

Issues in Nonprehensile Manipulation

Kevin M. Lynch, *Northwestern University, Evanston, IL, USA*

This paper outlines geometric and algorithmic issues common to various types of nonprehensile manipulation and gives some results for planar dynamic manipulation.

1 Overview

Nonprehensile manipulation is manipulation without a form- or force-closure grasp. Examples include pushing, throwing, juggling, tapping, batting, and rolling (Mason [24]; Higuchi [11]; Bühler and Koditschek [6]; Erdmann [8]; Huang *et al.* [12]; Zumel and Erdmann [40]; Aiyama *et al.* [1]; Trinkle and Zeng [37]). In each of these examples, the robot takes advantage of the natural task dynamics to help control the motion of the part. Nonprehensile manipulation occupies the majority of the manipulation spectrum, comprising everything between situations where the robot exerts complete control to situations where the natural dynamics exert complete control. During a baseball throw, the ball is at first held firmly in the hand, then is allowed to roll off the fingers, and finally follows a free-flight trajectory determined by gravity and air resistance. The nonprehensile manipulation problem is to arrange the rolling motion on the fingers such that the release state will allow the ball to reach the goal state.

Nonprehensile manipulation allows a simple, low degree-of-freedom robot to control more part freedoms through relative motion (slipping, rolling, and free flight) and to cause motion outside of its kinematic workspace. One cost of using nonprehensile manipulation, relative to pick-and-place manipulation, is that motion planning and control are harder. A nonprehensile manipulation planner must reason about dynamics.

Pick-and-place is usually concerned only with kinematics, and perhaps static friction for grasping (Lozano-Pérez *et al.* [16]). Also, it is not obvious what tasks are achievable by nonprehensile manipulation. With pick-and-place, if the initial and goal configurations of the part are within the robot's kinematic workspace and there is a free path connecting these configurations, then the goal configuration is achievable.

Interesting geometric and algorithmic problems in nonprehensile manipulation include

- generalizing the notion of a robot's kinematic workspace to consider the dynamically accessible workspace of the manipulated object;
- developing controllability tests based on the geometry, friction coefficient, and mass distribution of the object (and kinematic and dynamic descriptions of the manipulator when appropriate);
- characterizing classes of manipulator-object systems that are controllable;
- using properties of the system dynamics to derive efficient and complete motion planners; and
- determining feedback laws that stabilize the motion plans.

We have studied many of these issues in quasistatic pushing in (Lynch and Mason [20]; Lynch [17]). This paper summarizes our recent work on controllability and motion planning for planar dynamic nonprehensile manipulation, connects this work to control of an underactuated manipulator, and introduces the notion of dynamic workspaces in nonprehensile manipulation.

2 Example Systems

To motivate our discussion of controllability and motion planning, we will use the example of dynamic non-prehensile manipulation of a planar rigid body \mathcal{B} in zero gravity. The body \mathcal{B} moves like a puck on an air table. Depending on the control inputs used, \mathcal{B} represents

1. a hovercraft or planar spacecraft with unilateral thrusters;
2. an object being pushed with point contact along its boundary;
3. the third link of a 3R planar underactuated robot with a passive third joint;
4. an object being manipulated by a one joint robot.

The only constraints on the hovercraft thrusters are that they are fixed in the body frame and provide unilateral thrust. In the case of dynamic pushing, we assume a two-degree-of-freedom point robot P which can push \mathcal{B} at any point on its closed, piecewise smooth perimeter Γ . If we think of Γ as a pin-hole with P captured inside of it, we get the underactuated 3R robot, where the pin-hole is the third joint. In the last case, a one-degree-of-freedom revolute robot pushes \mathcal{B} with point contact. See Figure 1.

3 Controllability

The configuration space of \mathcal{B} is $\mathcal{C} = SE(2)$, the set of planar positions and orientations, and its state space is the tangent bundle $T\mathcal{C}$. A coordinate frame Σ_b is attached to the center of mass of \mathcal{B} , and its configuration in an inertial frame Σ_w is given by $\mathbf{q} = (x_w, y_w, \phi_w)^T$. The state of \mathcal{B} is written $(\mathbf{q}, \dot{\mathbf{q}}) \in T\mathcal{C}$. We define the zero velocity section Z as the three-dimensional space of zero velocity states $(\mathbf{q}, \mathbf{0})$.

The manipulation system can be written as the nonlinear control system

$$(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) = X_0(\mathbf{q}, \dot{\mathbf{q}}) + \sum_{i=1}^n u_i X_i(\mathbf{q}, \dot{\mathbf{q}}), \quad (1)$$




Figure 1: Control systems with the same dynamics but different control inputs. (a) A planar hovercraft with unilateral thrusters. (b) Dynamic pushing with a two-degree-of-freedom point robot. (c) The third link of a 3R planar robot with a passive third joint. (d) An object being manipulated by a single joint robot.

where $X_0(\mathbf{q}, \dot{\mathbf{q}}) = (\dot{x}_w, \dot{y}_w, \dot{\phi}_w, 0, 0, 0)^T$ is the drift vector field, $\mathbf{u} = (u_1, \dots, u_n)^T$ is the control input, and $X_i(\mathbf{q}, \dot{\mathbf{q}})$ is a control vector field. For systems 1–3, we choose the unit mass to be the mass of \mathcal{B} and the unit distance to be the radius of gyration of inertia of \mathcal{B} , and control forces are fixed in the body frame. For these systems, $X_i(\mathbf{q}, \dot{\mathbf{q}}) = (0, 0, 0, f_{xi} \cos \phi_w - f_{yi} \sin \phi_w, f_{xi} \sin \phi_w + f_{yi} \cos \phi_w, \tau_i)^T$, where (f_{xi}, f_{yi}) is a unit force expressed in the body frame Σ_b ($f_{xi}^2 + f_{yi}^2 = 1$) and τ_i is the torque about the center of mass. We write $\mathbf{f}_i = (f_{xi}, f_{yi}, \tau_i)^T$ and $\mathcal{F} = \bigcup_{i=1 \dots n} \mathbf{f}_i$.

We obtain each of the four systems (hovercraft, dynamic pushing, underactuated robot, and one joint robot) by placing different constraints on the control vector \mathbf{u} and the forces \mathcal{F} .

Hovercraft. The thruster forces must be unilateral. We will also place the constraints that only one thruster can be active at a time with unit thrust magnitude. The control set is $U_1 = \{\mathbf{0}, (1, 0, \dots, 0)^T, (0, 1, \dots, 0)^T, \dots, (0, 0, \dots, 1)^T\}$. The hovercraft has n thrusters providing forces \mathcal{F} fixed

in the body frame Σ_b .

Dynamic pushing. This case is the same as the hovercraft case, with the added restriction that each contact force $\mathbf{f}_i \in \mathcal{F}$ must be realizable from point contact somewhere on the closed, piecewise smooth perimeter Γ of \mathcal{B} .

Underactuated robot. We choose $n = 2$, with $f_1 = (1, 0, 0)^T$ and $f_2 = (0, 1, \tau)^T$ —the forces act along the x -axis of Σ_b through the center of mass of \mathcal{B} , and along the y -axis of Σ_b with nonzero torque, respectively. The control set is $U_3 = [-1, 1] \times [-1, 1]$. These controls are realizable at all joint angles away from the singularity $\sin \theta_2 = 0$, where θ_2 is the angle of the second link with respect to the first link.

One joint robot. We set $n = 1$ corresponding to the single input, the angular acceleration $\ddot{\theta}$ of the one joint robot. The force \mathbf{f}_1 is normal to the robot’s surface at the point contact with the object, and therefore is not fixed in the body frame Σ_b as in the previous three systems. By an input transformation we can treat the control input as $\mathbf{u} = u \in U_4 = \{0, 1\}$.

A *feasible trajectory* for the body \mathcal{B} is a solution of (1) for a control function $\mathbf{u} : [0, T] \rightarrow U_i$, $i = 1, \dots, 4$ for the systems 1–4 above.

Modifying notation from Nijmeijer and van der Schaft [28], we define $R^V(\mathbf{q}_0, \dot{\mathbf{q}}_0, T)$ to be the reachable set from $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ at time $T > 0$ by feasible trajectories remaining in the neighborhood V of $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ at times $t \in [0, T]$. Define $R^V(\mathbf{q}_0, \dot{\mathbf{q}}_0, \leq T) = \bigcup_{0 \leq t \leq T} R^V(\mathbf{q}_0, \dot{\mathbf{q}}_0, t)$. Then (1) (or simply \mathcal{B}) is *small-time accessible from* $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ if $R^V(\mathbf{q}_0, \dot{\mathbf{q}}_0, \leq T)$ has nonempty interior on $T\mathcal{C}$ for any neighborhood V of $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ and all $T > 0$. \mathcal{B} is *accessible from* $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ if there exists a finite time T such that $R^{T\mathcal{C}}(\mathbf{q}_0, \dot{\mathbf{q}}_0, \leq T)$ has nonempty interior. \mathcal{B} is *small-time locally controllable from* $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ if $R^V(\mathbf{q}_0, \dot{\mathbf{q}}_0, \leq T)$ contains a neighborhood of $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ for any neighborhood V and all $T > 0$. We define \mathcal{B} to be simply *locally controllable from* $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ if there exists a $T > 0$ such that $R^V(\mathbf{q}_0, \dot{\mathbf{q}}_0, \leq T)$ contains a neighborhood of $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ for any neighborhood V . \mathcal{B} is *controllable from* $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$

if, for any $(\mathbf{q}_1, \dot{\mathbf{q}}_1) \in T\mathcal{C}$, there exists a finite time T such that $(\mathbf{q}_1, \dot{\mathbf{q}}_1) \in R^{T\mathcal{C}}(\mathbf{q}_0, \dot{\mathbf{q}}_0, T)$. The phrase “from $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ ” can be eliminated from each of these definitions if the condition applies at all $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$.

A second-order system can be locally controllable only on the zero velocity section Z . (For any initial nonzero velocity state $(\mathbf{q}, \dot{\mathbf{q}})$ and any $\epsilon > 0$, there is an open subset V containing $(\mathbf{q}, \dot{\mathbf{q}})$ and $(\mathbf{q} - \epsilon\dot{\mathbf{q}}, \dot{\mathbf{q}})$ such that the system cannot attain $(\mathbf{q} - \epsilon\dot{\mathbf{q}}, \dot{\mathbf{q}})$ without leaving V .)

3.1 Hovercraft

Partial results for the hovercraft case have been reported previously by Manikonda and Krishnaprasad [22] and Lewis and Murray [15]. Manikonda and Krishnaprasad [22] observed that the Hamiltonian dynamics of the hovercraft on the cotangent bundle $T^*\mathcal{C}$ are invariant to the Lie group \mathcal{C} , suggesting the study of the reduced dynamics on the three-dimensional quotient manifold $T^*\mathcal{C}/\mathcal{C}$. They showed that the reduced dynamics of a hovercraft with a single bilateral thruster (or two opposing unilateral thrusters) are controllable on $T^*\mathcal{C}/\mathcal{C}$ —the hovercraft is controllable on its velocity space $T\mathcal{C}/\mathcal{C}$.

Lewis and Murray [15] studied the set of reachable configurations for mechanical control systems starting at rest. For an initial configuration \mathbf{q} at zero velocity and neighborhood $V_{\mathcal{C}}$ of \mathbf{q} on \mathcal{C} , they define $R_{\mathcal{C}}^{V_{\mathcal{C}}}(\mathbf{q}, T)$ to be the set of reachable configurations (with any velocity) at time T by trajectories remaining in the configuration neighborhood $V_{\mathcal{C}}$. They call a system *small-time locally configuration controllable* if $R_{\mathcal{C}}^{V_{\mathcal{C}}}(\mathbf{q}, \leq T)$ contains a neighborhood of \mathbf{q} on \mathcal{C} for any neighborhood $V_{\mathcal{C}}$ and all $T > 0$. (This is a weaker condition than small-time local controllability, which requires the locally reachable set to be a neighborhood of $(\mathbf{q}, \mathbf{0})$ on the full state space $T\mathcal{C}$, not just the configuration space.) They showed that a hovercraft with two bilateral thrusters (or four unilateral thrusters) is small-time locally configuration controllable.

Here we give stronger results for the unilateral thruster case (Lynch [21]). For a hovercraft with one to three unilateral thrusters, we get a new property with

each additional thruster: one thruster yields small-time accessibility on the hovercraft's state space \mathcal{TC} ; two thrusters yield global controllability on \mathcal{TC} ; and three thrusters yield small-time local controllability on the zero velocity section Z .

Proposition 1 *The hovercraft is small-time accessible with a single thruster ($n = 1$) if and only if the line of action of the thruster does not pass through the center of mass.*

Proof: To test for small-time accessibility we examine the Lie algebra of the system vector fields X_0 and X_1 . Without loss of generality, assume the thruster is aligned with the y -axis of Σ_b ($\mathbf{f}_1 = (0, 1, \tau)^T$), yielding $X_1 = (0, 0, 0, -\sin\phi_w, \cos\phi_w, \tau)^T$. We define the Lie bracket vector fields $X_2 = [X_0, X_1]$, $X_3 = [X_1, [X_0, X_1]]$, $X_4 = [X_1, [X_0, [X_0, X_1]]]$, $X_5 = [X_1, [X_1, [X_0, [X_0, X_1]]]]$, $X_6 = [X_0, [X_1, [X_1, [X_0, [X_0, X_1]]]]]$. We find that

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

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 $\tau \neq 0$ (the line of action of the thruster must not pass through the center of mass). Therefore the system satisfies the Lie Algebra Rank Condition (LARC) and the system is small-time accessible (Hermann and Krener [10]). \square

A single thruster is never sufficient for controllability, as the hovercraft's angular velocity is not controllable. It turns out that two thrusters are sufficient.

Proposition 2 *The hovercraft is controllable with two thrusters ($n = 2$) if and only if the two thrusters provide torque of opposite signs ($\tau_1 > 0, \tau_2 < 0$).*

The proof of this is given in (Lynch [21]). The approach is to find maneuvers that render the hovercraft controllable on its velocity space (the quotient space \mathcal{TC}/\mathcal{C}) and controllable on the zero velocity section Z . These properties are sufficient to show global controllability.

Proposition 3 *The hovercraft is small-time locally controllable on the zero velocity section Z with three*

thrusters ($n = 3$) if their lines of action intersect at a single point (which is not the center of mass) and the linear components of $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3$ positively span the plane.

Proof: Consider the system (Lewis and Murray [15])

$$(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) = X_0(\mathbf{q}, \dot{\mathbf{q}}) + u_1 X_1(\mathbf{q}, \dot{\mathbf{q}}) + u_2 X_2(\mathbf{q}, \dot{\mathbf{q}}), \quad |u_i| \leq 1. \quad (2)$$

We will also need the bracket terms $X_3 = [X_0, X_1]$, $X_4 = [X_0, X_2]$, $X_5 = [X_1, [X_0, X_2]]$, $X_6 = [X_0, [X_1, [X_0, X_2]]]$.

Now we give some definitions necessary to apply a version of Sussmann's [36] sufficient condition for small-time local controllability to a system such as (2). 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 $\sum_{i=0}^n \delta_i(B)$. B is called a "bad" bracket if $\delta_0(B)$ is odd and $\delta_i(B)$ is even for all $i \in \{1, \dots, n\}$, and B is 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 LARC at \mathbf{p} and all "bad" brackets evaluated at \mathbf{p} are neutralized, then the system is small-time locally controllable at \mathbf{p} .

Assume $X_1 = (0, 0, 0, \cos\phi_w, \sin\phi_w, 0)^T$ and $X_2 = (0, 0, 0, -\sin\phi_w, \cos\phi_w, \tau)^T$. The force \mathbf{f}_1 acts through the center of mass along the x -axis of Σ_b and \mathbf{f}_2 acts along the y -axis with torque τ about the center of mass. The force lines intersect at a point C not at the center of mass. Calculating the brackets above, we find that $\det(X_1 \ X_2 \ X_3 \ X_4 \ X_5 \ X_6) = \tau^4$; the LARC is satisfied provided $\tau \neq 0$. 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 $[X_2, [X_0, X_2]]$. We have

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

$$[X_2, [X_0, X_2]] = (0, 0, 0, -2\tau \cos\phi_w, -2\tau \sin\phi_w, 0)^T$$

which 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 span a body-fixed

force/torque plane \mathcal{P} which is neither the $\tau = 0$ plane nor orthogonal to the $\tau = 0$ plane, and the controls $|u_i| \leq 1$ define a compact, convex subset of \mathcal{P} containing the origin in the interior (relative to \mathcal{P}). By Sussmann’s [36] Proposition 2.3, small-time local controllability for this system implies small-time local controllability for the bang-bang system with the extremal 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 \mathcal{P} , and Sussmann’s proposition indicates that the extremal forces alone are sufficient. Therefore, any set of control forces that *positively* spans the plane \mathcal{P} also yields small-time local controllability. Any three unilateral forces which intersect at C and positively span the (x, y) plane also positively span the force plane \mathcal{P} . \square

3.1.1 Shape derivatives in grasping and motion derivatives in local controllability

We can draw a parallel between form-closure grasping of a planar object and small-time local controllability of a planar body with unilateral thrusters. In particular, a first-order analysis of form-closure grasping, considering only the contact normals at the finger locations, indicates that four fingers are necessary to constrain a planar object (Reuleaux [30]). Similarly, a first-order analysis can be used to show that four unilateral thrusters are sufficient for small-time local controllability at zero velocity states. In both cases, the forces (from the frictionless fingers or the thrusters) must positively span the body-fixed three-dimensional force-torque space (f_x, f_y, τ) . Four is the minimal number of vectors that can positively span a three-dimensional space.

Higher-order analyses can be used to reduce the number of fingers and thrusters needed. In Proposition 3 we proved that by using Lie brackets, *i.e.* motion derivatives, the number of thrusters required for small-time local controllability is reduced to three. Rimón and Burdick [31] have recently shown that contact curvature, *i.e.* shape derivative, may be used to decrease the number of fingers required for complete immobilization. In particular, they showed that three point fin-




Figure 2: Three finger grasps of two different objects. By a first-order mobility analysis, the only free motion for both objects is translation along the x -axis. (a) The bottom finger contacts the curve $y = x^4$ at $x = 0$. We have $y'(0) = y''(0) = y^{(3)}(0) = 0$ and $y^{(4)}(0) = 24$. A fourth-order test is required to determine that the object locally curves in the $+y$ direction, away from the finger, and motion along the x -axis is possible. (b) The bottom finger contacts the curve $y = -x^4$ at $x = 0$. A fourth-order test shows that motion along the x -axis is impossible, and the object is completely immobilized.

gers are sufficient for form-closure provided their contact normals positively span the plane and intersect at a single point C , and the centers of curvature of the perimeter of the body at the contacts do not all lie at C ([31], Proposition 4.1).

Higher-order shape derivatives (not just curvature) can be used in determining form-closure (Figure 2). Bracket terms of higher-order than those used in the proof of Proposition 3 do not create linearly independent vector fields.

3.2 Dynamic pushing

Dynamic pushing is similar to the hovercraft, except the control forces must be realizable by point contact with the perimeter Γ of \mathcal{B} . We assume a two-degree-of-freedom point pushing robot, and we place no restrictions on the motion of the pusher.

We can use Proposition 2 to show that any object is controllable by a point robot pushing at a single point of contact on the object, provided friction is nonzero. See Figure 3.

Proposition 4 For any planar body \mathcal{B} with a closed, piecewise smooth curve Γ of available contact points,




Figure 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.*

there exists a pushing contact point on Γ such that the object is controllable 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 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. \square

Proposition 5 *Any planar object \mathcal{B} with a closed, piecewise smooth curve Γ of available contact points is 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.*

Remark. This property ensures that the object can be dynamically pushed to follow any planar path arbitrarily closely (at sufficiently slow speeds). This requires the point robot to switch between distant con-

tacts. Because switching contact points takes finite time, \mathcal{B} is simply locally controllable, not small-time locally controllable. Note that the entire robot-object system does not have this property—the robot must make large motions.

Proof: If Γ is not a circle, the set of frictionless forces that can be applied through Γ (the normals to Γ) positively span the three-dimensional force-torque space (f_x, f_y, τ) (Mishra *et al.* [25]; Markenscoff *et al.* [23]), and the object is locally controllable by Proposition 3. If Γ is a circle, then the frictionless forces through Γ intersect at C , the center of the circle. If the center of mass of \mathcal{B} is not located at C , then \mathcal{B} is locally controllable by Proposition 3. \square

Dynamic pushing is, in one sense, a more complete primitive for planar manipulation than pick-and-place using a form- or force-closure grasp. No form- or force-closure grasp exists for a frictionless disk—the disk cannot be rotated by grasping and turning. On the other hand, any frictionless disk not centered at its center of mass is locally controllable by dynamic pushing.

3.3 Underactuated robot

We define R_A to be a fully-actuated 3R robot operating in a horizontal plane, and R_U to be the same robot with a passive third joint. The configuration of the robot is $\Theta = (\theta_1, \theta_2, \theta_3)^T \in M = T^3$. The workspace of the robot is populated with obstacles which include their boundaries, yielding the closed configuration obstacle set $CO \subset M$. We define the collision- and singularity-free configurations of the robot $M_{free} = \{\Theta \in M \mid \sin \theta_2 \neq 0, \Theta \notin CO\}$. The robot is confined to one of the two open sets $\theta_2 \in (0, \pi)$ or $\theta_2 \in (-\pi, 0)$, the RIGHTY and LEFTY configuration spaces.

Proposition 6 *If there exists a trajectory in M_{free} from $(\Theta_0, \mathbf{0})$ to $(\Theta_1, \mathbf{0})$ for the fully-actuated robot R_A , then there exists a trajectory in M_{free} from $(\Theta_0, \mathbf{0})$ to $(\Theta_1, \mathbf{0})$ for the robot R_U with the third joint passive.*

Proof: This result follows directly from Proposition 3. At any point of M_{free} , the robot can apply forces in any direction through the joint of the third link, the conditions of Proposition 3 are satisfied, and the third link

is small-time locally controllable at zero velocity. Any neighborhood of a configuration $\mathbf{q} \in \mathcal{C}$ of the third link maps to a neighborhood of $\Theta = K^{-1}(\mathbf{q}) \in M$, where K^{-1} is the inverse kinematic mapping. This mapping is smooth and one-to-one on the LEFTY and RIGHTY configuration spaces, and $K^{-1}(\mathbf{q}) \in \text{int}(K^{-1}(V(\mathbf{q})))$, where $V(\mathbf{q})$ is any neighborhood of \mathbf{q} , so small-time local controllability of the third link on Z implies small-time local controllability of the 3R robot on its zero velocity state space $\{(\Theta, \mathbf{0}) | \Theta \in M_{free}\}$. Therefore, if there is a trajectory for R_A from $(\Theta_0, \mathbf{0})$ to $(\Theta_1, \mathbf{0})$ which remains in M_{free} , then R_U can follow the associated path arbitrarily closely. The motion of R_U may be much slower, since in general it must remain near zero velocity to ensure that the path can be followed closely. \square

3.4 One joint robot

The results on dynamic pushing address the theoretical capabilities of dynamic nonprehensile manipulation from the viewpoint of the object alone. Just as important, however, are properties of the manipulator which 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.

Consider the system of Figure 4. The manipulator \mathcal{M} is a single revolute joint, and the object \mathcal{B} is a unit mass rod in point contact with \mathcal{M} . The distance from the contact to the rod's center of mass is r and the rod's radius of gyration of inertia is $\rho > 0$. The rod represents an arbitrary polygon in vertex contact with the manipulator. Because the contact is nonprehensile, in any neighborhood of a manipulator-object state, the best we can hope for is small-time accessibility.

The configuration of the manipulator-object system is $\mathbf{q}' = (\mathbf{q}, \theta) \in \mathcal{C}'$, where θ is the angle of the revolute joint. We constrain the manipulator to stay in contact with the rod endpoint at all times, applying zero force (simply "following" the rod) or applying a unit force. Assuming the single link is thin, the three-dimensional submanifold of contact configurations is given by $\{\mathbf{q}' \in \mathcal{C}' \mid F(\mathbf{q}') = \cos \theta (y_w - r \sin \phi_w) + \sin \theta (r \cos \phi_w - x_w) =$




Figure 4: A one-degree-of-freedom revolute robot manipulating a rod.

0}. The equations of motion are

$$\ddot{x}_w = -u \sin \theta \quad (3)$$

$$\ddot{y}_w = u \cos \theta \quad (4)$$

$$\ddot{\phi}_w = -\frac{r}{\rho^2} u \cos(\phi_w - \theta), \quad (5)$$

where the contact force u is

$$\begin{aligned} u = & \rho^2 (\ddot{\theta} (-r \cos(\phi_w - \theta) + x_w \cos \theta + y_w \sin \theta) + \\ & \dot{\phi}_w^2 (-r \sin(\phi_w - \theta)) + \\ & \dot{\theta}^2 (-r \sin(\phi_w - \theta) - x_w \sin \theta + y_w \cos \theta) + \\ & \dot{\phi}_w \dot{\theta} (2r \sin(\phi_w - \theta)) + \\ & \dot{\theta} \dot{x}_w (2 \cos \theta) + \\ & \dot{\theta} \dot{y}_w (2 \sin \theta)) / (r^2 \cos^2(\phi_w - \theta) + \rho^2). \end{aligned} \quad (6)$$

Rearranging Equation 6, we get the form $\ddot{\theta} = \ddot{\theta}_{drift} + u \ddot{\theta}_{control}$, 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. Treating the force u as the control input, the control system on the state space TC' is $(\dot{\mathbf{q}}', \ddot{\mathbf{q}}') = X_0(\mathbf{q}', \dot{\mathbf{q}}') + u X_1(\mathbf{q}', \dot{\mathbf{q}}')$, where the drift and control vector fields are written

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

By examining the Lie algebra of these vector fields, we conclude that the rod is small-time accessible from generic states provided $r \neq 0$. If the robot can also release the part, the entire robot-object system is small-time accessible. The rod is not small-time accessible

if the contact point coincides with the pivot of the manipulator—the robot can neither follow the motion of the drifting rod nor apply a force.

3.5 Dynamic workspaces

We would like to characterize the set of robot-object states reachable from another robot-object state. With pick-and-place manipulation, the reachable object states are limited by the reachable configurations of the end-effector, the “kinematic workspace.” When we include dynamics, however, the robot is capable of much more. Consider an object thrown to the robot from outside its kinematic workspace. Although the robot cannot initially affect the motion of the object, we know that at some point the robot will be able to contact it and, perhaps, throw it back outside of its kinematic workspace. Given an initial state of the robot-object system, we might define its “dynamic workspace” to be the set of reachable states. This dynamic workspace is a function of the robot, the object, and the initial state.

We illustrate the idea with an example. The two-degree-of-freedom robot is a point which can move freely inside a circle of radius R centered at $(0, 0)$. The object being manipulated is a frictionless disk of radius r which is not centered at its center of mass. For simplicity, the configuration of the disk $\mathbf{q} = (x_w, y_w, \phi_w)$ is measured at its center. By Proposition 1, we have small-time accessibility provided the robot can contact the disk and act as a single thruster. Therefore the disk’s small-time accessible workspace is defined by $(x_w^2 + y_w^2)^{1/2} < R + r$. For local controllability at zero velocity states, Proposition 3 implies that the robot must be able to contact the object at three points on the disk which are not confined to any closed half-circle. (Note that the robot does not have to be able to contact the disk at all points on its perimeter.) For $R > r$, this condition is satisfied by $(x_w^2 + y_w^2)^{1/2} < (R^2 - r^2)^{1/2}$, which defines the disk’s locally controllable workspace. The point robot can maneuver the disk to follow any path inside this workspace arbitrarily closely. If $r \geq R$, this workspace vanishes. See Figure 5.

Now assume a nonzero friction coefficient μ and friction angle α ($\mu = \tan \alpha$) between the robot and the

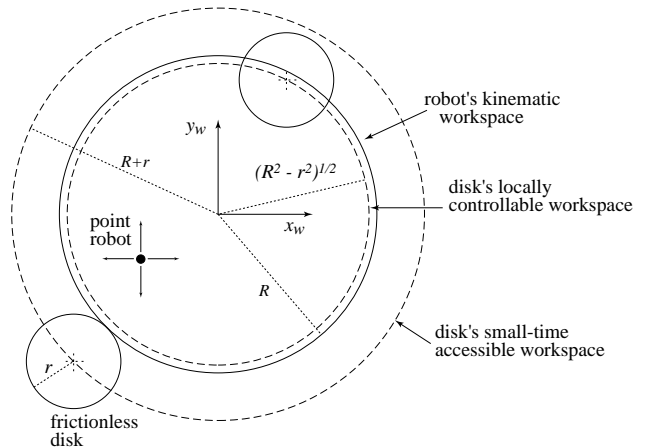


Figure 5: *The point robot, which can move freely in a circle of radius R , manipulates a frictionless disk (not centered at its center of mass) of radius r . The disk is small-time accessible provided its center stays in a circle of radius $R + r$, and it is locally controllable provided its center stays in a circle of radius $(R^2 - r^2)^{1/2}$. Disks are shown near the boundaries of the small-time accessible and locally controllable workspaces.*

disk. (The center of mass can now lie at the center of the disk.) With a nonzero friction coefficient, local controllability can be obtained from contacts confined to a closed half-circle (see Figure 6). As the friction angle α increases, the overlap between the robot’s workspace and the disk required for local controllability decreases—the disk is locally controllable if $(x_w^2 + y_w^2)^{1/2} < (R^2 - r^2 \cos \alpha)^{1/2} + r \sin \alpha$. As α approaches $\pi/2$ (infinite friction), the locally controllable workspace approaches the small-time accessible workspace, which is independent of the friction coefficient.¹

Finally, we have accessibility if the disk can reach a full-dimensional subset of its state space, not necessarily in small-time. The workspace from which the disk is accessible is the set of states that will carry the disk into its small-time accessible workspace. An example

¹We assume that the robot can apply any force in the friction cone. In some cases this may result in “pulling”—the robot moves in a direction with a negative component in the contact normal (Lynch and Mason [18]).

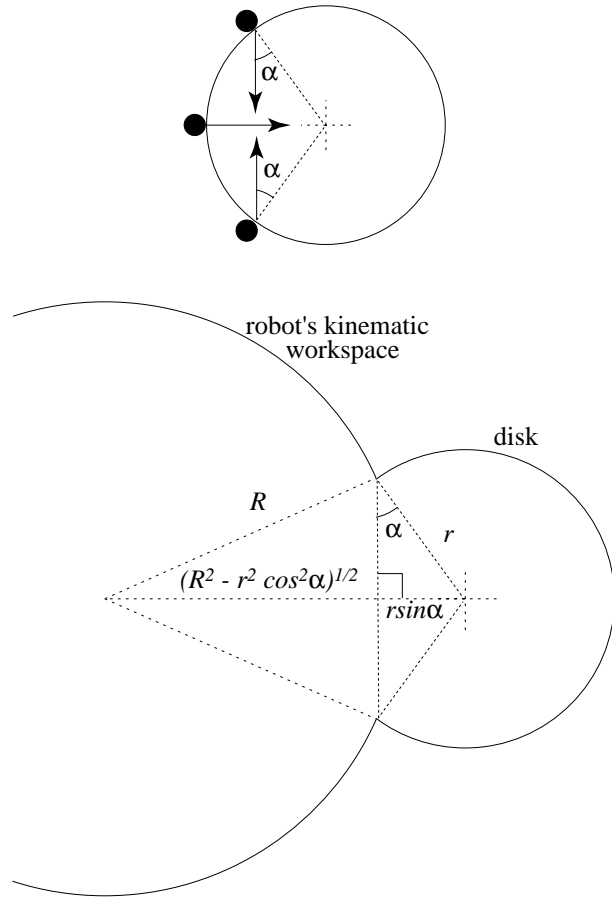


Figure 6: Top: At the friction angle α , forces through the three robot contact points pass through a single point and positively span a half-plane. At any greater value of friction, the condition of Proposition 3 is satisfied, and the disk is locally controllable. Bottom: Determining the minimum overlap of the disk and the robot workspace that allows the disk to be locally controllable by robot pushing with a friction angle α .

is shown in Figure 7.

We call a system *dynamically singular* at any state outside the small-time accessible workspace. The motion directions of the system are locally constrained to a lower-dimensional submanifold of the state space.

Obtaining a precise characterization of the accessible state space of an object during nonprehensile manipulation appears to be a very difficult problem in general.

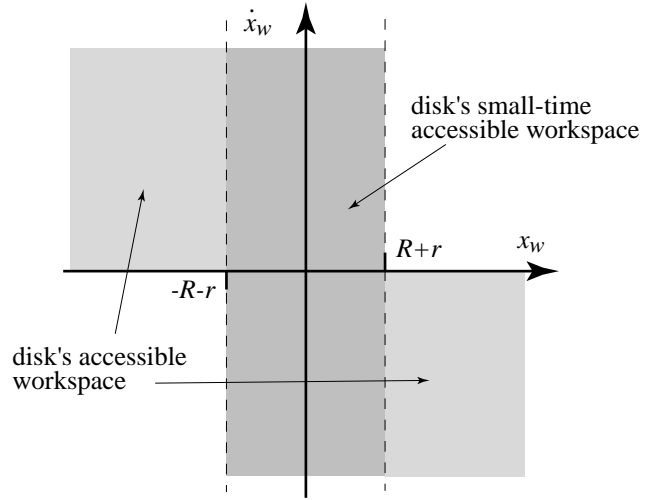


Figure 7: The (x_w, \dot{x}_w) state space for a disk centered at its center of mass, with $y_w = \dot{y}_w = 0$, ϕ_w and $\dot{\phi}_w$ arbitrary, and nonzero contact friction. The disk is small-time accessible in the band defined by $-R - r < x_w < R + r$, and the disk is accessible from the other shaded regions. The disk drifts into the small-time accessible band.

An approximate representation can be found by sampling the possible controls, simulating the system, and plotting the trajectories of the object.

4 Motion Planning

Nonprehensile manipulation is characterized by non-holonomic equality and inequality constraints. The former arise from using low degree-of-freedom robots to manipulate objects with more degrees-of-freedom; the latter arise from unilateral contact, friction constraints, and actuator constraints. We can distinguish between three basic approaches to the motion planning problem for nonholonomic systems (see also Popa and Wen [29]):

- 1) Control-theoretic algorithms make significant use of the structure of the system. Examples include systems without drift (Lafferiere and Sussmann [13]), systems which can be put into chained or differentially flat form (Murray *et al.* [26]; Murray *et al.* [27]), and Caplygin systems (Bloch *et al.* [4]). These approaches typically do not incorporate nonholonomic inequality

constraints (as arise in nonprehensile contact) or configuration constraints such as obstacles and joint limits.

2) Search-based algorithms sample the control space and can easily incorporate nonholonomic equality, inequality, and configuration constraints. Straightforward search can be quite general, requiring little information about the system, but as a result may suffer from high computational complexity. Information about the structure of the system can be used to perform a principled discretization of the control space (based on a critical decomposition or extremal controls), resulting in more efficient planners (*e.g.*, the shortest path information used in the mobile robot path planner of (Laumond *et al.* [14])).

3) Gradient-descent algorithms attempt to use smoothness in the input-output map to speed convergence to a solution. Typically a finite-dimensional parameterization of the control is chosen (for instance using Fourier bases, polynomials, splines, or piecewise constant functions), and the problem is cast as a constrained nonlinear optimization problem. This approach suffers from the same problems common to all gradient-descent optimization: local optima, possibly slow convergence, and numerical instabilities (particularly when using finite-difference gradients). Gradient descent for the control of nonholonomic systems has been explored by many (see, for example, (Divelbiss and Wen [7]; Sussmann [35]; Popa and Wen [29]; Sontag [34]; Fernandes *et al.* [9]; Žefran *et al.* [39])).

These methods can be used in concert; in particular, gradient descent can be used to locally optimize a solution found by another method.

Below we briefly describe a search-based planner for the underactuated robot and a gradient-descent approach to motion planning for a one joint robot (Lynch and Mason [19]).

4.1 A search-based planner for the underactuated robot

We would like to find fast, collision-free trajectories from $(\Theta_{init}, \mathbf{0})$ to $(\Theta_{goal}, \mathbf{0})$ for the underactuated 3R robot. Because the robot is small-time locally

controllable, it can follow any path closely. Arbitrary paths require the robot to stay near zero velocity, however. To find paths that can be executed quickly, we prefer to minimize the use of Lie bracket motions. We describe a planner that uses primitive motions corresponding to $\mathbf{f}_1 = (1, 0, 0)^T$ and $\mathbf{f}_2 = (0, 1, \tau)^T$ ($X_1 = (0, 0, 0, \cos \phi_w, \sin \phi_w, 0)^T$ and $X_2 = (0, 0, 0, -\sin \phi_w, \cos \phi_w, \tau)^T$) and minimizes the number of switches between them, implicitly minimizing the use of Lie bracket motions. By bringing the link velocity to zero at the switches, X_1 and X_2 yield velocity directions fixed in the frame of the passive link: pure translation along the link and pure rotation about the *center of percussion* of the link with respect to the joint. The vector fields X_1 and X_2 are unique in this respect, and for the purpose of path planning, this property allows us to treat the system as a first-order system with velocity vector fields $X'_1 = (\cos \phi_w, \sin \phi_w, 0)^T$ and $X'_2 = (-\sin \phi_w, \cos \phi_w, \tau)^T$. This allows us to decouple the full trajectory planning problem into the computationally simpler problems of path planning and time-scaling the path according to the manipulator dynamics and actuator constraints.

The planner is a modification of a path planning algorithm for mobile robots (Barraquand and Latombe [3]) and robotic pushing (Lynch and Mason [20]). It is a simple best-first search which minimizes the number of switches between vector fields. Starting from the initial link configuration \mathbf{q}_{init} (and corresponding initial joint configuration Θ_{init}), the planner integrates forward along each of $+X'_1$, $-X'_1$, $+X'_2$, and $-X'_2$ for a time δt , yielding four new link configurations. Each new collision-free configuration \mathbf{q}_{new} (and corresponding joint angles Θ_{new}) is added to a search tree T and to a sorted list $OPEN$ of configurations in T whose successors have not yet been generated. Configurations in $OPEN$ are sorted by the number of switches in their paths. The first configuration in $OPEN$ is then expanded. This process continues until a path is found to a user-specified goal neighborhood $\mathcal{G}(\mathbf{q}_{goal})$ or until $OPEN$ is empty (failure). Note that the planner is not exact, as it only finds a path to a goal neighborhood.

For a small enough δt , this planner will find a so-




Figure 8: The MEL underactuated manipulator.

lution when one exists. The planner is *resolution-complete*.

Once a path is found, we can apply the algorithms proposed by Shin and McKay [32] or Bobrow *et al.* [5] to find the time-optimal time-scaling of each motion segment of the path. The concatenation of these trajectories yields the time-optimal trajectory following the path returned by the planner.

We implemented planner results on the Mechanical Engineering Laboratory underactuated manipulator shown in Figure 8. Nonlinear feedback control is used to stabilize the motion of the third link during the rotational and translational motion segments (Shiroma *et al.* [33]). Time-optimal trajectories cause actuators to saturate, allowing no margin for error-recovery in feedback control. In our experiments, times for each motion segment were chosen experimentally for fast but robust performance.

Figure 9 shows an 11 segment path successfully implemented on the robot. The obstacles shown in the figure have been grown by 2 cm in each direction over the real experimental setup, providing room for error in execution. Joint limits are $-1.17 \text{ rad} < \theta_1 < 1.17 \text{ rad}$ (where 0 is vertical on the page) and $0.5 \text{ rad} < \theta_2 < 1.95 \text{ rad}$.

4.2 Gradient-descent motion planning for the one joint robot

A one joint robot such as that of Figure 4 can control more degrees-of-freedom of an object by sequencing

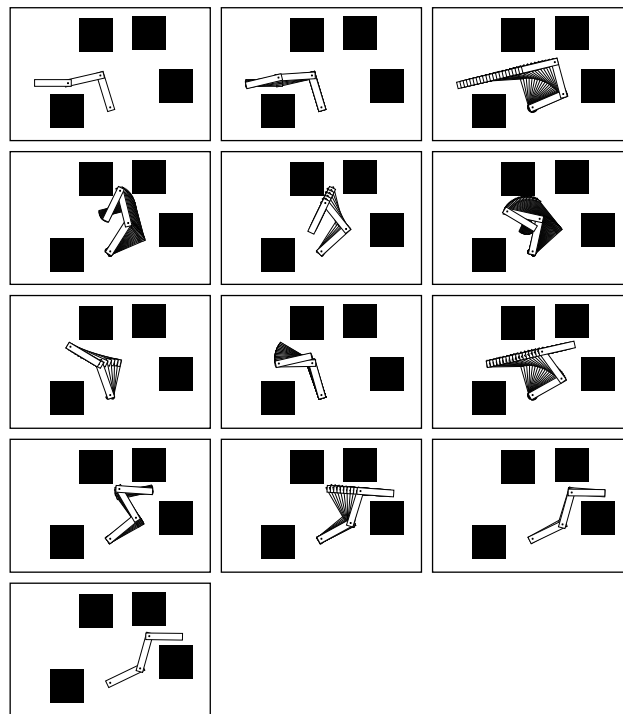


Figure 9: A path found by the planner using grown obstacles.

rolling, slipping, and free flight phases. We have previously described a motion planner for a 1 DOF robot performing rolling and throwing manipulation (Lynch and Mason [19]). This planner takes an initial guess for the trajectory of the robot and a sequence of manipulation phases (rolling, free flight, etc.) specified by the user, simulates the motion of the object, and using the results of the simulation, performs gradient-descent in the space of robot trajectories to find a motion which more nearly satisfies the goal state of the object and other constraints, such as friction inequalities and actuator limits. The manipulation problem is cast as a nonlinear optimization which is solved using sequential quadratic programming. The result is a robot motion, represented as a cubic spline, which takes the object to the goal state and locally optimizes a cost function. Several rolling and throwing tasks solved by the planner were successfully implemented on a real 1 DOF robot (see Figure 10). Žefran *et al.* [39] described a similar approach using nonlinear optimization for mo-




Figure 10: A 1 DOF robot motion to roll the square from one edge to another, found by the nonlinear optimization.

tion planning for systems with unilateral and changing dynamic constraints.

This system has much less apparent structure than the underactuated robot, so the gradient-descent planner assumes only smoothness in the input-output map. A major limitation of this planner relative to the planner for the underactuated robot is that it is not complete—it may not find a solution when one exists. Also, because the set of achievable tasks is not obvious, it is not clear when the planner has failed and when the desired goal state is unreachable. This problem points out the need for a better understanding of the dynamic workspace of the robot-object system. Also, the presence of inequality constraints due to unilateral contact and bounded friction exacerbates convergence and local minima problems. It appears that this gradient-

descent approach is better suited for local tuning of a rough solution than global motion planning.

5 Conclusion

Nonprehensile manipulation is a rich source of geometric problems. Examples include determining controllability (as discussed in this paper) and feedability (Akella *et al.* [2]; van der Stappen *et al.* [38]) based on the part geometry, contact friction, and mass distribution. These problems bear a resemblance to problems in robot grasping—characterizing graspable and ungraspable objects and determining the minimum number of fingers required for a grasp. The controllability and feedability properties are also closely bound up with characteristics of the manipulator, which may have fewer degrees-of-freedom than the object.

Nonprehensile manipulation also poses challenging problems for motion planning algorithms. We have described a gradient-descent motion planner which finds motion plans through a specified sequence of manipulation phases. Future motion planners should automatically determine a sequence of phases leading to the goal state. Also, the current motion planner makes very little use of the structure of the problem. Future work should be towards identifying and using this structure to derive complete motion planners. An example of using the structure of a nonprehensile manipulation problem is Erdmann’s [8] planner for two palm manipulation based on a critical decomposition of the control space.

Given the ubiquity of nonprehensile manipulation in everyday tasks and industrial applications such as vibratory parts feeding, a formal attack on the issues outlined in this paper should go a long way toward building a science of manipulation.

Acknowledgments

This paper describes joint work with Matt Mason at Carnegie Mellon, and Naoji Shiroma, Hirohiko Arai, and Kazuo Tanie at the Biorobotics Division, Mechanical Engineering Laboratory, Tsukuba, Japan. This work was supported by a grant from NSF at Carnegie

Mellon, a postdoctoral fellowship from the Science and Technology Agency of Japan, the Mechanical Engineering Laboratory, and Northwestern University.

References

- [1] Y. Aiyama, M. Inaba, and H. Inoue. Pivoting: A new method of graspless manipulation of object by robot fingers. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 136–143, Yokohama, Japan, 1993.
- [2] S. Akella, W. Huang, K. M. Lynch, and M. T. Mason. Planar manipulation on a conveyor with a one joint robot. In *International Symposium on Robotics Research*, pages 265–276, 1995.
- [3] J. Barraquand and J.-C. Latombe. Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles. *Algorithmica*, 10:121–155, 1993.
- [4] A. M. Bloch, M. Reyhanoglu, and N. H. McClamroch. Control and stabilization of nonholonomic dynamic systems. *IEEE Transactions on Automatic Control*, 37(11):1746–1757, 1992.
- [5] J. E. Bobrow, S. Dubowsky, and J. S. Gibson. Time-optimal control of robotic manipulators along specified paths. *International Journal of Robotics Research*, 4(3):3–17, Fall 1985.
- [6] M. Bühler and D. E. Koditschek. From stable to chaotic juggling: Theory, simulation, and experiments. In *IEEE International Conference on Robotics and Automation*, pages 1976–1981, Cincinnati, OH, 1990.
- [7] A. W. Divilbiss and J. Wen. Nonholonomic path planning with inequality constraints. In *IEEE International Conference on Decision and Control*, pages 2712–2717, 1993.
- [8] M. A. Erdmann. An exploration of nonprehensile two-palm manipulation using two zebras. In *Algorithms for Robotic Motion and Manipulation, WAFR 1996*, pages 239–254. A. K. Peters, Boston, MA, 1996.
- [9] C. Fernandes, L. Gurvits, and Z. Li. Near-optimal nonholonomic motion planning for a system of coupled rigid bodies. *IEEE Transactions on Automatic Control*, 30(3):450–463, Mar. 1994.
- [10] R. Hermann and A. J. Krener. Nonlinear controllability and observability. *IEEE Transactions on Automatic Control*, AC-22(5):728–740, Oct. 1977.
- [11] T. Higuchi. Application of electromagnetic impulsive force to precise positioning tools in robot systems. In *International Symposium on Robotics Research*, pages 281–285. Cambridge, Mass: MIT Press, 1985.
- [12] W. Huang, E. P. Krotkov, and M. T. Mason. Impulsive manipulation. In *IEEE International Conference on Robotics and Automation*, pages 120–125, 1995.
- [13] G. Lafferiere and H. Sussmann. Motion planning for controllable systems without drift. In *IEEE International Conference on Robotics and Automation*, pages 1148–1153, Sacramento, CA, 1991.
- [14] J.-P. Laumond, P. E. Jacobs, M. Taïx, and R. M. Murray. A motion planner for nonholonomic mobile robots. *IEEE Transactions on Robotics and Automation*, 10(5):577–593, Oct. 1994.
- [15] A. D. Lewis and R. M. Murray. Configuration controllability of simple mechanical control systems. *SIAM Journal on Control and Optimization*, 35(3):766–790, May 1997.
- [16] T. Lozano-Pérez, J. L. Jones, E. Mazer, and P. A. O’Donnell. *HANDEY: A Robot Task Planner*. The MIT Press, 1992.
- [17] K. M. Lynch. Locally controllable polygons by stable pushing. In *IEEE International Conference on Robotics and Automation*, pages 1442–1447, 1997.
- [18] K. M. Lynch and M. T. Mason. Pulling by pushing, slip with infinite friction, and perfectly rough surfaces. *International Journal of Robotics Research*, 14(2):174–183, Apr. 1995.
- [19] K. M. Lynch and M. T. Mason. Dynamic underactuated nonprehensile manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 889–896, 1996.
- [20] K. M. Lynch and M. T. Mason. Stable pushing: Mechanics, controllability, and planning. *International Journal of Robotics Research*, 15(6):533–556, Dec. 1996.
- [21] K. M. Lynch and M. T. Mason. Dynamic nonprehensile manipulation: Controllability, planning, and experiments, 1998. *International Journal of Robotics Research*, to appear.
- [22] V. Manikonda and P. S. Krishnaprasad. Controllability of Lie-Poisson reduced dynamics. Institute for Systems Research 57-59, University of Maryland, May 1997.

- [23] X. Markenscoff, L. Ni, and C. H. Papadimitriou. The geometry of grasping. *International Journal of Robotics Research*, 9(1):61–74, Feb. 1990.
- [24] M. T. Mason. Mechanics and planning of manipulator pushing operations. *International Journal of Robotics Research*, 5(3):53–71, Fall 1986.
- [25] B. Mishra, J. T. Schwartz, and M. Sharir. On the existence and synthesis of multifinger positive grips. *Algorithmica*, 2(4):541–558, 1987.
- [26] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.
- [27] R. M. Murray, M. Rathinam, and W. Sluis. Differential flatness of mechanical control systems: A catalog of prototype systems. In *ASME Int Mech Eng Congress and Expo*, 1995.
- [28] H. Nijmeijer and A. J. van der Schaft. *Nonlinear Dynamical Control Systems*. Springer-Verlag, 1990.
- [29] D. O. Popa and J. T. Wen. Nonholonomic path-planning with obstacle avoidance: A path-space approach. In *IEEE International Conference on Robotics and Automation*, pages 2662–2667, 1996.
- [30] F. Reuleaux. *The Kinematics of Machinery*. MacMillan, 1876. Reprinted by Dover, 1963.
- [31] E. Rimon and J. W. Burdick. New bounds on the number of frictionless fingers required to immobilize planar objects. *Journal of Robotic Systems*, 12(6):433–451, 1995.
- [32] K. G. Shin and N. D. McKay. Minimum-time control of robotic manipulators with geometric path constraints. *IEEE Transactions on Automatic Control*, 30(6):531–541, June 1985.
- [33] N. Shiroma, H. Arai, and K. Tanie. Nonlinear control of a planar free link under a nonholonomic constraint. In *International Conference on Advanced Robotics*, pages 103–109, 1997.
- [34] E. Sontag. Gradient techniques for systems with no drift: A classical idea revisited. In *IEEE International Conference on Decision and Control*, pages 2706–2711, 1993.
- [35] H. Sussmann. A continuation method for nonholonomic path-finding problems. In *IEEE International Conference on Decision and Control*, pages 2718–2723, 1993.
- [36] H. J. Sussmann. A general theorem on local controllability. *SIAM Journal on Control and Optimization*, 25(1):158–194, Jan. 1987.
- [37] J. C. Trinkle and D. C. Zeng. Prediction of the quasistatic planar motion of a contacted rigid body. *IEEE Transactions on Robotics and Automation*, 11(2):229–246, Apr. 1995.
- [38] A. F. van der Stappen, K. Goldberg, and M. H. Overmars. Geometric eccentricity and the complexity of manipulation plans. Submitted to *Algorithmica*, November 1996.
- [39] M. Žefran, J. Desai, and V. Kumar. Continuous motion plans for robotic systems with changing dynamic behavior. In J.-P. Laumond and M. Overmars, editors, *Algorithms for Robotic Motion and Manipulation*. A. K. Peters, Boston, MA, 1996.
- [40] N. B. Zumel and M. A. Erdmann. Nonprehensile two palm manipulation with non-equilibrium transitions between stable states. In *IEEE International Conference on Robotics and Automation*, pages 3317–3323, 1996.