaccuracies of the manipulators when executing a specified motion. All transitions are executed in two-point contact. However as the author himself notes, in two-point frictional contact, "equilibrium contact is possible if and only if one of...the sliding modes is possible"[33]. Hence, the reliance on two-point contact frictional stable support, rather than more robust three point contact first-order stable support, results in a general brittleness of the reorientation plans. Such brittleness can of course be offset by the use of sensors to detect slipping, which brings up a host of other complexities which are happily avoided by the system presented in the present work.

Paljug, *et.al.* ([66], [67]) present just such a system. Their emphasis is on force-closure grasps which are not form-closure, and their specific example is two planar 3-degree-of-freedom manipulators using palms to grasp a large object in a gravitational field. They are also interested in how these two robots can reposition the object in the palms without regrasping. The issue in these tasks is the trade-off between grasp (squeezing) force, which should be minimized, and grasp stability. Their method of determining stable grasps depends on the presence of friction (in that the contacts are expected not to slip), but is independent of the value of the coefficient of friction.

The arms are instrumented with tactile sensors on the palms to detect the contact point; one of the arms also has a force/torque sensor at its wrist. In order to maintain the stability of the grasp, they specifically disallow slipping: that is, they try to hold the object in such a way that the contact normals are always in the interior of the friction cones. In order to maintain the object in the aggregate friction cone without having to know the exact coefficient of friction, the planner (or grasp controller) constrains the normal of the palm to be aligned with the surface normal at the contact. This requires that the surface geometry be known, at least in the neighborhood of the nominal contact points; also, the resulting nonlinear equations which must be solved by the planner are generally not solvable in closed form for geometries more complicated than spheres or cylinders. A more significant problem is that, at least for the experimental setup described in [67], orientation error and local surface normal are not directly sensed, but derived from other sensory measurements. The input needed for feedback is hence rather noisy. Their experimental work suggests that in order for plans to succeed, the coefficient of friction must be large in order to tolerate large errors in palm orientation (with respect to the contact normal). If the errors in palm orientation exceed the angle of the friction cone, the contacts will slip, and the grasp can be lost. Conversely, the actual initial conditions of the object in the plans must be close to the nominal initial conditions in order to keep palm orientation error at a minimum. Again, part of the problem was that the object used in the experiment was a sphere; so manipulation was again performed using two-point rather than three point contact; the system would no doubt be more effective on polyhedral parts. The experiments with the sphere show that while sensors can get you out of some of the difficulties of frictional uncertainty, they are not a complete cure for the problem.

In his dissertation, Lynch [54] uses a technique somewhat complementary to ours in dealing with frictional uncertainty. Using a one-degree-of-freedom arm, he tries to find a manipulator trajectory which brings a part from a known initial state to a goal state by executing a sequence of dynamic and stable grasps, throws, rolls and catches. He implicitly models the contacts as having infinite friction; rolls are modeled as rotation about a pin joint. The surface of the arm is covered with a stiff foam to insure high friction contacts. The trajectory planning problem is formulated as a nonlinear constrained optimization problem over the space of manipulator acceleration profiles, with the initial and goal states of the part as input, and the coefficient of friction as a parameter to be minimized. The no-slip constraint is explicit in the problem formulation. The result is an open-loop manipulator trajectory which executes the desired reorientation and is maximally robust to friction, with respect to the no-slip assumption. Since the number of contacts is known at any state of the manipulation (two for a grasp, one for a roll, zero for a throw), Lynch was able to use the complete Coulomb model of friction in his planner. However, the planner had to directly simulate every trajectory which was evaluated during the optimization procedure, so the process of finding a reorientation plan was more computationally intensive than that used in our system. In empirical trials, the majority of problems seemed to be caused by the unmodeled compliance of the foam which lined the arm, and unmodeled motor frequency response. This was especially true of throws, where a small error in the release state can be exacerbated during the ballistic phase, leading to systematic errors in the final state. There did not seem to be serious errors due to slipping, although slipping did in fact sometimes occur.

From the examples of these three systems, in addition to our own system, one can see that there is a certain advantage to assuming an extreme situation with respect to the coefficient of friction — either no friction, or infinite friction. Of course, one must make sure that the environment is a reasonable approximation to whichever assumption one makes, and that the reorientation plans one uses are tolerant of minor violations of that assumption.

Another drawback to our friction model is that we do not specify whether we are in two or three point contact when we are looking at the (frictional) contacts. Because of this, we cannot identify configurations in a frictionally unstable equivalence region which would not roll to another state because the third contact prevents this. We also cannot identify configurations which could jam, due to friction. On the other hand, we do not know the exact contact configuration at all times; we do not know the exact palm trajectories until after we have found the plan (*i.e.*, the path through the graph), and even when we know the palm trajectories, we do not know the exact position of the part in the palms at all times during the execution of the plan. The part could be in two point contact when we assume it to be in three point contact, or vice versa. Rather than determining the part's behavior exactly under these circumstances, it seems better to identify *potential* violations of the all sliding assumption, when we can. We can also attempt to avoid potential jamming situations by minimizing the contact friction as much as we can, and perhaps putting a lower bound on the relative angle between the palms.

We close with another statement by Kao and Cutkosky:

> Errors in modeling contact conditions and friction limit surface are more serious than short term violations of quasistatic motion [namely, stick-slip, inertial and viscous effects, so long as those effects are small on average: "short in duration, zero-mean and uncorrelated over the duration of the trajectory"[43]]. Fortunately, the sensitivity of the *direction* of sliding motion is low except when... [the tangential contact force] $f_x \approx 0$. [43]

This is part of the conclusion of an exploration of the *quantitative* predictive power of the quasistatic model for dextrous manipulation (as the authors point out—and we have also pointed out, in Chapter 2—most quasistatic motion planning is primarily *qualitative* in its predictions). However, we believe that it is still a valid statement in our context, in that relying on an exact value for the coefficient of friction may not lead to reliable plans, if the estimated value is not correct. To the extent that we use imprecise knowledge of the coefficient of friction, or of the contact conditions, we will be unable to find optimal plans; perhaps we will not be able to find a plan at all. However, whatever plans we do find will be less likely to fail due to modelling error.

Chapter 6

Dynamics and Impact

There can be no space nor any part of space without gravitational potentials; for these confer upon space its metrical qualities, without which it cannot be imagined at all.

Albert Einstein
Ether and Relativity [31]

When the velocities of the system are larger than the quasistatic bounds, forces and velocities are no longer in as close correspondence as has been assumed in the previous sections. Centripetal and coriolis effects may begin to be significant. In addition, nonzero relative velocities generate impact forces as contacts are made. We would like conditions on the relative motion of the hand and object such that the hand aquires the object. We would further like the position of the object in the gripper to be known, that is, the part comes to rest in the palm in a predictable position, rather than tumbling to another position.

In this chapter, we will look at some conservative approximations which can be used to estimate bounds on the manipulator velocity, in order to minimize dynamic effects. In addition to assuming knowledge of the CG and radius of gyration, we will use a rigid body impulsive impact model, as described by [80] and used by [93], [94], and [98], with the mass of the hand being much greater than the mass of the object, so that the change in velocity due to collision takes place entirely in the motion of the object, and the motion of the hand is unaffected. We will also assume that at the moment of impact, the dominant force is the impact force.

6.1 Impact

We will show that taking into account the energy dissipation due to inelasticity of collisions can help us to bound the effects of impact forces on our system.

The object is parameterized by $(x_t, x_n, \rho\theta)$, where (x_t, x_n) are the coordinates of the center of gravity of the object with respect to the reference frame aligned with the contacting palm (see Figure 6.1), θ is the orientation of the object, and ρ is the radius of gyration of the object ([32]). For simplicity, at contact, we can abstract away the dimensions of the actual object, and think of the object as a rod, whose center of gravity is the same as the center of gravity of the object, located at length l from the contact point. The angle from the contact normal to the rod of length l in the counterclockwise direction will be called ζ.

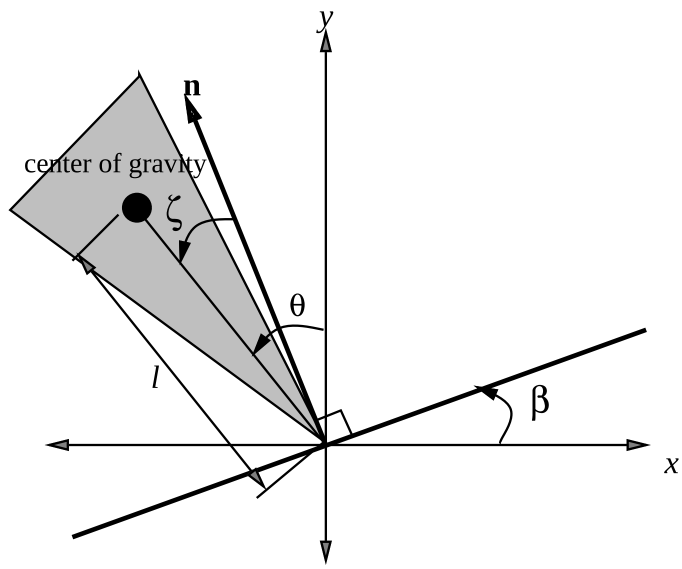

Figure 6.1: System Geometry

6.1.1 The Frictionless Case

We first look at the simplest case, where the contact friction and the support friction are both zero. In this case, the equation describing the change in velocity at impact in terms of the preimpact velocity $\mathbf{v}^- = [\dot{x}_t^-, \dot{x}_n^-, \rho\dot{\theta}^-]^T$ is

$$\Delta\mathbf{v} = -(1+e)(\hat{\mathbf{n}}^T\mathbf{v}^-)\hat{\mathbf{n}} \tag{6.1}$$

(see Chapter 2), where e is the coefficient of restitution and

$$\hat{\mathbf{n}} = \frac{\rho}{\sqrt{\rho^2 + l^2\sin^2\zeta}}\begin{bmatrix} 0 \\ 1 \\ \frac{l}{\rho}\sin\zeta \end{bmatrix}. \tag{6.2}$$

Upon impact, the contact point will bounce away from the hand, but we would like the object as a whole not to move away. One consideration is the energy constraints of the system. As a concrete example, let us look at what happens if we make a pure tilt transition from one equivalence region to another. As we make that transition, there is a brief interval of time during which the object loses stable contact with the palm, and falls into another stable contact configuration under the influence of gravity. We will model the instant it attains the new stable contact as an impulse. For the example in Figure (6.2), suppose that the cone has been rotating counterclockwise with an angular velocity ω_{man}, and up until the moment portrayed in (6.2a), the object has been tracking the movement of the cone. At the moment shown, the object is just about to lose stability, and will tilt into the preimage of the position shown in (6.2b). (The cone is at rest in (6.2b). It is possible that the object will continue to rotate beyond the desired position. In order for that to happen, it must have enough energy to escape the potential well of the desired final resting position: in other words, $E > mg\sqrt{13}$, where m is the mass of the object, g is the gravitational acceleration, and $\sqrt{13}$ is the distance from the CG to the impacting (in this case, right) vertex of the object. To gain a little leeway, we can take advantage of energy

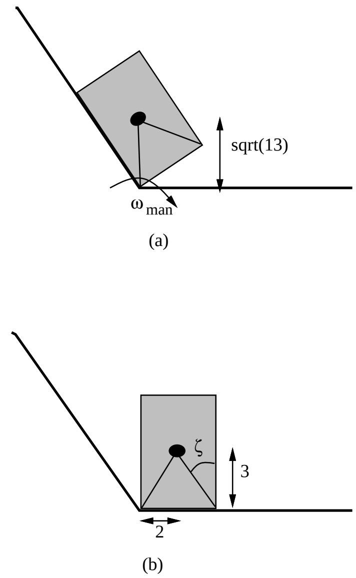

Figure 6.2: Impact example

lost due to collisions. What we would like, then, is for enough energy to be lost after the first bounce[1] for the object to be trapped in the correct energy valley.

6.1.2 Energy loss from collision: one bounce

Assuming no friction, we can look at the impact in configuration space. The impact and rebound will occur in a plane, and we can characterize the impact by the direction of the contact normal and its perpendicular in that plane. Define α to be the angle the velocity vector makes with the cspace normal at impact, CCW (actually, the negative of the velocity vector, so we get the small angle). Let $\mathbf{v}_n = \hat{\mathbf{n}}^T \mathbf{v}^- \hat{\mathbf{n}}$ be the cspace normal velocity, and v_n be its magnitude. Let $\mathbf{v_t}$ be the vector $\mathbf{v}^- - \mathbf{v}_n$, and v_t be its magnitude. $\mathbf{v_t}$ is perpendicular to $\mathbf{v_n}$, so we get

$$\alpha = \arctan(v_t, -v_n^-). \tag{6.3}$$

v_t will be unaffected by the impact, and ${v_n}^+ = -e{v_n}^-$. If we "dimensionalize" by $(v_n)^-$ (*i.e.* let it be of unit length) then ${v_t}^- = \tan\alpha$, so

$$\begin{aligned} ||\mathbf{v}^-||^2 &= 1 + \tan^2\alpha \\ ||\mathbf{v}^+||^2 &= e^2 + \tan^2\alpha. \end{aligned} \tag{6.4}$$

(See Figure 6.3.) Then

$$\begin{aligned} \frac{||\mathbf{v}^+||^2}{||\mathbf{v}^-||^2} &= \frac{e^2 + \tan^2\alpha}{1 + \tan^2\alpha} \\ &= \sin^2\alpha + e^2\cos^2\alpha. \end{aligned} \tag{6.5}$$

[1]In fact, the first impact will impart a counterclockwise impulse on the object, which will not cause it to go unstable. So we could actually look at the energy lost after the second impact, if we could predict what that impact would be. But it is easier, and more conservative, to simply look at the first impulse.

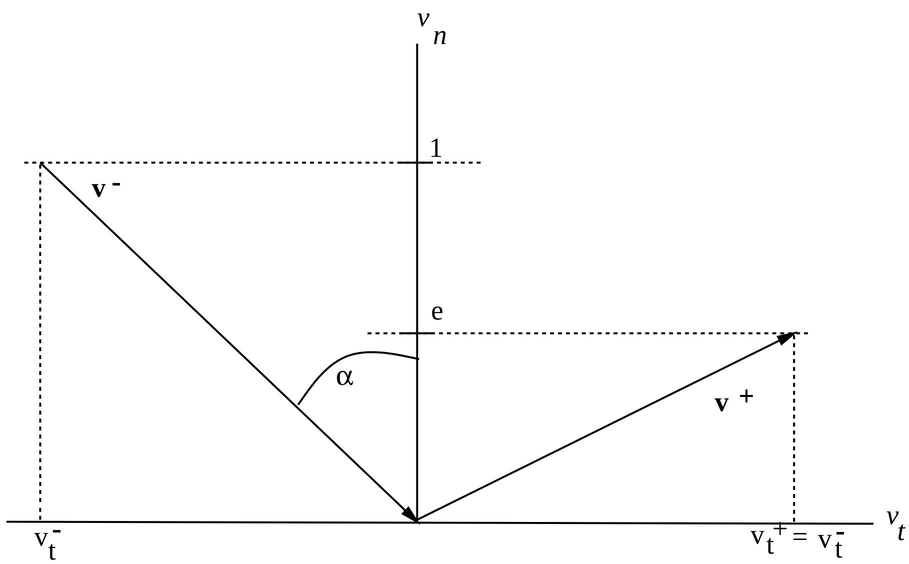

Figure 6.3: $\mathbf{v}^{-}$ and $\mathbf{v}^{+}$ in configuration space

or, in terms of the energy of the system, the proportion of kinetic energy retained by the system after impact (Figure 6.4) is

$$\mathrm{pK}(\alpha) = \sin^2\alpha + e^2\cos^2\alpha. \tag{6.6}$$

This relation depends only on α, independent of the actual cspace normal, or the orientation of the object at impact. The energy, however, depends on the configuration of object. The CG of the object will be at a specific height (and thus, a specfic potential energy) for a given impact configuration ($\mathbf{x} = x, y, \rho\theta$) (in world coordinates). Call the potential energy function $\mathrm{P}(\mathbf{x})$. Then

$$\begin{aligned} \mathrm{E}^{+}(\mathbf{x}) &= \mathrm{P}(\mathbf{x}) + \mathrm{pK}(\alpha)\mathrm{K}(\mathbf{x}) \\ &= \mathrm{P}(\mathbf{x}) + \mathrm{pK}(\alpha)(\mathrm{E}^{-}(\mathbf{x}) - \mathrm{P}(\mathbf{x})) \\ &= (1 - \mathrm{pK}(\alpha))\mathrm{P}(\mathbf{x}) + \mathrm{pK}(\alpha)\mathrm{E}^{-}(\mathbf{x}) \end{aligned} \tag{6.7}$$

where α may be a function of $\mathbf{x}$.

To return to the example in Figure (6.2), let us assume that after the object loses stability, the manipulator no longer imparts velocity to the object. This is reasonable, since the only contact is the object vertex to the cone vertex. Let us also assume, for the purposes of emphasizing the impact analysis, that centrifugal forces are negligible, in other words that the object will not go sliding up the left palm due to centrifugal forces. This is equivalent to supposing that the initial state of the cone and object were just slightly counterclockwise of the position shown in Figure (6.2a), with zero initial velocity. The energy of the object as it loses stability is

$$\mathrm{E}_o = mg\sqrt{13} + \frac{1}{2}13m{\omega^2}_{man} + \frac{1}{2}m\rho^2{\omega^2}_{man} \tag{6.8}$$

The first term is the potential energy, the second term is the translational kinetic energy imparted to the object by the rotating left palm, and the third term is the rotational kinetic

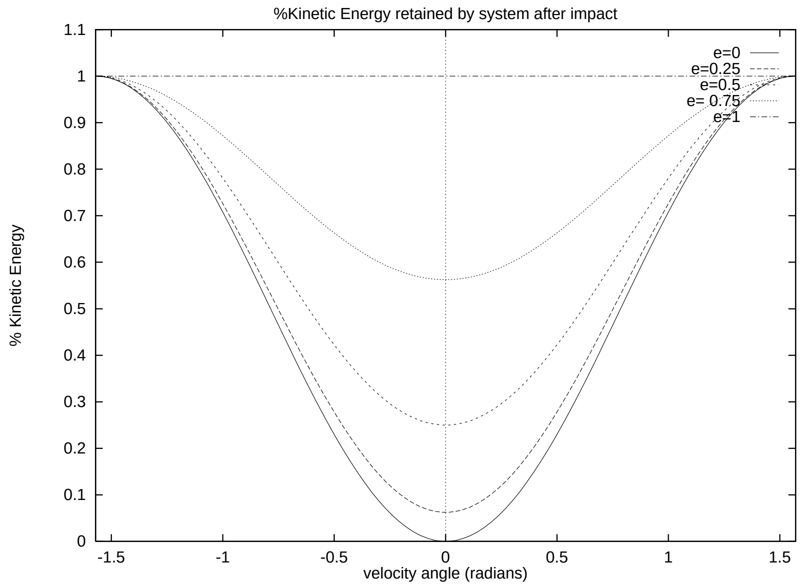

Figure 6.4: Percent of Kinetic Energy retained after collision

energy. ρ is the radius of gyration of the object. We will assume the object's mass to be uniformly distributed; then $\rho^2 = 4.33$

Upon making contact (Figure (6.2b)), the energy is

$$\mathrm{E}_f{}^- = 3mg + (\frac{1}{2}13m\omega^2{}_{man} + \frac{1}{2}m\rho^2\omega^2{}_{man} + mg(\sqrt{13} - 3)) \tag{6.9}$$

The first term is again the potential energy. The second term is the kinetic energy, represented here as the initial kinetic energy plus the difference in the final and initial potential energies. From equation (6.6), the energy remaining immediately after impact is then

$$\mathrm{E}_f{}^+ = 3mg + (\sin^2\alpha + e^2\cos^2\alpha)(\frac{1}{2}m(13 + \rho^2)\omega^2{}_{man} + mg(\sqrt{13} - 3)). \tag{6.10}$$

In order for the object to be caught in the correct energy valley, we want $\mathrm{E}_f{}^+ < \sqrt{13}mg$.

The answer to the question "how slow is quasistatic?" (neglecting centrifugal and momentum effects) is then determined by the equation

$$\omega^2{}_{man} < \frac{(1 - e^2)\cos^2\alpha}{\sin^2\alpha + e^2\cos^2\alpha}\frac{2(\sqrt{13} - 3)g}{13 + \rho^2} \tag{6.11}$$

For this system, $\cos\alpha = 0.148839$, so $\alpha = 1.4214$. Figure 6.5 shows the plot of the square root of Equation (6.11) as a function of e.

In principle, this analysis can be extended to form the "n-bounce preimages" of each stable resting state, for arbitrary n. To extend it in this way, beyond $n = 2$, however, will cost more complex computation, of decreasing additional utility.

6.2 Other Dynamic Effects

Other velocity effects can arise from centrifugal or coriolis forces, such as the object being thrown from the palms. As a simple example, suppose we have a point mass resting on one of

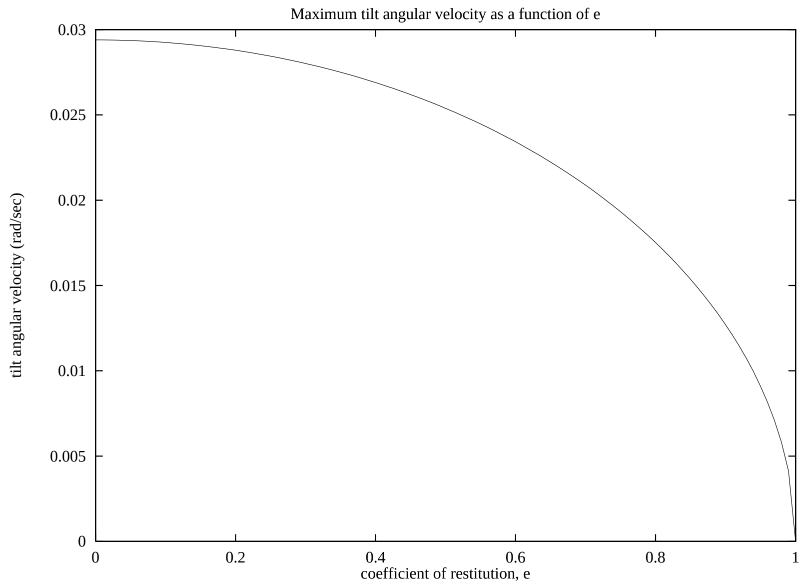

Figure 6.5: Graph of maximum ω_{man} as a function of e.

the (frictionless) palms. Set a frame fixed to the palm, with the origin at the hinge point, and the x axis along the palm. The y axis points normal to the palm, towards the point mass. The palm rotates with a constant velocity, ω. At time $t = 0$, the palm is in orientation θ_0. Using the formulas for Lagrange's equations in a rotating frame [83], we can derive the equations of the point mass.

6.2.1 Lagrange's equations in a rotating frame

For a system with N independent generalized coordinates and M constraints, Lagrange's equations of motion are given by the N equations, each of the form

$$\frac{d}{dt}\frac{\partial T}{\partial \dot{q}_k} - \frac{\partial T}{\partial q_k} = Q_k + \sum_j^M \lambda_j a_{jk} \tag{6.12}$$

where q_k are the N generalized coordinates, T is the kinetic energy of the system, Q_k are the generalized forces on the system, λ_j are the M (unknown) Lagrange multipliers, and a_{jk} are the coefficients of the M constraint relations

$$\sum_i^N a_{ji}\dot{q}_i = 0, \text{ for } j = 1 \text{ to } M. \tag{6.13}$$

If the system coordinates are in a frame which is rotating with respect to a fixed world frame, there are an additional three degrees of freedom associated with the motion of the rotating frame: (x_0, y_0, θ). There are also an additional three equations of motion [83]:

$$\begin{aligned} \frac{d}{dt}\frac{\partial T}{\partial \dot{x}_0} - \dot{\theta}\frac{\partial T}{\partial \dot{y}_0} &= \sum F_x \\ \frac{d}{dt}\frac{\partial T}{\partial \dot{y}_0} + \dot{\theta}\frac{\partial T}{\partial \dot{x}_0} &= \sum F_y \end{aligned} \tag{6.14}$$

$$\frac{d}{dt}\frac{\partial T}{\partial \dot{\theta}} + \dot{x}_0\frac{\partial T}{\partial \dot{y}_0} - \dot{y}_0\frac{\partial T}{\partial \dot{x}_0} = \sum M$$

where $(\sum F_x, \sum F_y, \sum M)$ are the total forces and moments acting on the origin of the rotating axis.

For a point mass in the cone frame, there are two generalized coordinates, (q_x, q_y), taken with respect to the cone's center of rotation, which are the position of the mass radially out on the palm and out along the normal to the palm, respectively (In other words, the palm is always parallel to the q_x axis). There is also one constraint,

$$q_y \geq 0. \tag{6.15}$$

The kinetic energy of the point mass (expressed in world coordinates) is

$$T = \frac{1}{2}m(\dot{x}^2 + \dot{y}^2). \tag{6.16}$$

The conversion from cone frame coordinates to world frame coordinates is given by

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \cos(\theta_0 + \omega t) & -\sin(\theta_0 + \omega t) \\ \sin(\theta_0 + \omega t) & \cos(\theta_0 + \omega t) \end{pmatrix} \begin{pmatrix} q_x \\ q_y \end{pmatrix} \tag{6.17}$$

$$\begin{aligned} \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} &= \begin{pmatrix} \cos(\theta_0 + \omega t) & -\sin(\theta_0 + \omega t) \\ \sin(\theta_0 + \omega t) & \cos(\theta_0 + \omega t) \end{pmatrix} \begin{pmatrix} \dot{q_x} \\ \dot{q_y} \end{pmatrix} \\ &+ \omega \begin{pmatrix} -\sin(\theta_0 + \omega t) & -\cos(\theta_0 + \omega t) \\ \cos(\theta_0 + \omega t) & -\sin(\theta_0 + \omega t) \end{pmatrix} \begin{pmatrix} q_x \\ q_y \end{pmatrix}. \end{aligned} \tag{6.18}$$

The generalized forces on the system are given by

$$\mathbf{Q} = -mg \begin{pmatrix} \sin(\theta_0 + \omega t) \\ \cos(\theta_0 + \omega t) \end{pmatrix} \tag{6.19}$$

$$\mathbf{F} = \mathbf{0} \tag{6.20}$$

Plugging Equations 6.17 and 6.18 into the expression for T, Equation 6.16, and differentiating that expression according to Equations 6.12 and 6.14, results in an expression for the motion of the point mass in the rotating cone frame:

$$\ddot{q_x} = q_x\omega^2 + 2\dot{q_y}\omega - g\sin(\omega t + \theta_0) \tag{6.21}$$

$$\ddot{q_y} = -2\dot{q_x}\omega + q_y\omega^2 - g\cos(\omega t + \theta_0) + \lambda \tag{6.22}$$

λ represents the contact force exerted on the point mass by the palm. If contact breaks, $\lambda = 0$. Otherwise, $q_y = \dot{q_y} = \ddot{q_y} = 0$

Suppose, for example, the mass begins in a stationary state resting on one of the palms at location $(q_x, 0)$, as the palms begin to rotate. By Equation 6.21, there is a force $q_x\omega^2$ which moves the mass radially outward. If either q_x or ω is high enough, the mass can slide completely off the palm, as shown in Figure 6.6a. If the object does not slide off the palm, gravity can then pull the object back radially inward towards the vertex. The object will then slide all the way down, transforming its initial potential energy into kinetic energy. This can cause it to continue sliding when it contacts the other palm, sometimes sliding off the palm, as shown in Figure 6.6b. More often, however, (in the case of objects with a rotational degree of freedom) the impact with the other palm can transfer the kinetic energy from translational to rotational motion of the object, and the object may tumble and rotate in a way not predicted by the planner.

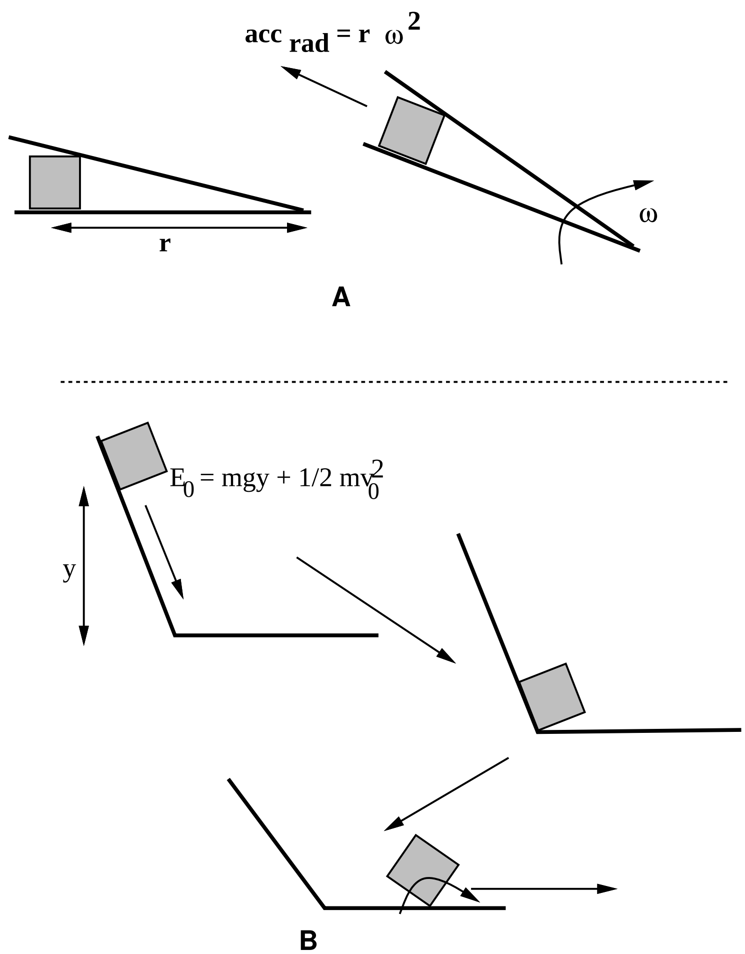

Figure 6.6: A: In the cone frame, the mass experiences a force radially outward as the cone rotates. If the radially outward component of the force on the mass dominates any forces radially inward (such as gravity), the part may fly off the palm. B: If the energy of the mass at the end of the motion is too high, when the part slides down the palm, much of its energy (minus what is lost in the impact and what is converted to rotational velocity) will be converted into sliding motion, and the part may slide off the palm. Part of the energy will be converted to rotational motion as well, and the part may rotate away from the desired edge.

6.2.2 An Example with a Complete Transition Motion

We will extend the example from Section 6.1.2, again using the configuration in Figure 6.2. This time we will analyze a complete but simplified transition motion, starting with the object at rest on the left palm. The palm then tilts with constant rotational velocity until the right palm is horizontal, stops, and waits for one second. The object then completes any further motion under only the influence of gravity. Because we are only looking for an estimation of the gross behavior of the part, we will treat the object as a point mass on the left palm during its centrifugal motion, and ignore any flight phase due to the sudden deceleration of the palms (in other words, we assume that the object is always in contact with the palm).

The cone opening is $\phi = \pi/2 + \arctan 2/3$. This cone opening was chosen so that with the object resting on the left palm as in Figure 6.2, and the right palm horizontal, the contact normal with the right palm passes through the CG of the object. This will simplify the equations.

Let the units of measurement for this example be centimeters, and the origin of the world frame be the vertex of the cone. Imagine that we start with the object at rest on the horizontal left palm, five centimeters from the vertex of the cone ($q_x = -5$ cm, $q_y = 0$). The coefficient of restitution is $e = 0.1$. The gravitational acceleration is $g = 98$ cm/s^2, or one-tenth of the usual gravitational acceleration. We do not want the object to fly off the edge of the left palm from centrifugal forces. We do want the object to transit to the configuration shown in Figure 6.2(b), without tumbling away. We wish to know how fast we can rotate the palms and still meet these two constraints. We will assume that the angular velocity of the palms, ω, is constant throughout the tilting motion.

If we take the friction of the palm into account, with coefficient of friction μ, the equation for the radial acceleration of a point mass (assuming $q_y = \dot{q_y} = 0$) is given by

$$\ddot{q_x} = \begin{cases} 0, & \text{if } \left|\frac{q_x\omega^2 - g\sin(\omega t)}{2\dot{q_x}\omega + g\cos(\omega t)}\right| < \mu \text{ and } \dot{q_x} = 0 \\ q_x\omega^2 - g\sin(\omega t) - \lambda\mu \operatorname{sgn}\dot{q_x}, & \text{otherwise.} \\ \lambda = 2\dot{q_x}\omega + g\cos(\omega t) & \end{cases} \tag{6.23}$$

The term λ comes from the force normal to the palm which the point mass experiences due to gravity and Coriolis forces.

We can integrate[2] Equation 6.23, assuming that the point mass never breaks contact with the palm. We will use a coefficient of friction $\mu = 0.2$, and a coefficient of restitution of 0.1, to estimate how much velocity is lost when the point mass hits the right palm. Using various values of ω, we find, for instance, that using a value of $\omega = -2$ rad/sec, that the mass will come to rest almost 6 centimeters out on the right palm (ω is negative, because the rotation is clockwise). For $\omega = -1$ rad/sec, the mass comes to rest about 5 centimeters out on the right palm, and for $\omega = -0.5$ rad/sec, the mass comes to rest about 1 centimeter out on the right palm. See Figure 6.7.

Returning to the problem of our example rectangle from Section 6.1.2, we shall assume that our point mass analysis conservatively estimates the sliding behavior of the rectangle. For the an object with extent, some of the energy will go into rotational motion, as well as sliding. We still need to check that the object does not tumble away from the desired orientation, shown in Figure 6.2b. Let us take $\omega = -0.5$ rad/sec. Suppose the rectangle as oriented in

[2] All numerical integrations in this section were done using the second and third order Runge-Kutta routine in MATLAB (ODE23).

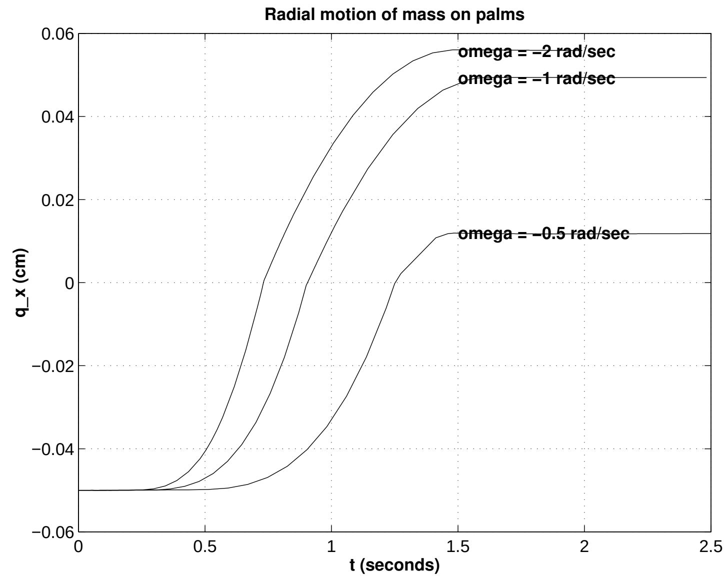

Figure 6.7: Radial motion of a point mass on a cone with $\phi = 2.1588$, tilting clockwise at -2, -1, and -0.5 rad/sec respectively. In each case the mass starts at rest on a horizontal left palm, 5 centimeters from the cone vertex.

Figure 6.2b is 7.5 cm tall, 5 cm wide; this is in the same proportions as the rectangle in the figure. The rectangle will make contact with the right palm when $q_x = -3.75$ cm. This occurs approximately at $t_{hit} \approx 1.4$ sec. Since this is less than the time it takes for the right palm to reach the horizontal, the cone is still rotating. From our numerical integration, we can determine q_x and $\dot{q_x}$. Using Equations (6.17) and (6.18), assuming that $q_y = 2.5$ cm., and $\dot{q_y} = 0$, we can determine the position and velocity of the rectangle's center of gravity (in configuration space):

$$\begin{pmatrix} x \\ y \\ \theta_0 + \omega t_{hit} \end{pmatrix} = \begin{pmatrix} -1.18 \\ 4.35 \\ -0.72 \end{pmatrix} \text{cm} \tag{6.24}$$

$$\begin{pmatrix} \dot{x} \\ \dot{y} \\ \rho\omega \end{pmatrix} = \begin{pmatrix} 16.18 \\ -13.28 \\ -1.3 \end{pmatrix} \text{cm/sec.} \tag{6.25}$$

The potential energy of the rectangle at impact $P = mg\,y = 426\,m$, in units of g-cm^2/sec^2. The kinetic energy $K = \frac{1}{2}m\dot{\mathbf{x}}^T\dot{\mathbf{x}} = 220\,m$. The right palm at the moment of impact is oriented at $\theta = 0.265$ with respect to the world frame. Because we have chosen ϕ so that the contact normal of the right palm on impact passes through the center of gravity, as in Figure 6.2a, the contact normal is given by $\hat{\mathbf{n}} = [-0.262\,0.965\,0]^T$. Hence, the impact angle $\gamma = 2.5$. The energy remaining after impact (using a coefficient of restitution $e = 0.1$) is given by:

$$E^+ = P + \mathrm{pK}(\gamma)K, \tag{6.26}$$

where $\mathrm{pK}(\gamma)$ (see Equation (6.6)) gives the percent of kinetic energy remaining after impact as a function of impact angle and coefficient of restitution. Plugging in all the numbers from our example into Equation (6.26), we determine that $E^+ = 502\,m$. The energy barrier which must be overcome to tumble to a new resting configuration is $E_{barrier} = mg\sqrt{20.3} = 442\ m$. Therefore, using a cone rotation velocity of -0.5 rad/sec, it is still possible that the the part can tumble away from the desired orientation shown in Figure 6.2b.

These estimates are quite conservative, and when we try this rotation on the air table, we find that we can use a manipulator velocity of up to about 2 rad/sec in magnitude, and the reorientation succeeds. There are other reorientations which will not succeed with a manipulator velocity this high. In practice, we have found approximately 0.4 to 0.5 rad/sec to be a fairly safe range of rotational velocity for most of the trajectories we have tested.

Chapter 7

Theoretical Results

> ...out of the pictures which are all that we can really see, we imagine a world of solid things...this world is constructed so as to fulfill a certain code of rules, some called axioms, and some called definitions, and some called postulates, and some assumed in the course of demonstration....

William Kingdon Clifford
The Postulates of the Science of Space, 1872 [29]

In this chapter, we present the proofs of results used in previous chapters. All of these proofs refer to the frictionless version of the planner.

7.1 Connectedness of Equivalence Regions

First, we present some results on the connectedness of the equivalence regions, and on the decomposition of manipulator trajectories. Once a high level path from start to goal has been found by searching the state transition graph, this high level path must be instantiated as a specific trajectory for the palms to follow. We will show that all manipulator trajectories can be composed from two simple types of motions.

By a *connected region* we mean a region X such that if $x_1 \in X$ and $x_2 \in X$, there is a curve segment whose endpoints are x_1 and x_2, and which is completely contained in X. By a *simply connected region*, we mean a connected region X without holes. More formally, a simply connected region is a connected region X that can be continuously shrunk to a single point.

A *stable edge* of an object is an edge on which the object can rest on a horizontal palm, unsupported by the other palm, without tipping over, in the face of small disturbances. In other words, a stable edge is an edge e_s such that, if one were to project the center of gravity onto e_s, this projection would lie in the interior of the line segment e_s.

In this chapter, except where noted, we will assume that we have palms of infinite length; in other words, that the palms are much longer than the object, so that arbitrarily small cone openings are feasible. We will also assume frictionless contact.

Let $\mathcal{S}$ be an equivalence region corresponding to a stable edge, and p, q be points in $\mathcal{S}$.

Theorem 1 *$\mathcal{S}$ and its projection into the (ϕ, β) plane are simply connected.*

First we recall the description of $\mathcal{S}$ from Chapter 3:

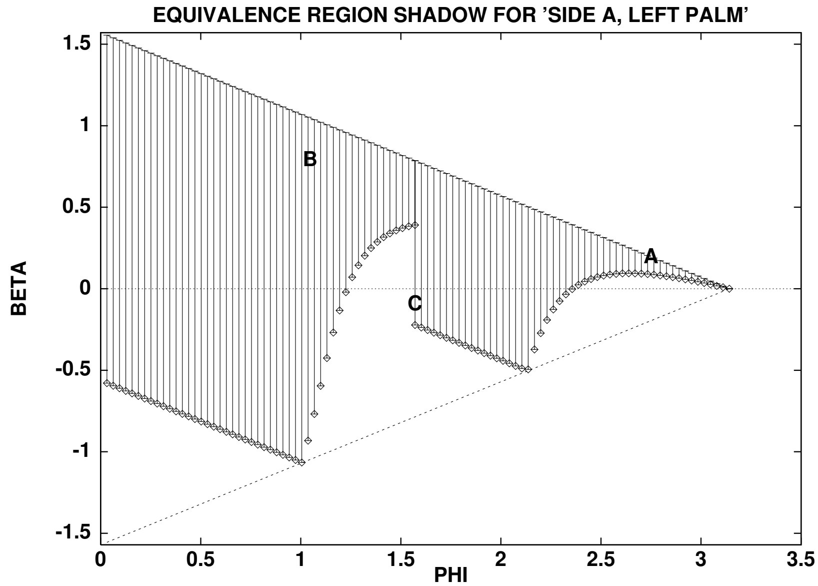

Figure 7.1: An example $\mathcal{S}$

$$\begin{aligned} \partial\theta &= 1 \cdot \partial\beta - \frac{\varepsilon}{2} \cdot \partial\phi \\ \varepsilon &= -1, \text{ resting on left palm} \\ & \quad\ 1, \text{ resting on right palm} \end{aligned} \tag{7.1}$$

over the range $\mathcal{P}_{\mathcal{S}}$ of (ϕ, β) for which the object remains stably supported.

The expression for θ describes a continuous, one to one and onto mapping from $\mathcal{P}_{\mathcal{S}}$ to $\mathcal{S}$. Its inverse is simply the projection operation $\pi(\phi,\, \beta,\, \theta) = (\phi,\, \beta)$, which is also continuous. Hence $\mathcal{S}$ and $\mathcal{P}_{\mathcal{S}}$ are homeomorphic to each other ([11]). Homeomorphisms preserve topological properties such as simple connectedness, so $\mathcal{S}$ and $\mathcal{P}_{\mathcal{S}}$ are topologically indistinguishable. In the following, we will use the notation $\mathcal{S}$ to refer either to an equivalence region corresponding to a stable edge, or its projection into the $(\phi,\, \beta)$ plane.

In order to prove Theorem 1, we will need the following lemmas. First, we need to show that if a planar set is made up of cross sections which are themselves connected sets (line intervals), and all of the cross sections are connected to each other, then that set is simply connected.

Lemma 2 *Let $\mathcal{S} \subset \Re^2$, and let $\mathcal{S}_\phi = \{(x, y) \in \mathcal{S} | x = \phi\} \subset \Re$ be the ϕ-cross sections of $\mathcal{S}$. Let $\mathcal{I}$ be an interval in $\Re$. If:*

- *There exists a function $f(\cdot) : \mathcal{I} \mapsto \Re^2$ such that $f(\phi) \in \mathcal{S}_\phi$ and $f(\cdot)$ is continuous.*
- *All $\mathcal{S}_\phi$ are themselves connected (line intervals);*
- $\mathcal{S} \cap \mathcal{I} \times \Re = \mathcal{S}$

then $\mathcal{S}$ is simply connected.

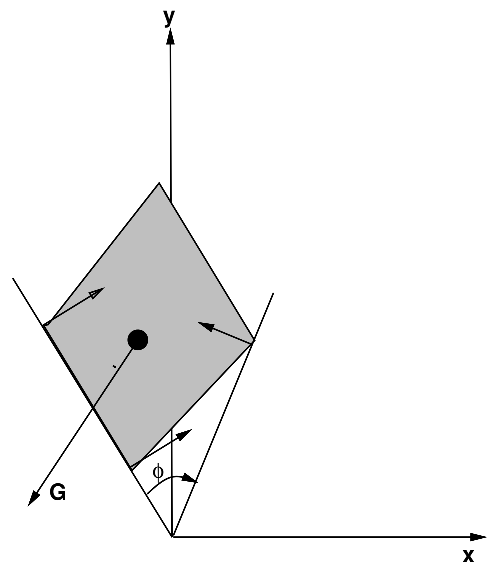

Figure 7.2: Contact normals and gravity vector in cone frame.

Given a feasible value for ϕ, the function f picks a point from the ϕ-cross section, $\mathcal{S}_\phi$, and does this in a continuous manner as ϕ is varied. The last condition of the theorem simply guarantees that we pick our interval $\mathcal{I}$ long enough that f exists in all the ϕ-cross sections of $\mathcal{S}$.

Proof: (Lemma 2) First, we prove connectedness. Let $p = (p_x, p_y), q = (q_x, q_y) \in \mathcal{S}$. We wish to show that there is a path between p and q, entirely contained in $\mathcal{S}$. If $p_x = q_x$, then by hypothesis, we are done. Otherwise, we look at the case $p_x \neq q_x$. We know that there exists a connected path from $f(p_x)$ to $f(q_x)$; namely, the f-image of the interval $[p_x,\, q_x]$. By composing this path with the path from p_x to $f(p_x)$, and the path from $f(q_x)$ to q_x, we have a complete path satisfying our requirements, and so $\mathcal{S}$ is connected.

Next, we will prove simple connectedness. The mapping

$$g(x, y, t) = (1 - t)(x, y) + t\, f(x)$$

continuously deforms the region $\mathcal{S}$ into the curve $f(x)$ as t varies from 0 to 1. By the continuity of f, this curve can be shrunk to a single point, and we are done. □

We now need to show that our set $\mathcal{S}$ is, in fact, made up of connected ϕ-cross sections, in this case cross sections formed by constant cone openings. Then we need to show the existence of an appropriate function f.

Lemma 3 *For a particular equivalence region $\mathcal{S}$, and the cone $Cone(\phi_0, \beta)$, with fixed ϕ_0, the range of β for which $(\phi_0,\, \beta) \in \mathcal{S}$ is an interval in $\Re$.*

Phrased more simply: If we hold the opening of the cone fixed, with the object in a stable position within the cone, we can tilt the cone in a certain range of β, and the object will stay in the same stable fixed position relative to the cone frame. The range of β for which the object is stable in that particular orientation is an interval. Hence, the ϕ-cross sections of $\mathcal{S}$ are also intervals.

Proof: (Lemma 3) With respect to the cone frame, the contact normals in configuration space are fixed, and the negative of the applied force due to gravity, $-G(\beta)$, varies as

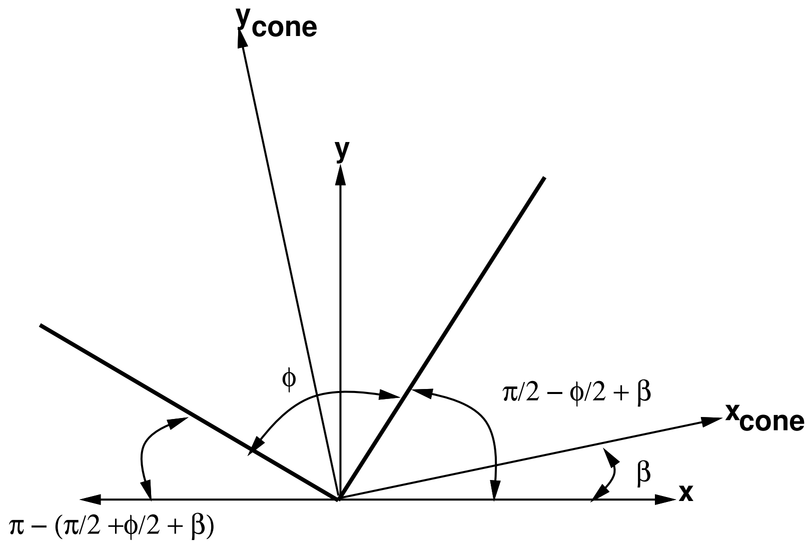

Figure 7.3: Orientation of left and right palms

$[\sin\beta \ \cos\beta \ 0]$ (Figure 7.2). Note that β is constrained to lie between $-\pi/2$ and $\pi/2$. Hence,G is a one-to-one function of β. The condition for stability is then that $-G(\beta)$ be expressible as a nonnegative combination of the contact normals, *i.e.* that $-G(\beta)$ be contained in the cone in force-torque space formed by the contact normals. The trajectory of $-G(\beta)$ is an arc of a circle, and so its intersection with the cone is closed and connected. Hence, the range of stable β is an interval, by the continuity of G. $\square$

Lemma 4 *The function* $f(\phi) = (\phi, -\varepsilon\frac{\pi-\phi}{2})$, *where* $\varepsilon = -1$ *if the object is resting on the left palm,* $\varepsilon = 1$ *if the object is resting on the right palm, is defined on the half-open interval* $(0, \pi]$, *and satisfies the conditions in Lemma 2.*

Proof: (Lemma 4) If a cone is in the configuration (ϕ, β), it is easy to show (see Figure 7.3) that the left palm makes an angle $\frac{\pi+\phi}{2} + \beta$ with the horizontal axis, and the right palm makes an angle $\frac{\pi-\phi}{2} + \beta$. If $\varepsilon = -1$, then at the point f(ϕ), the left palm is at an angle $\frac{\pi+\phi}{2} + \frac{\pi-\phi}{2} = \pi$. If $\varepsilon = 1$, then at the point f(ϕ), the right palm is at an angle $\frac{\pi-\phi}{2} - \frac{\pi-\phi}{2} = 0$. In either case, the palm on which the object is resting is horizontal, and since the resting edge is, by hypothesis, stable, $f(\phi) \in \mathcal{S}_\phi$, over the interval $0 < \phi \leq \pi$, and $f(\phi)$ is obviously continuous. $\square$

We note that the cone opening ϕ must lie in the half-open interval $(0, \pi]$; hence if we pick this as our interval $\mathcal{I}$, we have satisfied the third condition of Theorem 1. The proof that $\mathcal{S}$ is simply connected is then a straightforward application of Lemmas 2,3, and 4.

7.2 Sufficiency of pure tilts and fixed-θ squeezes for plan execution

We will assume that the palms have at least sufficient freedom to open flat (Cone($\pi, 0$)), and that they have symmetric rotational freedom. We will also assume that the palms have infinite extent.

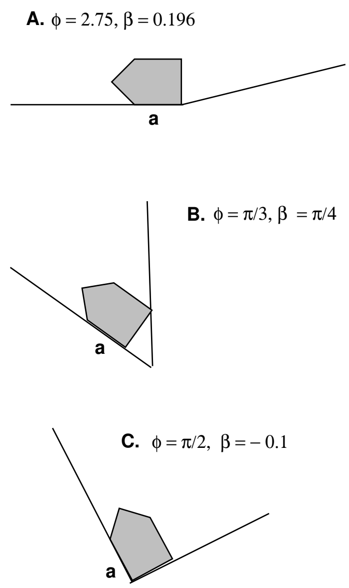

Figure 7.4: Different contact formations which are in the same equivalent region

In the following, let $\varepsilon = 1$ if $\mathcal{S}$ is an equivalence region corresponding to the object resting on the right palm Let $\varepsilon = -1$ if the object is resting on the left palm.

Theorem 5 (Pure Tilts) *If $p = (p_x, p_y), q = (q_x, q_y) \in \mathcal{S}$, and $p_x = q_x = \phi$, then the line segment $\overline{pq}$ is completely contained in $\mathcal{S}$.*

In other words, if p, q are both in $\mathcal{S}$, and both correspond to points with the same cone opening ϕ, then the pure tilt motion between them preserves the stable configuration of the part. This follows directly from Lemma 3.

Unfortunately, equivalence regions are not in general convex: if β is held constant and ϕ varied, the intersection of that line with the shadow of $\mathcal{S}$ will not in general be an interval (see Figure 7.1).

In addition to pure tilts (ϕ constant, β varying) and pure squeezes (β constant, ϕ varying), there is another type of simple movement of the palms related to pure squeezes: a *fixed-θ squeeze.* In a fixed-θ squeeze, the palm making edge contact with the object stays fixed, and the other palm opens and closes, resulting in the cone configuration changing, but the object orientation relative to the world frame remaining fixed as long as the state is stable. In the parameter space, this corresponds to $\frac{\partial \beta}{\partial \phi} = \varepsilon \frac{1}{2}$, where $\varepsilon = -1$ for the left palm held fixed, 1 for the right palm held fixed. Like a pure tilt, a fixed-θ squeeze between two (appropriately collinear) points in $\mathcal{S}$ will be completely contained in the equivalence region, subject to the following constraint.

Recall that an equivalence region is defined by the stable edge contact between a particular edge and a given palm, without considering the third contact with the other palm. Hence, an equivalence region can further be separated into smaller vertex-equivalent subregions, where in each subregion the contact formation is the same. For example, for an equivalence region corresponding to the right palm making contact with edge $\overline{ab}$, one subregion could correspond to the left palm making (one point) contact with vertex c, and another subregion could correspond to the left palm making contact with vertex d, etc. In Figure 7.4, the top and middle

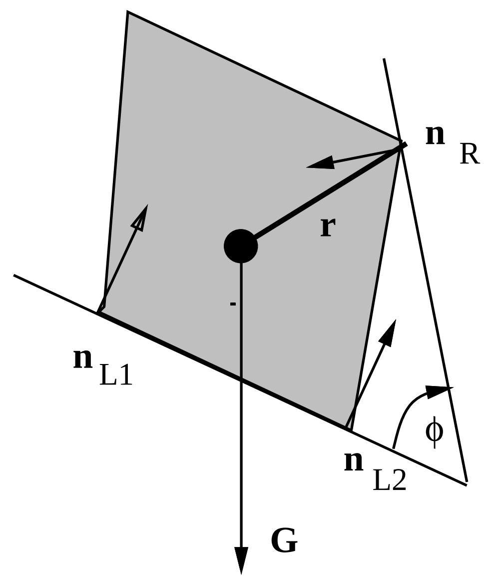

Figure 7.5: Contact normals and gravity vector in realspace.

configurations correspond to points in two different subregions of the region represented in Figure 7.1. (In the top pose, the object vertex which is in the vertex of the cone is making two point contact: contact simultaneously with both palms.) We will show that a fixed-θ squeeze is completely contained in $\mathcal{S}$ if its endpoints are in the same subregion of $\mathcal{S}$. Note that a point in $\mathcal{S}$ of four-point contact (as in the bottom configuration of figure 3.5) is in the closure of two different subregions (possibly even two different regions), but will not be considered to be a member of either region for the purposes of the next theorem.

Theorem 6 (Fixed-θ Squeezes) *If $p, q \in \mathcal{S}$, p and q are vertex-equivalent, and $\overline{p\,q}$ has slope $\frac{\partial\beta}{\partial\phi} = \varepsilon\frac{1}{2}$, then $\overline{p\,q}$ is completely contained in $\mathcal{S}$.*

Proof: Without loss of generality, assume that the left palm is making edge contact with a particular edge of the object, and the right palm is making one point contact with a vertex. Let $\hat{n}_{L1}, \hat{n}_{L2}$ be the (force-torque) contact normals of the left palm in configuration space, and $\hat{n}_R$ be the contact normal of the right palm. Let G be the force vector corresponding to gravity. See Figure (7.5). If $\hat{n}_{L1}, \hat{n}_{L2}, G$ are coplanar (which is true when the left palm is horizontal), the stability of the object is independent of the right palm, and the theorem is immediately true. Therefore, we will only consider the case where $\hat{n}_{L1}, \hat{n}_{L2}, G$ are in general position. By the definition of a fixed-θ squeeze, $\hat{n}_{L1}, \hat{n}_{L2}$, and G are fixed, and $\hat{n}_R$ varies with ϕ. The stability criterion is that $-G$ is in the positive span of $\hat{n}_{L1}, \hat{n}_{L2}$, and $\hat{n}_R$:

$$\begin{aligned} -G &= k_1\ \hat{n}_{L1} + k_2\hat{n}_{L2} + k_3\hat{n}_R, \\ k_1, k_2 &\geq 0 \\ k_3 &> 0 \end{aligned} \tag{7.2}$$

k_3 is strictly positive because $\hat{n}_{L1}, \hat{n}_{L2}$, and G are linearly independent. This equation can be rewritten:

$$
\begin{aligned}
(-\hat{n}_R) &= \frac{1}{k_3}G + \frac{k_1}{k_3}\hat{n}_{L1} + \frac{k_2}{k_3}\hat{n}_{L2}, \text{ or} \\
(-\hat{n}_R) &= k_1'\hat{n}_{L1} + k_2'\hat{n}_{L2} + k_3'G, \\
k_1', k_2' &\geq 0 \\
k_3' &> 0
\end{aligned} \tag{7.3}
$$

So, the fixed-θ squeeze will correspond to a line segment in (ϕ, β) space which is entirely contained in (the projection of) $\mathcal{S}$ if and only if the trajectory swept out by $-\hat{n}_R$ only enters and leaves the cone formed by $\hat{n}_{L1}, \hat{n}_{L2}$, and G at most once each.

In the frame with the left palm aligned with the x axis, and the object center of gravity as the origin, $\hat{n}_R = [-\sin\phi,\ -\cos\phi,\ -r_x\cos\phi + r_y\sin\phi]$, as ϕ varies from 0 to π. $[r_x, r_y]$ is the (fixed) vector from the CG to the vertex in contact with the right palm. The condition for stability is then

$$
\begin{aligned}
-\begin{bmatrix} -\sin\phi \\ -\cos\phi \\ -r_x\cos\phi + r_y\sin\phi \end{bmatrix} &= [\hat{n}_{L1}\ \hat{n}_{L2}\ G]\begin{bmatrix} k_1' \\ k_2' \\ k_3' \end{bmatrix} \\
k_1', k_2' &\geq 0 \\
k_3' &> 0
\end{aligned} \tag{7.4}
$$

Notice that as ϕ varies, the curve swept out by $-\hat{n}_R$ lies in a plane $\mathcal{P}$ in force-torque space which goes through the origin. The intersection of this plane with the cone described by $[\hat{n}_{L1}\,\hat{n}_{L2}\,G]$ is either only the origin, in which case we are done, or else the intersection is a two dimensional cone in $\mathcal{P}$. We wish to show that $-\hat{n}_R$ only enters and leaves this cone at most once each.

The plane $\mathcal{P}$ can be rotated into the (x, y) plane by the full rank transformation

$$
\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ r_y & -r_x & 1 \end{bmatrix} \tag{7.5}
$$

which functions as a projection into the (x, y) plane for all points in $\mathcal{P}$. Therefore the stability criterion is that the circular arc $[\sin\phi\ \cos\phi]$ only pass in and out of a cone at most once each. Since the trajectory of $-\hat{n}_R$ does not make a full revolution, this is clearly true, and we are done. □

Since subregions of a given equivalence region will have as the boundary between them a region of four point contact which is in the closure of both subregions, we can transit from one subregion to another by doing a fixed-θ squeeze to take us to the four point contact region, and then another fixed-θ squeeze into the new subregion.

Theorem 7 *For the equivalence regions of cone with palms of equal finite length, but a full 180 degrees of rotational freedom for each palm, Theorems 5 and 6 still hold.*

Proof: For a cone, finite palm length corresponds to a lower bound on possible ϕ, but not of possible β for a given possible ϕ. In other words, the new boundary of the equivalence region

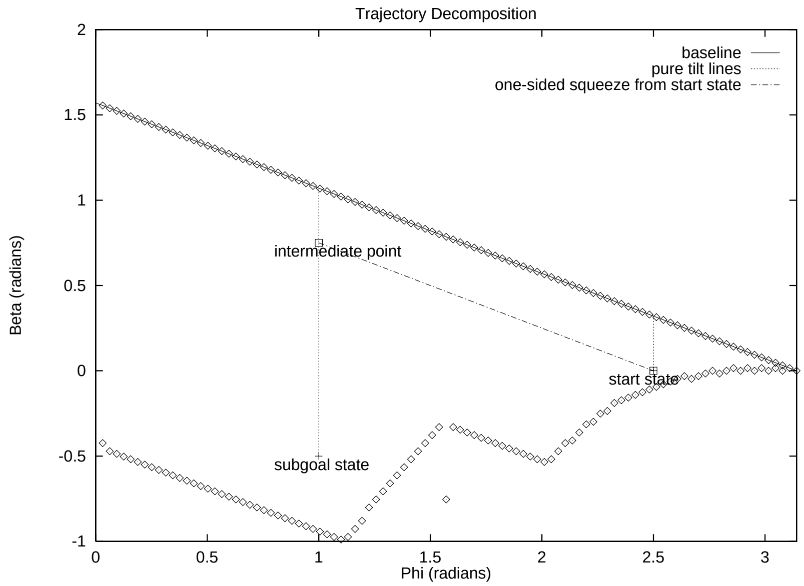

Figure 7.6: Trajectory decomposition

is a line of the form $\phi =$ constant. Therefore, the pure tilt and pure squeeze properties still hold. $\square$

Armed with Theorems 5, 6, and 7, we can now see that for our manipulator, we can plan motions from one boundary point to another by finding piecewise linear (in (ϕ, β) space) motions composed of pure tilts and fixed-θ squeezes, without having to keep track of anything except the valid range of β which is contained in $\mathcal{S}$ for each ϕ.

In fact, for an object which has all statically stable sides, we know that all the equivalence regions contain the point corresponding to the cone $\mathrm{Cone}(\pi, 0)$: in other words, the entire line (call it the "baseline") corresponding to the bottom palm being horizontal is contained in each equivalence region. The baseline is the edge of the equivalence region which corresponds to the line $\beta(\phi) = \varepsilon\frac{\pi-\phi}{2}$, where $\varepsilon = 1$ for the object resting on the left palm, -1 for the object resting on the right palm. So, for any start point and any subgoal point in an equivalence region corresponding to a stable edge, one of the following is true:

- There is a pure tilt connecting the two points, or there is a fixed-θ squeeze connecting the two points. Then there is a single operation that will bring you from the start state to the subgoal state.

- If there is not a pure tilt or fixed-θ squeeze connecting the two points, then draw the pure tilt line from the start state to the baseline, and the pure tilt line from the goal state to the baseline. These are both guaranteed to be contained in the shadow of the equivalence region by Theorem 5. There will be some line (a fixed-θ line) contained in $\mathcal{S}$ which is parallel to the baseline and which passes through these two pure tilt lines. (At the very least, the baseline will satisfy this requirement). See Figure 7.6. Therefore, the subgoal can be achieved from the start state in at most two tilts and one fixed-θ squeeze.

For an unstable edge, equivalence regions as we have defined them are no longer connected. However, in each component of the equivalence region, the above results about pure tilts and fixed-θ squeezes still hold.

The results are summarized below.

Theorem 8 (Path Planning) *Given any two points in an equivalence region $\mathcal{S}$ corresponding to a stable edge of a polygon P, there exists a path entirely contained in $\mathcal{S}$. which can be decomposed into at most two tilts and one fixed-θ squeeze.*

If the equivalence region does not correspond to a stable edge, but the two points are in the same component of the equivalence region, there still exists a path between the two points composed only of pure tilts and fixed-θ squeezes which is entirely contained in that component.

Therefore, if there exists a path from an arbitrary start state to an arbitrary end state, where the transitions between equivalence regions are given by pure tilt motions, the entire path between the start state and the goal state can be decomposed naturally into pure tilts and fixed-θ squeezes.

7.3 Existence of Plans

For planning paths from known start states, breadth-first search over the transition graph associated with a part will find a minimal length reorientation plan, assuming the plan exists. For a part with all stable edges, a reorientation plan between any two stable states exists if and only if the transition graph is *strongly connected.* A strongly connected graph is one where there is a directed path from any node to any other node. For a part with some unstable edges we would like there to be a directed path from any stable edge on either palm to any other stable edge on either palm. We can say in this case, that the subgraph of nodes of the transition graph corresponding to the stable edges of the part is a *strongly connected component* of the transition graph. In either case, if the stable states of a polygon P form a strongly connected component of the transition graph, we will call P *orientable.* An arbitrary polygon P is not always orientable. Consider the triangle in Figure 7.7. If the sliding transfer motion is possible (*i.e.* bidirectional arcs from the equivalence region ($edge$, left) to ($edge$, right)), then the graph is strongly connected. If these sliding motions are not possible, then the graph is disconnected into three clique-pairs. We present some conditions which are sufficient to determine whether a part P is orientable.

Let us orient the edges of the polygon so that edge e_i[1] is a ray whose origin is vertex v_i, and the vertices (and edges) are ordered counterclockwise with respect to the center of gravity. When the polygon is in a fixed position (say $\theta = 0$), each edge e_i then forms an angle α_i with the horizontal. Then the *relative angle* between edges e_i, and e_j can be defined to be $\alpha_{ij} = \alpha_j - \alpha_i$. See Figure 7.8.

Lemma 9 *For a cone with infinite length palms, suppose polygon P is resting stably on edge e_0 on the left palm. Let the right palm make contact with vertex v_k. Let the vector from the center of gravity of P to vertex v_k be $\mathbf{v}_k$, and ${\phi^\perp}_k$ be the cone opening such that the right palm is perpendicular to $\mathbf{v}_k$. Then:*

- *For all cone openings such that ${\phi^\perp}_{k+1} < \phi < {\phi^\perp}_k$, P will rotate to edge e_k on the right palm, upon executing a pure clockwise tilt. If e_k is stable, the part will stay in that configuration.*

[1] In the subscripts for polygon edges or vertices, or of the nodes of the transition graphs G or G' (which will be presented later), $i \pm k$ is shorthand for $(i \pm k)$ modulo N.

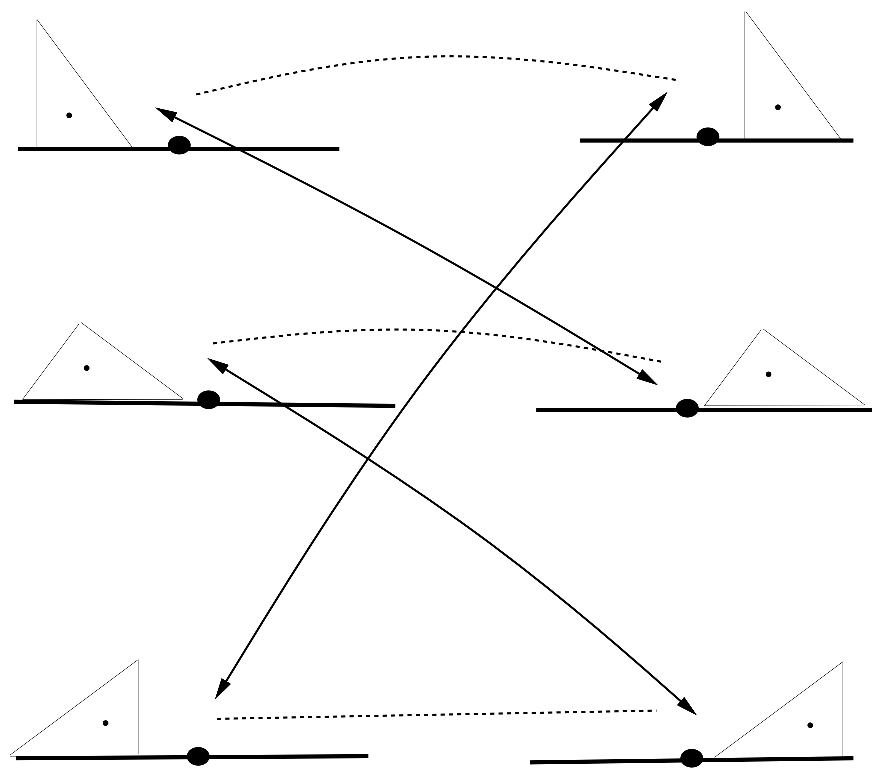

Figure 7.7: If sliding transfer motions (shown as dashed lines) are physically feasible, then the transition graph for this object is strongly connected. Otherwise, it is disconnected into three cliques

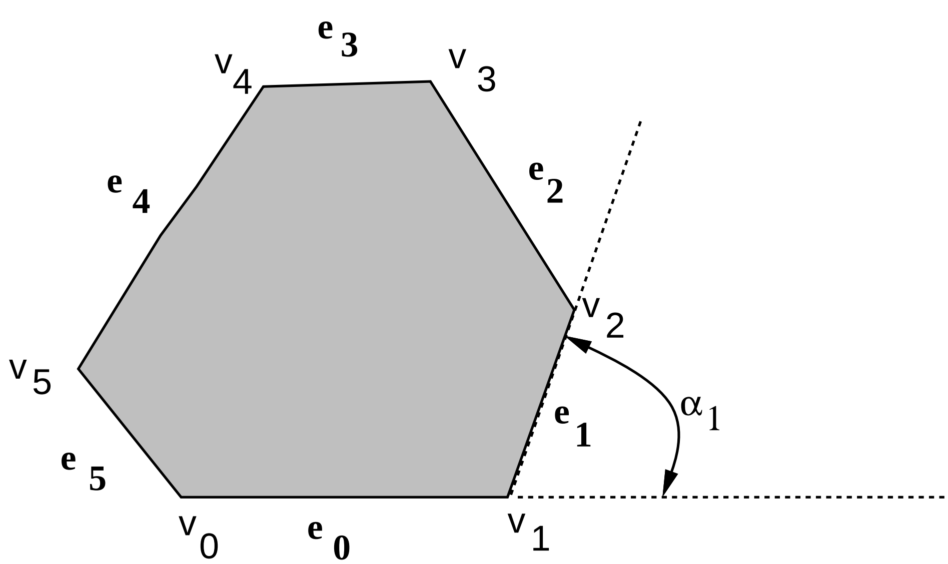

Figure 7.8: Edge and vertex labeling conventions.

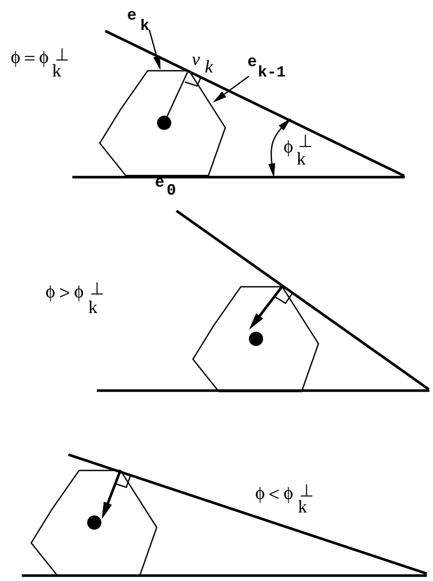

Figure 7.9: For $\phi > \phi^{\perp}{}_k$, object will transit to edge e_{k-1} on the right palm. For $\phi < \phi^{\perp}{}_k$, object will transit to edge e_k on the right palm.

- *For all cone openings such that $\phi^{\perp}{}_k < \phi < \phi^{\perp}{}_{k-1}$, P will rotate to edge e_{k-1} on the right palm, upon executing a pure clockwise tilt. If the e_{k-1} is stable, the part will stay in that configuration.*

In general, the object may also slide as well as rotate as the palms rotate, but since equivalence regions are determined solely by edge-palm contact — in other words, by orientation in the palm frame — only the rotational motion of the object is significant for our purposes. First we will show that the contact with the right palm determines the feasible rotation sense of the part. We will use the following terminology. Suppose we have a vector $\mathbf{n}$ at a point $\mathbf{p}$. Let $\mathbf{t}$ be another vector at $\mathbf{p}$, perpendicular to $\mathbf{n}$, such that $(\mathbf{t}, \mathbf{n})$ form a right handed frame at $\mathbf{p}$. We will say that a point $\mathbf{x}$ is *to the right of* $\mathbf{n}$ if $\mathbf{t}^T(\mathbf{x} - \mathbf{p}) > 0$. Similarly, $\mathbf{x}$ is *to the left of* $\mathbf{n}$ if $\mathbf{t}^T(\mathbf{x} - \mathbf{p}) < 0$. Note that any point on the line which is defined by the vector $\mathbf{n}$ and the point $\mathbf{p}$ may be used as the origin; the left and right regions will remain the same.

Lemma 10 *A part P is resting in equilibrium on edge e_0 on the (not necessarily horizontal) left palm. The right palm is also making contact with the part, and its contact normal is $\mathbf{n}_3$, emanating from the contact point between P and the right palm. If the center of gravity of P is to the left of $\mathbf{n}_3$, the part can only remain stable or make a positive rotation as the cone rotates clockwise. If the center of gravity is to the right of $\mathbf{n}_3$, the part can only remain stable or make a negative rotation as the cone rotates clockwise.*

We are ignoring the case where $\mathbf{n}_3$ goes through the center of gravity, as it is not needed for the proof of Lemma 9.

Proof: Referring to Figure 7.10, let n_1, n_2 be the lines through the contact normals on the left palm (with n_1 to the left of n_2, as in the figure), and let n_3 be the line through $\mathbf{n}_3$, the contact normal on the right palm. In the cone frame, these lines are held fixed, and as the cone tilts, the gravity vector $\mathbf{G}$ sweeps counterclockwise in a region contained in the region which starts

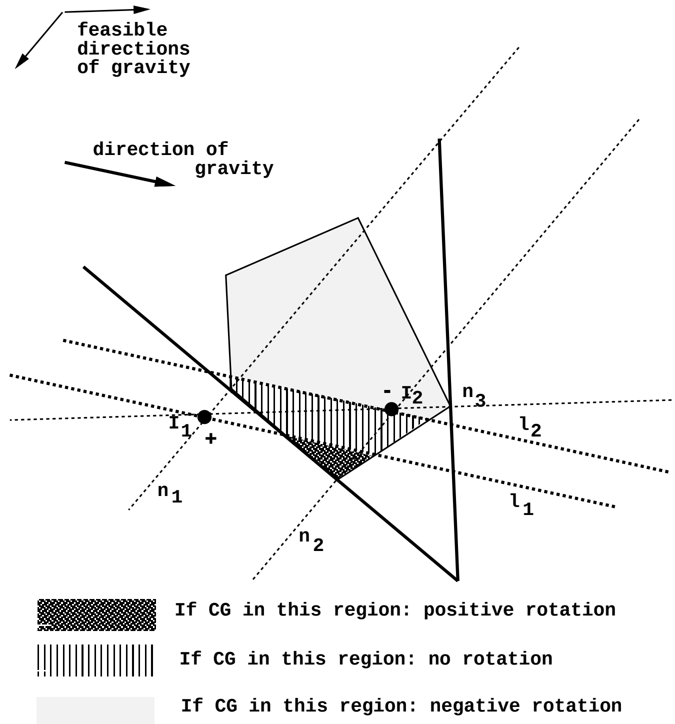

Figure 7.10: Determining the rotation sense of a part

with the left palm horizontal and ends with the right palm horizontal. $\mathbf{G}$ is anti-parallel to $\mathbf{n}_{1,2}$ when the left palm is horizontal, and anti-parallel to $\mathbf{n}_3$ when the right palm is horizontal. As shown in Lemma 3, some arc (possibly empty) of gravity's trajectory corresponds to the object being in equilibrium; by hypothesis, this arc, β_I, is non-empty and contains our start state. If edge e_0 is stable, then the cone orientation where $\mathbf{G}$ is anti-parallel to $\mathbf{n}_{1,2}$ is contained in β_I, and without loss of generality, we may consider the left palm horizontal to be our start state. Otherwise, all of β_I must be counterclockwise of this orientation, and since we are sweeping counterclockwise, we need not consider it.

As we rotate $\mathbf{G}$, at some point force balance will no longer be feasible, and the object must rotate to another stable state (on the right palm). When equilibrium is no longer feasible, the part will move in a direction which minimizes potential energy: the instantaneous velocity of the center of gravity must make a positive dot product with $\mathbf{G}$. If $\mathbf{G}$ is anti-parallel to $\mathbf{n}_3$, then it is obvious that if the center of gravity is to the left of n_3, the part can only minimize its potential energy by falling to the left: a positive rotation. If the center of gravity is to the right, the object can only make a negative rotation. In either case, the terms of the lemma are satisfied, and we are done. We now consider the case when $\mathbf{G}$ is in the interior of the feasible range of gravity directions.

By doing a kinematic analysis, we find that there are only two feasible centers of rotation at the moment force balance is lost. One is the point of intersection of n_1 and n_3 (labeled I_1 in Figure 7.10), which corresponds to contact 2 breaking. I_1 must have a positive rotation sense to prevent contact 2 from penetrating the palm. The other feasible center of rotation is the intersection of n_2 and n_3 (labeled I_2). I_2 must have a negative rotation sense.

Drawing lines l_1 and l_2 parallel to $\mathbf{G}$ and passing through I_1 and I_2, respectively, we make the following observations. Taking $-\mathbf{G}$ as a directional reference (and any point on l_1 as the origin), the instantaneous motion vector of every point to the left[2] of l_1 will make a positive dot product with $\mathbf{G}$ only if I_1 is the center of rotation. Hence, if the center of gravity is to the left of l_1, the part will make a positive rotation. The instantaneous motion vector of every point to the right of l_2 (now using a point on l_2 as the origin) will only make a positive dot product with $\mathbf{G}$ if I_2 is the center of rotation. Therefore if the center of gravity is to the right of l_2, the part will make a negative rotation. No point in between l_1 and l_2 can ever make a positive dot product with $\mathbf{G}$. Hence, if the center of gravity is between l_1 and l_2, the part will remain in equilibrium.

Making this construction for the (equilibrium) starting configuration of the object and the cone, we observe that the center of gravity must lie in the region between this original l_1 and l_2. Some subset of this original region is to the left of n_3, and some subset to the right. We will call the interior of the left region, plus the segment of l_1 which intersects P, L_0. We will call the interior of the right region, plus the segment of l_2 which intersects P, R_0. Recall that we are explicitly ignoring the case where n_3 passes through the center of gravity. Also, note that L_0 and R_0 are fixed regions defined with respect to the start state; they do not change as the gravity direction rotates.

As the cone tilts, l_1 and l_2 will pivot about I_1 and I_2, remaining parallel to $\mathbf{G}$. It is easy to see that if the center of gravity is in L_0, it must always be either between l_1 and l_2 or to the left of l_1 (using $-\mathbf{G}$ as the directional reference). If the center of gravity is in R_0, it must always be between l_1 and l_2 or to the right of l_2 (using $-\mathbf{G}$ as the directional reference). The lemma is therefore proved. □

We can now prove Lemma 9.

Proof: (Lemma 9) We use the same reference frame as in Lemma 10. Using an argument similar to the one in the first paragraph of the proof of Lemma 10, we first argue that at some point during the cone tilt, force balance will in fact no longer be feasible. At the very least, force balance (for the initial pose of the object in the palms) will not be feasible when the right palm is horizontal; this is when $\mathbf{G}$ is anti-parallel to $\mathbf{n}_3$. Notice from Figure 7.9 that if $\phi > \phi^{\perp}{}_k$, then the center of gravity is to the left of $\mathbf{n_3}$. By Lemma 10, when force balance is no longer feasible, the part will rotate counterclockwise, to edge e_{k-1} on the right palm.

It is possible that as the part rotates counterclockwise, the part will make edge palm contact between the left palm and the edge of the part clockwise of e_0 (edge e_{N-1}, by our labeling convention) before the right palm makes contact with edge e_{k-1}. If this happens, the left palm loses contact with the vertex at n_2, but maintains contact with the vertex at n_1, which is common to both edges e_0 and e_{N-1}. Let us assume that when the part loses stability, it rotates instantaneously from edge e_0 on the left palm to edge e_{N-1} on the left palm, so that gravity maintains the same orientation during the transition. The part will rotate counterclockwise about I_1 in such a way that the motion of the center of gravity makes a positive dot product with gravity — the center of gravity is on the left of the line l_1 (with $-G$ as the directional reference). As the part rotates, the line l_1 (and the rotation point I_1) will move, but maintain the same orientation. Recall from Chapter 3 that if a part is in two point contact with the palms (one contact per palm), with gravity in a fixed direction with respect to the palms, the potential energy curve of the part is either concave down, with a single maximum in the interior of the range of feasible orientations for those two contacts, or it has no maximum in the interior of this

[2] Note that l_1 is always to the left of l_2 with respect to $-\mathbf{G}$ for all valid directions of $-\mathbf{G}$.

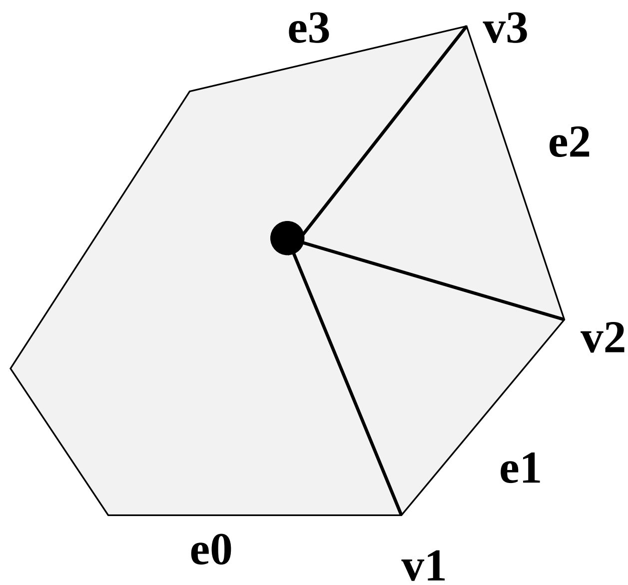

Figure 7.11: If edge $e0$ is resting on the left palm, then the vectors $\mathbf{v_i}$ formed at vertices $v1$, $v2$, and $v3$ form strictly acute angles with $e0$. Hence, there are pure tilt motions to transfer the part to one of edges $e1$, $e2$, or $e3$ on the right palm.

range. The unique local maximum (unstable equilibrium) will correspond to a configuration where the center of gravity lies on the line l_1. Since the potential energy cannot increase, the center of gravity must remain strictly to the left of the moving l_1. When edge e_{N-1} makes contact with the left palm, three point contact will be reestablished, with the former n_1 now playing the role of n_2. Hence, the center of gravity must be strictly to the left of the new l_2. It must either be between the new l_1 and l_2, or to the left of both of them.

If the center of gravity is in the closed region between the new l_1 and l_2 we are again in the situation described by Lemma 10, which we can apply again to get the same answer (counterclockwise rotation). If the center of gravity is still to the left of the new l_1, the part will continue to rotate counterclockwise, until in fact it reaches edge e_{k-1} on the right palm.

If $\phi < {\phi^{\perp}}_k$, then the center of gravity is to the right of $\mathbf{n_3}$. By Lemma 10, when force balance is no longer feasible, the part will rotate clockwise; by a similar argument as above, it will rotate to edge e_k on the right palm. □

If edge e_k is not stable, the part may not remain on edge e_k, but may transfer to the nearest (in the sense of having the nearest local minima of potential energy) stable edge. In fact, the part can be made to transfer to that nearest edge by tilting until the right palm is horizontal, and opening both palms flat.

In the case of only stable edges, however, it is easy to show using Lemma 9 that one can transfer from, *e.g.* edge e_0 to any edge e_i such that the vector $\mathbf{v_i}$ forms a strictly acute angle with edge e_0 (see Figure 7.11). From this follows the (weaker) conclusion that one can hence transfer from edge e_0 to any stable edge e_i which has a positive relative angle less than or equal to π. If e_0 is resting initially on the right palm, a similar result holds for all edges e_i with relative angle $\alpha_i - \alpha_0$ negative and greater than $-\pi$. Hence, if there is an arc of the transition graph connecting two stable edges (e_i, left) with (e_j, right), and $\alpha_{ij} \leq \pi$, this arc is bidirectional: that is, it corresponds to two directed arcs in opposite directions[3].

Lemma 11 *For a cone with infinite length palms, if a polygon P is resting stably on edge e_0 on the left palm, there is a cone trajectory that will transfer P to rest on any stable edge e_i on the right palm if the relative (CCW) angle between e_i and e_0 is nonzero and less than or equal to π. The corresponding arc of the transition graph is bidirectional.*

Using this fact, we can find our first set of sufficient conditions for P to be orientable.

Theorem 12 *Let P be a convex polygon with N all stable edges, G the corresponding transition graph. In order for P to be orientable, it is sufficient that either:*

1. *Sliding transfers be possible, or*
2. *For every edge e_i of P, there be at least two edges e_j, e_k such that $\alpha_{ij}, \alpha_{ik} \leq \pi$.*

If either of the above two cases is true, then the length of any reorientation plan is bounded by N. This plan can be found in $\mathcal{O}(N^2)$ operations.

Proof: By Lemma 11, every node (e_0, left) has at least one (bidirectional) arc, to (e_1, right), since two adjacent edges must have relative angle less than π. We can then collapse the graph G into a smaller graph G', with N nodes. Each node p_i of G' corresponds to the node pair $((e_i, \text{left}), (e_{i+1}, \text{right}))$. Suppose that case 1 of the theorem is true. Then (e_i, left) is connected to (e_i, right). With respect to G', this means that every node p_i is connected bidirectionally to node p_{i-1}. Suppose case 2 of the theorem is true. Then edges e_{i+1} and e_{i+2} both have relative angles with edge e_i less than or equal to π. Hence, by Lemma 11, (e_i, left) is connected bidirectionally to (e_{i+2}, right). This means that in G', node p_i is connected to node p_{i+1}. In either case, G', and hence G, contains a cycle, and is therefore strongly connected. The length of the cycle in G is $2N$, and the length of the shortest path between any two nodes of G is no more than N. If we search G using breadth-first search, this minimum length path will be found in $k * 2N$ operations, where k is the maximum number of children of any given node in G. Since k can be no more than N, a path can be found in $\mathcal{O}(N^2)$ operations. □

For a polygon P with unstable edges, we can construct a sort of subpolygon, P', which has all stable edges. For each edge e_i consider the half plane formed by all points in the plane which make a nonnegative dot product with the outward facing edge normal. Call that half plane H_i.

Every vertex of P which is connected to two stable edges is also in P'. For each unstable edge e_u take the edges e_{CW} and e_{CCW} which are the closest stable edges to e_u going clockwise and counterclockwise, respectively. Let the intersection of the extensions of e_{CW} and e_{CCW} be x. If x is in H_u, then x is a new vertex of P'. Otherwise, P' is unbounded between e_{CW} and e_{CCW}. If none of the edges of P' are unbounded, then P' is *bounded*. Figure 7.12 shows an example of a bounded and an unbounded subpolygon. Make the one to one correspondence between the edges of P' and the stable edges of P. Any bidirectional arcs in the transition graph of P' which are gotten by the application of Lemma 11 will also be in the transition graph of P, since the operation of forming P' preserves the relative angles between the stable edges. However, an unbounded P' no longer has the property that adjacent edges must have relative angle less than π. In fact, if an edge e_0 of P' has as its counterclockwise adjacent edge

[3] It is more succinctly represented as an undirected arc, of course, but we will choose to represent it as two directed arcs to distinguish it from arcs which are truly directed, that is, one-way only.

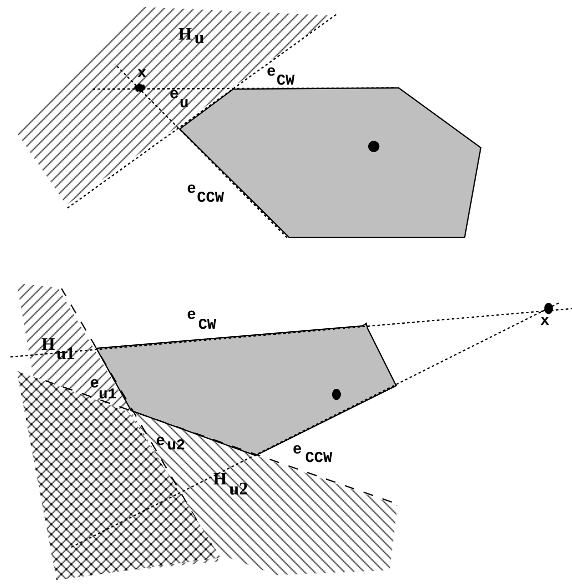

Figure 7.12: The top polygon has a bounded subpolygon. The bottom polygon has an unbounded subpolygon

an unbounded edge, then it is true that the next counterclockwise edge after that will make a relative angle with e_0 of greater than or equal to π. Nonetheless, we can still extend Theorem 12 to any polygon whose subpolygon P' is bounded and satisfies the the conditions of the theorem.

Corollary 13 *Any polygon P whose subpolygon P' is bounded and satisfies the conditions of Theorem 12 is orientable. The length of any path is bounded by M, the number of stable states, and can be found in $\mathcal{O}(M^2)$ operations, after P' has been constructed.*

If P has N sides, M of which are stable, constructing P' takes time $\mathcal{O}(N)$.

Theorem 12 and its Corollary are useful for "eyeballing" polygons and determining quickly whether their transition graphs might be strongly connected. In fact, if condition 1 of the theorem (namely, that sliding transfer are possible) is satisfied, we will show later that any polygon is orientable. Sliding transfers may not always be practical, however, since the palms may not actually make contact at the cone vertex, or there may be an obstacle such as a motor or hinge near the vertex. In the case where sliding transfers are disallowed, Theorem 12 is a fairly weak theorem, and there are many orientable polygons which do not satisfy its preconditions.

In order to tell which edges can be reached from a given edge e_i on the left palm, one can make use of the following construction (see Figure 7.13). Color each stable edge of the polygon P a different color, and color any unstable edges with the color of its nearest stable edge (in the potential energy sense, as before). Then draw a circle around P, with its center of gravity as the center of the circle. Each colored set of edges of P can then be projected up into an arc of this circle, which will be considered to be also colored appropriately. Label each arc α_i; each arc α_i contains within itself the projection of exactly one stable edge, e_{α_i}. We will call this circle $\mathcal{C}_P$. Now draw a line l_i through the center of the circle parallel to the inward pointing normal of edge e_{α_i}, $\mathbf{n}_i$. Let $\mathbf{t}_i$ be a vector perpendicular to $\mathbf{n}_i$ such that the frame $(\mathbf{t}_i, \mathbf{n}_i)$ is a right-handed frame. Set this frame so that its origin is at the intersection of l_i and e_{α_i}. Then when $\mathcal{C}_P$ is rotated so that this frame is aligned with the world frame (that frame where gravity

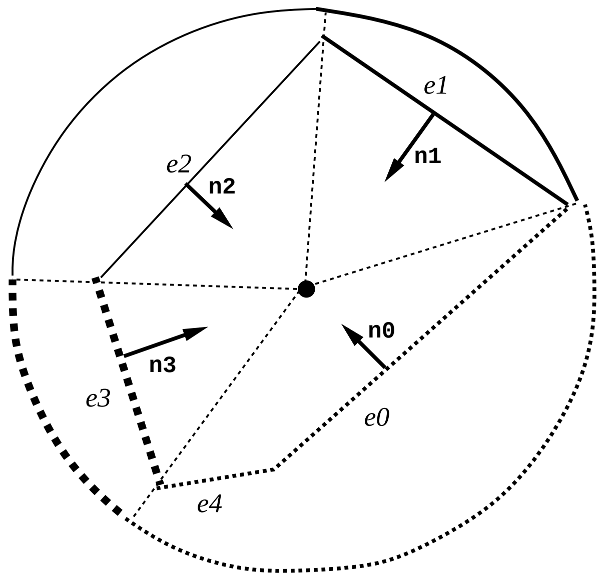

Figure 7.13: Construction for determining reachable edges. Note that edges $e0$ and $e4$ correspond to the same arc on the circle. This is because edge $e0$ is the stable edge to which the unstable edge $e4$ will fall.

is anti-parallel to the vertical axis), we have a description of the orientation of P when it is resting on the stable edge e_{α_i}. We can use the frame $(\mathbf{t}_i, \mathbf{n}_i)$ to orient the plane with respect to n_i.

Definition 7.1 (Open right half plane) *The line l_i divides the plane into two half planes. The open half plane of directions which make a positive dot product with $\mathbf{t}_i$ is the* open right half plane with respect to n_i*)*

Any arc of the circle (other than α_i) which reaches into open right half plane with respect to n_i corresponds to a stable edge of P which is reachable using a clockwise tilt starting from the state (e_{α_i}, left palm). This follows from Lemma 9; it is a generalization of the observation that from an edge e_0 one can reach any edge e_i such that the vector from the center of gravity to the vertex v_i makes a strictly acute angle with e_0. See Figure 7.14. Similarly, any arc of the circle (other than α_i) in the open left half of the plane corresponds to a stable edge of P which is reachable from edge e_{α_i} using a counterclockwise tilt starting state (e_{α_i}, right palm). We can then build a transition graph G for the arcs α_i of $\mathcal{C}_P$. Suppose that the state (e_{α_0}, left palm) can reach the state (e_{α_k}, right palm). The corresponding transition is bidirectional if the right endpoint of α_0 is in the open left half plane with respect to $\mathbf{n}_k$.

Theorem 14 *If sliding transfers are possible, any polygon P is orientable. The path length from any initial state to any goal state is bounded by $2M$, where M is the number of stable edges. It can be found in $\mathcal{O}(M^2)$ operations.*

Proof: Recall that the time to find a path through a graph by breadth-first search is bounded by Nk, where N is the number of nodes and k is the maximum number of children of any node. Every node of G ($2M$ of them) has at most M children, so we only need to prove

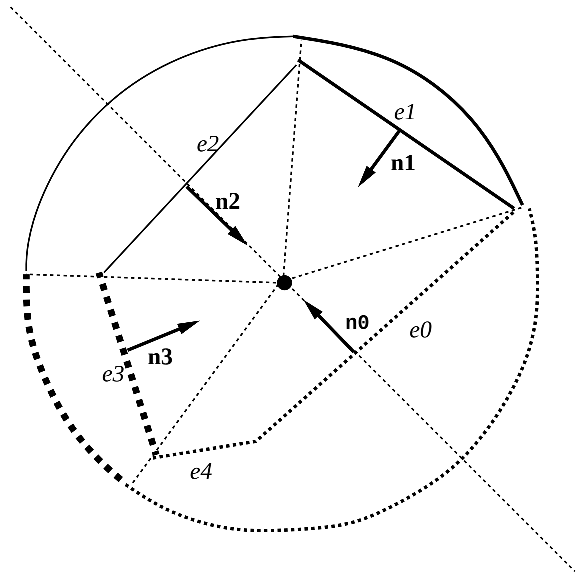

Figure 7.14: If P is resting on edge $e0$ on the left palm, it can reach edges $e1$ or $e2$ on the right palm by a clockwise tilt. If P is resting on edge $e1$ on the right palm, it can reach edges $e3$ or $e2$ on the left palm by a counterclockwise tilt.

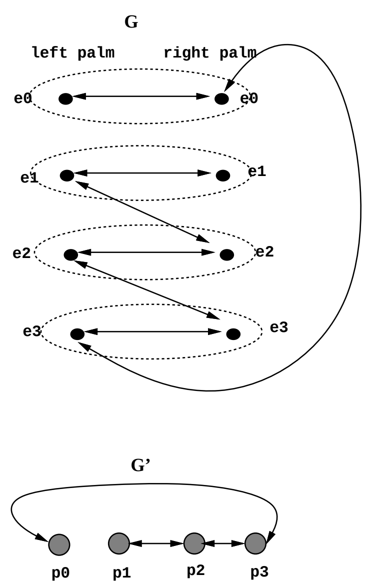

Figure 7.15: Worst case scenario for Theorem 14

that G is strongly connected Construct the circle diagram $\mathcal{C}_P$. There are M colored arcs α_i in $\mathcal{C}_P$, each one containing the projection of one stable edge, e_{α_i}. Construct the lines l_i and the frames $(\mathbf{t}_i, \mathbf{n_i})$ as described previously. Since we know that the nodes $(e_{\alpha_i}, \text{left})$ and $(e_{\alpha_i}, \text{right})$ have a bidirectional arc between them (sliding transfer), we an collapse G into a smaller graph G'. A node p_i of G' corresponds to the node pair $((e_{\alpha_i}, \text{left}), (e_{\alpha_i}, \text{right}))$, G will be strongly connected if and only if G' is strongly connected. If every arc of $\mathcal{C}_P$ can reach its neighboring counterclockwise arc bidirectionally, then we are in a situation analogous to Case 1 of Theorem 12, and we are done.

Otherwise, suppose that arc α_0 cannot reach its neighboring counterclockwise arc, bidirectionally. This case also breaks down into two cases. It may be the case that P has only one stable edge; in that case, G' is a single node, and hence trivially strongly connected, and we are done.

Suppose that P has more than one stable edge. If α_0 cannot reach any other counterclockwise arc bidirectionally, then there are no normals $\mathbf{n}_i$ in the open right half plane of $\mathbf{n}_0$. Hence, all the normals $\mathbf{n}_i$ are in the same closed half plane: the closed left half plane of $\mathbf{n}_0$. $\mathbf{n}_0$ will be the normal the furthest in the counterclockwise direction. Hence, it is the case that $(e_{\alpha_i}$, left) can be brought to $(e_{\alpha_{i+1}}$, right) bidirectionally for all i greater than 0. Therefore (see Figure 7.15), for the graph G', every node p_i is connected to p_{i+1} for $i > 0$. The graph G' forms a path of length M, and G forms a path of length $2M$, and we are done. □

Theorem 15 *If sliding transfers are not possible, in order for P to be orientable it is sufficient that every arc α_i of $\mathcal{C}_P$ can reach two other arcs α_{i+1} and α_{i+2} going counterclockwise. That is, α_{i+1} and α_{i+2} both reach into the open right half plane of $\mathbf{n}_i$. The path from any initial state to any goal is bounded by $2M$ and can be found in at most $\mathcal{O}(M^2)$ operations.*

This theorem can be restated as in Chapter 3:

> *If sliding transfers are not possible, in order for P to be orientable it is sufficient that every stable edge e_{α_i}, when resting on the left palm, can reach two other stable edges $e_{\alpha_{i+1}}$ and $e_{\alpha_{i+2}}$ on the right palm by a clockwise tilt. If P has M stable edges, the path from any initial state to any goal is bounded by $2M$ and can be found in at most $\mathcal{O}(M^2)$ operations.*

Proof: If α_{i+1} and α_{i+2} both reach into the open right half plane of $\mathbf{n}_i$, then α_{i+1} must be completely contained in that open half right plane. Hence $\mathbf{n}_{i+1}$ is in the open right half plane, and $(e_{\alpha_i}$, left) can be brought to $(e_{\alpha_{i+1}}$, right) bidirectionally. So we can again collapse G into a smaller graph G', where the nodes p_i correspond to the node pairs $((e_{\alpha_i}$, left), $(e_{\alpha_{i+1}}$, right)). We also know that $(e_{\alpha_i}$, left) can reach $(e_{\alpha_{i+2}}$, right) at least unidirectionally. Hence, node p_i of G' can reach node p_{i+1} at least unidirectionally, for all i. G' therefore forms a directed cycle of length M, as seen in Figure 7.16, and G forms a directed cycle of length $2M$. □

An analogous result naturally holds in the clockwise case, as well.

7.4 Homing Sequences

As described in Chapter 5, if we wish to find a plan to orient a part to a known final state from an arbitrary unknown initial state, we can construct a larger transition graph whose nodes are the elements of the power set of the set of equivalence regions, and whose arcs are cone motions.

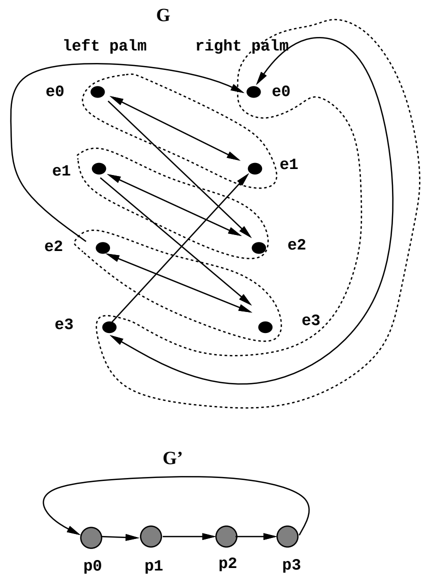

Figure 7.16: Worst case scenario for Theorem 15

One of the questions which we wish to address is the discretization of cone motion space. How many different cone motions, or how finely do we have to sample the cone motion space, if we want to guarantee that we will find an orientation plan, if one exists? We would also like to find conditions for which a polygon P can be oriented to a particular goal from any arbitrary initial state.

We recall Lemma 9:

> For a cone with infinite length palms, suppose polygon P is resting stably on edge e_0 on the left palm. Let the right palm make contact with vertex v_k. Let the vector from the center of gravity of P to vertex v_k be $\mathbf{v}_k$, and ${\phi^\perp}_k$ be the cone opening such that the right palm is perpendicular to $\mathbf{v}_k$. Then:
>
> - For all cone openings such that ${\phi^\perp}_{k+1} < \phi < {\phi^\perp}_k$, P will rotate to edge e_k on the right palm, upon executing a pure clockwise tilt. If e_k is stable, the part will stay in that configuration.
> - For all cone openings such that ${\phi^\perp}_k < \phi < {\phi^\perp}_{k-1}$, P will rotate to edge e_{k-1} on the right palm, upon executing a pure clockwise tilt. If the e_{k-1} is stable, the part will stay in that configuration.

We can construct such intervals for a polygon with unstable edges by performing a similar procedure on $\mathcal{C}_P$. Orient $\mathcal{C}_P$ so that $\mathbf{n}_0$ is parallel to the y axis. Label the endpoints of each arc in a manner analogous to the labelling of polygon vertices: every arc α_i has two "vertices" (endpoints) v_{α_i} and $v_{\alpha_{i+1}}$, going counterclockwise. $\mathbf{v}_{\alpha_i}$ is the vector from the center of $\mathcal{C}_P$ to v_{α_i}. If edge e_{α_0} is resting on the left palm, then let ${\phi^\perp}_i$ be the cone opening such that the right palm is perpendicular to $\mathbf{v}_{\alpha_i}$, for $i > 0$. Then again,

1. For all cone openings such that ${\phi^\perp}_{i+1} < \phi < {\phi^\perp}_i$, upon executing a clockwise tilt, P will rotate to arc α_i, and hence to edge e_{α_i} as the unique stable edge within that arc.

2. For all cone openings such that $\phi^{\perp}{}_i < \phi < \phi^{\perp}{}_{i-1}$, upon executing a pure clockwise tilt, P will rotate to edge $e_{\alpha_{i-1}}$ on the right palm.

3. If a sliding transfer of edge e_{α_0} to the right palm is feasible, it will occur upon executing a clockwise tilt, for $\phi > \phi_1^{\perp}$. This follows from item 2.

This follows from Lemma 9. We can now determine how finely we must sample cone motion space in order to guarantee finding a plan, if one exists.

Theorem 16 *Suppose we have a polygon P, and a cone with infinite length palms. If transitions between equivalence regions are executed using pure tilts, then there is a finite set of cone motions such that a (frictionless) homing sequence exists if and only if it can be constructed from this finite set of cone motions.*

Proof: Assume we determined all the intervals of cone openings for all the stable edges of the object on the left palm, as described above. Sort the endpoints (both upper and lower) of all the intervals into a sequence $\{\zeta_i\}$. Let the interval I_i be the interval from ζ_i to ζ_{i+1}. Do the same for the right palm, and call those intervals J_i. Pick a sequence of cone openings, one out of each interval, $\gamma_i \in I_i$ and $\rho_i \in J_i$.

Let $\beta_{max}(\phi) = (\pi - \phi)/2.0$. $\beta_{max}(\phi)$ corresponds to a cone with cone opening ϕ tilted until the left palm is horizontal; $-\beta_{max}(\phi)$ corresponds to the right palm horizontal. By using $\pm\beta_{max}(\phi)$, we guarantee that when we tilt, for example, right, that at the end of the tilt the object must be resting in edge-palm contact with the right palm, and vice-versa if we tilt left. If a homing sequence exists for P, it can be found by searching only over the set of cone motions

$$\begin{array}{rcl}
(\gamma_i, \beta_{max}(\gamma_i)) & \longrightarrow & (\gamma_i, -\beta_{max}(\gamma_i)) \\
 & & \text{(tilts from left palm to right palm)} \\
(\rho_i, -\beta_{max}(\rho_i)) & \longrightarrow & (\rho_i, \beta_{max}(\rho_i)) \\
 & & \text{(tilts from right palm to left palm)}
\end{array}$$

plus the appropriate fixed-θ squeezes to move from any $(\gamma_i, \beta_{max}(\gamma_i))$ to any $(\gamma_j, \beta_{max}(\gamma_j))$, and from any $(\rho_i, -\beta_{max}(\rho_i))$ to any $(\rho_j, -\beta_{max}(\rho_j))$.

Each value in the sequence $\{\zeta_i\}$ represents a cone opening where some initial equivalence region $\mathcal{E}_i$ stops transiting to equivalence region $\mathcal{E}_j$ upon executing a pure tilt (for cone openings less than ζ_i), and begins transiting to equivalence region $\mathcal{E}_k$ (for cone opening greater than ζ_i). Hence, the endpoints of each interval I_i (or J_i) represent critical cone openings where the behavior of the system will change for some sets of equivalence regions. No such critical points will occur inside any interval I_i (J_i), by construction. Each interval I_i (J_i) represents an interval of cone motions whose effect is identical on any given set of equivalence regions. Since the number of intervals is finite, there are only a finite number of equivalent sets of cone motions. By our construction, we have found them all, and if there exists a homing sequence, it will have to be composed of this finite set of motions. □

We can construct our power set transition graph using the set of cone motions as obtained in Theorem 16. Then by Theorem 16, a breadth-first search through the hypergraph is guaranteed to find the shortest homing sequence for our object, if one exists. Breadth-first backchaining search from a particular desired goal state back to the maximum hyperset is likewise guaranteed to find the shortest path to a desired state, if one exists.

Chapter 8

Conclusion

> Visualizing vividly the condition as a whole may mean an essential advance; and separating the condition into appropriate parts may be an important step forward. When we have found a figure that we can easily imagine, or a notation that we can easily retain, we can reasonably believe that we have made some progress.
>
> G. Polya
> *How to Solve It* [70]

8.1 Future Work

8.1.1 Existence of Homing Sequences

Clearly, in order for a part to be homed to any arbitrary goal state, it must be orientable. However, not all orientable parts can be homed to any arbitrary state. As an obvious example, a regular polygon of uniform mass distribution cannot be homed to any arbitrary edge. Any edge of a regular polygon will behave in exactly the same way for any given cone motion, and hence there is no cone motion which will bring two initial states both to a single end state. Without such merging of the states, no homing sequence can exist.

A non-square rectangle with uniform mass distribution is another orientable part which cannot be homed to any arbitrary goal. According to our planner, however, it can, be homed to either palm, to any stable orientation modulo π. Intuitively, this is because there exists an interval of cone openings which will bring both the long edge and the short edge on a given palm to a long edge on the opposite palm. Hence there is some reduction of state; but we cannot tell which long edge we are resting on, because we cannot mechanically distinguish one from the other. If the rectangle does not have uniform mass distribution, it is possible that it may again be homed to any arbitrary state, since now a given edge may not look exactly like its opposite counterpart.

Definition 8.1 (Group) *[81] A* group *is an ordered pair $(G, \circ)$ such that G is a set, $\circ$ is an associative binary operation on G, and there exists $e \in G$ such that*

- *if $a \in G$, then $a \circ e = e \circ a = a$*
- *if $a \in G$, then there exists $a^{-1} \in G$ such than $a \circ a^{-1} = a^{-1} \circ a = e$*

The element e is known as the *identity* of G. The set of in-the-plane, symmetry preserving rotations on a polygon, together with the operation which composes rotations, forms a group, whose identity is the 0 (or equivalently, 2π) rotation. When we say symmetry preserving, we are also including the location of the center of gravity. For example, a regular n-gon with uniform mass distribution has associated with it the group $\{\pi/n, 2\pi/n, ..., 2\pi\}$. The example polygon we have been using, depicted in Figure 3.2 of Chapter 3, has associated with it the group $\{2\pi\}$: there is no symmetry preserving rotation of this polygon which remains in the plane, other than the identity. These groups are called *cyclic groups*, because all of their elements can be generated by repeated composition of the smallest rotation in the group.

We make the following conjecture:

Conjecture *An orientable polygon can be homed to any arbitrary stable state if its associated cyclic group is only the identity. Otherwise, it can be homed at best to any given orientation modulo the smallest rotation of the associated cyclic group.*

Informally, we conjecture that any orientable polygon can be homed up to symmetry.

8.1.2 Minimizing Dynamic Instability

Up to this point, the analysis used by the planner has assumed that the state of the system can be described entirely by positional parameters: the orientation of the palms, and object, and which palm the object is resting upon (or was resting upon last). In other words, the state of the object is described only by its potential energy and by the (possibly frictional) contact forces applied to the object by the palms. Under these assumptions, the only way to change the state of the object is by changing the contact forces upon the object. State changes due to velocity effects have been ignored. Velocity effects can be effects due to momentum or kinetic energy, such as tumbling or impact. They can also be effects due to centrifugal or coriolis forces, such as the object being thrown from the palms or sliding off a palm due to centrifugal motion.

Unpredicted state changes due to centrifugal or coriolis forces can be minimized by keeping the rotational velocity of the palms sufficiently low, and by avoiding abrupt accelerations and decelerations. For the experimental system used to test the plans, the rotational velocity of the palms has been approximately constant in magnitude, in the range of about 0.4 to 0.5 radians per second. Although this leads to abrupt decelerations at the end of a motion, the dynamic effects due to these decelerations have not generally been observed to cause plan failure.

A more significant cause of plan failure has been excess kinetic energy of the object. In some cases, after a tilt has been executed to bring the object to a new resting state on a new palm, the translational energy given to the object by the palm motions will cause the object to continue to slide tangentially on the resting palm, until the object slides off the palm. In other cases, the rotational energy imparted to the object by the palm motions will cause it to tumble from its nominal resting state to another resting state not predicted by the planner.

The most common cause of this excess kinetic energy is a palm trajectory where a motion corresponding to a very small cone opening is followed by a tilt. This pinching of the object by the palms moves the objects radially outward on its resting palm. If the palms then execute a tilt, the effect is similar to that of a long tether which is anchored at one end, and has a mass attached to the other end. If the tether is then swung in a circle about its fixed end, it exerts a force on the mass radially inward proportional to $r\omega^2$, where r is the length of the tether and ω is the velocity of the rotation. The mass exerts a force of equal magnitude and

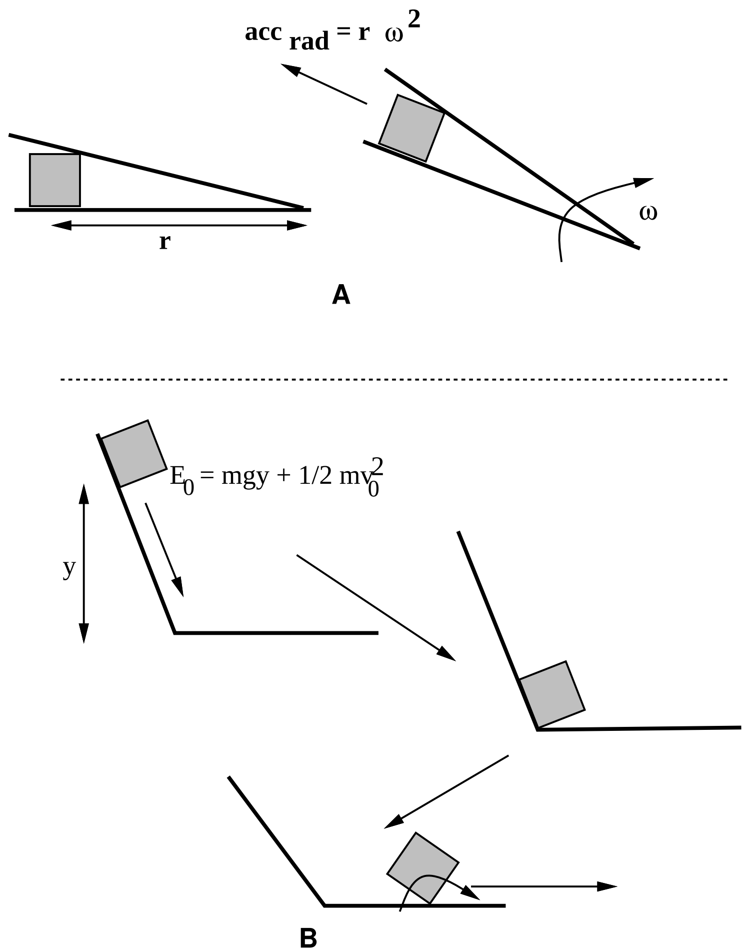

Figure 8.1: A: In the cone frame, if the radially outward component of the force on the mass dominates any forces radially inward, the part may fly off the palm. B: If the energy of the mass at the end of the motion is too high, when the part slides down the palm, much of its energy (minus what is lost in the impact and what is converted to rotational velocity) will be converted into sliding motion, and the part may slide off the palm.

opposite direction (radially outward) on the tether. If the tether were then to stretch or break in response to that force, the mass would then move radially outward. In the same way, the object on the palms wants to slide radially outward on its resting palm, and may possibly slide entirely off the palm. Even after the tilt is finished, the object will have some residual, radially outward velocity left.

Even if the tilt is successfully executed, and the object does not slide off the palms, it is sometimes the case that the next motion corresponds to the palms opening very wide. In this case, the object starts out radially far from the vertex of the cone being formed by the two palms, and gravity pulls the object towards the vertex, usually so that the object can be transferred to the other palm. The object will then slide all the way down, transforming its initial potential energy into kinetic energy. This can cause it to continue sliding when it contacts the other palm, sometimes off the palm. More often, however, the impact with the other palm can transfer the kinetic energy from translational to rotational motion of the object, and the object will tumble and rotate in a way not predicted by the planner. An algorithm for designing palm trajectories to minimize dynamic effects at higher joint velocities would be useful. Alternatively, a method of predicting dynamically unstable moves, much as we can identify frictionally unstable moves, would enable us to easily expand the planner framework which we already have.

We have already considered one way of trying to minimize dynamic effects for the plans which the existing planner generates. Recall that by Lemma 9, there are intervals of cone openings, I_i, each with an associated tilt direction, such that pure tilts for any cone opening in I_i can cause the same transition for a particular initial state, or set of possible initial states. Consider the two scenarios mentioned previously; in order to minimize the centrifugal effects in the first scenario, one must maximize the cone opening of the pinch, to keep the radial distance of the object from the cone vertex as small as possible. To minimize the dynamic effects in the second scenario, one must minimize the cone opening which follows the pinch (or open the palms very slowly) in order to keep the kinetic energy of the object as small as possible. Both of these objectives can be met simultaneously by minimizing the difference between successive cone openings in a given trajectory.

It follows from Lemma 9 that every arc of the power set transition graph represents a set of motions, or an interval of cone openings with an associated rotation direction. If we wish to minimize the difference between successive cone openings in a palm trajectory, we can do so by picking the cone openings at the endpoints of the cone opening intervals, or at least near the endpoints, if we wish to leave a little slack to account for the inaccuracies in palm positioning of the physical system. Which endpoint we pick depends on the cone openings required by the tilts preceeding and succeeding a given tilt. In principle, minimizing the difference between successive cone openings should minimize the dynamic effects in the scenarios above.

We took a few homing sequences which the existing planner produced, but which failed due to dynamic effects when run at our customary joint velocity (0.5 radians per second), and optimized them as above. Unfortunately, this first attempt at optimization was not successful. Plans which tended to fail due to dynamic effects when run at our default joint velocity still failed after optimization, at the same joint velocity. However, it is still plausible that with some additional modifications, some version of this optimization can be made to work.

8.1.3 Nonpolygonal objects

The generalization of our planar planner to nonpolygonal objects should be relatively straightforward in the case of a part whose set of equilibrium resting poses on a horizontal plane is discrete. In the following, we will discuss the frictionless case. Kriegman [46] gives the conditions under which a particular pose of a planar, piecewise smooth part will be stable in one or two point contact on a horizontal (infinite length)[1] palm.

For a piecewise smooth planar part, there are three types of faces with which a palm can make contact: a flat edge; a boundary point which is smooth (that is, a point at which the curve describing the boundary of the part is differentiable) — we will refer to this as a *smooth point*; and a boundary point which is not smooth — we will refer to this as a *vertex*. A palm can only make contact with an edge if the edge is aligned with the palm. A palm can only make contact with a smooth point if the tangent to the part's boundary at the smooth point is aligned with the palm. A palm can make contact at a vertex in some range of relative orientations. There are four types of two point contact: with an edge, with two smooth points, with two vertices, or with a smooth point and a vertex. Clearly, all these two-point contacts are analogous to the edge contact we have examined throught this dissertation, with similar conditions for stability. The conditions for stable one point contact are as follows:

1. The tangent of the part's boundary curve at the contact point is aligned with the palm.
2. The line from the contact point to the center of gravity lies on the normal of the palm at the contact point.
3. The length of this line is strictly less than the radius of curvature of the part at the contact point.

Condition 1 says that a stable one point contact can only occur at a smooth point. A smooth point which satisfies Conditions 1 and 2 but not Condition 3 corresponds to an unstable equilibrium contact. For brevity, we will refer to smooth points which satisfy Conditions 1 and 2 as *smooth equilibrium points*. Note that an actual stable edge has some point in its interior which satisfies all three conditions, since a straight line has an infinite radius of curvature.

One attempt at extending the planner is to generalize the concept of "edge" to edges or points on the boundary of the part corresponding to any equilibrium one or two point contact with the palm. Those which are stable correspond to stable edges, those which are not correspond to unstable edges. If we can then build concisely described equivalence regions, analogous to those in the polygonal case, for a piecewise smooth object, we can extend the planner to handle more general parts.

Unfortunately, the nonpolygonal case differs from the polygonal case in that, for the polygonal case, all equivalence regions of stable configurations of the part in the two palms correspond directly to a given edge making contact with a given palm. This is no longer true in the nonpolygonal case. The part can make either two or three point contact with the two palms. If it is making three point contact, one of the palms is making two point contact, and hence is making "edge" contact. It seems feasible that there is a corresponding equivalence region of stable states, with the same nice properties as in the polygonal case: simple connectedess if the "edge" is stable, piecewise linear stable paths through the equivalence region, etc.

However, the analogy is not be true in the two point contact case (that is, one point contact with each palm). If the part is making two point contact with the cone, either both of the palms are making vertex contact, or at least one of the palms is making contact with a smooth

[1]We use the convention of an infinite palm simply to indicate that we are ignoring contact between a part and the end of the palm

point. By the results in Appendix A, two point vertex contact is never stable. Hence, in order for a two point contact configuration to be stable, at least one of the contacts must be with a smooth point.

For a two point contact pose in the palms, the center of rotation will be at the intersection of the lines through the contact normals. If one draws a line parallel to gravity through this point, then in order for the part to be in equilibrium, the center of gravity must lie on this line. Hence, it is possible for a nonpolygonal piecewise smooth part to be in stable two point contact, even if neither of the contact points are smooth stable equilibrium points. For example, consider a circular disc with an offset center of gravity. This part has only one smooth stable equilibrium point, on the radius which passes through the center of gravity. Call this radial line r_{CG}. Whatever the opening or orientation of the cone, it is clear that the center of rotation is always the center of the disc, and is is equally clear that the part will always rotate in the palms until r_{CG} is parallel to the direction of gravity, and remain there, stably.

It may be possible to draw an analogy to the polygonal case if one considers unstable smooth equilibrium to be a generalization of vertices. Then an "edge" is the segment of the part boundary between two "vertices", a "vertex" being either an actual vertex or an unstable smooth equilibrium point. If the "edge" is an actual stable edge, or contains a smooth equilibrium point, then it is stable; otherwise, it is unstable. An ellipse, for example, would have two "vertices" and two stable "edges". The circle in the example above would have one "vertex" and one stable "edge". Stable "edges", like stable polygonal edges, have the property that the orientation of the part is completely determined on a horizontal palm by specifying which "edge" is in contact with the palm. There would still be some scenarios which are not seen in the polygonal case, such as a single "edge" making contact with both palms, but it seems feasible that the polygonal planner can be extended to this case.

We have not yet addressed the case where there is an continuum of smooth equilibrium points, for example an arc of a circle where the location of the center of gravity happens to coincide with the radius of the circle. It may be desirable to consider this arc as a special case "edge", since in some sense, none of the points on this arc can be distinguished from one another.

8.1.4 Three dimensional case

Kriegman has also presented a method of finding resting poses for three dimensional parts resting on a plane perpendicular to gravity which have a piecewise smooth convex hull [46], as well as the stable resting poses and capture regions for smooth parts [47]. Once a workable model of a "three dimensional cone" is determined, one could extend these algorithms to find the stable resting poses of a three dimensional part in this three dimensional cone.

There are a number of possibilities for extending our principles to the three dimensional case, including the use of both palms and fingers as constraints. If one uses fingers, however, then the convexity or nonconvexity of the part becomes an issue, whereas in the case of only palms, only the convex hull of the part matters.

In our planar system, we constrain two degrees of freedom, so the object state can be described simply by the part orientation, plus which palm is making edge contact with the part. The analogy to three dimensions would seem to be three constraints, leaving the object state to be described by its three orientatation parameters, and its relationship to the three palms (or fingers). It may be possible to use four constraints. One could fix one of the degrees of freedom of the part to be rotation about the z axis, and then orient the part up to rotation about z.

Some possibilities for three dimensional systems include

- Three or four fingers. Abell and Erdmann [1] presented a two finger stable support reorientation planner for convex polygonal planar parts. A three dimensional analogue is no doubt possible.

- Four palms arranged as two two-dimensional cones set in perpendicular frames. A variation might be one pair of palms and two fingers.

- One tiltable support plane (like the bottom of a tray tilter) and two or three palms, rather like the corner of a box with no top whose sides are hinged.

A problem with these higher dimensional formulations is that it is much harder to determine the configuration space constraint surfaces and the potential energy surfaces or curves in higher dimension. A possibility which may be computationally more tractable is to treat the problem as successive planar problems. Suppose for the sake of discussion that the convex hull of a part is polyhedral. If we have a support plane (tilted so that there a gravitional force in the support plane) and two palms in the plane, then there is one face of the polyhedron which is contacting the support plane, and the projection of the part onto the support plane which forms the polygon which the palms manipulate. There may be configurations of this polygon for which the part will switch which face contacts the support plane, or we might have some mechanism, such as a wiper to actively change the contacting face. At any rate, the orientation problems in three dimensions are reduced to a succession of orientation problems in the plane.

8.1.5 Part Singulation and Pipelining

In order for a manipulator such as the one we have discussed to be practical in assembly line situations, we must consider part singulation: how to separate the parts sufficiently so that they do not interfere with each other in the manipulator.

We must also consider feed rate. If an assembly line requires a feed rate of, for example, a part every five seconds, then our parts orienter must be able to reorient a part in less than five seconds, or we must be able to orient multiple parts in parallel so that there are always parts available to the rest of the line while the orienter is working on the current batch. One can imagine having a line of these manipulators operating in tandem to orient multiple parts. However, the problem of making sure only a single part lands in each manipulator remains.

8.1.6 Sensors

While we have modeled our orientation procedures as openloop and sensorless, it may be desirable to have some sensory feedback in order to determine when an error in the plan has occurred. One would like this sensory data to be minimal, in keeping with the criteria of simplicity and speed. Determining what sort of sensory feedback, and how to incorporate it, is a task worth addressing in the future.

8.2 Conclusion

We have presented a model of nonprehensile manipulation, using two one degree of freedom palms, and developed a planning method for part reorientation with our model. Our method finds feasible paths through the space of equivalent state configurations of the object in the palms, without requiring that the palms maintain stable support of the object over the entire path. We have shown that such a device can reliably orient parts in the plane. This device has demonstrated a number of points.

First, simple low degree of freedom devices can be used for reliable, fast manipulation of objects. In parts orienting scenarios, one would like to avoid complex mechanisms and sensors, which may break down or need careful recalibration. In addition, integrating sensor input into a manipulation algorithm increases the computational complexity, and may slow the action of the device. Sensory feedback is of course necessary in unknown or unstructured environments. However, for tasks in a structured environment, where the same action is repeated continuously, manipulators such as the one studied in this dissertation have a distinct advantage.

Second, because the devices are mechanically simple, the analysis of their mechanics is also relatively simple. To change the task from one object to another, or to change the goal state for the same object, requires only a simple software modification. APOS trays or bowl feeders, which have the same strengths of quick, reliable performance for a given task, must be custom designed to each task, whereas devices such as this one can be used for a variety of tasks. The planner we have designed is efficient, and flexible. Planning reorientations for a variety of objects requires the geometric descriptions of the objects: vertices, center of gravity, radius of gyration, and an upper bound estimate on the coefficient of friction between the part and the palms.

Third, by not relying on force closure grasps, we can exploit gravitational forces to guide the object into the correct state, without excessively precise control over the manipulator motions. Nor do we need extremely precise knowledge of frictional or restitutional coefficients. Rough estimates are sufficient. The primary mechanical analysis used by the planner is frictionless and quasistatic. Knowledge of frictional and dynamic forces is only approximate, yet the resulting plans are robust to initial conditions, friction and to small errors in the calibration of the manipulator.

Other issues such as parts singulation, higher throughput and more general object shape must be addressed in order to make such a device as we have described truly practical. However, we believe that the continued development of devices such as the one presented in this dissertation is necessary to meet the demands of modern industrial automation.

Appendix A

Properties of the Potential Energy Curve

> There is the story of how he informed Halley of one of his most fundamental discoveries of planetary motion. 'Yes,' replied Halley, 'but how do you know that? Have you proved it?' Newton was taken aback— 'Why, I've known it for years,' he replied. 'If you give me a few days, I'll certainly find a proof for it'—as in due course he did.
>
> John Maynard Keynes
> *Newton, the Man* [45]

In Chapter 3, we asserted that the potential curve of a part in two point contact with the palms (one point per palm) would be either monotonic or have a single maximum in the interior of the range of part orientations for which this two point contact is maintained. We will now prove this assertion.

Assume we have a part P in two point contact with a pair of infinite length palms, one contact per palm. The palms have a relative angle ϕ. The cone frame is defined with the origin at the vertex of the cone formed by the two palms, such that the y axis is along the cone bisector, pointing into the cone's interior. The cone frame is oriented at an angle β in the world frame (where gravity is antiparallel to the y axis). This two point contact is maintained over some range of part orientations, I. There is also an orientation of the part, θ_0, such that the two contacts have the same y coordinate, in the cone frame. We say the center of gravity is above the vertex of the cone if the y coordinate of the center of gravity in the cone frame is positive.

Theorem 17 *For any location of the center of gravity such that the center of gravity is contained inside the cone when P is in orientation θ_0, the potential energy of P as a function of orientation θ will either be monotonic or have a single maximum within the interior of I.*

For Theorem 17, the range of valid center of gravity locations (excluding issues of collision) is shown as the shaded portion in Figure A.1. This valid region also includes the palms themselves.

Suppose that at $\theta = \theta_0$, the center of gravity is not inside the cone, but is still above the cone vertex, as in Figure A.2. There may still be orientations of the object contained in the interval $[-(\pi-\phi)/2+\theta_0, (\pi-\phi)/2+\theta_0]$ for which the object is completely contained in the cone. Suppose, without loss of generality, that $\theta_0 = 0$. Then specifically, there may be an interval $I^+ = [\theta_1{}^+, \theta_2{}^+]$ such that all θ in I^+ satisfy $0 < \theta_1{}^+ < \theta < \theta_2{}^+ \leq \frac{\pi-\phi}{2}$, and for all θ in

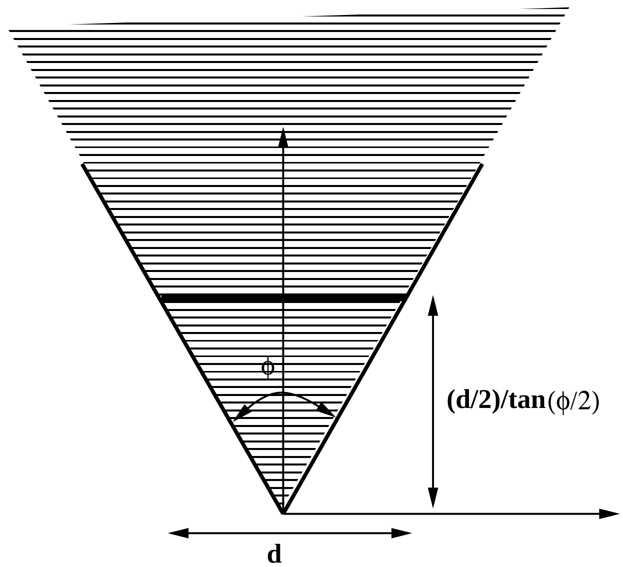

Figure A.1: The shaded portion shows the range of locations of the center of gravity satisfying the conditions of Theorem 17.

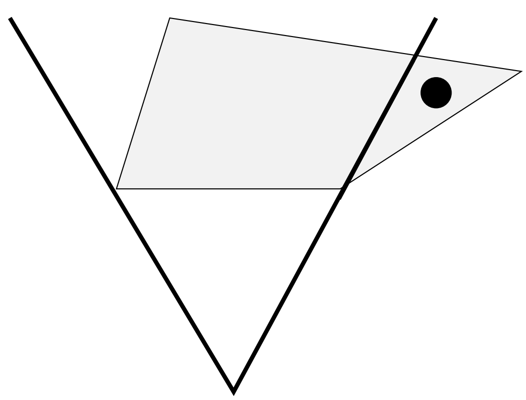

Figure A.2: The center of gravity is not inside the cone at $\theta = \theta_0$.

I^+, the center of gravity is contained in the cone. Or, there may be an interval $I^- = [\theta_1{}^-, \theta_2{}^-]$ such that all θ in I^- satisfy $\frac{-(\pi-\phi)}{2} \leq \theta_1{}^- < \theta < \theta_2{}^- < 0$, and for all θ in I^-, the center of gravity is contained in the cone. If either interval exists, then any orientation of the part in the existing interval is a *valid orientation*.

Theorem 18 *For any location of the center of gravity such that the center of gravity is outside the cone but above its vertex when P is in orientation θ_0, either*

1. *neither I^+ nor I^- exist, or*

2. *the potential energy of P as a function of orientation θ will either be monotonic or have a single maximum within the interior of the interval of valid orientations.*

We will prove Theorem 17, and then sketch the proof of Theorem 18. Before proving Theorem 17, we will give a set of simple assertions about trigonometric functions.

Assertion 1 *If either* $\sin x$ *or* $\cos x$ *is taken over an interval I of x of length at most π, then it will either be monotonic or have exactly one extremum in the interior of I.*

Assertion 2 *Let* $\arctan_2(y,x)$ *be the four quadrant arctangent of the point (x, y) (in other words, the range of* $\arctan_2$ *is* $[-\pi,\pi]$*). Then*

$$\begin{aligned} A\cos x + B\sin x &= \sqrt{A^2+B^2}\cos(x - \arctan_2(B,A)) \\ &= \sqrt{A^2+B^2}\sin(x + \arctan_2(A,B)) \end{aligned}$$

Assertion 3 *Consider the function*

$$f(x) = N\cos x + M\sin x.$$

Its derivative is

$$f'(x) = -N\sin x + M\cos x.$$

If $N > 0$, and $0 \leq x < \pi/2$, then $f'(-x) \geq f'(x)$. Therefore, by Assertions 1 and 2, either f is monotonic (if both derivatives have the same sign), or it has a single maximum between $-x$ and x. If $N = 0$, then f is monotonic between $-x$ and x.

Now let us consider a rod of length d in the interior of a cone with opening ϕ. The cone frame is as described above. Let the orientation of the rod be θ_0, when the rod is horizontal (in the cone frame). Without loss of generality, let $\theta_0 = 0$. The endpoints of the rod are given by $\mathbf{p}_l = (p_{lx}, p_{ly})$ and $\mathbf{p}_r = (p_{rx}, p_{ry})$. It is clear that the range of valid orientations of the rod such that the left endpoint $\mathbf{p}_l$ stays on the left palm, and the right endpoint $\mathbf{p}_r$ stays on the right palm is given by $I = [-(\pi-\phi)/2, (\pi-\phi)/2]$. The locations of the endpoints in the cone frame as a function of orientation are given by

$$\begin{aligned} \mathbf{p}_l(\theta) &= \kappa(\theta)\begin{pmatrix} -\sin\frac{\phi}{2} \\ \cos\frac{\phi}{2} \end{pmatrix}, \qquad \text{(A.1)} \\ & \kappa(\theta) \geq 0 \\ \mathbf{p}_r(\theta) &= \mathbf{p}_l + d\begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix} \qquad \text{(A.2)} \end{aligned}$$

where the function $\kappa(\theta)$ is to be determined.

The contact constraints are expressed by the equations

$$\frac{p_{ly}}{p_{lx}} = \frac{-1}{\tan\frac{\phi}{2}} = -\frac{p_{ry}}{p_{rx}}. \tag{A.3}$$

Combining Equations A.1, A.2, and A.3, we derive an expression for the left endpoint:

$$\mathbf{p}_l = \frac{d}{2}\frac{1}{\sin\frac{\phi}{2}}(\cos\theta - \tan\frac{\phi}{2}\sin\theta)\begin{pmatrix} -\sin\frac{\phi}{2} \\ \cos\frac{\phi}{2} \end{pmatrix}. \tag{A.4}$$

We next consider the endpoints of this rod to be the contacting vertices of our part P. Again, let $\theta = 0$ when the rod is horizontal in the cone frame. The location of the left vertex with respect to the center of gravity is

$$\mathbf{r}_l = r_l \begin{pmatrix} \cos(\theta+\theta_l) \\ \sin(\theta+\theta_l) \end{pmatrix}$$

where r_l is the distance to the left vertex, and θ_l is the angle to the left vertex, both with respect to the center of gravity.

The center of gravity location in cone coordinates is then given by

$$CG(\theta) = \frac{d}{2}\frac{1}{\sin\frac{\phi}{2}}(\cos\theta - \tan\frac{\phi}{2}\sin\theta)\begin{pmatrix} -\sin\frac{\phi}{2} \\ \cos\frac{\phi}{2} \end{pmatrix} - r_l\begin{pmatrix} \cos(\theta+\theta_l) \\ \sin(\theta+\theta_l) \end{pmatrix}. \tag{A.5}$$

The direction of gravity in the cone frame is $G = (-\sin\beta, -\cos\beta)^T$. The potential energy as a function of θ for a fixed β is then given by $\mathrm{P}_\beta(\theta) = -G^T CG(\theta)$.

Theorem 19 *Over the range $I = [-(\pi-\phi)/2, (\pi-\phi)/2]$, $\mathrm{P}_\beta(\theta)$ is either monotonic or has a single extremum in the interior of I.*

Proof: By cranking through the algebra, the potential energy function can be written

$$\begin{aligned}
\mathrm{P}_\beta(\theta) &= N\cos\theta + M\sin\theta \\
N &= -r_l\sin(\theta_l+\beta) + \frac{d}{2}\frac{1}{\sin\frac{\phi}{2}}\cos(\frac{\phi}{2}+\beta) \\
M &= -r_l\cos(\theta_l+\beta) - \frac{d}{2}\frac{1}{\cos\frac{\phi}{2}}\cos(\frac{\phi}{2}+\beta).
\end{aligned} \tag{A.6}$$

By Assertion 2, P_β can be written as a pure sine or cosine. The range of valid cone openings is $(0,\pi]$; hence, the interval of valid orientations I has length strictly less than π. By Assertion 1, P_β is therefore either monotonic or has a single extremum in the interior of I. □

Assertion 4

$$\begin{gathered}
0 < \frac{\phi}{2} \leq \pi/2 \\
-(\pi-\phi)/2 \leq \theta \leq (\pi-\phi)/2 \\
-(\pi-\phi)/2 \leq \beta \leq (\pi-\phi)/2
\end{gathered}$$

The endpoints of the range of valid β correspond to cone orientations such that the right and left palms, respectively, are horizontal. We can now prove Theorem 17.

Proof: (Theorem 17) First, we note that the orientations of the left and right palms in the world frame are given by

$$\begin{aligned} \psi_{left} &= (\pi + \phi)/2 + \beta \\ \psi_{right} &= (\pi - \phi)/2 + \beta \end{aligned}$$

By Assertion 4, the palms are therefore always in the closed Quadrants I or II of the plane; that is, $\sin \psi_{left} \geq 0$ and $\sin \psi_{right} \geq 0$. By hypothesis, the center of gravity is inside the cone at the orientation $\theta = 0$; hence it lies in the closed Quadrants I or II of the plane. Therefore, the y component of the center of gravity location, CG_y, is always greater than or equal to zero. But CG_y is also proportional to the potential energy of the part, with a positive constant of proportionality:

$$\begin{aligned} CG_y(\theta = 0) &\propto \mathrm{P}_\beta(\theta = 0) \\ &= N \cos(0) + M \sin(0) \\ &= N \end{aligned}$$

Therefore, N is nonnegative for any valid value of β. By Assertions 3 and 4, it follows that either P_β is monotonic or it has a single maximum in the interior of the range of valid part orientations. □

The sketch of the proof of Theorem 18 is as follows. Suppose that I^+ is valid. If we reparameterize P_β by $\alpha = \theta - {\theta_1}^+$, then we get a new function

$${\mathrm{P}_\beta}^+(\alpha) = N \cos(\alpha + {\theta_1}^+) + M \sin(\alpha + {\theta_1}^+),$$

where N and M are as defined in the proof of Theorem 19. This can be rewritten

$${\mathrm{P}_\beta}^+(\alpha) = N^+ \cos \alpha + M^+ \sin \alpha,$$

valid over the interval $0 \leq \alpha \leq ({\theta_2}^+ - {\theta_1}^+)$, which is of length strictly less than $\pi/2$. All values of α in this interval satisfy the conditions of Assertion 3. By hypothesis, at the orientation $\alpha = 0$ the center of gravity is inside the cone, and hence in either Quadrant I or II of the world frame. Using the same reasoning as in the proof of Theorem 17, we can say

$$\begin{aligned} CG_y(\alpha = 0) &\propto {\mathrm{P}_\beta}^+(\alpha = 0) \\ &= N^+ \cos(0) + M^+ \sin(0) \\ &= N^+ \geq 0 \end{aligned}$$

Therefore, ${\mathrm{P}_\beta}^+$ is either monotonic or has a maximum in the interior of I^+.

If I^- is valid, then we can reparameterize by $\alpha = \theta - {\theta_1}^-$, and derive an expression for the potential energy valid over the range $0 \leq \alpha \leq {\theta_2}^- - {\theta_1}^-$, and the same result follows.

Bibliography

[1] T. Abell and M. Erdmann. "Stably Supported Rotations of a Planar Polygon with Two Frictionless Contacts", *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1995.

[2] Y. Aiyama, M. Inaba, and H. Inoue. "Pivoting: A New Method of Graspless Manipulation of Object by Robot Fingers", *Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1993.

[3] S. Akella. *Robotic Parts Transfer and Orienting for Flexible Assembly.* Ph.D. Thesis, Robotics Institute, Carnegie Mellon University, 1996.

[4] S. Akella, W. Huang, K. M. Lynch, and M. T. Mason. "Planar manipulation on a conveyor with a one joint robot", presented at the *1995 International Symposium on Robotics Research*, 1995.

[5] S. Akella and M. T. Mason. "Posing Polygonal Objects in the Plane by Pushing", *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pp. 2255-2262, 1992.

[6] R. L. Andersson. "Understanding and applying a robot ping-pong player's expert controller", *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*, 1989.

[7] H. Arai and O. Khatib; "Experiments with Dynamic Skills", *Proceedings of the 1994 Japan-USA Symposium on Flexible Automation*, 1994.

[8] H. Arai and S. Tachi. "Position Control of a Manipulator with passive joints using Dynamic Coupling", *IEEE Transactions on Robotics and Automation*, vol. 7, no. 4, pp. 528-534, 1991.

[9] Aristotle. *Physics*, Book VII. From *Aristotle Selected Works*, H. G. Apostle and L. P. Gerson, trans. Peripatetic Press, Grinnell, Iowa, 1986.

[10] D. Baraff. "Issues in Computing Contact Forces for Non-penetrating Rigid Bodies", *Algorithmica*, vol. 10, pp. 292-352, 1993.

[11] D. J. Bell. *Mathematics of Linear and Nonlinear Systems*, Clarendon Press, Oxford, 1990.

[12] A. Bicchi, C. Melchiorri, D. Balluchi. "On the Mobility and Manipulability of General Multiple Limb Robots", *IEEE Transactions on Robotics and Automation*, vol. 11, no. 2, pp. 215-228, 1995.

[13] K. F. Böhringer, B. R. Donald, N. C. MacDonald. "What Programmable Vector Fields Can and Cannot Do: Force Field Algorithms for MEMS and Vibratory Plate Parts Feeders", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, vol. 1, pp. 822-829, 1996.

[14] G. Boothroyd, C. Poli and L. E. Murch. *Automatic Assembly*, M. Dekker, New York, 1982.

[15] G. Boothroyd, A. H. Redford, C. Poli, L. E. Murch. "Statistical Distributions of Natural Resting Aspects of Parts for Automatic Handling", *Manufacturing Engineering Transactions*, vol. 1, pp. 93-105, 1972.

[16] J. L. Borges. "The Garden of Forking Paths", *Labyrinths*, New Directions, New York, 1964.

[17] Raymond M. Brach. *Mechanical Impact Dynamics*, John Wiley & Sons, New York, 1991.

[18] D. L. Brock. "Enhancing the Dexterity of a Robot Hand using Controlled Slip", *Proceedings of the 1988 IEEE Conference on Robotics and Automation*, 1988.

[19] D. L. Brock. *Enhancing the Dexterity of a Robot Hand using Controlled Slip* M.S. Thesis, Massachusetts Institute of Technology, 1987. Available as Technical Report AI-TR 992.

[20] M. Brokowski, M. Peshkin, and K. Goldberg. "Curved Fences for Part Alignment", *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, 1993.

[21] R. C. Brost. "Automatic Grasp Planning in the Presence of Uncertainty", *International Journal of Robotics Research*, vol. 7, no. 1, 1988.

[22] R. C. Brost. *Analysis and Planning of Planar Manipulation Tasks.* Ph.D. Thesis, School of Computer Science, Carnegie Mellon University, January 1991. Available as CMU-CS-91-149.

[23] R. C. Brost and K. Y. Goldberg. "A Complete Algorithm for Designing Planar Fixtures Using Modular Components", *IEEE Transactions on Robotics and Automation*, vol. 12, no. 1, pp. 31-46, 1996.

[24] R. C. Brost and M. T. Mason. "Graphical Analysis of planar rigid-body dynamics with multiple frictional contacts". In H. Muira and S. Arimoto, editors, *Robotics Research: The Fifth International Symposium*, pp. 293-300. Cambridge, Massachusetts, 1990, MIT Press.

[25] M. Bühler. *Robotic Tasks with Intermittent Dynamics*, Ph.D. Thesis, Yale University, 1990.

[26] M. Caine. "The Design of Shape Interactions Using Motion Constraints", *Proceedings of the 1994 IEEE Conference on Robotics and Automation*, vol. 1, pp. 366-371, 1994.

[27] I. Calvino. "The model of models", *Mr. Palomar*, Harcourt Brace Jovanovich, San Diego, 1985.

[28] A. Christiansen. *Automatic Aquisition of Task Theories for Robotic Manipulation*, Ph.D. Thesis, Computer Science, Carnegie Mellon University, 1992. Available as Technical Report CMU-CS-92-111.

[29] W. K. Clifford. "The Postulates of the Science of Space", *The World of Mathematics*, Vol. 1, J. R. Newman, ed. Simon and Schuster, New York, 1956.

[30] J. P. Den Hartog. *Mechanics*, Dover Publications, New York, 1948.

[31] A. Einstein. "Ether and Relativity", *Sidelights on Relativity*, Dover Publications, New York, 1983.

[32] M. A. Erdmann. "On a Representation of Friction in Configuration Space", *International Journal of Robotics Research*, vol. 13, no. 3, June 1994.

[33] M. A. Erdmann. "An Exploration of Nonprehensile Two-Palm Manipulation: Planning and Execution", presented at *The International Symposium of Robotics Research*, 1995.

[34] M. A. Erdmann and M. T. Mason. "An Exploration of Sensorless Manipulation", *IEEE Journal of Robotics & Automation*, vol. 4, no. 4, 1988.

[35] M. Erdmann, M. T. Mason, and G. Vaněček, Jr. "Mechanical Parts Orienting: The Case of a Polyhedron on a Table", *Algorithmica*, vol. 10, pp. 226-247, 1993.

[36] A. O. Farahat and J. C. Trinkle. "An Investigation of Dextrous Manipulation Planning using C-Functions and Stability Functions", *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1992.

[37] K. Y. Goldberg. "Orienting Polygonal Parts without Sensors", *Algorithmica*, vol. 10, pp. 201-225, 1993.

[38] S. Goyal, A. Ruina and J. Papadopoulos. "Planar Sliding with Dry Friction 1: Limit Surface and Moment Function", *Wear*, vol. 143 no. 2, 1991.

[39] S. Goyal, A. Ruina and J. Papadopoulos. "Planar Sliding with Dry Friction 2: Dynamics of Motion", *Wear*, vol. 143 no. 2, 1991.

[40] H. Hitakawa. "Advanced parts orientation system has wide application", *Assembly Automation*, vol. 8, no. 3, 1988.

[41] L. Joskowicz and E. Sacks. *Unifying Kinematics and Dynammics for the Automatic Analysis of Machines.* IBM Technical Report RC 15573, 1990.

[42] I. Kao and M. R. Cutkosky. "Quasistatic Manipulation with Compliance and Sliding", *International Journal of Robotics Research*, vol 11, No. 1, February 1992.

[43] I. Kao and M. R. Cutkosky. "Comparison of Theoretical and Experimental Force/Motion Trajectories for Dextrous Manipulation With Sliding", *International Journal of Robotics Research*, vol. 12, No. 6, December 1993.

[44] J. Kepler. *Gesammelte Werke*, heraus gegeben unter der Leitung von Walther von Duck und Max Caspar. München, C. H. Beck, 1937. cited in R. Torretti. *Philosophy of Geometry from Riemann to Poincaré*, D. Reidel Publishing, Dordrecht, 1984.

[45] J. M. Keynes. "Newton, the Man", exerpted in *The World of Mathematics*, Vol. 1, J. R. Newman, ed. Simon and Schuster, New York, 1956.

[46] D. J. Kriegman. "Computing Stable Poses of Piecewise Smooth Objects", *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 55, no. 2, pp. 109-118, March 1992.

[47] D. J. Kriegman. "Let Them Fall Where They May: Capture Regions of Curved 3D Objects", *Proceedings, IEEE International Conference of Robotics and Automation*, 1994.

[48] J. Krishnasamy, M. J. Jakiela, D. E. Whitney, "Mechanics of Vibration-assisted Entrapment with Application to Design", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, pp. 838-845, 1996.

[49] D, Koditschek and M. Bühler. "Analysis of a Simplified Hopping Robot", *International Journal of Robotics Research*, vol. 10, no. 6, December 1991.

[50] J. C. Latombe. *Robot Motion Planning*, Klewer Academic Publishers, Norwell, 1991.

[51] T. Lozano-Pérez. "Spatial Planning: A Configuration Space Approach", *IEEE Transactions on Computers*, vol. C-32, no. 2, pp. 108-120, 1983.

[52] T. Lozano-Pérez, M. T. Mason and R. H. Taylor. "Automatic Synthesis of Fine-Motion Strategies for Robots", *International Journal of Robotics Research*, vol. 3, no. 1, 1984.

[53] K. M. Lynch. "The Mechanics of Fine Manipulation by Pushing", *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, 1992.

[54] K. M. Lynch. *Nonprehensile robotic manipulation: Controllability and planning.* Ph.D. Thesis, Robotics Institute, Carnegie Mellon University, March 1996. Available as CMU-RI-TR-96-05.

[55] K. M. Lynch and M. T. Mason. "Pulling by Pushing, slip with Infinite Friction, and Perfectly Rough Surfaces", *International Journal of Robotics Research*, vol. 14, no. 2, pp. 174-183, 1995.

[56] K. M. Lynch and M. T. Mason. "Stable Pushing: Mechanics, Controllability, and Planning", *The First Workshop on the Algorithmic Foundations of Robotics*, AK Peters, Boston, 1995.

[57] M. Mani and W. Wilson. "A Programmable Orienting System for Flat Parts", *Proceedings, North American Research Institute Conference XIII*, 1985.

[58] M. Mason. "Mechanics and Planning of Manipulator Pushing Operations", *International Journal of Robotics Research*, vol. 5, no. 3, Fall 1986.

[59] M. Mason and K. Lynch. "Dynamic Manipulation", *Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1993.

[60] M. Mason and K. Lynch. "Throwing a Club: Early Results", presented at *The International Symposium of Robotics Research*, Hidden Valley, PA; Oct. 2-5, 1993.

[61] R. Mattikalli, D. Baraff, P. Khosla, and B. Repetto. "Gravitational Stability of Frictionless Assemblies", *IEEE Transactions on Robotics and Automation*, vol 11, no. 3, pp. 374-388, 1995.

[62] R. Mattikalli, D. Baraff, and P. Khosla. "Finding All Stable Orientations of Assemblies with Friction", *IEEE Transactions on Robotics and Automation*, vol 12, no. 2, pp. 331-342, 1996.

[63] B. Mirtich, Y. Zhuang, K. Goldberg, J. Craig, R. Zanutta, B. Carlisle, and J. Canny. "Estimating Pose Statistics for Robotic Part Feeders", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, vol 2, pp. 1140-1146, 1996.

[64] B. K. Natarajan. "Some Paradigms for the Automated Design of Parts Feeders", *International Journal of Robotics Research*, vol. 8, no. 6, pp. 98-109, 1989.

[65] V. Nguyen. "Constructing Force-Closure Grasps", *International Journal of Robotics Research*, vol. 7, no. 3, pp. 3-16, June 1988.

[66] E. Paljug, X. Yun. "Experimental Results of Two Robot Arms Manipulating Large Objects" *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, 1993.

[67] E. Paljug, X. Yun, and V. Kumar. "Control of Rolling Contacts in Multi-Arm Manipulation", *IEEE Transactions on Robotics & Automation*, vol. 10, no. 4, August 1994.

[68] M. A. Peshkin and A. C. Sanderson. "Planning Robotic Manipulation Strategies", *IEEE Journal of Robotics & Automation*, vol. 4, no. 5, October 1988.

[69] M. A. Peshkin and A. C. Sanderson. "The Motion of a Pushed, Sliding Workpiece", *IEEE Journal of Robotics & Automation*, vol. 4, no. 6, December 1988.

[70] G. Polya. *How to Solve It*, Princeton University Press, Princeton, 1957.

[71] J. Ponce and B. Faverjon. "On Computing Three-Finger Force-Closure Grasps of Polygonal Objects", *IEEE Transactions on Robotics and Automation*, vol. 11, no. 6, 1995.

[72] J. Ponce. "On Planning Immobilizing Fixtures for Three-Dimensional Polyhedral Parts", *Proceedings of the 1996 IEEE Conference on Robotics and Automation*, vol. 2, pp. 509-514, 1996.

[73] M. H. Raibert. *Legged Robots that Balance*, The MIT Press, Cambridge, MA, 1986.

[74] A. S, Rao and K. Y. Goldberg. "Manipulating Algebraic Parts in the Plane", *IEEE Transactions on Robotics & Automation*, vol. 11, no. 4, pp. 598-602, August 1995.

[75] A. Rao, D. J. Kriegman, and K. Y. Goldberg. "Complete Algorithms for Feeding Polyhedral Parts Using Pivot Grasps", *IEEE Transactions on Robotics & Automation*, vol. 12, no. 2, pp. 331-342, April 1996.

[76] F. Reuleaux. *The Kinematics of Machinery*, Dover Publications, New York, 1963.

[77] E. Rimon and A. Blake. "Caging 2D Bodies by 1-Parameter Two-Finger Gripping Systems", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, vol. 2, pp. 1458-1464, 1996.

[78] E. Rimon and J. Burdick. "Mobility of Bodies in Contact I: A new 2nd order mobility index for multifingered grasps", *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 2329-2335, 1994.

[79] A. A. Rizzi and D. E. Koditschek "Progress in Spatial Robot Juggling", *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, 1992.

[80] E. J. Routh. *Dynamics of a System of Rigid Bodies, Vol. 1*, Dover Publications, New York, 1960.

[81] W. R. Scott. *Group Theory*, Dover Publications, New York, 1987.

[82] W. J. Stronge. "Rigid body collisions with friction", *Proceedings of the Royal Society of London,* A 431, 1990.

[83] W. T. Thomson. *Introduction to Space Dynamics*, Dover Publications, New York, 1986.

[84] J. C. Trinkle. "On the Stability and Instantaneous Velocity of Grasped Frictionless Objects", *IEEE Transactions on Robotics & Automation*, vol. 8, no. 5, pp. 560-572, October 1992.

[85] J. C. Trinkle, A. O. Farahat, and P. F. Stiller. "First Order Stability Cells of Active Frictionless Rigid Body Systems", *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, 1993.

[86] J. C. Trinkle, A. O. Farahat, and P. F. Stiller, P. F. "Second Order Stability Cells of Frictionless Rigid Body Systems", *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, 1994.

[87] J. C. Trinkle, A. O. Farahat, and P. F. Stiller. "First Order Stability Cells of Active Multi-Rigid-Body Systems" *IEEE Transactions on Robotics and Automation*, vol. 11, no. 4, pp. 545-557, 1995.

[88] J. C. Trinkle and J. J. Hunter. "A Framework for Planning Dextrous Manipulation", *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, 1991.

[89] J. C. Trinkle, R. P. Paul. "Planning for Dexterous Manipulation with Sliding Contacts", *International Journal of Robotics Research*, vol. 9, no. 3, June 1990.

[90] J. C. Trinkle, R. C. Ram, A. O. Farahat, and P. F. Stiller. "Dextrous Manipulation Planning and Execution of an Enveloped Slippery Workpiece", *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, 1993.

[91] J. C. Trinkle and D. C. Zeng. "Prediction of the Quasistatic Planar Motion of a Contacted Rigid Body", *IEEE Transactions on Robotics & Automation*, vol. 11 no. 2, April, 1995.

[92] A. S. Wallack. "Generic Fixture Design Algorithms for Minimal Modular Fixture Toolkits", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, 1996.

[93] Y. Wang. *On Impact Dynamics of Robotic Operations*, Technical Report CMU-RI-TR-86-14, Robotics Institute, Carnegie Mellon University, 1986.

[94] Y. Wang. *Dynamic Analysis and Simulation of Mechanical Systems with Intermittent Constraints*, Ph.D. Thesis, Mechanical Engineering, Carnegie Mellon University, 1989.

[95] J. Weigley, K. Y. Goldberg, M. Peshkin, and M. Brokowski. "A Complete Algorithm for Designing Passive Fences to Orient Parts", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, 1996.

[96] T. Yoshikawa, Y. Yokokohji and A. Nagayama. "Object Handling bu Three-Fingered Hands using Slip Motion", *Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 99-105, 1993.

[97] X. Yun. "Object Handling Using Two Arms without Grasping", *International Journal of Robotics Research*, vol. 12, no. 1, Feb. 1993.

[98] N. B. Zumel and M. A. Erdmann. "Balancing of a Planar Bouncing Object", *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, 1994.

[99] N. B. Zumel and M. A. Erdmann. "Nonprehensile Two Palm Manipulation with Non-Equilibrium Transitions between Stable States", *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, 1996.