Dynamic Nonprehensile Manipulation: Controllability, Planning, and Experiments

Kevin M. Lynch and Matthew T. Mason

The International Journal of Robotics Research 1999; 18; 64
DOI: 10.1177/027836499901800105

Kevin M. Lynch

Department of Mechanical Engineering
Northwestern University
Evanston, Illinois 60208-3111, USA
kmlynch@nwu.edu

Matthew T. Mason

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213, USA

Dynamic Nonprehensile Manipulation: Controllability, Planning, and Experiments

Abstract

We are interested in using low-degree-of-freedom robots to perform complex tasks by nonprehensile manipulation (manipulation without a form- or force-closure grasp). By not grasping, the robot can use gravitational, centrifugal, and Coriolis forces as virtual motors to control more degrees of freedom of the part. The part s extra motion freedoms are exhibited as rolling, slipping, and free flight.

This paper describes controllability, motion planning, and implementation of planar dynamic nonprehensile manipulation. We show that almost any planar object is controllable by point contact, and the controlling robot requires only two degrees of freedom (a point translating in the plane). We then focus on a one-joint manipulator (with a two-dimensional state space), and show that even this simplest of robots, by using slipping and rolling, can control a planar object to a full-dimensional subset of its six-dimensional state space. We have developed a one-joint robot to perform a variety of dynamic tasks, including snatching an object from a table, rolling an object on the surface of the arm, and throwing and catching. Nonlinear optimization is used to plan robot trajectories that achieve the desired object motion via coupling forces through the nonprehensile contact.

1. Introduction

We are interested in using low-degree-of-freedom robots to perform complex tasks by nonprehensile manipulation (manipulation without a form- or force-closure grasp). By not grasping, the robot can use gravitational, centrifugal, and Coriolis forces as virtual motors to control more degrees of freedom of a part. The part's extra motion freedoms are exhibited as rolling, slipping, and free flight. An example of dynamic nonprehensile manipulation is shooting a basketball: the ball is sent to the basket using rolling and free flight (Fig. 1). The motion of the ball relative to the hand is induced by the forces due to the acceleration of the contact constraint.

Dynamic nonprehensile manipulation offers several potential benefits:

- *New robot primitives.* Nonprehensile manipulation provides an option for a robot when it must manipulate an object too large or too heavy to be grasped and lifted.
- *Simpler manipulators.* The structure of the manipulator can be simplified by eliminating the gripper.
- *Flexibility.* Grippers are often designed to grasp parts of particular shapes and sizes. With nonprehensile manipulation, if the robot can apply a force to the part, it can manipulate the part. Nonprehensile manipulation allows a manipulator to control multiple parts simultaneously, using whatever surfaces of the manipulator are available.
- *Increased workspace size.* We usually define the workspace of a robot as the kinematic workspace of its end effector. In some situations, a more useful concept is the set of reachable states for an object manipulated by the robot. If the robot can throw the object to points outside its kinematic workspace, the size of the robot's workspace is effectively increased.
- *Increased workspace dimensionality.* When an object is carried with a grasp, the manipulator must have at least as many degrees of freedom as has the object we want to control. If we allow the object to move relative to the manipulator, we can control more degrees of freedom than the manipulator itself has. This is *underactuated* manipulation. It is not always necessary to attach a motor to every degree of freedom we would like to control.

The International Journal of Robotics Research
Vol. 18, No. 1, January 1999, pp. 64-92,

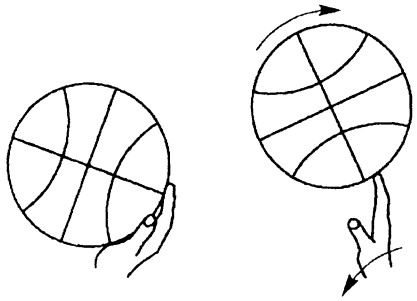

Fig. 1. Shooting a basketball using rolling contact.

These benefits come at the expense of increased complexity in planning and control. For example, planning for pick-and-place manipulation requires only a kinematic model of the world; dynamic manipulation requires a dynamic model. Dynamic nonprehensile manipulation transfers some of the complexity of a robot system from hardware (grippers, joints, and actuators) to planning and control.

In this paper, we study *controllability* and *motion planning* for planar nonprehensile manipulation with a dynamic model. The controllability problem is to characterize the object's accessible state space during nonprehensile manipulation. We begin by assuming no constraints on the motion of the robot, and we show that almost any planar object is controllable by dynamic pushing with point contact. The controlling robot requires only two degrees of freedom (a point translating in the plane). We then focus on a one-joint manipulator (with a two-dimensional state space), and show that even this simplest of robots, by using slipping and rolling, can control an object to a full-dimensional subset of its six-dimensional state space.

The motion-planning problem is to find a manipulator trajectory to transfer the object to the goal state using frictional, gravitational, and dynamic forces. This is a kind of "dynamic pick-and-place." By using a sequence of different manipulation phases, such as rolling and free flight, the robot can control more degrees of freedom of the object. We use nonlinear optimization to solve for a manipulator trajectory that satisfies dynamic constraints and force-inequality constraints at the nonprehensile contact. The resulting trajectory takes the object to the goal state while locally minimizing a given objective function. We have successfully implemented trajectories found by the motion planner on a one-joint direct-drive arm to perform a variety of dynamic tasks, such as snatching an object from a table, rolling an object on the surface of the arm, and throwing and catching.

This work pursues a minimalist approach to robotic manipulation. We are motivated by the academic interest to understand the simplest mechanisms capable of performing a given task and the economic motive to construct simpler, cheaper robots. Simple robots employing dynamic nonprehensile manipulation have been successfully used for parts-feeding tasks in industry (Boothroyd, Poli, and Murch 1982; Hitakawa 1988). Dynamic manipulation may also be useful in space, where dynamic effects dominate.

1.1. Relation to Previous Work

Our previous work on underactuated nonprehensile manipulation has focused on quasi-static pushing. By analyzing the mechanics of pushing, we have shown:

- A two-degrees-of-freedom robot (a point translating in the plane) can push an object arbitrarily closely along any path in the object's three-dimensional configuration space, unless the object is a frictionless disk centered at its center of mass (Lynch and Mason 1996).
- A single revolute joint, operating above a fixed-speed conveyor, can move any polygon from any initial configuration upstream of the joint to a specified goal configuration on the conveyor by using pushing and conveyor drift (Akella et al. 1995).

During quasi-static pushing, the state of a pushed object is simply its configuration, and its motion is subject to a set of nonholonomic velocity constraints arising from limits on the forces that can be applied by pushing. In this sense, pushing an object is similar to more familiar kinematic nonholonomic systems with rolling constraints. Examples include dextrous manipulation and wheeled mobile robots.

Dynamic nonprehensile manipulation resembles pushing, but now we have second-order dynamics and second-order nonholonomic constraints. There is also a drift term corresponding to the object's motion when no control force is applied. The problems studied in this paper can be thought of as dynamic pushing where support friction is negligible.

1.2. Overview

The work described in this paper draws on previous work in nonprehensile manipulation, motion planning and control for nonholonomic systems, optimal trajectory planning, and minimalism in robotics. We touch on some related work in these areas in Section 2. Section 3 provides some definitions. The reachable state space of an object during dynamic nonprehensile manipulation is studied in Section 4. Section 5 describes a dynamic motion planner, and Section 6 presents some experiments in dynamic manipulation with a one-joint robot.

This paper presents early work on dynamic underactuated nonprehensile manipulation, to begin to understand its basic properties and to test experimental feasibility. In the conclusion, we discuss future directions.

2. Related Work

2.1. Nonprehensile Manipulation

Koditschek (1993) has provided a thoughtful review of seminal works in the area of dynamic manipulation. Some examples of dynamic nonprehensile manipulation are discussed below.

2.1.1. Parts Reorienting

Using a dynamic simulator that takes careful account of friction (Erdmann 1994), Erdmann (1995) designed the acceleration of a flat palm to "stand up" a block resting on it. Similarly, Arai and Khatib (1994) used dynamic forces to roll a cube sitting on a paddle held by a PUMA robot arm. The paddle has three-degrees-of-motion freedom in the vertical plane. The paddle first imparts an angular velocity to the object, causing it to begin to roll about the desired edge, then accelerates downward to decrease the impact at the end of the roll. The approach assumes no slip at the rolling contact.

In this paper, we address the problem of planning such trajectories automatically, explicitly considering friction constraints and manipulator kinematic and dynamic constraints, especially for low-degree-of-freedom manipulators.

2.1.2. Throwing

Learning approaches have been applied to improve a robot's ability to throw (Aboaf, Atkeson, and Reinkensmeyer 1987; Schneider and Brown 1993). Throwing a club using a *dynamic grasp* was previously reported by Mason and Lynch (1993a, 1993b). (A dynamic grasp is defined as a manipulator acceleration such that the dynamic load keeps the object fixed to the manipulator as it moves.) Positioning an object allowed to slide freely on a supporting surface has been studied by Huang, Krotkov, and Mason (1995) and Zhu and colleagues (1996).

2.1.3. Catching

Burridge, Rizzi, and Koditschek (1995) described three different "mirror laws" for controlled juggling, catching, and palming a ball with a planar paddle. By switching control laws from juggling to catching to palming, the robot transitions from juggling the ball to balancing it on the paddle. Hove and Slotine's (1991) catching robot uses visual data to estimate the flight of a ball, matches the trajectory of the end effector to that of the ball, and then grabs the ball.

In this paper, the catching surface is motionless, and the object's arrival state at the end of flight is chosen to make the catch robust.

2.1.4. Batting

Batting combines catching and throwing into a single collision. Robot juggling by batting has been demonstrated for one or two pucks with a planar juggler (Bühler and Koditschek 1990), a ball in space (Aboaf, Drucker, and Atkeson 1989; Rizzi and Koditschek 1992), and two balls in space (Rizzi and Koditschek 1993). The problem of batting a planar polygon to a desired stable orientation was addressed by Zumel and Erdmann (1994). Schaal and Atkeson (1993) used batting in their "devil sticking" robot. Andersson (1989) built a machine to play Ping-Pong, which is an adversarial form of batting.

2.1.5. Pushing

Mason (1986) discovered a simple rule for determining the rotation sense of a pushed object that depends on the center of friction of the object, not its precise support distribution. This result has been used and extended by many to plan parallel-jaw grasps, to construct parts feeders, and to plan pushing paths among obstacles. See the work by Lynch and Mason (1996) for an extensive bibliography of work on pushing.

2.1.6. Others

Other examples of nonprehensile manipulation include tray-tilting to position polygonal parts sliding in a tray (Erdmann and Mason 1988) and to orient polyhedral parts rolling on a rough surface (Erdmann, Mason, and Vánek 1993); using robot fingers to rotate a polyhedral part about an edge (Sawasaki, Inaba, and Inoue 1989) or pivot it about a vertex (Aiyama, Inaba, and Inoue 1993) in contact with a floor; cooperative manipulation by two nonprehensile palms (Erdmann 1995; Zumel and Erdmann 1996) or two frictionless pins (Abell and Erdmann 1995); positioning and orienting parts on a vibrating support surface (Böhringer et al. 1995; Reznik and Canny 1998); and various forms of whole-arm manipulation (Salisbury 1987; Trinkle et al. 1993).

2.2. Motion Planning and Control of Underactuated Systems

Results from nonlinear control theory are useful in characterizing the controllability of nonholonomically constrained robot systems. Good introductions to the field of nonlinear geometric control theory are provided by Isidori (1989) and Nijmeijer and van der Schaft (1990), and the text by Boothby (1986) is an excellent introduction to many of the differential geometric concepts in nonlinear control. The Lie algebra rank condition, an important nonlinear analog to the Kalman controllability rank condition of linear control theory, and its implications are discussed by Brockett (1976), Haynes and Hermes (1970), Hermann and Krener (1977), Jurdjevic (1972), Sussmann and Jurdjevic (1972), and Sussmann (1983). Although there is no general necessary and sufficient condition for the controllability of nonlinear systems, Sussmann (1987) derived a general sufficient condition for small-time local controllability. Even if a nonholonomic system is fully controllable, however, Brockett (1983) showed that it may not

stabilizable to a single equilibrium point using time-invariant smooth state feedback. This suggests the use of algorithmic or nonsmooth control. Introductions to nonholonomic motion planning and the use of nonlinear control theory in robotics are given by Latombe (1991), Murray, Li, and Sastry (1994), and Li and Canny (1993).

Dynamic nonprehensile manipulation is closely related to the problem of controlling the position and orientation of a hovercraft with thrusters (Manikonda and Krishnaprasad 1997) and controlling the motion of an unactuated joint of a robot arm. As shown by Oriolo and Nakamura (1991), unactuated robot joints typically result in nonintegrable second-order constraints. Research on controlling unactuated joints by dynamic coupling includes controlling the swinging motion of a high-bar robot (Takashima 1991) or the Acrobot (Hauser and Murray 1990; Spong 1994); the control of running and hopping robots (Berkemeier and Fearing 1992); controlling the base of a mobile robot by motions of the end effector (De Luca, Mattone, and Oriolo 1996); a series of papers by Arai and colleagues (Arai and Tachi 1991; Arai, Tanie, and Tachi 1993), and related work by Bergerman, Lee, and Xu (1995), examining controllability and describing control algorithms for manipulators with some joints equipped only with brakes. Control of unactuated joints without brakes was studied by Suzuki, Koinuma, and Nakamura (1996), Arai (1996), and Lynch and colleagues (1998).

One distinguishing feature of nonprehensile manipulation is that contact forces are unilateral, while most controllability analyses assume bidirectional controls. One exception is the work by Goodwine and Burdick (1996), which is based on Sussmann's (1987) general theorem. In this paper, we demonstrate controllability using unilateral contact forces.

Other researchers have studied underactuated manipulation using kinematic coupling instead of dynamic coupling. Examples include pivoting an object within a grasp (Brock 1988; Carlisle et al. 1994; Rao, Kriegman, and Goldberg 1995); rolling an object between two palms (Bicchi and Sorrentino 1995); rolling a ball on a plane or another ball (Li and Canny 1990); and controlling an n-link planar manipulator with only two motors using nonholonomic gears (Sørdalen, Nakamura, and Chung 1994). Perhaps the most common example of underactuated manipulation is assembly, where robots possessing just a few degrees of freedom are required to assemble a set of parts with a larger number of degrees of freedom (Koditschek 1991; Ferbach and Rit 1996). Wheeled mobile robots are perhaps the most heavily studied type of kinematically constrained nonholonomic system (Laumond 1986; Barraquand and Latombe 1993; Laumond et al. 1994).

The *snakeboard* is an interesting nonholonomic system that is similar to a skateboard, except the directions of the front and rear wheels can be controlled independently. By shifting his or her weight, the rider produces motion of the snakeboard through rolling of the wheels. A simplified model of the snakeboard has been studied by Ostrowski and colleagues (1995).

A common property of underactuated systems is the inability to independently drive all state variables to their desired values using time-invariant smooth state feedback (Brockett 1983). Thus some sort of time-varying or nonsmooth feedback law, or open-loop control, is required. Strategies for generating open-loop controls for systems with nonholonomic constraints have been proposed by Murray and Sastry (1993), Lafferiere and Sussmann (1991), Sussmann (1993), Sontag (1993), and Leonard and Krishnaprasad (1995). The problem is simplified when the nonholonomic system possesses a certain structure, such as differential flatness (Fliess et al. 1995; Murray, Rathinam, and Sluis 1995). Divelbiss and Wen (1993) and Fernandes, Gurvits, and Li (1994) proposed an iterative optimization approach akin to Newton's method, where the controls are given by finite Fourier series. At each step, the coefficients of the Fourier terms are adjusted to minimize a merit function. In the work of Divelbiss and Wen, the merit function was given by the error at the final state and the potential functions representing obstacles. They illustrated the approach by planning the motion of a tractor pulling a trailer. Fernandes, Gurvits, and Li (1994) used a merit-function quadratic in the Fourier coefficients and the error at the final state. They presented examples of controlling the attitude of a falling cat and a space platform/manipulator system.

The approach to trajectory planning for dynamic nonprehensile manipulation employed in this paper is an iterative optimization technique similar to these works. An important characteristic of the trajectory-planning problem for dynamic nonprehensile manipulation is the inequality constraints imposed by unilateral frictional contact. A closely related approach that handles inequality constraints was recently proposed by Zefran, Desai, and Kumar (1996).

2.3. Optimal Trajectory Planning

The problem of finding open-loop controls for nonholonomically constrained systems is related to the optimal trajectory-planning problem for fully actuated systems. To increase productivity, much work has focused on time-optimal trajectory planning (Shin and McKay 1985; Bobrow, Dubowsky, and Gibson 1985; Shiller and Dubowsky 1991; Canny et al. 1988; Donald and Xavier 1989). Objective functions other than time have also been considered (Witkin and Kass 1988; Yen and Nagurka 1988; De Luca, Lanari, and Oriolo 1991; Chen 1991; Martin and Bobrow 1997). In these works, robot motions are typically represented as finite-dimensional polynomials, Fourier series, or splines, and the problem is converted into a nonlinear optimization.

2.4. Minimalism

Canny and Goldberg (1994) argue that RISC (reduced intricacy in sensing and control) robotics using simple devices results in cheaper, more flexible systems. The minimalist approach to

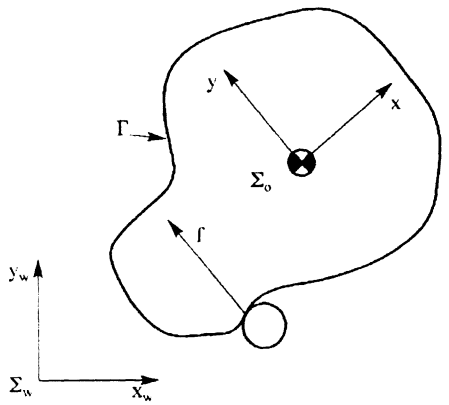

Fig. 2. Notation for dynamic nonprehensile manipulation. The object can be pushed at any point on its perimeter Γ.

robotics is the attempt to find the simplest systems capable of performing a given task or class of tasks, where the complexity of a system may be measured in terms of the sensors, effectors, computation, communication, and so forth (Böhringer et al. 1995). This paper explores minimal actuator solutions to dynamic manipulation tasks.

3. Definitions and Assumptions

All problems considered in this paper are planar. The planar object $\mathcal{O}$ can be contacted anywhere along its closed, piecewise-smooth perimeter Γ (Fig. 2). Friction between the robot and the object conforms to Coulomb's law.

The configuration space of the object is $\mathcal{C} = SE(2) = \mathbf{R}^2 \times SO(2)$. An object frame $\Sigma_\mathcal{O}$ is fixed to the center of mass of the object. Coordinates in this frame are $(x, y, \phi)^T$. The configuration of $\Sigma_\mathcal{O}$ in the inertial world frame $\Sigma_\mathcal{W}$ is written $\mathbf{q} = (x_w, y_w, \phi_w)^T$. The state space of the object is the tangent bundle $T\mathcal{C}$, and the object's state is given by $(\mathbf{q}, \dot{\mathbf{q}})$. The tangent space at $(\mathbf{q}, \dot{\mathbf{q}})$ is written $T_{(\mathbf{q},\dot{\mathbf{q}})}T\mathcal{C}$.

Generalized forces $\mathbf{f} = (f_x, f_y, \tau)^T$ are written in the object frame $\Sigma_\mathcal{O}$. A *pure force* is a force with a zero torque component ($\tau = 0$), and a *pure torque* is a force with a zero linear component ($f_x = f_y = 0$). A *force direction* $\hat{\mathbf{f}} = \left(\hat{f_x}, \hat{f_y}, \hat{\tau}\right)^T$ is defined as $\mathbf{f}/|\mathbf{f}|$. The *force sphere* is the two-dimensional unit sphere S^2 of all force directions.

4. Accessibility

This section studies the accessible state space of a planar object $\mathcal{O}$ during dynamic nonprehensile manipulation. We consider two kinds of control forces: unilateral contact forces fixed in the object frame (Section 4.1) and forces arising from contact with a one-joint robot (Section 4.2).

Using notation from Sussmann (1983) and Nijmeijer and van der Schaft (1990), we define $R^W(\mathbf{q}, \dot{\mathbf{q}}, T)$ to be the object's reachable state space from $(\mathbf{q}, \dot{\mathbf{q}})$ at time $T > 0$ by object trajectories remaining in the neighborhood W of $(\mathbf{q}, \dot{\mathbf{q}})$ at times $t \in [0, T]$. Define $R^W(\mathbf{q}, \dot{\mathbf{q}}, \leq T) = \bigcup_{0 \leq t \leq T} R^W(\mathbf{q}, \dot{\mathbf{q}}, t)$. Then the system (or simply the object) is *small-time accessible from* $(\mathbf{q}, \dot{\mathbf{q}})$ if $R^W(\mathbf{q}, \dot{\mathbf{q}}, \leq T)$ contains a nonempty open set of $T\mathcal{C}$ for any neighborhood W of $(\mathbf{q}, \dot{\mathbf{q}})$ and all $T > 0$. The object is *small-time locally controllable from* $(\mathbf{q}, \dot{\mathbf{q}})$ if $R^W(\mathbf{q}, \dot{\mathbf{q}}, \leq T)$ contains a neighborhood of $(\mathbf{q}, \dot{\mathbf{q}})$ for any neighborhood W and all $T > 0$. The object is *controllable from* $(\mathbf{q}, \dot{\mathbf{q}})$ if, for any $(\mathbf{q}_1, \dot{\mathbf{q}}_1) \in T\mathcal{C}$, there exists a finite time T such that $(\mathbf{q}_1, \dot{\mathbf{q}}_1) \in R^{T\mathcal{C}}(\mathbf{q}, \dot{\mathbf{q}}, T)$. The phrase "from $(\mathbf{q}, \dot{\mathbf{q}})$" can be eliminated from each of these definitions if the condition is satisfied for all $(\mathbf{q}, \dot{\mathbf{q}})$.

We begin (in Section 4.1) by examining the case of no constraints on the motion of the manipulator $\mathcal{M}$, which can contact any point on the object's perimeter Γ. With this assumption, we demonstrate sufficient conditions for the controllability of the object by pushing and batting. The control forces for this system correspond to unilateral, body-fixed forces from point contact with the object. These forces are similar to those obtained with thrusters.

Because of the difficulty of breaking contact and re-contacting a moving object, we then include manipulator motion constraints in the analysis (Section 4.2). We study the simplest possible case: a single-degree-of-freedom robot that maintains point contact with the object as it moves. Control forces arise from the frictionless contact. We show that a one-degree-of-freedom revolute robot, with just a two-dimensional state space, can take a planar object to a six-dimensional subset of its six-dimensional state space by using slipping and rolling. In other words, the equality constraints on the state of the manipulator (the pivot remains fixed) usually do not translate to equality constraints on the state of the object.

4.1. No Manipulator Constraints

To visualize the control system with no manipulator constraints, imagine an object floating on an air table that may be tilted to yield a gravitational acceleration in the plane of the table. The object is pushed or batted by a manipulator. Alternatively, the object can be considered to be a planar free-flying rigid body with gas jets attached to its perimeter. The angle that the gas jets can take with respect to the normal of the perimeter is determined by the friction coefficient μ.

The control system is written

$$(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) = X_0(\mathbf{q}, \dot{\mathbf{q}}) + uX_i(\mathbf{q}, \dot{\mathbf{q}}), \quad i \in \{1, \ldots, n\},$$

where X_0 is a drift vector field, u is a nonnegative scalar control, and X_i is the control vector field, where i chooses which of the n control vector fields is used. (Note that only one control force is applied at a time, and u must be nonnegative due to unilateral contact. The control is determined by both u and i, but usually we refer to u as the control input.) The vector field X_i, $i \in \{1, \ldots, n\}$, corresponds to a unit force $\hat{\mathbf{f}}_i$ fixed in the object frame $\Sigma_{\mathcal{O}}$, arising from point contact on the perimeter Γ. The set of force directions $\cup_i \hat{\mathbf{f}}_i$ is denoted $\hat{\mathcal{F}}$. For simplicity, the mass m and radius of gyration ρ (where the inertia is $m\rho^2$) of $\mathcal{O}$ are assumed to be unit.

We consider two control sets: U_u, the set of nonnegative inputs $[0, \infty)$, and U_d, the discrete set $\{0, 1\}$. While we are typically interested in the case $u \in U_u$, which corresponds to the usual assumption in grasping that grasp forces can be as large as necessary to maintain the grasp, several of the results also apply to the case $u \in U_d$. To state these results in full generality, we note when they apply to $u \in U_d$. Unless otherwise noted, the control set U_u is assumed.

The state of $\mathcal{O}$ is $(\mathbf{q}, \dot{\mathbf{q}}) = (x_w, y_w, \phi_w, \dot{x}_w, \dot{y}_w, \dot{\phi}_w)^T$, and the tangent vector is $(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) = (\dot{x}_w, \dot{y}_w, \dot{\phi}_w, \ddot{x}_w, \ddot{y}_w, \ddot{\phi}_w)^T$. The drift field X_0 is written $(\dot{x}_w, \dot{y}_w, \dot{\phi}_w, 0, g, 0)^T$, where g is the gravitational acceleration (possibly zero).

We begin by checking accessibility for a single control vector field, $n = 1$. To test for small-time accessibility, we examine the Lie algebra of the system's vector fields. If $\mathcal{V}$ is a family of vector fields (corresponding to constant controls) on a manifold M, then the accessibility Lie algebra $L(\mathcal{V})$ is the smallest subalgebra of vector fields on M containing $\mathcal{V}$. (For a finite family $\mathcal{V}$, defining $B_0(\mathcal{V}) = \mathcal{V}$ and $B_{k+1}(\mathcal{V}) = B_k(\mathcal{V}) \cup \{[V_i, V_j] \text{ for all } V_i, V_j \in B_k(\mathcal{V})\}$, where the Lie bracket $[V_i, V_j]$ is given at each $\mathbf{p} \in M$ as

$$[V_i, V_j](\mathbf{p}) = \frac{\partial V_j(\mathbf{p})}{\partial \mathbf{p}} V_i(\mathbf{p}) - \frac{\partial V_i(\mathbf{p})}{\partial \mathbf{p}} V_j(\mathbf{p}),$$

recall that the Lie algebra $L(\mathcal{V})$ is spanned by elements of $B_\infty(\mathcal{V})$.) The tangent vectors of $L(\mathcal{V})$ at $\mathbf{p}$ are $L(\mathcal{V})(\mathbf{p})$. Then the system satisfies the *Lie algebra rank condition* at $\mathbf{p}$, and therefore is small-time accessible from $\mathbf{p}$, if $L(\mathcal{V})(\mathbf{p})$ is the tangent space $T_{\mathbf{p}}M$ (Hermann and Krener 1977; Sussmann 1987). Note that $\mathcal{V}$ need not be symmetric for small-time accessibility; in particular, if V is an element of $\mathcal{V}$, it is not necessary that $-V$ also belong to $\mathcal{V}$. Such symmetries are often required for small-time local controllability.

For the case of $n = 1$, we study the Lie algebra of the vector fields X_0 and X_1. Without loss of generality, assume the unit-control force $\hat{\mathbf{f}}_1$ is $(0, f_y, \tau)^T$ in the object frame $\Sigma_{\mathcal{O}}$, so the control vector field X_1 is written $(0, 0, 0, -f_y \sin\phi_w, f_y \cos\phi_w, \tau)^T$. We define the Lie bracket vector fields:

$$\begin{aligned} X_2 &= [X_0, X_1] \ , \\ X_3 &= [X_1, [X_0, X_1]\ ] \ , \\ X_4 &= [X_1, [X_0, [X_0, X_1]\ ]\ ] \ , \\ X_5 &= [X_1, [X_1, [X_0, [X_0, X_1]\ ]\ ]\ ] \ , \\ X_6 &= [X_0, [X_1, [X_1, [X_0, [X_0, X_1]\ ]\ ]\ ]\ ] \ . \end{aligned}$$

We find that

$$\det(X_1\ X_2\ X_3\ X_4\ X_5\ X_6) = -16 f_y^4 \tau^8,$$

indicating that these six vector fields span the tangent space $T(\mathbf{q}, \dot{\mathbf{q}})\mathcal{TC}$ at any state $(\mathbf{q}, \dot{\mathbf{q}})$, provided $f_y \neq 0$ (the control must not be a pure torque) and $\tau \neq 0$ (the control must not be a pure force through the center of mass). Note that a pure torque is not possible by frictional contact with the perimeter of a bounded object.

The tangent vectors $X_1(\mathbf{q}, \dot{\mathbf{q}})$, $X_3(\mathbf{q}, \dot{\mathbf{q}})$, and $X_5(\mathbf{q}, \dot{\mathbf{q}})$ span the acceleration space at $(\mathbf{q}, \dot{\mathbf{q}})$, and $X_2(\mathbf{q}, \dot{\mathbf{q}})$, $X_4(\mathbf{q}, \dot{\mathbf{q}})$, and $X_6(\mathbf{q}, \dot{\mathbf{q}})$ span the velocity space.

PROPOSITION 1. The planar object $\mathcal{O}$ is small-time accessible for the control set U_d, with or without gravity, if and only if $\hat{\mathcal{F}}$ contains a force direction that is neither a pure force nor a pure torque.

Proposition 1 implies that all planar objects are small-time accessible by point contact, except for a frictionless disk centered at its center of mass. For such an object, all contacts with its perimeter produce zero torque.

Clearly, $n = 1$ is never sufficient for controllability; the angular velocity of the object can only change in one direction. It can be shown that $n = 2$ is sufficient for controllability, provided the signs of $\hat{\tau}_1$ and $\hat{\tau}_2$ are opposite.

PROPOSITION 2. The planar object $\mathcal{O}$ is controllable, in zero gravity (for the control set U_d) or nonzero gravity (for the control set U_u), if and only if $\hat{\mathcal{F}}$ contains force directions $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$ such that $\hat{\tau}_1 > 0$, $\hat{\tau}_2 < 0$, and at least one force direction $\hat{\mathbf{f}}_i \in \hat{\mathcal{F}}$ has a nonzero linear component.

Proof. The conditions of the proposition are clearly necessary. The conditions are shown to be sufficient in the appendix. □

Proposition 2 implies that any object is controllable by point contact with its perimeter Γ, except for a frictionless disk centered at its center of mass. In fact, if friction is nonzero, Proposition 2 implies that the object is controllable from a single point of contact (Fig. 3).

THEOREM 1. For any planar object $\mathcal{O}$ with a closed, piecewise-smooth curve Γ of available contact points, there exists a pushing contact point on Γ such that the object is controllable in zero gravity (for the control set U_d) or nonzero gravity (for the control set U_u), provided the friction coefficient at the contact is nonzero.

Proof. The radius function $r : \Gamma \rightarrow \mathbf{R}$ measures the distance from the center of mass to points on the object's perimeter

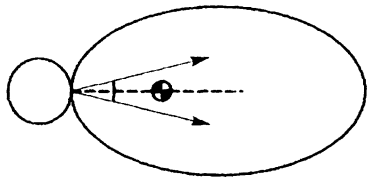

Fig. 3. The object is controllable by pushing at the contact point shown for any nonzero friction coefficient. An example friction cone is illustrated. Both positive and negative torques can be applied through the contact, making the object controllable by Proposition 2.

Γ. Assume the curve Γ is parameterized by s. At each point $\Gamma(s)$ where $dr(\Gamma(s))/ds = 0$, the contact normal of Γ passes through the center of mass. (If $dr(\Gamma(s))/ds$ is discontinuous at s, such as at a vertex of a polygon, the contact normal can be chosen as any value in the range defined by the normals as we approach s from both directions.) There are at least two such points, because Γ is closed and $r(\Gamma)$ attains at least one local maximum and one local minimum. If Γ is not a single point, there is at least one point $\Gamma(s)$ at which $dr(\Gamma(s))/ds = 0$ and $r(\Gamma(s)) \neq 0$ for any center-of-mass location. If $r(\Gamma(s)) \neq 0$, $dr(\Gamma(s))/ds = 0$, and the friction coefficient is nonzero, then the center of mass is in the interior of the friction cone and positive and negative torques can be applied from $\Gamma(s)$. Applying Proposition 2, the proof is complete. □

We now consider conditions for small-time local controllability at a state $(\mathbf{q}, \mathbf{0})$. The object $\mathcal{O}$ is small-time locally controllable at $(\mathbf{q}, \mathbf{0})$ if, given any neighborhood W of $(\mathbf{q}, \mathbf{0})$, $(\mathbf{q}, \mathbf{0})$ is interior to the set of states reachable by trajectories remaining in W. One consequence of this property is that the object can follow any path in its configuration space arbitrarily closely by staying sufficiently close to zero velocity.

PROPOSITION 3. In the presence of gravity or other disturbance forces, the planar object $\mathcal{O}$ is small-time locally controllable at all states $(\mathbf{q}, \mathbf{0})$ for the control set U_u if and only if the set of force directions $\hat{\mathcal{F}}$ positively spans the force sphere (i.e., there is no closed hemisphere of the force sphere containing all force directions in $\hat{\mathcal{F}}$). A minimum of four force directions ($n \geq 4$) is necessary.

The proof of Proposition 3 is straightforward and can be found in the work of Lynch (1996). This is the familiar condition for a force-closure grasp of a planar object; the difference is that for a grasp, all contacts are simultaneously active. We can directly apply various theorems regarding the existence of positive grasps.

THEOREM 2. Any planar object $\mathcal{O}$ with a closed, piecewise-smooth curve Γ of available contact points is small-time locally controllable at all states $(\mathbf{q}, \mathbf{0})$ for the control set U_u, with or without gravity or other disturbance forces, unless the contact is frictionless and Γ is a circle.

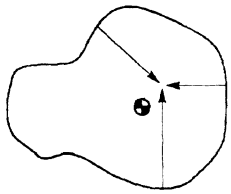

Fig. 4. Without gravity, the object is small-time locally controllable at any state $(\mathbf{q}, \mathbf{0})$ by the three unilateral forces shown.

Proof. See the works of Mishra, Schwartz, and Sharir (1987) and Markenscoff, Ni, and Papadimitriou (1990). □

Tighter sufficient conditions on the set of force directions $\hat{\mathcal{F}}$ can be found for small-time local controllability in zero gravity (Proposition 5) and for subsets of the zero-velocity space in the presence of gravity (Proposition 4).

PROPOSITION 4. The planar object $\mathcal{O}$ is small-time locally controllable at a state $(\mathbf{q}, \mathbf{0})$ for the control set U_u if the negated gravitational force direction, expressed in the object's frame $\Sigma_{\mathcal{O}}$, is in the interior of $CH_{S^2}(\hat{\mathcal{F}})$, the convex hull of the force directions $\hat{\mathcal{F}}$ on the force sphere.

If the condition of Proposition 4 is satisfied, $n \geq 3$ and the object is small-time locally controllable on the simply connected three-dimensional subset of its configuration space $\{\mathbf{q} \in \mathcal{C} \mid \phi_{min} < \phi_w < \phi_{max}\}$, for a suitably defined world frame $\Sigma_{\mathcal{W}}$. This is just the angle range for stable equilibrium if all contacts were acting simultaneously.

Proposition 5 addresses small-time local controllability in zero gravity, and is relevant to controlling the position and attitude of a free-flying planar robot with gas jets, a hovercraft with a single rotating thruster (Manikonda and Krishnaprasad 1997), or an unactuated joint of an underactuated manipulator (Arai 1996; Lynch et al. 1998).

PROPOSITION 5. In the absence of gravity, with the control set U_d, the object $\mathcal{O}$ is small-time locally controllable at any state $(\mathbf{q}, \mathbf{0})$ if the set of force directions $\hat{\mathcal{F}}$ positively spans a great circle of the force sphere that does not lie in the $\tau = 0$ plane.

REMARK. The condition of Proposition 5 is satisfied by any set of three or more force lines that intersect at a single point, provided the force lines positively span the plane and the intersection point is not at the object's center of mass (Fig. 4).

Proof. Consider the system (Lewis and Murray 1997):

$$(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) = X_0(\mathbf{q}, \dot{\mathbf{q}}) + u_1 X_1(\mathbf{q}, \dot{\mathbf{q}}) + u_2 X_2(\mathbf{q}, \dot{\mathbf{q}}),$$

where $u_1, u_2 \in [-1, 1]$, and the bracket terms

$$\begin{aligned} X_3 &= [X_0, X_1] \ , \\ X_4 &= [X_0, X_2] \ , \\ X_5 &= [X_1, [X_0, X_2] \ ] \ , \\ X_6 &= [X_0, [X_1, [X_0, X_2] \ ] \ ] \ . \end{aligned}$$

Now we give some definitions that are necessary to apply Sussmann's (1987) sufficient condition for small-time local controllability. For a bracket term B, we define $\delta_i(B)$ as the number of times X_i appears in B, and the degree of B is $\sum_{i=0}^{n} \delta_i(B)$. We call B a "bad" bracket if $\delta_0(B)$ is odd and $\delta_i(B)$ is even for all $i \in \{1, \ldots, n\}$; we call B a "good" bracket otherwise. A "bad" bracket B is "neutralized" at a state $\mathbf{p}$ if B, evaluated at $\mathbf{p}$, is the linear combination of "good" brackets of lower degree evaluated at $\mathbf{p}$. Sussmann proved that if the system satisfies the Lie algebra rank condition at $\mathbf{p}$, and all "bad" brackets evaluated at $\mathbf{p}$ are neutralized, then the system is small-time locally controllable at $\mathbf{p}$.

Consider the control vector fields $X_1 = (0, 0, 0, \cos\phi_w, \sin\phi_w, 0)^T$ and $X_2 = (0, 0, 0, 0, 0, 1)^T$. The force $\mathbf{f}_1$ acts through the center of mass along the x-axis of the object frame $\Sigma_{\mathcal{O}}$ and $\mathbf{f}_2$ is a pure torque. Calculating the brackets defined above, we find that $\det(X_1 \ X_2 \ X_3 \ X_4 \ X_5 \ X_6) = 1$; the Lie algebra rank condition is satisfied. Because we only use brackets up to degree four, the only "bad" brackets to be neutralized are the drift field (which vanishes at $\dot{\mathbf{q}} = \mathbf{0}$) and the "bad" brackets of degree three: $[X_1, [X_0, X_1]]$ and $[X_2, [X_0, X_2]]$. In this case,

$$[X_1, [X_0, X_1]] = [X_2, [X_0, X_2]] = (0, 0, 0, 0, 0, 0)^T ,$$

and the system is small-time locally controllable at states $(\mathbf{q}, \mathbf{0})$.

The forces $\mathbf{f}_1$ and $\mathbf{f}_2$ define a plane orthogonal to the $\tau = 0$ plane in the three-dimensional body-fixed force space, and the controls $u_1, u_2 \in [-1, 1]$ define a compact, convex subset of this plane containing the origin in the interior (relative to the plane). By Sussmann's (1987) Proposition 2.3, small-time local controllability for this system implies small-time local controllability for the bang-bang system with controls $u_1, u_2 \in \{-1, 1\}$. Scaling, small-time local controllability holds for any compact, convex set of control forces that contains a neighborhood of the origin in this plane, and Sussmann's proposition indicates that the extremal forces alone are sufficient. Therefore, any set of control forces that *positively* spans a plane orthogonal to the $\tau = 0$ plane (equivalently, any set of force directions that positively spans a great circle of the force sphere orthogonal to the $\tau = 0$ plane) also yields small-time local controllability.

Now consider a force $\mathbf{f}_2$ in the y-direction of the object frame $\Sigma_{\mathcal{O}}$ with some torque about the center of mass, and its corresponding vector field $X_2 = (0, 0, 0, -\sin\phi_w, \cos\phi_w, \tau)^T$. Then X_i, $i = 1, \ldots, 6$, satisfy the Lie algebra rank condition, provided τ is not zero. The "bad" brackets

$$\begin{aligned} [X_1, [X_0, X_1]] &= (0, 0, 0, 0, 0, 0)^T , \\ [X_2, [X_0, X_2]] &= (0, 0, 0, -2\tau\cos\phi_w, -2\tau\sin\phi_w, 0)^T , \end{aligned}$$

are clearly neutralized (the latter being a multiple of X_1), and the system is small-time locally controllable. The two forces $\mathbf{f}_1$ and $\mathbf{f}_2$ can span any plane that is neither the $\tau = 0$ plane nor orthogonal to the $\tau = 0$ plane. As above, any set of forces that positively spans the same plane also yields small-time local controllability.

Taking the two cases together, we see that small-time local controllability holds, provided the set of force directions $\hat{\mathcal{F}}$ positively spans a great circle of the force sphere that does not lie in the $\tau = 0$ plane. □

If the object is a frictionless disk not centered at its center of mass, contact with the perimeter Γ defines a great circle of force directions that does not lie in the $\tau = 0$ plane. Proposition 5 therefore allows us to strengthen Theorem 2 for the case of zero gravity.

THEOREM 3. In the absence of gravity, with the control set U_d, any planar object $\mathcal{O}$ with a closed, piecewise-smooth curve Γ of available contact points is small-time locally controllable at all states $(\mathbf{q}, \mathbf{0})$, unless the contact is frictionless and Γ is a circle centered at the object's center of mass.

Proposition 5 implies that three force directions are sufficient for small-time local controllability in the absence of gravity.[1] In contrast, a force-closure grasp requires at least four unilateral force directions. Furthermore, we cannot simply grasp and rotate a frictionless disk, but if the disk is not centered at its center of mass, then Theorem 3 indicates that the position *and* orientation of the object are locally controllable by point-contact pushing in zero gravity. In this sense, dynamic pushing is a more complete primitive for planar manipulation.

Theorem 3 implies that a two-degrees-of-freedom point robot, which can translate freely in the plane, can dynamically push an object to closely follow any path in its three-dimensional configuration space. In general, the robot will have to move between distant contacts on the object to make it follow the desired path. Because transit between distant contacts takes time, the object will no longer be *small-time* locally controllable, but it can still be pushed to closely follow any path. In other words, there exists a time $T > 0$ such that $R^W(\mathbf{q}, 0, \leq T)$ contains a neighborhood of $(\mathbf{q}, 0)$ for all $(\mathbf{q}, 0)$ and any neighborhood W of $(\mathbf{q}, 0)$.

1. Recent work by Lewis (1997) and Manikonda and Krishnaprasad (1997) indicates that a single bidirectional force (equivalently, two opposing unilateral forces) is insufficient, meaning that three unilateral force directions are also necessary.

4.2. With Manipulator Constraints

The results of the previous section address the theoretical capabilities of dynamic nonprehensile manipulation from the viewpoint of the object alone. Just as important, however, are properties of the manipulator that is controlling the object. While an object may be controllable by point contact, the manipulator may not be able to achieve the contacts and motions necessary to bring the object to the desired state.

We would like to understand the global reachability properties of a manipulator/object system—given the kinematic and dynamic specifications of a manipulator $\mathcal{M}$, the mass, inertia, and shape of an object $\mathcal{O}$, the friction between them, and their initial state, where can the manipulator take the object? Unfortunately, such a question appears to be very difficult to answer, so we are forced to look locally. Because the contact is nonprehensile, in any neighborhood of a manipulator/object state, the best we can hope for is small-time accessibility in the absence of gravity.

Proposition 1 says that a single control force, fixed in the object frame, is sufficient for small-time accessibility. This indicates that a one-degree-of-freedom manipulator may be sufficient for small-time accessibility. This would imply that even the simplest of robots is capable of performing interesting dynamic manipulation in the plane—the object's reachable state space is a full-dimensional subset of its six-dimensional state space.

We examine this possibility for a single prismatic joint and a single revolute joint.[2] Note that Proposition 1 cannot be directly applied to these systems, as the control forces are not fixed in the object frame.

4.2.1. Example 1: A Single Prismatic Joint

Consider the system of Figure 5. The manipulator $\mathcal{M}$ is a single prismatic joint, and the object $\mathcal{O}$ is a unit-mass rod in point contact with $\mathcal{M}$ at an angle $\phi_w \in (0, \pi)$. The distance from the contact to the rod's center of mass is r, and the rod's radius of gyration is $\rho \in (0, \infty)$. The rod represents an arbitrary polygon in vertex contact with the manipulator.

The configuration of the system is $\mathbf{q}' = (\mathbf{q}, y_m) \in \mathcal{C}'$, where y_m is the position of $\mathcal{M}$ in the world frame $\Sigma_{\mathcal{W}}$. The state space of the system is the eight-dimensional manifold $T\mathcal{C}' = \mathbf{R}^3 \times SO(2) \times \mathbf{R}^4$. We assume that the manipulator stays in contact with the rod endpoint at all times, and it may apply zero force (simply "following" the rod) or a nonzero force. The three-dimensional submanifold of contact configurations is $\{\mathbf{q}' \in \mathcal{C}' | F(\mathbf{q}') = y_w - y_m - r \sin \phi_w = 0\}$.

The conditions that $\mathcal{O}$ remain in contact with $\mathcal{M}$ are:

$$\frac{dF(\mathbf{q}'(t))}{dt} = 0,$$

2. Sign errors were propagated through the equations in this section in earlier work (Lynch 1996).

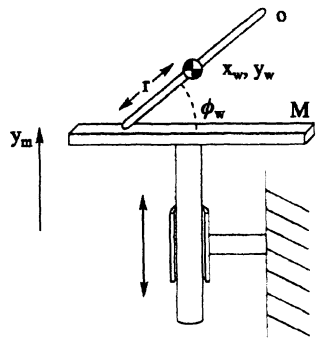

Fig. 5. A one-degree-of-freedom prismatic robot manipulating a rod. The rod represents an arbitrary polygon in vertex contact.

$$\frac{d^2F(\mathbf{q}'(t))}{dt^2} = 0.$$

Erdmann (1984, 1994) refers to these constraints as the *first* and *second variation constraints*, respectively. These constraints state that the velocity and acceleration of the system normal to the constraint surface must be zero. For the system of Figure 5, these constraints are written:

$$\frac{dF(\mathbf{q}'(t))}{dt} = \dot{y}_w - \dot{y}_m - r\dot{\phi}_w \cos\phi_w = 0,$$

$$\frac{d^2F(\mathbf{q}'(t))}{dt^2} = \ddot{y}_w - \ddot{y}_m + r(\dot{\phi}_w^2 \sin\phi_w - \ddot{\phi}_w \cos\phi_w) = 0. \quad (1)$$

(The state variables' dependence on time is omitted for clarity.) Assuming the contact is frictionless, the acceleration of $\mathcal{O}$ must satisfy the constraints

$$\ddot{x}_w = 0, \quad (2)$$

$$\ddot{y}_w r\cos\phi_w + \rho^2\ddot{\phi}_w = 0. \quad (3)$$

Equation (2) constrains the direction of the contact force, and eq. (3) constrains the force to pass through the contact point. Using Equations (1)–(3), we solve for the acceleration of $\mathcal{O}$ as a function of the manipulator control $\ddot{y}_m$ and the system state $(\mathbf{q}', \dot{\mathbf{q}}')$:

$$\begin{aligned} \ddot{x}_w &= 0, \\ \ddot{y}_w &= K, \\ \ddot{\phi}_w &= \frac{-r}{\rho^2} K \cos\phi_w, \end{aligned}$$

where the contact force K is given by

$$K = \frac{\rho^2(\ddot{y}_m - \dot{\phi}_w^2 r \sin\phi_w)}{\rho^2 + r^2\cos^2\phi_w}.$$

Treating K as the control input, the control system is $(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) = X_0(\mathbf{q}, \dot{\mathbf{q}}) + KX_1(\mathbf{q}, \dot{\mathbf{q}})$, where the drift vector field $X_0(\mathbf{q}, \dot{\mathbf{q}})$ and the control vector field $X_1(\mathbf{q}, \dot{\mathbf{q}})$ (which are the projections of the vector fields on the system state space TC' to vector fields on the object's state space TC) are as follows:

$$\begin{aligned} X_0(\mathbf{q}, \dot{\mathbf{q}}) &= (\dot{x}_w, \dot{y}_w, \dot{\phi}_w, 0, 0, 0)^T, \\ X_1(\mathbf{q}, \dot{\mathbf{q}}) &= (0, 0, 0, 0, 1, -\frac{r}{\rho^2}\cos\phi_w)^T. \end{aligned}$$

The drift field is written without a gravity term, but one can be included without changing the results here.

It is clear that $\mathcal{O}$ is not accessible as eq. (2) integrates to yield the velocity constraint $\dot{x}_w = c_1$ and the position constraint $x_w = c_1 t + c_2$. The situation is (locally) similar with slipping contact and nonzero friction, except the applied force is constrained to act along a friction-cone edge instead of the contact normal.

We might ask instead if the rod is small-time accessible on its reduced state space $(y_w, \phi_w, \dot{y}_w, \dot{\phi}_w)$. Constructing the vector fields $X_2 = [X_0, X_1]$, $X_3 = [X_1, [X_0, X_1]]$, and $X_4 = [X_1, [X_0, [X_0, X_1]]]$, and projecting to the reduced tangent space, we see that these vector fields span if $-4r^4 \cos^2\phi_w \sin^2\phi_w/\rho^8 \neq 0$. This determinant is zero if $r = 0$ (the contact is coincident with the rod's center of mass) or $\cos\phi_w = 0$ or $\sin\phi_w = 0$. Taking higher-order brackets shows that the rod is small-time accessible on its reduced state space (provided $r \neq 0$) unless $\dot{\phi}_w = 0$ and $\cos\phi_w = 0$ (the rod is perpendicular to the robot's surface). In this case, the rod cannot be rotated.

If there is nonzero friction at the contact between the rod and the manipulator, and the contact is not initially slipping, then in general the rod is small-time accessible on its full state space (Lynch 1996). This derives from the fact that the linear direction of the force applied to the rod, within the friction cone, is a function of the manipulator's acceleration (and the state of the system). This gives control over the linear direction of the applied force, not just the magnitude as in the slipping case. Of course, the rod must undergo both sticking and slipping contact phases for it to be small-time accessible; otherwise, the configuration of the rod is confined to a four-dimensional set where the endpoint is pinned to the manipulator.

4.2.2. *Example 2: A Single Revolute Joint*

Now consider the system of Figure 6. The manipulator $\mathcal{M}$ is a single revolute joint, and the object $\mathcal{O}$ is a rod as before. The configuration of the system is $\mathbf{q}' = (\mathbf{q}, \theta) \in \mathcal{C}'$, where θ is the angle of the revolute joint. Assuming the single link is thin, the three-dimensional submanifold of contact configurations is given by $\{\mathbf{q}' \in \mathcal{C}' \mid F(\mathbf{q}') = \cos\theta(y_w - r\sin\phi_w) + \sin\theta(r\cos\phi_w - x_w) = 0\}$. After collecting acceleration, centrifugal, and Coriolis terms, the second variation constraint is written:

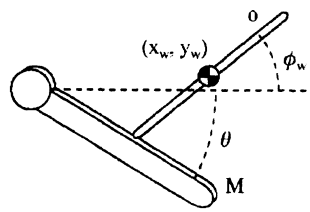

Fig. 6. A one-degree-of-freedom revolute robot manipulating a rod.

$$\begin{aligned} \frac{d^2F(\mathbf{q}'(t))}{dt^2} = 0 = \\ &\ddot{x}_w(-\sin\theta) + \ddot{y}_w(\cos\theta) + \ddot{\phi}_w(-r\cos(\phi_w - \theta)) + \\ &\ddot{\theta}(r\cos(\phi_w - \theta) - x_w\cos\theta - y_w\sin\theta) + \\ &\dot{\phi}_w^2(r\sin(\phi_w - \theta)) + \\ &\dot{\theta}^2(r\sin(\phi_w - \theta) + x_w\sin\theta - y_w\cos\theta) + \\ &\dot{\phi}_w\dot{\theta}(-2r\sin(\phi_w - \theta)) + \\ &\dot{\theta}\dot{x}_w(-2\cos\theta) + \dot{\theta}\dot{y}_w(-2\sin\theta). \end{aligned} \tag{4}$$

Assuming the contact is frictionless, we have constraints on the linear direction (eq. (5)) and the point of application (eq. (6)) of the contact force:

$$\ddot{x}_w\cos\theta + \ddot{y}_w\sin\theta = 0, \tag{5}$$

$$\ddot{y}_w r\cos(\phi_w - \theta) + \ddot{\phi}_w\rho^2\cos\theta = 0. \tag{6}$$

Solving eqs. (4)–(6) for the acceleration of the object in terms of the control $\ddot{\theta}$ and the system state $(\mathbf{q}', \dot{\mathbf{q}}')$, we get

$$\ddot{x}_w = -K\sin\theta, \tag{7}$$

$$\ddot{y}_w = K\cos\theta, \tag{8}$$

$$\ddot{\phi}_w = -\frac{r}{\rho^2}K\cos(\phi_w - \theta), \tag{9}$$

where the contact force K is

$$\begin{aligned} K = &\rho^2(\ddot{\theta}(-r\cos(\phi_w - \theta) + x_w\cos\theta + y_w\sin\theta) + \\ &\dot{\phi}_w^2(-r\sin(\phi_w - \theta)) + \\ &\dot{\theta}^2(-r\sin(\phi_w - \theta) - x_w\sin\theta + y_w\cos\theta) + \\ &\dot{\phi}_w\dot{\theta}(2r\sin(\phi_w - \theta)) + \\ &\dot{\theta}\dot{x}_w(2\cos\theta) + \\ &\dot{\theta}\dot{y}_w(2\sin\theta))/(r^2\cos^2(\phi_w - \theta) + \rho^2). \end{aligned} \tag{10}$$

Equations (7)–(9) have the structure we expect—the force is normal to the manipulator link, and the torque about the center of mass of $\mathcal{O}$ depends on the angle between the object and

the manipulator. The equations also show that even for this simplest of systems, the force K is a complex function of the state and control $\ddot{\theta}$.

Rearranging eq. (10), we get the form $\ddot{\theta} = \ddot{\theta}_{drift} + K\ddot{\theta}_{control}$, where $\ddot{\theta}_{drift}$ is the acceleration of the robot needed to stay in contact with the rod while applying zero force, and $\ddot{\theta}_{control}$ is the additional acceleration of the robot required to apply a unit force to the rod:

$$\begin{aligned}
\ddot{\theta}_{drift} &= (\dot{\phi}_w^2(-r\sin(\phi_w-\theta)) + \\
&\quad \dot{\theta}^2(-r\sin(\phi_w-\theta) - x_w\sin\theta + y_w\cos\theta) + \\
&\quad \dot{\phi}_w\dot{\theta}(2r\sin(\phi_w-\theta)) + \dot{\theta}\dot{x}_w(2\cos\theta) + \\
&\quad \dot{\theta}\dot{y}_w(2\sin\theta))/(r\cos(\phi_w-\theta) - x_w\cos\theta - \\
&\quad y_w\sin\theta), \\
\ddot{\theta}_{control} &= \frac{\rho^2 + r^2\cos^2(\phi_w-\theta)}{\rho^2(-r\cos(\phi_w-\theta) + x_w\cos\theta + y_w\sin\theta)}.
\end{aligned}$$

The denominator of these expressions is zero when the contact point of the rod is at the robot's pivot point. At this singularity, the robot cannot apply a force ($\ddot{\theta}_{control}$) or maintain contact with a moving rod ($\ddot{\theta}_{drift}$).

Treating the force K as the control input, the control system on the state space TC' is $(\dot{\mathbf{q}}', \ddot{\mathbf{q}}') = X_0(\mathbf{q}', \dot{\mathbf{q}}') + KX_1(\mathbf{q}', \dot{\mathbf{q}}')$, where the drift and control vector fields are written

$$\begin{aligned}
X_0(\mathbf{q}', \dot{\mathbf{q}}') &= (\dot{x}_w, \dot{y}_w, \dot{\phi}_w, \dot{\theta}, 0, 0, 0, \ddot{\theta}_{drift})^T, \\
X_1(\mathbf{q}', \dot{\mathbf{q}}') &= (0, 0, 0, 0, -\sin\theta, \cos\theta, -\frac{r}{\rho^2}\cos(\phi_w-\theta), \\
&\quad \ddot{\theta}_{control})^T.
\end{aligned}$$

Because the acceleration of the object is a function of the manipulator state $(\theta, \dot{\theta})$, and not just its acceleration $\ddot{\theta}$, we cannot immediately project the vector fields to vector fields on the object's state space TC as we did with the prismatic joint. We must look at the vector fields X_0 and X_1, and their Lie brackets, on the full state manifold TC'. After we have constructed the Lie brackets, we can look at their projection to TC to determine the accessibility of the object $\mathcal{O}$.

As in the proof of Proposition 1, we construct the vector fields

$$\begin{aligned}
X_2 &= [X_0, X_1], \\
X_3 &= [X_1, [X_0, X_1]], \\
X_4 &= [X_1, [X_0, [X_0, X_1]]], \\
X_5 &= [X_1, [X_1, [X_0, [X_0, X_1]]]], \\
X_6 &= [X_0, [X_1, [X_1, [X_0, [X_0, X_1]]]]].
\end{aligned}$$

These vector fields are highly complex trigonometric functions. To completely answer the question of small-time accessibility, we must consider even higher-order bracket terms. Owing to the complexity of the bracket terms, we focus on the particular case of zero-velocity states $(\mathbf{q}', \mathbf{0})$.

Evaluated at zero velocity and projected to TC, the vector fields $X_1 \ldots X_6$ take the form $(\mathbf{0}, \mathbf{a})^T$, $(-\mathbf{a}, \mathbf{0})^T$, $(\mathbf{0}, \mathbf{b})^T$, $(-\mathbf{b}, \mathbf{0})^T$, $(\mathbf{0}, \mathbf{c})^T$, $(-\mathbf{c}, \mathbf{0})^T$, where $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$ are three-vectors and $\mathbf{0}$ is the zero three-vector. Because of this form, it is sufficient to look at $\det(\mathbf{a}\ \mathbf{b}\ \mathbf{c})$ to determine accessibility. The determinant is $4r^2h$, where h is a complex trigonometric function of the configuration of the system. Thus, from a zero-velocity state, these vector fields span unless $r = 0$ (the contact is coincident with the rod's center of mass) or $h = 0$. Because $h = 0$ defines a lower-dimensional manifold of the configuration space, the rod is small-time accessible from a generic configuration.

For the particular case $\phi_w = \pi/2$ and $\theta = 0$ (the rod is perpendicular to the robot), we have $h = -2/\rho^4 x_w^2$, which is undefined at $x_w = 0$. This corresponds to a contact point at the pivot of the robot, so the robot cannot apply a force and the rod is not small-time accessible. Otherwise, we have small-time accessibility provided $r \neq 0$.

We conjecture that by taking higher-order Lie brackets, a rod with $r \neq 0$ can be shown to be small-time accessible at all states such that the contact point does not coincide with the pivot of the manipulator. Unlike the case of a frictionless prismatic joint, we have accessibility with a frictionless revolute joint because the force angle varies with the joint angle θ, giving the robot some control over the direction of the applied force.

4.2.3. *Discussion*

The example above shows that even a one-degree-of-freedom revolute robot can take an object to a six-dimensional subset of its state space by using slipping and rolling between the robot and the object. We should therefore be able to do interesting planar dynamic manipulation with even the simplest of robots. The four state-equality constraints of the robot link (pivot point is fixed) do not translate to state-equality constraints for the object.

The examples above only consider polygonal edge-vertex contacts. Smooth objects and manipulators can be handled by modifying the constraints. Smooth surfaces may result in other hindrances to small-time accessibility; for example, a frictionless disk centered at its center of mass can never be small-time accessible by any type of contact.

The system is *dynamically singular* at a system state $(\mathbf{q}', \dot{\mathbf{q}}')$ if the set of motion directions of $\mathcal{O}$ loses rank on $T_{(\mathbf{q},\dot{\mathbf{q}})}TC$. (The object may still be accessible if the system can break the dynamic singularity at some time T.) Figure 7 gives examples of dynamically singular systems.

In short, the object is not small-time accessible if (1) the linear force direction is fixed in the world frame (as with a frictionless prismatic joint); (2) the force always passes

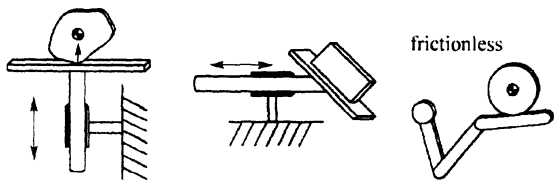

Fig. 7. Dynamically singular nonprehensile systems. (a) The system is initially at rest. At this configuration, any acceleration of the prismatic joint either results in zero force applied to the object, or a force through the center of mass. This dynamic singularity is unstable. (b) If the object is initially at rest in line contact with the manipulator, no acceleration of the manipulator can cause the object to rotate. Even if the robot has two or more prismatic joints, the system may be singular with line contact. (c) If the disk is frictionless and centered at its center of mass, the system is dynamically singular at every state. The angular velocity of the disk cannot be changed.

through the object's center of mass (as with a frictionless disk, or if the center of mass of the object is coincident with the contact point); or (3) the robot cannot apply a force (as when the contact point coincides with the robot pivot).

Although we have shown that the accessible state space may be full-dimensional, it appears extremely difficult to calculate the shape of the accessible state space from a given state. Ideally, we would have a representation similar to a robot's kinematic workspace.

4.3. *Algorithmic Control*

The differential geometric approach to reachability yields little insight into how to generate controls to reach a goal state. Here we describe an intuitive approach to controlling the degrees of freedom of the object using *algorithmic* control, which is the basis of the planner described in Section 5.

The idea behind algorithmic control is to control a subset of the state variables at any given time, but to switch between subsets so the goal state is reached. For example, a unicycle can be driven to the zero state by first reorienting it so that it points toward the origin, then driving it to the origin, and finally reorienting it to the zero angle. In general, we must account for drift in the variables that are not being directly controlled.

In the context of dynamic nonprehensile manipulation, we can define the following control phases (Fig. 8):

1. *Dynamic grasp.* An object is in a dynamic grasp if it makes line contact with the manipulator and the manipulator accelerates such that the object remains fixed against it. With a dynamic grasp, up to $\min(2n, 6)$ of the object's state variables can be directly controlled by an n-joint manipulator.

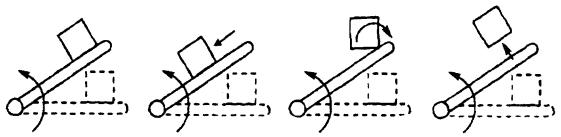

Fig. 8. Manipulation phases (left to right): dynamic grasp, slip, roll, and free flight.

2. *Slip.* Controlled slip provides control of two more state variables: the slipping distance and the slipping velocity.
3. *Roll.* Rolling provides control of two more state variables: the rolling angle and angular velocity.
4. *Free flight.* After the object is released, it follows a one-dimensional path through its state space, parameterized by its time of flight.

A control algorithm can sequence these phases. The dimension of the accessible state space is the sum of the independent freedoms of each phase, up to a maximum of six.

Example 1. An n-degree-of-freedom manipulator carries an object with a dynamic grasp, allows it to begin rolling, and then releases it. The dimension of the accessible state space of the object is upper-bounded by $\min(2n+2+1, 6)$. The "controls" are the state of the robot and the roll angle and velocity at release, and the time of flight (assuming the arrival time at the goal state can be chosen freely).

5. Motion Planning

For a given initial state of the object and the manipulator, the planning problem is to find a manipulator trajectory to take the object to the goal state using frictional, gravitational, and dynamic forces. We are especially interested in the following dynamic tasks:

1. *Snatch:* transfer an object initially at rest on a table to rest on the manipulator. The manipulator accelerates into the object, transferring control of the object from the table to the manipulator.
2. *Throw:* throw the object to a desired goal state. The object is carried with a dynamic grasp and released instantaneously (no slipping or rolling) at a state where the free-flight dynamics will take the object to the goal state (possibly a catch).
3. *Roll:* roll a polygonal object on the manipulator from one statically stable edge to another statically stable edge.
4. *Rolling throw:* allow the object to begin rolling before throwing it. By controlling the roll angle and velocity before the release, the dimension of the object's accessible state space is increased by two.

The manipulation phases in these tasks are dynamic grasp, roll, and free flight. Slipping contact is not used. The object is a polygon in point or line contact with the manipulator.

To solve these problems, we cast trajectory planning as a constrained nonlinear optimization problem, where the system's initial state and goal state (or state manifold) are specified as constraints to the optimization. The trajectory is also subject to a set of nonlinear equality and inequality constraints arising from constraints on the manipulator motion and the dynamics governing the object's motion relative to the manipulator. Because dynamic nonprehensile manipulation relies on friction between the object and the manipulator, and friction coefficients are often uncertain and varying, the optimization is usually asked to minimize the required friction coefficient for successful manipulation. Unlike other work on optimizing the time or energy of a robot's motion, we are more concerned with making the manipulation maximally robust to variations in the friction coefficient.

5.1. Problem Specification

Every task is assumed to consist of a sequence of manipulation phases made up of one or more of the following: a (dynamic) grasp phase g, a roll phase r, and a flight phase f. The nonlinear program is written to handle any sequence of manipulation phases S from the simple finite-state machine:

$$S \rightarrow ABC, \quad A \rightarrow g|\varepsilon, \; B \rightarrow r|rg|\varepsilon, \; C \rightarrow f|\varepsilon,$$

where ε is the empty string. With this notation, a throw (as defined above) is denoted gf, a roll is denoted grg, and a rolling throw is denoted grf. A snatch can be either g or rg. We assume that there is no rebound from the impact at the end of a roll (the transition from r to g). The instant the new edge contacts, if the dynamic grasp conditions are met, then the object is assumed to be in a dynamic grasp.

The times of the manipulation phases are t_{g1} for the first dynamic-grasp phase g, t_{roll} for the rolling phase r, t_{g2} for the second g phase, and t_{flight} for the flight phase f. If a phase is omitted, its corresponding time duration is zero. We also define the cumulative times $T_{g1} = t_{g1}, T_{roll} = T_{g1} + t_{roll}, T_{g2} = T_{roll} + t_{g2}$, and $T_{flight} = T_{g2} + t_{flight}$. Finally, we define $T = T_{g2}$, where T is the total time the manipulator is in contact with the object during the manipulation.

The manipulation primitives—the snatch, the roll, the throw, and the rolling throw—can be composed or "glued" together. The glue between primitives is a static equilibrium carry of the object.

In the next three subsections, we describe the three fundamental elements of the nonlinear program: the design variables, the constraints, and the objective function.

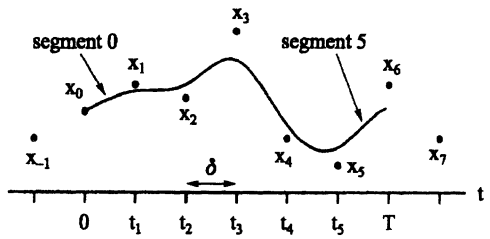

Fig. 9. A cubic B-spline trajectory with nine knot points and six segments.

5.1.1. Design Variables

The design variables consist of the variables $\mathbf{x}$ specifying the trajectory of the manipulator over the interval $[0, T]$; the times of each phase of the manipulation, t_{g1}, t_{roll}, t_{g2}, and t_{flight} (some subset of these will be applicable based on the problem specification); and the required friction coefficient μ between the manipulator and the object. These variables are not independent; the trajectory of the manipulator implicitly defines the time of each phase. However, the problem formulation is much simpler if we make each of these variables explicit and constrain them to be dynamically consistent. Although we cannot control the friction coefficient μ, it is convenient to represent μ as a design variable and enforce the resulting friction constraints.

Many different finite parameterizations of manipulator trajectories have been explored, including polynomials, Fourier bases, summed Fourier and polynomial functions, splines, and piecewise-constant acceleration segments. After some experimentation, we decided to represent trajectories as uniform cubic B-splines (Bartels, Beatty, and Barsky 1987; Chen 1991).

For an n-joint robot, $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2, \ldots \mathbf{x}^n)$, where $\mathbf{x}^i$ is the vector of knot points for the cubic B-spline position history of joint i. The time of each knot point x_j^i is given by t_j, and the knot points are evenly spaced in time. The position of the joint passes "near" the knot points; the actual position at each time is obtained by taking a weighted sum of the four knot points that are closest in time. The weighting basis functions are cubic polynomials of time. Therefore, the position is C^2 and piecewise cubic, the velocity is C^1 and piecewise quadratic, and the acceleration is C^0 and piecewise linear (constant jerk segments); see Figure 9.

5.1.2. Constraints

Constraints in the optimization are determined by limitations on the motion of the manipulator and constraints on the object's motion. Object motion constraints include inequality constraints due to Coulomb friction, and equality constraints due to Newton's laws. To simplify the notation, the dependencies of the constraints on the design variables is omitted.

Manipulator Constraints Constraints on the manipulator's motion include the following:

1. Position constraints: $t \in [0, T]$, and

$$\mathbf{p}(\Theta(t)) \leq \mathbf{0},$$

where Θ is the arm configuration and $\mathbf{p}$ is a vector-valued function representing joint limits and obstacles.

2. Joint velocity constraints: $t \in [0, T]$, and

$$\dot{\Theta}_{min} \leq \dot{\Theta}(t) \leq \dot{\Theta}_{max} .$$

3. Joint torque constraints: $t \in [0, T]$, and

$$\tau_{min} \leq \tau(t) \leq \tau_{max} ,$$

where $\tau(t)$ is the torque to move the arm and the object along the trajectory.

4. Initial state constraints:

$$\Theta(0) = \Theta_0, \ \dot{\Theta}(0) = \dot{\Theta}_0.$$

Object Constraints During a dynamic grasp or rolling phase, contact friction constraints must be enforced to prevent slipping or breaking contact. These force constraints encode the unilateral nature of contact (forces can only be applied into the object) and the finite friction coefficient μ.

To maintain a dynamic grasp, the sum of the negated gravitational vector $-\mathbf{g}$ and the manipulator's acceleration, measured in the object frame, must fall inside an acceleration cone $\mathcal{A}$. The cone $\mathcal{A}$ is determined by the line contact and the friction coefficient μ, and it lives in the three-dimensional space of body-centered accelerations (two linear components and one angular component).[3] The cone is bounded by four edges, $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$, and $\mathbf{a}_4$, numbered so that the interior of the cone lies to the left as we move from $\mathbf{a}_1$ to $\mathbf{a}_2$, and so forth. (These edges correspond to the accelerations of the object from forces through the endpoints of the line contact and on the boundaries of the friction cone; see Figure 10.) The magnitude of these vectors is not important, but for concreteness, assume they are unit vectors. We now form the 3×4 matrix

$$\mathbf{A}_g = (\mathbf{a}_2 \times \mathbf{a}_1 \mid \mathbf{a}_3 \times \mathbf{a}_2 \mid \mathbf{a}_4 \times \mathbf{a}_3 \mid \mathbf{a}_1 \times \mathbf{a}_4),$$

where each column of $\mathbf{A}_g$ is an outward-pointing normal to a face of the acceleration cone $\mathcal{A}$. The acceleration of the manipulator must have a nonpositive dot product with each column vector of $\mathbf{A}_g$.

3. More generally, we can talk about a force cone in the body-fixed force space, but for the particular case of a planar object with a coordinate frame attached to the center of mass, the two representations are equivalent.

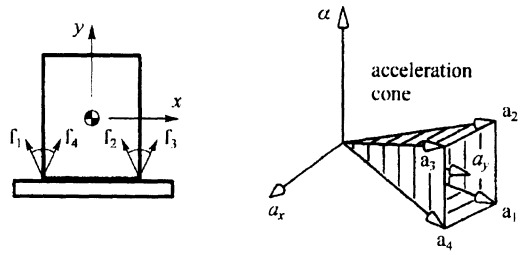

Fig. 10. The contact forces on the left map to the acceleration cone on the right, expressed in the object frame $\sum\mathcal{O}$. To maintain the dynamic grasp, the manipulator acceleration (measured in $\sum\mathcal{O}$) must lie inside the acceleration cone.

We recognize five types of object constraints.

1. Dynamic grasp constraints (dynamic grasp phase):

$$\mathbf{A}_{gi}^T \left(\dot{\mathbf{J}}_{gi}(\Theta(t))\dot{\Theta}(t) + \mathbf{J}_{gi}(\Theta(t))\ddot{\Theta}(t) - \mathbf{g} \right) \leq 0,$$

where $\mathbf{g}$ is the gravitational acceleration, and i is 1 or 2, depending on the current dynamic grasp phase. During a dynamic grasp phase, $\mathbf{J}_{gi}$ and $\dot{\mathbf{J}}_{gi}$, the manipulator Jacobian and its time derivative, are measured at the center of mass of the object, so that all accelerations can be represented in the object frame $\sum_{\mathcal{O}}$.

During rolling contact, a single point of the object is in contact with the manipulator, so $\mathcal{A}$ is a planar acceleration cone (now measured in the world frame $\Sigma_{\mathcal{W}}$), bounded by the two edges $\mathbf{a}_l$ and $\mathbf{a}_r$ (from forces on the left and right edges of the friction cone, respectively). Define the vector $\mathbf{a}_\perp$ normal to the plane of this cone: $\mathbf{a}_\perp = \mathbf{a}_r \times \mathbf{a}_l$. We form the matrix $\mathbf{A}_r = (\mathbf{a}_\perp \times \mathbf{a}_l \mid \mathbf{a}_r \times \mathbf{a}_\perp)$, where each column of $\mathbf{A}_r$ lies in the $(\mathbf{a}_l, \mathbf{a}_r)$ plane and is an outward-pointing normal to an edge of the acceleration cone. The acceleration of the object must have a nonpositive dot product with each column vector of $\mathbf{A}_r$.

The object acceleration $\mathbf{a}_{roll} = (a_{roll,x}, a_{roll,y}, \alpha_{roll})^T$ (measured at its center of mass) required to maintain the rolling contact is determined by assuming a pin joint at the contact and finding the object acceleration consistent with the motion of the manipulator. The acceleration $\mathbf{a} = (a_x, a_y, \alpha)^T$ at the contact point on the robot (including negated gravity) is given by

$$\mathbf{a} = \dot{\mathbf{J}}_r(\Theta(t))\dot{\Theta}(t) + \mathbf{J}_r(\Theta(t))\ddot{\Theta}(t) - \mathbf{g},$$

where $\mathbf{J}_r$ and $\dot{\mathbf{J}}_r$ are measured at the contact point on the manipulator. The constraint that the linear acceleration (a_x, a_y) matches the acceleration of the contact point on the object is expressed

$$(a_x, a_y) = -\omega^2\mathbf{r} + (-r_y\alpha_{roll}, r_x\alpha_{roll}) + (a_{roll,x}, a_{roll,y}),$$

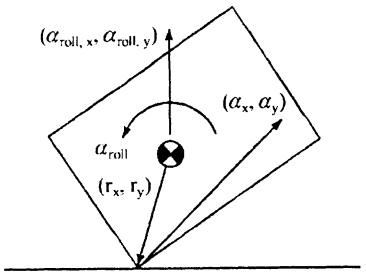

Fig. 11. Rolling notation.

where ω is the angular velocity of the object and $\mathbf{r} = (r_x, r_y)$ is the vector from the center of mass of the object to the contact point in the world frame (Fig. 11). This gives us two constraints on the object acceleration $\mathbf{a}_{roll}$; the third is that $\mathbf{a}_{roll}$ must result from a force through the contact point:

$$\alpha_{roll} = \frac{1}{\rho^2}(r_x a_{roll,y} - r_y a_{roll,x}),$$

where ρ is the object's radius of gyration. After a little manipulation, we get

$$a_{roll,x} = \frac{a_x(\rho^2 + r_x^2) + a_y r_x r_y + r_x \omega^2(\rho^2 + r_x^2 + r_y^2)}{\rho^2 + r_x^2 + r_y^2},$$

$$a_{roll,y} = \frac{a_y(\rho^2 + r_y^2) + a_x r_x r_y + r_y \omega^2(\rho^2 + r_x^2 + r_y^2)}{\rho^2 + r_x^2 + r_y^2},$$

$$\alpha_{roll} = \frac{r_x a_{roll,y} - r_y a_{roll,x}}{\rho^2}.$$

Now we can write the rolling friction constraints:

2. Rolling-friction constraints (rolling phase only):

$$\mathbf{A}_r^T \mathbf{a}_{roll} \leq 0.$$

3. Roll-angle constraints (rolling phase only):

$$\psi_{min} \leq \psi(t) \leq \psi_{max}$$

where ψ is the angle of the object relative to the manipulator. This constraint prevents the object from penetrating the manipulator during the rolling phase.

4. Roll-completed constraint (for rolls only):

$$\psi(T_{roll}) = \psi_{goal}.$$

5. Release-state constraints (for throws only):

$$s(\mathbf{q}, \dot{\mathbf{q}}, t_{flight}) = 0.$$

These constraints specify that the object reaches the goal submanifold by free flight, where $(\mathbf{q}, \dot{\mathbf{q}})$ is the release state and t_{flight} is the time of flight. The goal submanifold is usually specified by goal values of some subset of the state variables.

In principle, the manipulator constraints (1)–(3) and object constraints (1)–(3) should be satisfied at all times during their domain of applicability. In practice, the constraints are only enforced at p uniformly sampled points during each manipulation phase. In the examples here, p is chosen between 20 and 50.

5.1.3. Objective Function

In most of our problems, we minimize the required friction coefficient μ between the object and manipulator to make the manipulation maximally robust to variations in friction. (An exception is the roll of Section 6.3, which minimizes the impact velocity at the end of the roll.) If friction in the actual system is high, we can fix the value of the friction coefficient in the optimization and instead minimize functions such as energy or time.

5.2. Sequential Quadratic Programming

Sequential quadratic programming (SQP) is used to solve the nonlinear program. A generalization of Newton's method for unconstrained optimization, SQP finds a step away from the current iterate by minimizing a quadratic model of the problem. At each iteration, SQP determines the direction to step by solving a quadratic subprogram, where the objective function is a quadratic approximation at the current point and nonlinear constraints are linearized.

The constrained nonlinear program can be written

$$\min f(x)$$

$$\begin{aligned} \text{subject to } c_i(x) &\leq 0, \quad i \in \mathcal{I}, \\ c_i(x) &= 0, \quad i \in \mathcal{E}, \end{aligned}$$

where $x \in \mathbf{R}^m$ is the iterate, f is the objective function, each c_i is a constraint mapping $\mathbf{R}^m$ to $\mathbf{R}$, and $\mathcal{I}$ and $\mathcal{E}$ are index sets for inequality and equality constraints, respectively. The Lagrangian function is defined as

$$L(x, \lambda) = f(x) + \sum_{i \in \mathcal{I} \cup \mathcal{E}} \lambda_i c_i(x),$$

where the λ_i are the Lagrange multipliers. At each iterate, the direction of the step is computed by a quadratic programming subproblem of the form

$$\min \nabla f(x_k)^T d + \frac{1}{2} d^T H_k d$$

$$\text{subject to } c_i(x_k) + \nabla c_i(x_k)^T d \leq 0, \quad i \in \mathcal{I},$$
$$c_i(x_k) + \nabla c_i(x_k)^T d = 0, \quad i \in \mathcal{E},$$

where x_k is the current iterate, the constraints c_i are local linear approximations, and H_k is a positive definite estimate of the Hessian of the Lagrangian ($\nabla^2_{xx} L(x_k, \lambda_k)$). The solution d_k to this subproblem defines the direction of descent. The distance moved in the step direction is determined by a line search to minimize a merit function consisting of the objective function and a penalty function based on the violation of the constraints.

To solve SQP problems, we used CFSQP (C code for Feasible Sequential Quadratic Programming, as described by Lawrence, Zhou, and Tits (1994), using the code due to Schittkowski (1986) to solve each quadratic programming subproblem. CFSQP is a variant of the general approach described above, and it is unique in that it maintains "feasible" iterates during the optimization—once an iterate is found that satisfies all linear constraints and nonlinear inequality constraints, all subsequent iterates will also satisfy these constraints. For problems without an *rg* subsequence, this property ensures that each iterate in the solution process corresponds to a physically valid motion, though the goal state may not be achieved.

As with all iterative optimization routines, SQP finds a local optimum which is not necessarily the global optimum. In addition, the finite-dimensional parameterization of the manipulator trajectory artificially limits the space of possible trajectories. The particular local optimum achieved by SQP depends on the shape of the feasible space and the initial guess. This problem can be alleviated by solving with many different initial guesses and choosing the best solution.

5.3. *Putting It Together*

A dynamic task is specified by a geometric description of the polygonal object, along with its mass, center of mass, and radius of gyration ρ; the initial state and the desired goal state; and the sequence of manipulation phases to use (e.g., *g*, *rg*, *gf*, *grg*, or *grf*). A guess is also required to initialize the optimization. The motion of the object during rolling phases is simulated using fourth-order Runge-Kutta. The SQP strategy requires gradients of the constraints with respect to the design variables, and these are calculated using finite differences. All functions are implemented in C on a Sun SPARC 20.

6. Experiments

Although the planner of Section 5 can be applied to any robot, motivated by the results of Section 4, our experiments have been with a single-joint robot. We built a one-degree-of-freedom arm powered by an NSK direct-drive motor (Fig. 12).

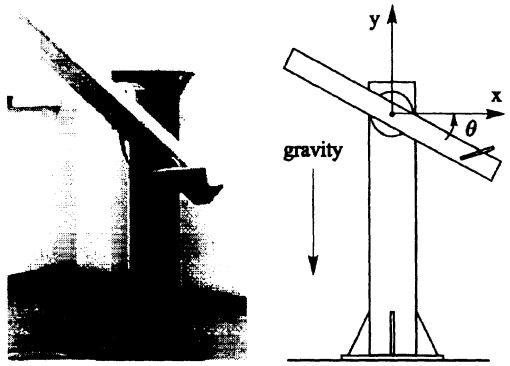

Fig. 12. The NSK direct-drive arm.

The hollow aluminum arm is centrally mounted and is 122 cm long, with a 10-cm-square cross section. It also has a "palm" mounted at a 45-degree angle. The top surfaces of the arm and the palm are used as manipulation surfaces. To increase friction and damping on these surfaces, they have been covered with a soft, 5-mm-thick foam.

Because dynamic manipulation requires precise trajectory following, we have carefully modeled the response of the NSK motor. As a result, we obtain good open-loop tracking of the planned arm trajectories (Lynch 1996). Small feedback corrections are also used.

Trajectories specified by the planner are implemented directly on the robot, without modification. In some cases, however, effects that are not modeled in the planner cause the plans to fail when they are implemented on the robot. For example, the planner assumes that if the dynamic grasp constraints are satisfied at the end of a roll (an *rg* sequence), then the object is immediately in a dynamic grasp. Impact is not modeled. Also, the planner's rigid-body assumption is violated by the soft foam, which has the effect of slightly rounding a rolling vertex. When these unmodeled effects cause a plan to fail, the problem specification can be modified to compensate. An example of this is described in Section 6.3.

To execute a throw, the arm is maximally decelerated at the release point. This causes the object to be released nearly instantaneously. If the arm is also to catch the object, it follows a bang-bang trajectory to reach the catching configuration. Catches are made robust by the soft foam (which minimizes rebound) and by appropriately choosing the object's impact state with the immobile arm. It is possible to choose the object's arrival state such that the impact immediately cancels the object's pre-impact velocity (Lynch 1996).

We now describe a snatch, throw, roll, and rolling throw. The test objects are lightweight, and made of wood. More details can be found in an earlier work (Lynch 1996), and a video of some of the experiments is available (Lynch and Mason 1997).

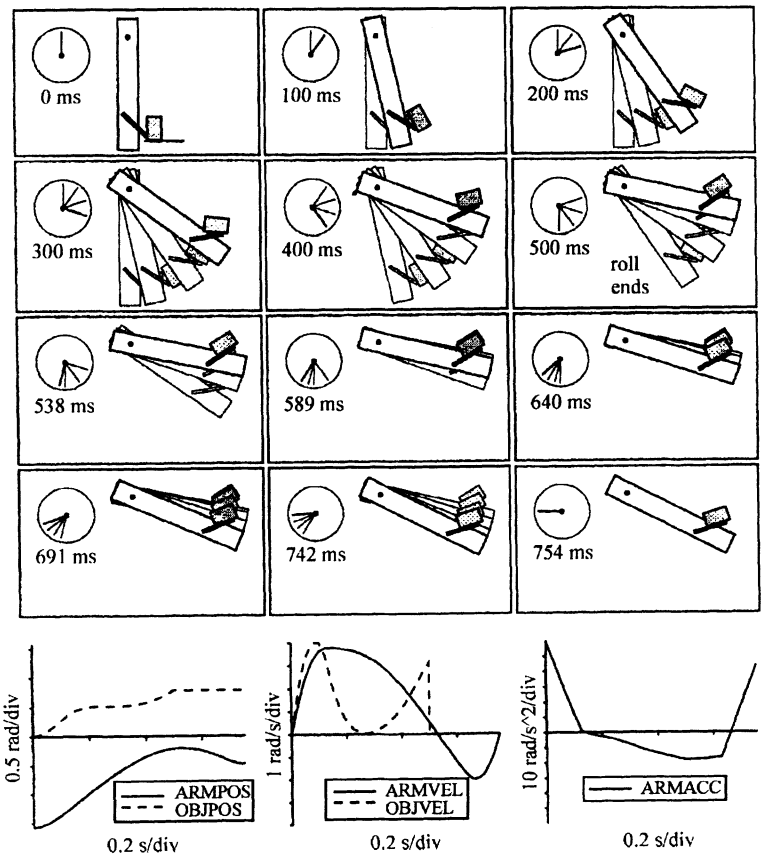

Fig. 13. A snatch: the trajectory found by the optimization. The required friction coefficient is 0.609. The plots show the angle of the arm (ARMPOS) and the angle of the object (OBJPOS) relative to it. The time between frames is not constant, so an equal number of frames of the roll and dynamic grasp phases can be seen. The clock indicates the time of each frame, and the previous three frames create a motion blur.

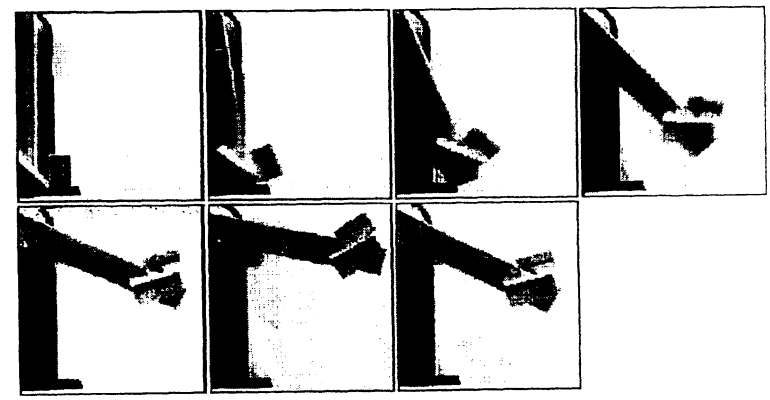

Fig. 14. A snatch: implementation on the robot.

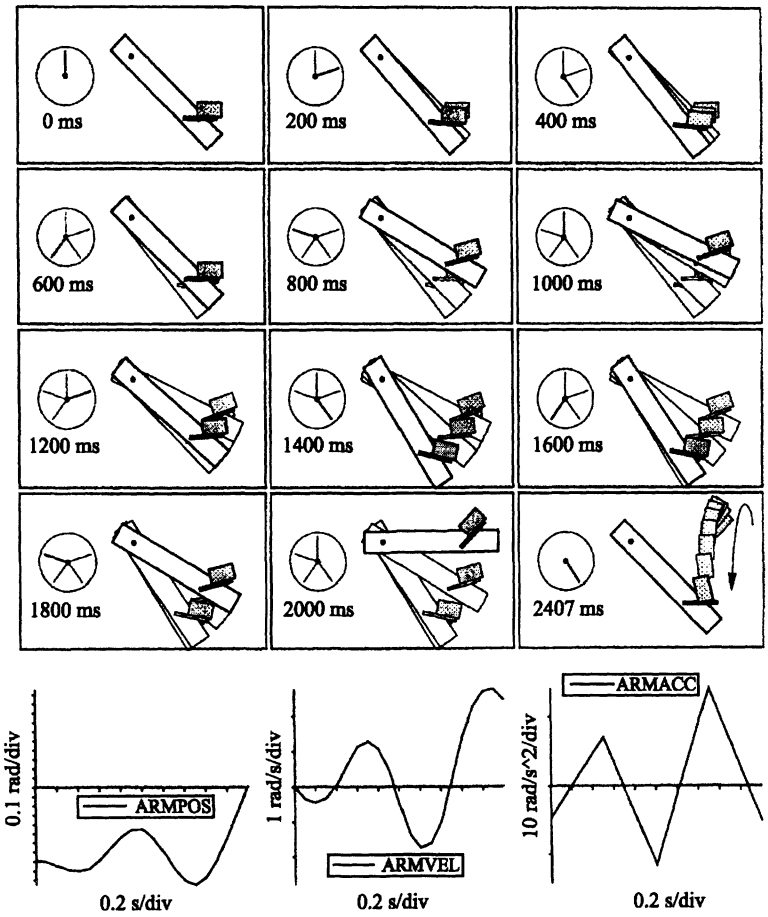

Fig. 15. A throw: the required friction coefficient is 0.156.

6.1. A Snatch

Using the phase sequence *rg*, the optimization finds the snatching trajectory shown in Figure 13 after 200 iterations and 120 sec. The trajectory consists of nine knot points, and the required friction coefficient μ is 0.609. The initial guess is for the arm to remain motionless, but the constraints of the optimization pull it toward a solution where the palm accelerates into the wooden block. The goal is any statically stable configuration after the roll has been completed.

The implementation on the arm works consistently (Fig. 14). If the same trajectory is slowed down too much, the block is simply pushed off the table, and if it is sped up too much, the block is thrown.

6.2. A Throw

After snatching the wooden block, the arm reorients it on the palm by throwing and catching it. Using the phase sequence *gf*, the optimization finds the trajectory shown in Figures 15 and 16 after 98 iterations and 13 sec. The trajectory consists of seven knot points and the required friction coefficient is 0.156. The goal state is a catching configuration where the block is over-rotated to counteract its angular velocity. The catching position of the arm (palm is horizontal) is specified.

Interestingly, the solution is to "double-pump" before throwing. Path reversals (spiraling through the state space) are often necessary to minimize the required friction. In this example, the throw could be accomplished without oscillation, but a larger friction coefficient would be required to prevent the block from slipping off before reaching the release state.

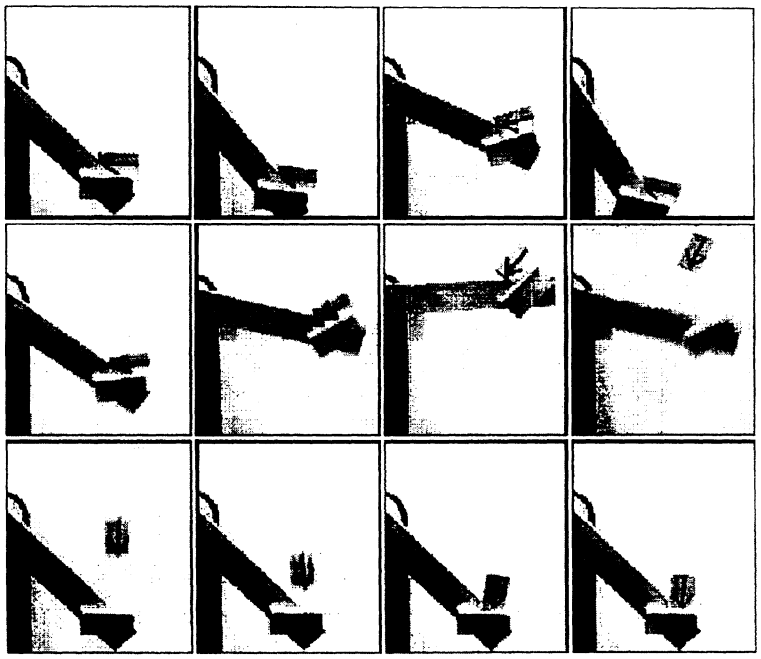

Fig. 16. A throw: implementation on the robot.

This is an example of composing a snatch and a throw. The throw and catch is experimentally robust, but occasionally the catch fails if the block tilts slightly out of the plane during the snatch.

6.3. A Roll

The manipulation phase sequence is grg, where the duration of the final g phase is zero. In other words, the arm stops moving as soon as the object has finished its roll, and this final configuration is statically stable.

The nine-knot rolling trajectory of Figures 17, 18, and 19 takes 32 sec and 55 iterations to find. The object is a 27-cm-square frame. Note the windup before the roll, which has the effect of "throwing" the square over its rolling vertex. In this example, the contact friction μ is set to 1.5 and the objective is to minimize the squared impact velocity at the end of the roll. The solution yields an impact velocity of 3.1 rad/sec. The implementation is robust, and the rebound on impact is small.

We limit the angle of the arm at the end of the roll to be greater than or equal to -0.3 rad. This constraint is active in the solution. Without this constraint, the solution is to end the roll with the arm at $-\pi/4$ rad and the center of mass of the square balanced over the rolling vertex. While this solution minimizes the impact velocity (zero), it is not experimentally robust.

6.4. A Rolling Throw

In this example, a wooden cube (7.6 cm on each side) is thrown with an angular velocity different from that of the arm at release. This is only possible with a rolling throw. Here the block rotates half a revolution clockwise before landing on the arm in the same position. Notice that the block does not begin to roll when the arm is at its nadir; centrifugal and gravitational forces combine to begin the roll. The nine-knot trajectory in Figure 20 takes 54 iterations and 35 sec to find. The required friction coefficient is 1.011. Examining the geometry of the object, we see that the friction coefficient must be greater than 1.0 to apply a clockwise torque to the object through the rolling vertex at the beginning of the roll, indicating that the solution is nearly optimal.

The rolling throw is the dynamic task most sensitive to trajectory error, as any error in the roll is propagated to the flight phase. While rolling throws are fairly repeatable, the accuracy is generally less than that of a throw. The implementation is shown in Figure 21.

7. Conclusion and Future Work

Dynamic underactuated nonprehensile manipulation exploits dynamic effects to achieve interesting behaviors with simple robots. This paper has studied its theoretical properties and

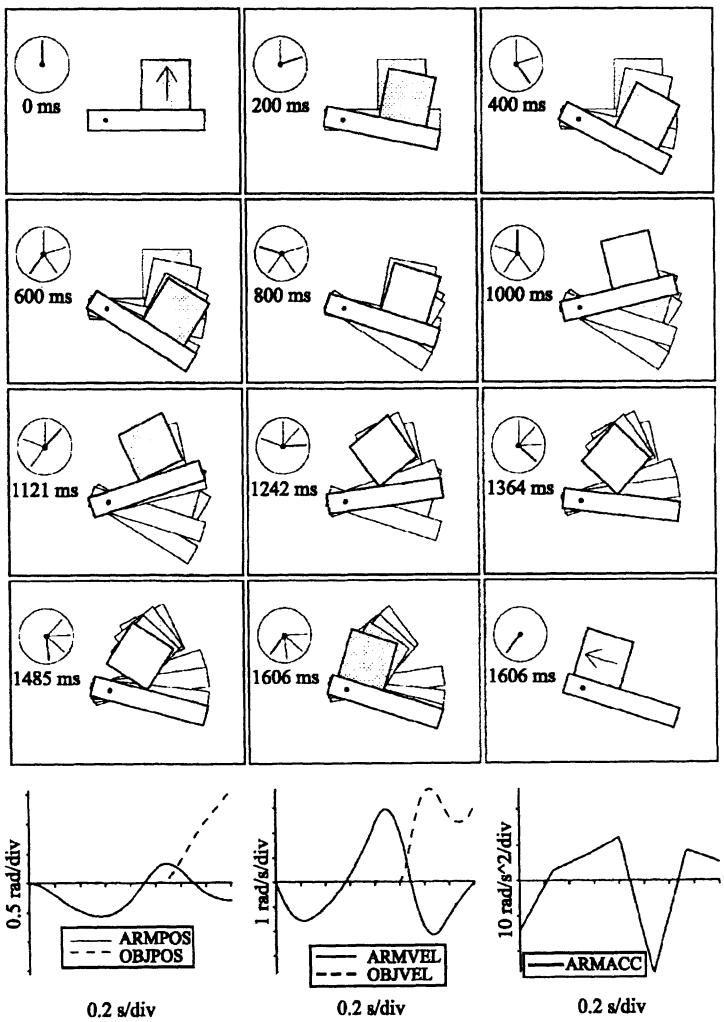

Fig. 17. A roll: the trajectory found by the optimization.

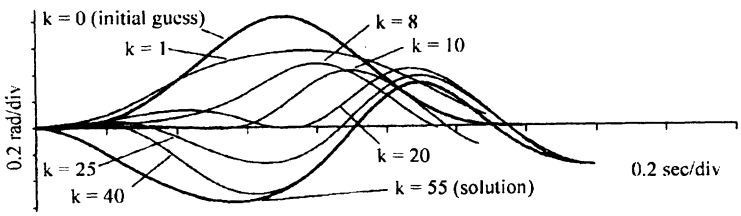

Fig. 18. A roll: the initial guess, solution, and intermediate iterates.

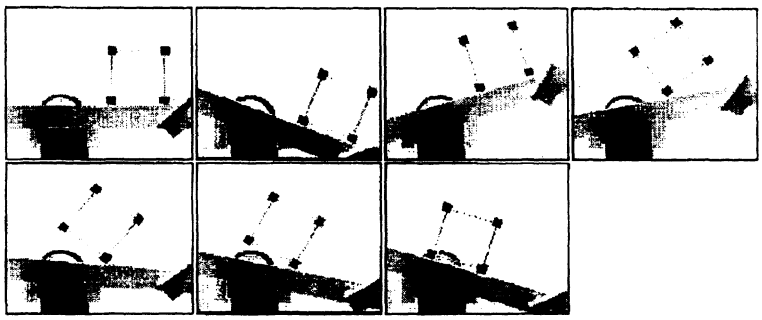

Fig. 19. A roll: implementation on the robot.

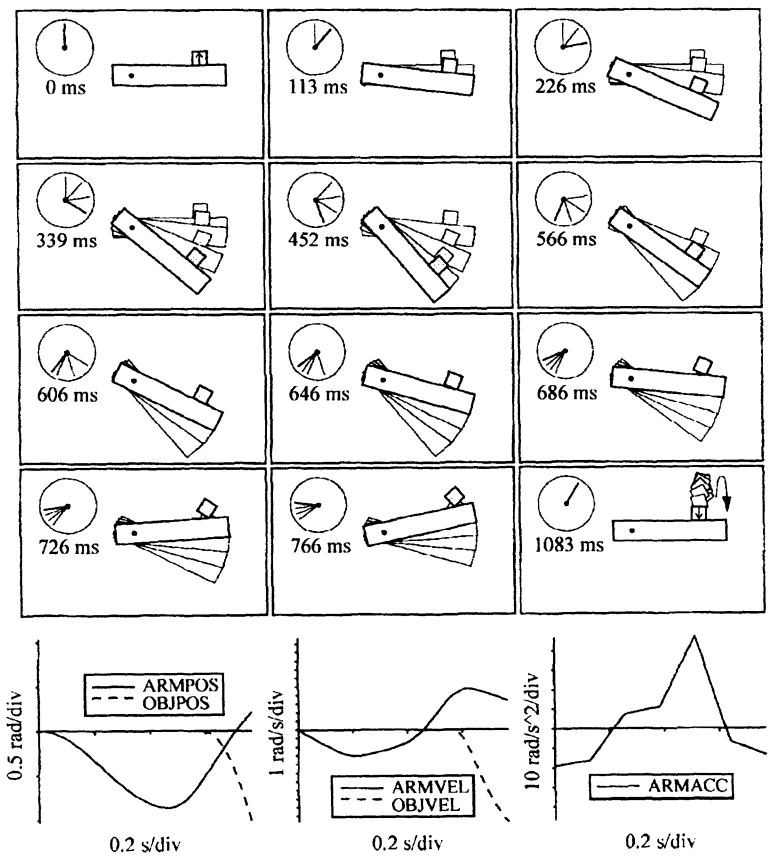

Fig. 20. A rolling throw: the trajectory found by the optimization.

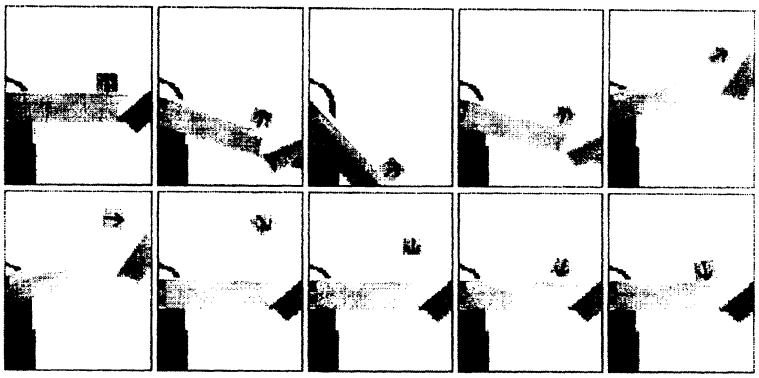

Fig. 21. A rolling throw: implementation on the robot.

presented a dynamic manipulation planner that has been successfully implemented on a real robot to perform a variety of dynamic tasks. Although the one-joint robot is a very simple system, it is a flexible test bed for research on dynamic nonholonomic manipulation using rolling, slipping, and free-flight manipulation phases.

There is much more work to be done. Future work should address feedback stabilization of the dynamic manipulation trajectories, possibly using vision or force feedback; automatic planning of manipulation-phase sequences; a more detailed analysis of the geometry of accessible states; using repeated contacts (batting) to increase the accessible state space; nonprehensile manipulation with more complex contact geometry and spatial parts; and tractable approaches to complete and globally optimal manipulation planning.

Appendix: Controls with Two Force Directions

In this appendix, we prove the sufficiency of Proposition 2 by adapting the results of Lynch (1998). We consider the minimum number of force directions (two) that satisfies the conditions of the proposition: $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$, not both pure torques, such that $\hat{\tau}_1 > 0, \hat{\tau}_2 < 0$. (Alternatively, we could consider three force directions, $\hat{\mathbf{f}}_1 = (0,0,1)^T$, $\hat{\mathbf{f}}_2 = (0,0,-1)^T$, and $\hat{\mathbf{f}}_3$ any unit force with a nonzero linear component, but this system can approximate the minimal system, so it suffices to prove sufficiency with the minimal system.) Without loss of generality, let $\mathbf{f}_1 = (f_1,0,1)^T$ ($f_1 = 1$) and $\mathbf{f}_2 = (f_2\cos\beta, f_2\sin\beta, -1)^T$. (Simply define the x-axis of the object frame $\Sigma_{\mathcal{O}}$ to align with the linear component of $\mathbf{f}_1$ and choose the unit length appropriately.) If $f_2 \neq 0$, the angle of the linear component of $\mathbf{f}_2$ is β in the object frame $\Sigma_{\mathcal{O}}$. The corresponding unit forces are $\hat{\mathbf{f}}_1, \hat{\mathbf{f}}_2$, with linear components $\hat{f}_1, \hat{f}_2$.

We first prove the sufficiency of Proposition 2 for the zero-gravity case with the control set U_d (Section A.1). We then modify the proof for the case of nonzero gravity with the control set U_u (Section A.2).

We will find the following lemma useful.

LEMMA 1. Consider a body-fixed force $\mathbf{f}$, with a linear magnitude $f = \sqrt{f_x^2 + f_y^2} > 0$ and torque $\tau \neq 0$, applied to $\mathcal{O}$ for a time t. Define the smooth function $m_{\mathbf{f}}(t, \omega_0)$ that maps the application time t and the initial angular velocity ω_0 of $\mathcal{O}$ to the magnitude of the total linear impulse delivered during the application (linear forces integrated over t). Then for any range of initial angular velocities $(-\omega_{max}, \omega_{max}), \omega_{max} > 0$, there exists a time $T > 0$ such that $\partial m_{\mathbf{f}}(t,\omega_0)/\partial t > 0$ for $m_{\mathbf{f}}(t,\omega_0)$ restricted to $t \in (0,T), \omega_0 \in (-\omega_{max}, \omega_{max})$.

Proof. The idea behind the lemma is simple. If the object does not rotate, the total linear impulse is ft. Rotation of the object, and therefore rotation of the force in the world frame $\Sigma_{\mathcal{W}}$, leads to some cancellation in the linear force components, giving a total linear impulse less than ft. If the rotation of the object during the application is less than $\pi/2$, however, clearly $m_{\mathbf{f}}(t,\omega_0)$ is monotonic with t (the cosine of the angle between any two instantaneous forces is positive). The magnitude of the angular velocity during the force application is upper-bounded by $|\omega_0| + T|\tau|$, so there always exists a $T > 0$ found by

$$(|\omega_0| + T|\tau|)T = \frac{\pi}{2} \quad \Rightarrow \quad T = \frac{-|\omega_0| + \sqrt{\omega_0^2 + 2\pi|\tau|}}{2|\tau|},$$

such that $\partial m_{\mathbf{f}}(t,\omega_0)/\partial t > 0$ for all $t \in (0,T)$.

More formally, from a state $(0,0,0,v_x,v_y,\omega_0)^T$, the equations of motion during the application of a force $(f,0,\tau)^T$ are given by

$$\ddot{\phi}_w(t) = \tau,$$

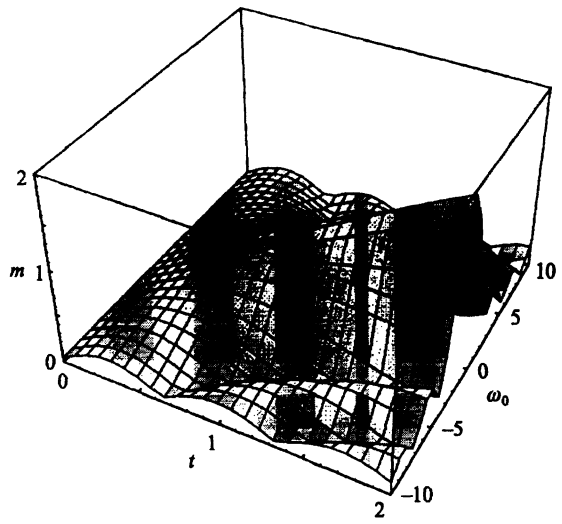

Fig. 22. The total linear impulse delivered to $\mathcal{O}$ as a function of its initial angular velocity ω_0 and the time of application t of the constant force $(f_x, f_y, 1)^T$, where $f_x^2 + f_y^2 = 1$.

$$\ddot{x}_w(t) = f\cos\phi_w(t) = f\cos(\frac{\tau t^2}{2} + \omega_0 t),$$
$$\ddot{y}_w(t) = f\sin\phi_w(t) = f\sin(\frac{\tau t^2}{2} + \omega_0 t).$$

Integrating, we obtain expressions for $\dot{x}_w(t)$ and $\dot{y}_w(t)$ in terms of Fresnel integrals, and $m_{\mathbf{f}}(t,\omega_0) = ((\dot{x}_w(t) - \dot{x}_w(0))^2 + (\dot{y}_w(t) - \dot{y}_w(0))^2)^{\frac{1}{2}}$. We find that $\partial m_{\mathbf{f}}(t,\omega_0)/\partial t \to f$ as t approaches 0 from above, for any ω_0, verifying the lemma for any given $(-\omega_{max}, \omega_{max})$. An example plot of $m_{\mathbf{f}}(t,\omega_0)$ for $f = \tau = 1$ is shown in Figure 22. □

A.1. Zero Gravity

Now we are ready to prove that the object is controllable in zero gravity using the unit forces $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$ (the control set U_d). It is sufficient to show the following:

1. The object is controllable on its velocity space, the quotient space $TC/\mathcal{C}$. Any velocity is attainable from any other velocity.
2. The object is controllable on its three-dimensional zero velocity section $Z = \{(\mathbf{q}, 0) | \mathbf{q} \in C\}$. The object can be moved from any $(\mathbf{q}_1, \mathbf{0})$ to any other $(\mathbf{q}_2, \mathbf{0})$.

Using these properties, the object $\mathcal{O}$ can be brought to a zero-velocity state $(\mathbf{q}_1, \mathbf{0})$ (Property 1), moved to an arbitrary zerovelocity state $(\mathbf{q}_2, \mathbf{0})$ (Property 2), and accelerated to the goal $(\mathbf{q}_3, \dot{\mathbf{q}}_3)$. This last step is possible because, by Property 1, there is some control $u(\cdot)$ to take $\mathcal{O}$ from zero velocity to the goal velocity $\dot{\mathbf{q}}_3$. Applying the control inverse from $(\mathbf{q}_3, \dot{\mathbf{q}}_3)$, we obtain the zero-velocity state $(\mathbf{q}_2, \mathbf{0})$.

Property 1: The object is controllable on its velocity space, the quotient space $TC/\mathcal{C}$. Any velocity is attainable from any other velocity.

We decouple the proof of this property by showing that the object's angular velocity can be changed arbitrarily without changing the linear velocity, and the linear velocity can be changed arbitrarily without changing the angular velocity. We begin with the former.

There exists a $\Delta\omega > 0$, for any initial angular velocity $\dot{\phi}_w$, such that two applications of the unit force $\hat{\mathbf{f}}_1$ (respectively $\hat{\mathbf{f}}_2$), for times t_1 and t_2, can change the angular velocity to any value in the open interval $(\dot{\phi}_w, \dot{\phi}_w + \Delta\omega)$ (respectively, $(\dot{\phi}_w - \Delta\omega, \dot{\phi}_w)$) without changing the linear velocity. By patching together open sets, $\mathcal{O}$ can be moved from $(\dot{x}_{w1}, \dot{y}_{w1}, \dot{\phi}_{w1})$ to $(\dot{x}_{w1}, \dot{y}_{w1}, \dot{\phi}_{w2})$ for any $\dot{x}_{w1}, \dot{y}_{w1}, \dot{\phi}_{w1}, \dot{\phi}_{w2}$.

Consider $\hat{\mathbf{f}}_1$ (similar reasoning holds for $\hat{\mathbf{f}}_2$) applied for a time t_1 at an initial angular velocity ω_1. By Lemma 1, there exists a neighborhood W of $(0, \omega_1)$ such that for all $(t,\omega) \in W$, $\partial m_{\hat{\mathbf{f}}_1}(t,\omega)/\partial t > 0$ ($m_{\hat{\mathbf{f}}_1}$ does not achieve a local maximum or minimum, and the constant $m_{\hat{\mathbf{f}}_1}$ contours never become parallel to the t-axis). Therefore, we can choose t_1 from an open interval $(0, T)$ such that $(t_1, \omega_1) \in W$ and there exists a $(t_2, \omega_1 + t_1\hat{\tau}_1) \in W$ (where $\omega_1 + t_1\hat{\tau}_1$ is the angular velocity after application of $\hat{\mathbf{f}}_1$ for time t_1) such that $m_{\hat{\mathbf{f}}_1}(t_1, \omega_1) = m_{\hat{\mathbf{f}}_1}(t_2, \omega_1 + t_1\hat{\tau}_1)$. In other words, the total linear impulse delivered by the first application can be exactly matched by the linear impulse of the second application. By allowing the object to rotate sufficiently between applications, the linear impulses cancel, restoring the original linear velocity, while the angular velocity is transferred to a point of an open interval $(\omega_1, \omega_1 + \Delta\omega)$ parameterized by $t_1 \in (0, T)$.

Note that if t_1 is chosen to exactly zero the angular velocity after the first application ($t_1 = -\omega_1/\hat{\tau}_1$, $\omega_1 < 0$), $\mathcal{O}$ cannot rotate to the angle for the second application. We simply avoid such a value of t_1, possibly creating the two open t_1 intervals $(0, -\omega_1/\hat{\tau}_1)$ and $(-\omega_1/\hat{\tau}_1, T)$.

To see that the linear velocity can be changed arbitrarily without changing the angular velocity, first assume the object is always rotating. The unit forces $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$ can be applied in pairs such that their total torque impulses cancel and their linear components sum to yield a net change of velocity in the desired direction. Because the object is rotating, the linear components of $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$ can take any direction in the world frame $\Sigma_{\mathcal{W}}$.

Finally, the object can be maneuvered to any desired velocity by first transferring it to the desired linear velocity (possibly after giving the object an initial angular velocity) and then transferring it to the desired angular velocity.

Property 2: The object is controllable on its three-dimensional space Z of zero-velocity states. The object can be moved from any zero-velocity state $(\mathbf{q}_1, \mathbf{0})$ to any other $(\mathbf{q}_2, \mathbf{0})$.

This can be proven by demonstrating that $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$ are sufficient to steer $\mathcal{O}$ from any zero-velocity state $(\mathbf{q}, \mathbf{0})$ to a neighborhood of $\mathbf{q}$ on the three-dimensional zero-velocity section Z. By patching together neighborhoods, $\mathcal{O}$ can be moved from any $(\mathbf{q}_1, \mathbf{0})$ to any $(\mathbf{q}_2, \mathbf{0})$. (Construct any curve in Z connecting these two states. Each point on the curve is in the interior of its open accessible set. These open sets comprise an open cover of the curve, and because the curve is compact, there is a finite subcover.)

To greatly simplify the discussion, we first consider the limiting case where control forces are applied in short bursts (impulses) such that the motion of the object during the application of a control force is zero. We then relax this assumption. The magnitudes of the linear impulses delivered during applications of $\hat{\mathbf{f}}_1$ and $\hat{\mathbf{f}}_2$ are ϵ and $f_2\epsilon$, respectively, and the angular impulses are ϵ and $-\epsilon$, respectively, where ϵ is a small positive value. Consider the following four-step sequence beginning with the object at $(0, 0)$:

1. Apply an $\hat{\mathbf{f}}_1$ impulse and allow the object to drift for time π/ϵ. The new state is $(\pi, 0, \pi, \epsilon, 0, \epsilon)^T$.
2. Apply an $\hat{\mathbf{f}}_1$ impulse, canceling the object's linear velocity and doubling its angular velocity. Allow the object to rotate in place an angle ψ_1, where ψ_1 can be chosen arbitrarily. The new state is $(\pi, 0, \pi + \psi_1, 0, 0, 2\epsilon)^T$.
3. Apply an $\hat{\mathbf{f}}_2$ impulse and allow the object to drift for time π/ϵ. The new state is $(\pi + \pi f_2 \cos(\pi + \psi_1 + \beta), \pi f_2 \sin(\pi + \psi_1 + \beta), 2\pi + \psi_1, f_2\epsilon \cos(\pi + \psi_1 + \beta), f_2\epsilon \sin(\pi + \psi_1 + \beta), \epsilon)^T$.
4. Apply an $\hat{\mathbf{f}}_2$ impulse, stopping the object's motion. The final state is $(\pi + \pi f_2 \cos(\pi + \psi_1 + \beta), \pi f_2 \sin(\pi + \psi_1 + \beta), 2\pi + \psi_1, 0, 0, 0)^T$.

The final set of configurations R_1 reachable from the zero configuration $\mathbf{0}$ is a one-dimensional curve on Z, parameterized by ψ_1, given by $f(\psi_1, \mathbf{0}) = (\pi(1 - f_2 \cos(\psi_1 + \beta)), -\pi f_2 \sin(\psi_1 + \beta), \psi_1)^T$. The mapping f is independent of ϵ; $-\epsilon$ determines only the time of motion. We consider ψ_1 to be the control. Repeating the sequence from each point of R_1, we obtain a reachable set R_{12}, and repeating again from each point of R_{12}, we obtain a set of reachable configurations $R_{123} = \{f(\psi_3, f(\psi_2, f(\psi_1, \mathbf{0})))|\psi_1, \psi_2, \psi_3 \in S^1\}$ on Z:

$$
\begin{aligned}
&(\pi(1 - f_2 \cos(\psi_1 + \beta) + \cos\psi_1(1 - f_2 \cos(\psi_2 + \beta)) \\
&+ f_2 \sin\psi_1 \sin(\psi_2 + \beta) + \cos(\psi_1 + \psi_2) \\
&\times (1 - f_2 \cos(\psi_3 + \beta)) + f_2 \sin(\psi_1 + \psi_2) \sin(\psi_3 + \beta)),
\end{aligned}
$$

$$
\begin{aligned}
&\pi(-f_2 \sin(\psi_1 + \beta) + \sin\psi_1(1 - f_2 \cos(\psi_2 + \beta)) \\
&- f_2 \cos\psi_1 \sin(\psi_2 + \beta) + \sin(\psi_1 + \psi_2)(1 - f_2 \\
&\cos(\psi_3 + \beta)) - f_2 \cos(\psi_1 + \psi_2) \sin(\psi_3 + \beta)),
\end{aligned}
$$

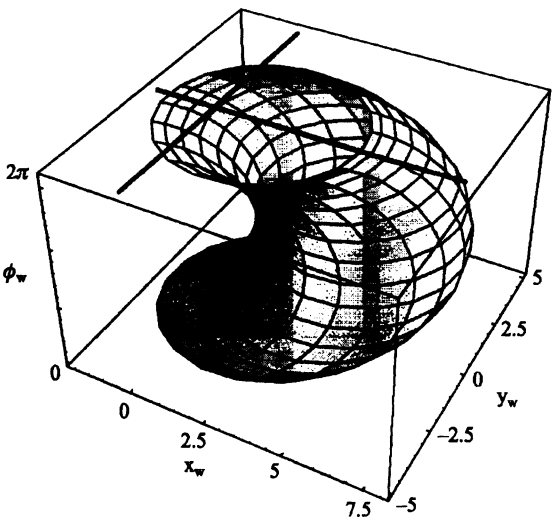

Fig. 23. The reachable set R_{123} for $\mathbf{f}_1 = (1, 0, 1)^T$, $\mathbf{f}_2 = (0.5, 0, -1)^T$. Note that the twisted tube wraps around at $\phi_w = 2\pi$ and that $(0, 0, 0)^T$ is in the interior.

$$
\psi_1 + \psi_2 + \psi_3)^T.
$$

(Note that angles are modulo 2π.)

We wish to show that this set contains a neighborhood of the origin on Z. First we observe that by choosing $\psi_1 = \psi_2 = \psi_3 = 2\pi/3$, the object returns to the zero configuration $(0, 0, 0)^T$. We will define this point in the three-dimensional control space as $\Psi_0 = (2\pi/3, 2\pi/3, 2\pi/3)^T$. To show that $(0, 0, 0)^T$ is in the interior of R_{123} on Z, we look at the determinant of the Jacobian matrix of R_{123},

$$
\det \begin{pmatrix} \frac{\partial x_w}{\partial \psi_1} & \frac{\partial x_w}{\partial \psi_2} & \frac{\partial x_w}{\partial \psi_3} \\ \frac{\partial y_w}{\partial \psi_1} & \frac{\partial y_w}{\partial \psi_2} & \frac{\partial y_w}{\partial \psi_3} \\ \frac{\partial \phi_w}{\partial \psi_1} & \frac{\partial \phi_w}{\partial \psi_2} & \frac{\partial \phi_w}{\partial \psi_3} \end{pmatrix} = \pi^2(1 + f_2^2 - 2f_2 \cos\beta) \sin\psi_2.
$$

The determinant is nonzero at Ψ_0 unless $1 + f_2^2 - 2f_2 \cos\beta = 0$. This equation holds only when $f_2 = 1, \beta = 0$; that is, $\mathbf{f}_2 = (1, 0, -1)^T$. For any other choice of $\mathbf{f}_2$, the set R_{123} gives a neighborhood of reachable configurations of $(0, 0, 0)^T$ on Z. Figure 23 shows the reachable set R_{123} for $f_2 = 0.5, \beta = 0$.

If $\mathbf{f}_2 = (1, 0, -1)^T$, the reachable configuration space by the control sequence above is a one-dimensional curve, regardless of the number of applications of the sequence. Therefore we modify the control sequence above by prepending it with the following two steps (assume the object begins at $(0, 0)$):

- Apply an $\hat{\mathbf{f}}_1$ impulse and allow the object to drift for time π/ϵ. The new state is $(\pi, 0, \pi, \epsilon, 0, \epsilon)^T$.

- Apply an $\hat{\mathbf{f}}_2$ impulse, stopping the object's motion. The new state is $(\pi, 0, \pi, 0, 0, 0)^T$.

Then, continuing with the four steps described previously, the final reachable curve is $f(\psi_1, 0) = (\pi \cos\psi_1, \pi \sin\psi_1, \psi_1 + \pi)^T$. Applying the complete six-step control sequence two more times, we get $R_{123} = \{f(\psi_3, f(\psi_2, f(\psi_1, \mathbf{0})))| \psi_1, \psi_2, \psi_3 \in S^1\}$ on Z:

$$(\pi(\cos\psi_1(1 - \cos\psi_2) + \sin\psi_1 \sin\psi_2 + \cos(\psi_1 + \psi_2)\cos\psi_3 - \sin(\psi_1 + \psi_2)\sin\psi_3),$$

$$\pi(\sin\psi_1(1 - \cos\psi_2) - \cos\psi_1 \sin\psi_2 + \sin(\psi_1 + \psi_2)\cos\psi_3 + \cos(\psi_1 + \psi_2)\sin\psi_3),$$

$$\psi_1 + \psi_2 + \psi_3 + \pi)^T.$$

The control $\Psi_0 = (5\pi/3, 5\pi/3, 5\pi/3)^T$ returns the object to $(0, 0, 0)^T$. The determinant of the Jacobian is $-\pi^2 \sin\psi_2$, indicating that this control is nonsingular—R_{123} contains a neighborhood of $(0, 0, 0)^T$ on Z.

We have shown that by applying controls in a neighborhood V of Ψ_0 in the control space, the object $\mathcal{O}$, starting from the zero-velocity configuration $(0, 0, 0)^T$, can reach a neighborhood W of $(0, 0, 0)^T$ on Z using impulses. W is invariant to $\mathbf{q}$ when expressed in the object frame $\Sigma_{\mathcal{O}}$, so from any configuration $\mathbf{q}$ on Z, $\mathcal{O}$ can reach a neighborhood of configurations of $\mathbf{q}$ on Z.

Finally, we replace the impulses with the unit-control forces $\hat{\mathbf{f}}_1, \hat{\mathbf{f}}_2$. Consider the first two steps of the four-step control sequence. The impulse of the first step is approximated by an application of $\hat{\mathbf{f}}_1$ for a time $T/2$ (linear impulse is $m_1 = m_{\hat{\mathbf{f}}_1}(T/2, 0)$), free rotation of the object by an angle δ, and another application for a time $T/2$ (linear impulse is $m_2 = m_{\hat{\mathbf{f}}_1}(T/2, T\hat{\tau}_1/2)$). The second step is approximated by an application of $\hat{\mathbf{f}}_1$ for a time T (linear impulse is $m_3 = m_{\hat{\mathbf{f}}_1}(T, T\hat{\tau}_1)$). The total torque impulses of the two steps are equal. For small values of T, δ can be chosen so that the total linear impulse of the two steps are equal, thus zeroing the linear velocity after the two steps. (Note that $m_1 + m_2 > m_3 > m_1 > m_2$, and δ allows partial cancellation between m_1 and m_2. As $T \to 0$, $\delta \to 0$.) Thus the final velocity is $(0, 0, 2T\hat{\tau}_1)^T$, identical to the impulse case where $\epsilon = T\hat{\tau}_1$. The configuration-error vector (due to motion of the object during the force application and the linear velocity error during the motion between steps 1 and 2) is $e_1(\psi_1, T)$, smooth in ψ_1 and T, and $e_1(\psi_1, T) \to \mathbf{0}$ as $T \to 0$.

Continuing in an analogous manner for steps 3 and 4 (the object returns to zero velocity after step 4), and repeating the sequence twice more, we get a smooth final configuration error $e_{123}(\psi_1, \psi_2, \psi_3, T)$ such that $e_{123}(\psi_1, \psi_2, \psi_3, T) \to 0$ as $T \to 0$. For any $W \in R_{123}$ above, where W is a neighborhood of $(0, 0, 0)^T$ on Z, we can define the set $W' = \{w + e_{123}(\psi_1, \psi_2, \psi_3, T) | w \in W\}$, where ψ_1, ψ_2, ψ_3 are the controls (for the impulse case) that take the object to w. By choosing T small enough, W' also contains a neighborhood of $(0, 0, 0)^T$ on Z. Therefore, the object $\mathcal{O}$ can reach a neighborhood of any initial configuration $\mathbf{q}$ on Z with unit-control forces, and $\mathcal{O}$ is controllable on its zero-velocity section.

This completes the proof for the case of zero gravity and the control set U_d.

A.2. Nonzero Gravity

We now introduce nonzero gravity with the control set U_u (the control forces are $u\mathbf{f}_1$ and $u\mathbf{f}_2$, where $u \in [0, \infty)$). Scaling, we can choose gravity $g > 0$ to be as small as desired. The proof for the zero-gravity case is easily modified.

Property 1 holds simply by adding applications of $\mathbf{f}_1$ and $\mathbf{f}_2$ during the control sequence such that their torque impulses cancel while their linear impulses sum to exactly cancel the change in linear velocity due to gravity (gT, where T is the total time of the maneuver). Similarly, Property 2 holds by adding applications of $\mathbf{f}_1$ and $\mathbf{f}_2$ during each control sequence to cancel the change in velocity due to gravity. This introduces a position error at the end of the sequence that goes to zero as g goes to zero. By choosing g sufficiently small, from any $\mathbf{q}$ on Z, the object can reach a neighborhood of $\mathbf{q}$ on Z. Finally, Properties 1 and 2 are sufficient to prove that the reachable state space from any $(\mathbf{q}, \dot{\mathbf{q}})$ is the entire state space $T\mathcal{C}$ in zero gravity. Because the effect of gravity can be made arbitrarily small, this also holds with nonzero gravity and the control set U_u.

Acknowledgments

This work was funded by NSF under grant IRI-9318496. Thanks to Pradeep Goel and NSK for providing the motor used in the experiments and André Tits and Craig Lawrence for providing the CFSQP software. We also thank Richard Murray, Andrew Lewis, P. S. Krishnaprasad, and Yan-Bin Jia for helpful discussions regarding the controllability of a planar object, and an anonymous reviewer for suggestions to improve the readability of the paper.

References

Abell, T., and Erdmann, M. A. 1995. Stably supported rotations of a planar polygon with two frictionless contacts. *IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Washington, DC: IEEE, pp. 411–418.

Aboaf, E. W., Atkeson, C. G., and Reinkensmeyer, D. J. 1987. Task-level robot learning: Ball throwing. AI Memo 1006, Massachusetts Institute of Technology, Cambridge, MA.

Aboaf, E. W., Drucker, S. M., and Atkeson, C. G. 1989 (Scottsdale, AZ). Task-level robot learning: Juggling a tennis ball more accurately. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 1290–1295.

Aiyama, Y., Inaba, M., and Inoue, H. 1993 (Yokohama, Japan). Pivoting: A new method of graspless manipulation of object by robot fingers. *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Los Alamitos, CA: IEEE, pp. 136–143.

Akella, S., Huang, W., Lynch, K. M., and Mason, M. T. 1995. Planar manipulation on a conveyor with a one-joint robot. *Proc. of the Int. Symp. on Robot. Res.* Cambridge, MA: MIT Press.

Andersson, R. L. 1989 (Scottsdale, AZ). Understanding and applying a robot Ping-Pong player's expert controller. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 1284–1289.

Arai, H. 1996. Controllability of a 3-DOF manipulator with a passive joint under a nonholonomic constraint. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 3707–3713.

Arai, H., and Khatib, O. 1994. Experiments with dynamic skills. *Proc. of the 1994 Japan-USA Symp. on Flexible Automat.*, pp. 81–84.

Arai, H., and Tachi, S. 1991. Position control of a manipulator with passive joints using dynamic coupling. *IEEE Trans. Robot. Automat.* 7(4):528–534.

Arai, H., Tanie, K., and Tachi, S. 1993. Dynamic control of a manipulator with passive joints in operational space. *IEEE Trans. Robot. and Automat.* 9(1):85–93.

Barraquand, J., and Latombe, J.-C. 1993. Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles. *Algorithmica* 10:121–155.

Bartels, R. H., Beatty, J. C., and Barsky, B. A. 1987. *An Introduction to Splines for Use in Computer Graphics and Geometric Modeling.* San Francisco, CA: Morgan Kaufmann.

Bergerman, M., Lee, C., and Xu, Y. 1995. Experimental study of an underactuated manipulator. *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Washington, DC: IEEE, pp. 317–322.

Berkemeier, M. D., and Fearing, R. S. 1992 (Nice, France). Control of a two-link robot to achieve sliding and hopping gaits. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 286–291.

Bicchi, A., and Sorrentino, R. 1995. Dexterous manipulation through rolling. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 452–457.

Bobrow, J. E., Dubowsky, S., and Gibson, J. S. 1985. Time-optimal control of robotic manipulators along specified paths. *Int. J. Robot. Res.* 4(3):3–17.

Böhringer, K., Brown, R., Donald, B., Jennings, J., and Rus, D. 1995. Distributed robotic manipulation: Experiments in minimalism. *Proc. of the Int. Symp. on Exp. Robot.*

Boothby, W. M. 1986. *An Introduction to Differentiable Manifolds and Riemannian Geometry.* New York: Academic Press.

Boothroyd, G., Poli, C., and Murch, L. E. 1982. *Automatic Assembly.* New York: Marcel Dekker.

Brock, D. L. 1988. Enhancing the dexterity of a robot hand using controlled slip. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 249–251.

Brockett, R. W. 1976. Nonlinear systems and differential geometry. *Proc. of the IEEE* 64(1):61–72.

Brockett, R. W. 1983. Asymptotic stability and feedback stabilization. In Brockett, R. W., Millman, R. S., and Sussmann, H. J. (eds): *Differential Geometric Control Theory.* Boston, MA: Birkhauser.

Bühler, M., and Koditschek, D. E. 1990 (Cincinnati, OH). From stable to chaotic juggling: Theory, simulation, and experiments. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 1976–1981.

Burridge, R. R., Rizzi, A. A., and Koditschek, D. E. 1995. Toward a dynamical pick and place. *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Washington, DC: IEEE, pp. 292–297.

Canny, J. F., and Goldberg, K. Y. 1994. "RISC" industrial robotics: Recent results and open problems. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 1951–1958.

Canny, J., Reif, J., Donald, B., and Xavier, P. 1988 (White Plains, NY). On the complexity of kinodynamic planning. *Proc. of the IEEE Symp. on the Found. of Comp. Sci.* Los Alamitos, CA: IEEE, pp. 306–316.

Carlisle, B., Goldberg, K., Rao, A., and Wiegley, J. 1994. A pivoting gripper for feeding industrial parts. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 1650–1655.

Chen, Y.-C. 1991. Solving robot trajectory planning problems with uniform cubic B-splines. *Optimal Control Appl. Methods* 12:247–262.

De Luca, A., Lanari, L., and Oriolo, G. 1991. A sensitivity approach to optimal spline robot trajectories. *Automatica* 27(3):535–539.

De Luca, A., Mattone, R., and Oriolo, G. 1996. Dynamic mobility of redundant robots using end-effector commands. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 1760–1767.

Divelbiss, A. W., and Wen, J. 1993. Nonholonomic path planning with inequality constraints. *Proc. of the IEEE Int. Conf. on Decision and Control.* Los Alamitos, CA: IEEE, pp. 2712–2717.

Donald, B., and Xavier, P. 1989 (Scottsdale, AZ). A provably good approximation algorithm for optimal-time trajectory planning. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 958–963.

Erdmann, M. A. 1984 (August). On motion planning with uncertainty. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Erdmann, M. A. 1994. On a representation of friction in configuration space. *Int. J. Robot. Res.* 13(3):240–271.

Erdmann, M. A. 1995. An exploration of nonprehensile two-palm manipulation: Planning and execution. *Proc. of the Int. Symp. on Robot. Res.* Cambridge, MA: MIT Press.

Erdmann, M. A., and Mason, M. T. 1988. An exploration of sensorless manipulation. *IEEE Trans. Robot. Automat.* 4(4):369–379.

Erdmann, M. A., Mason, M. T., and Vaněček, Jr., G. 1993. Mechanical parts orienting: The case of a polyhedron on a table. *Algorithmica* 10:226–247.

Ferbach, P., and Rit, J.-F. 1996. Planning nonholonomic motions for manipulated objects. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 2935–2942.

Fernandes, C., Gurvits, L., and Li, Z. 1994. Near-optimal nonholonomic motion planning for a system of coupled rigid bodies. *IEEE Trans. Automat. Control* 30(3):450–463.

Fliess, M., Lévine, J., Martin, P., and Rouchon, P. 1995. On differentially flat nonlinear systems. *Int. J. Control* 61(6):1327–1361.

Goodwine, B., and Burdick, J. 1996. Controllability with unilateral control inputs. *Proc. of the Conf. on Decision and Control*, pp. 3394–3399.

Hauser, J., and Murray, R. M. 1990. Nonlinear controllers for nonintegrable systems: The acrobot example. *Am. Control Conf.*. Los Alamitos, CA: IEEE, pp. 669–671.

Haynes, G. W., and Hermes, H. 1970. Nonlinear controllability via Lie theory. *SIAM J. Control* 8(4):450–460.

Hermann, R., and Krener, A. J. 1977. Nonlinear controllability and observability. *IEEE Trans. Automat. Control* 22(5): 728–740.

Hitakawa, H. 1988. Advanced parts orientation system has wide application. *Assembly Automation* 8(3):147–150.

Hove, B. M., and Slotine, J.-J. E. 1991. Experiments in robotic catching. *Proc. of the Am. Control Conf.* Los Alamitos, CA: IEEE, pp. 380–385.

Huang, W., Krotkov, E. P., and Mason, M. T. 1995. Impulsive manipulation. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 120–125.

Isidori, A. 1989. *Nonlinear Control Systems: An Introduction.* Berlin: Springer-Verlag.

Jurdjevic, V. 1972. Certain controllability properties of analytic control systems. *SIAM J. Control* 10(2):354–360.

Koditschek, D. E. 1991. Robot assembly: Another source of nonholonomic control problems. *Am. Control Conf.* Los Alamitos, CA: IEEE, pp. 1627–1632.

Koditschek, D. E. 1993. Dynamically dexterous robots. In Spong, M. W., Lewis, F. L., and Abdallah, C. T. (eds.): *Robot Control: Dynamics, Motion Planning, and Analysis.* New York: IEEE Press.

Lafferiere, G, and Sussmann, H. 1991 (Sacramento, CA). Motion planning for controllable systems without drift. *IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 1148–1153.

Latombe, J.-C. 1991. *Robot Motion Planning.* Boston, MA: Kluwer Academic.

Laumond, J.-P. 1986. Feasible trajectories for mobile robots with kinematic and environment constraints. *Int. Conf. on Intell. Autonomous Sys.*, pp. 346–354.

Laumond, J.-P., Jacobs, P. E., Taïx, M., and Murray, R. M. 1994. A motion planner for nonholonomic mobile robots. *IEEE Trans. Robot. Automat.* 10(5):577–593.

Lawrence, C., Zhou, J. L., and Tits, A. L. 1994. User's guide for CFSQP version 2.3. Technical Report 94-16, Institute for Systems Research, University of Maryland.

Leonard, N. E., and Krishnaprasad, P. S. 1995. Motion control of drift-free, left-invariant systems on Lie groups. *IEEE Trans. Automat. Control* 40(9):1539–1554.

Lewis, A. D. 1997 (Brussels). Local configuration controllability for a class of mechanical systems with a single input. *1997 European Control Conference* (CD-ROM).

Lewis, A. D., and Murray, R. M. 1997. Configuration controllability of simple mechanical control systems. *SIAM J. Control Optimization* 35(3):766-790.

Li, Z., and Canny, J. 1990. Motion of two rigid bodies with rolling constraint. *IEEE Trans. Robot. Automat.* 6(1):62–72.

Li, Z., and Canny, J. 1993. *Nonholonomic Motion Planning.* Boston, MA: Kluwer Academic.

Lynch, K. M. 1996 (March). Nonprehensile Robotic Manipulation: Controllability and Planning. PhD thesis, The Robotics Institute, Carnegie Mellon University. Available on the Web as Technical Report CMU-RI-TR-96-05.

Lynch, K. M., and Mason, M. T. 1996. Stable pushing: Mechanics, controllability, and planning. *Int. J. Robot. Res.* 15(6):533–556.

Lynch, K. M., and Mason, M. T. 1997. Nonprehensile manipulation. *1997 IEEE Int. Conf. on Robot. and Automat. Video Proc.* Washington, DC: IEEE.

Lynch, K. M., Shiroma, N., Arai, H., and Tanie, K. 1998. Motion planning for a 3-DOF robot with a passive joint. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE.

Lynch, K. M. Forthcoming. Controllability of a planar body with unilateral thrusters; IEEE Transactions on Automatic Control.

Manikonda, V., and Krishnaprasad, P. S. 1997. Controllability of Lie-Poisson reduced dynamics. *Proc. of the Am. Control Conf.* Washington, DC: IEEE.

Markenscoff, X., Ni, L., and Papadimitriou, C. H. 1990. The geometry of grasping. *Int. J. Robot. Res.* 9(1):61–74.

Martin, B. J., and Bobrow, J. E. 1997. Minimum effort motions for open-chain manipulators with task-dependent end-effector constraints. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 2044–2049.

Mason, M. T. 1986. Mechanics and planning of manipulator-pushing operations. *Int. J. Robot. Res.* 5(3):53–71.

Mason, M. T., and Lynch, K. M. 1993a (Yokohama, Japan). Dynamic manipulation. *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Los Alamitos, CA: IEEE, pp. 152–159.

Mason, M. T., and Lynch, K. M. 1993b (October, Hidden Valley, PA). Throwing a club: Early results. *Proc. of the Int. Symp. on Robot. Res.* Cambridge, MA: MIT Press.

Mishra, B., Schwartz, J. T., and Sharir, M. 1987. On the existence and synthesis of multifinger positive grips. *Algorithmica* 2(4):541–558.

Murray, R. M., Li, Z., and Sastry, S. S. 1994. *A Mathematical Introduction to Robotic Manipulation.* Boca Raton, FL: CRC Press.

Murray, R. M., Rathinam, M., and Sluis, W. 1995. Differential flatness of mechanical control systems: A catalog of prototype systems. *Proc. of the ASME Int. Mech. Eng. Congress and Expo.* New York: ASME.

Murray, R. M., and Sastry, S. S. 1993. Nonholonomic motion planning: Steering using sinusoids. *IEEE Trans. Automat. Control* 38(5):700–716.

Nijmeijer, H., and van der Schaft, A. J. 1990. *Nonlinear Dynamical Control Systems.* Berlin: Springer-Verlag.

Oriolo, G., and Nakamura, Y. 1991. Control of mechanical systems with second-order nonholonomic constraints: Underactuated manipulators. *Proc. of the IEEE Int. Conf. on Decision and Control.* Los Alamitos, CA: IEEE, pp. 2398–2403.

Ostrowski, J., Burdick, J., Lewis, A. D., and Murray, R. M. 1995. The mechanics of undulatory locomotion: The mixed kinematic and dynamic case. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 1945–1951.

Rao, A., Kriegman, D., and Goldberg, K. Y. 1995. Complete algorithms for reorienting polyhedral parts using a pivoting gripper. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 2242–2248.

Reznik, D., and Canny, J. 1998. A flat rigid plate is a universal planar manipulator. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE.

Rizzi, A. A., and Koditschek, D. E. 1992 (Nice, France). Progress in spatial robot juggling. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 775–780.

Rizzi, A. A., and Koditschek, D. E. 1993 (Atlanta, GA). Further progress in robot juggling: The spatial two-juggle. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 919–924.

Salisbury, J. K. 1987 (August, Santa Cruz, CA). Whole-arm manipulation. *Proc. of the Int. Symp. on Robot. Res.* Cambridge, MA: MIT Press.

Sawasaki, N., Inaba, M., and Inoue, H. 1989 (Tokyo). Tumbling objects using a multifingered robot. *Proc. of the 20th Int. Symp. on Industrial Robots and Robot Exhibition*, pp. 609–616.

Schaal, S., and Atkeson, C. G. 1993 (Atlanta, GA). Open-loop stable-control strategies for robot juggling. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 913–918.

Schittkowski, K. 1986. *QLD: A Fortran Code for Quadratic Programming Users Guide.* Universität Bayreuth, Germany: Mathematisches Institut.

Schneider, J. G., and Brown, C. M. 1993 (Atlanta, GA). Robot skill learning, basis functions, and control regimes. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 403–410.

Shiller, Z., and Dubowsky, S. 1991. On computing the global time-optimal motions of robotic manipulators in the presence of obstacles. *IEEE Trans. Robot. and Automat.* 7(6):785–797.

Shin, K. G., and McKay, N. D. 1985. Minimum-time control of robotic manipulators with geometric path constraints. *IEEE Trans. Automat. Control* 30(6):531–541.

Sontag, E. 1993. Gradient techniques for systems with no drift: A classical idea revisited. *Proc. of the IEEE Int. Conf. on Decision and Control.* Los Alamitos, CA: IEEE, pp. 2706–2711.

Sørdalen, O. J., Nakamura, Y., and Chung, W. J. 1994. Design of a nonholonomic manipulator. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 8–13.

Spong, M. W. 1994. Swing-up control of the acrobot. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 2356–2361.

Sussmann, H. J. 1983. Lie brackets, real analyticity, and geometric control. In Brockett, R. W., Millman, R. S., and Sussmann, H. J. (eds.): *Differential Geometric Control Theory.* Boston, MA: Birkhauser.

Sussmann, H. J. 1987. A general theorem on local controllability. *SIAM J. Control Optimization* 25(1):158–194.

Sussmann, H. 1993. A continuation method for nonholonomic path-finding problems. *Proc. of the IEEE Int. Conf. on Decision and Control.* Los Alamitos, CA: IEEE, pp. 2718–2723.

Sussmann, H. J., and Jurdjevic, V. 1972. Controllability of nonlinear systems. *J. Differential Equations* 12:95–116.

Suzuki, T., Koinuma, M., and Nakamura, Y. 1996. Chaos and nonlinear control of a nonholonomic free-joint manipulator. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 2668–2675.

Takashima, S. 1991 (Osaka). Control of gymnast on a high bar. *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Los Alamitos, CA: IEEE, pp. 1424–1429.

Trinkle, J. C., Ram, R. C., Farahat, A. O., and Stiller, P. F. 1993. Dexterous manipulation planning and execution of an enveloped slippery workpiece. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Los Alamitos, CA: IEEE, pp. 442–448.

Witkin, A., and Kass, M. 1988. Spa[illegible]ne constraints. *Comp. Graphics* 22(4):159–168.

Yen, V., and Nagurka, M. L. 1988. [illegible] suboptimal trajectory-planning problem for robotic [illegible]pulators. *ISA Trans.* 27(1):51–59.

Žefran, M., Desai, J., and Kumar, V. 1996. Continuous motion plans for robotic systems with changing dynamic behavior. In Laumond, J.-P., and Overmars, M. (eds.): *Algorithms for Robotic Motion and Manipulation*. Boston, MA: A. K. Peters.

Zhu, C., Aiyama, Y., Chawanya, T., and Arai, T. 1996. Releasing manipulation. *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Sys.* Washington, DC: IEEE, pp. 911–916.

Zumel, N. B., and Erdmann, M. A. 1994 (San Diego, CA). Balancing of a planar bouncing object. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 2949–2954.

Zumel, N. B., and Erdmann, M. A. 1996. Nonprehensile two-palm manipulation with nonequilibrium transitions between stable states. *Proc. of the IEEE Int. Conf. on Robot. and Automat.* Washington, DC: IEEE, pp. 3317–3323.