Control of Nonprehensile Manipulation

Kevin M. Lynch   Todd D. Murphey

Mechanical Engineering Dept.
Northwestern University
Evanston, IL 60208 USA

Abstract. *Nonprehensile manipulation* is the process of manipulating a part without a form- or force-closure grasp. Without such a grasp, the part is free to roll, slide, or break contact with the robot(s) manipulating it. Controlling the motion of a part using nonprehensile manipulation becomes the challenging problem of controlling a dynamic system with equations of motion incorporating the robot and part geometry, friction and restitution laws, and changing dynamics due to changing contact states. Drawing from the work of others and our own previous work, in this paper we pose several open problems in the control of nonprehensile manipulation.

1 Introduction

Nonprehensile manipulation is the process of manipulating a part without a form- or force-closure grasp. Without such a grasp, the part is free to roll, slide, or break contact with the robot(s) manipulating it. Examples of nonprehensile manipulation are shown in Fig. 1.

Many nonprehensile manipulation tasks share common characteristics. First, nonprehensile manipulation often involves *unilateral controls*. The robot manipulates the part by contact forces. The controls may be unilateral because the robot can only push the part, not pull it. There are no kinematic equality constraints guaranteeing the actuators' connection to the part at all times. Indeed, this is part of the power of nonprehensile manipulation — the nature of the connection of the actuators to the part can be changed. Second, nonprehensile manipulation exhibits *nonsmooth dynamics*. The robot manipulates the part through a set of contacts. When one or more of the contacts changes its mode, say from sticking contact to slipping, or from slipping to no contact, the dynamics of the system changes in a nonsmooth manner. Some of these switches can be controlled, as when a finger pushing a part moves away and recontacts it at another location or when an impact is used for the purpose of control. Other switches are uncontrolled, as when the friction force between an object and a surface it is sliding on changes due to indeterminacy in the support force distribution. This kind of uncertainty is inherent in many manipulation systems. Finally, both *part geometry* and *robot geometry* are important. Since forces are applied to the part through contact, the set of available contacts afforded by the geometry of the part is critical in determining the controllability of the part and in devising a control law. The

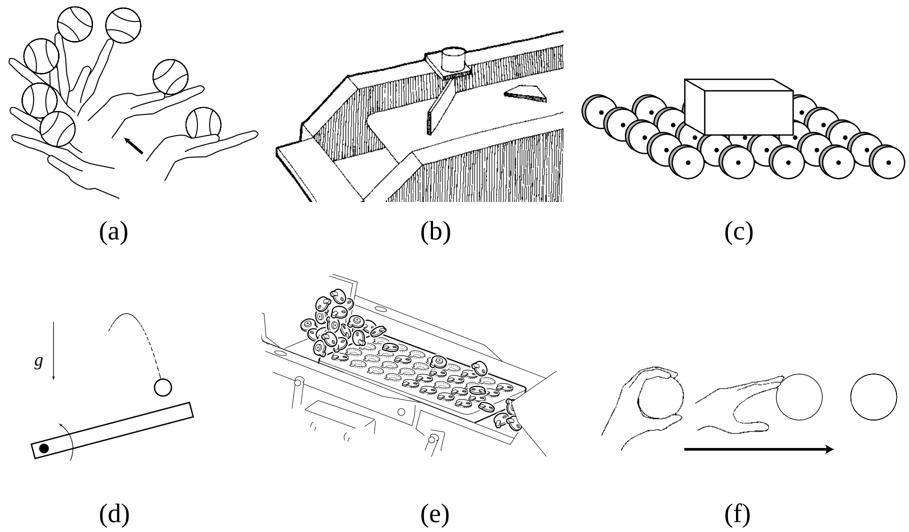

Fig. 1. Examples of nonprehensile manipulation. (a) A human juggler performs a "butterfly," dynamically rolling a ball from the palm to the back of the hand. (b) A part is carried on a constant-speed conveyor, and a rotating fence pushes the part back upstream while rotating it. By a series of pushing motions and conveyor drift, the part can be positioned and oriented. (c) A box is carried on an array of rolling wheels. The wheels may change orientation and speed. (d) A single-joint robot bat-juggles a puck in a gravity field. (e) The Sony APOS parts feeder orients parts in a tray of part-shaped depression. Parts wash over the vibrating tray, and parts which fall in a hole in the right orientation are trapped, while parts in a wrong orientation bounce out and continue into a recirculation bin. The vibration and hole shapes are designed to "juggle" parts into place in an open-loop fashion. (Figure taken from [33] and used with permission from MIT Press.) (f) Throwing a ball. Note that the initial carry phase is prehensile.

robot's geometry (and kinematics) is also important, as it determines how the robot can contact the part.

Some nonprehensile manipulation systems are *underactuated* and others are *overactuated*. In the former case, there are fewer robot actuators than part degrees-of-freedom. Examples include robot pushing and juggling. The ability to break contact and recontact leads to favorable controllability properties despite the underactuation. The latter case includes distributed manipulation by actuator arrays, where many actuators interact with a single part. In this case, the dimension of the input space can be quite large, while the dimension of the output space is small (typically three or six). For such a system the mechanics may be nonsmooth due to stick-slip phenomena associated with friction. Additionally, the mechanics of these systems are typically not well posed, in the sense that for a given set of inputs, solutions are not unique.

Control problems in nonprehensile manipulation include:

- *Defining sensible and testable notions of controllability.* The state space of a manipulation system naturally decomposes into a robot state space and a part state space, and we are typically interested in the local and global controllability (or "feedability" in the case of parts feeding) properties of the part. Some progress has been made for stratified systems in [15]. Ideally we would be able to derive tests for controllability in terms of the kinematics and dynamics of the robot, the geometry and mass properties of the part, and friction and restitution coefficients.
- *Reduction of quasistatic and dynamic nonprehensile manipulation control systems to kinematic systems for motion planning purposes.* Recent progress has been made in characterizing when (smooth) mechanical control systems can be treated as kinematic systems in [19,10,9]. Initial results are given in [28,34] for manipulation systems with uncertainty.
- *Trajectory generation.* Given the nonsmooth dynamics of the manipulation system, the problem is to find a set of controls to take the system to the goal state. This includes finding a sequence of contact states as well as continuous motion plans within each contact state [43,29].
- *Stabilizing a planned trajectory or equilibrium.* The problem of positioning and orienting a part (parts feeding) is to find a control law to stabilize an equilibrium.

Another important issue is reducing the sensing requirements needed to implement a successful control law. Full state sensing may not be required if we can identify equivalence classes of states where the same control will take the system closer to the goal state [14].

In this paper, we will consider a few different nonprehensile manipulation systems and list some open questions in control related to the issues raised above. The formulation in this paper is heavily influenced by previous work on rolling manipulation [5,17,31,32,12], juggling [6,27,7,8,39,40,11,27], pushing [28,26,1], control of nonsmooth manipulation systems [43,29,15,34,36], and distributed manipulation [3,22,34,36].

2 Definitions

In this section we propose some common definitions for nonprehensile systems. Nonprehensile manipulation systems come in so many shapes and flavors, however, that the utility of these definitions is limited. Nonetheless, these definitions provide a starting point for our discussion.

A nonprehensile manipulation system consists of a robot R and a part P. (The definition can easily be extended to multiple parts.) The configuration of the robot is written $q_R \in \mathcal{Q}_R$, and its dimension is m_R. Let z_R be the state of the robot, where $z_R = (q_R, \dot{q}_R)$ if the system is dynamic, and $z_R = q_R$ if the system is quasistatic. Let $\mathcal{Z}_R$ be the robot state space, and $\dim(\mathcal{Z}_R) = n_R$. The control applied to the robot is $u \in \mathcal{U}$.

Analogously, we can define the configuration of the part $q_P \in \mathcal{Q}_P$, its dimension $m_P = \dim(\mathcal{Q}_P)$, and its state $z_P \in \mathcal{Z}_P$ of dimension n_P. The configuration of the entire nonprehensile manipulation system is then written $q = (q_R, q_P) \in \mathcal{Q} = \mathcal{Q}_R \times \mathcal{Q}_P$, and the state is written $z = (z_R, z_P) \in \mathcal{Z} = \mathcal{Z}_R \times \mathcal{Z}_P$, where $\dim(\mathcal{Z}) = n_R + n_P = n$. Although the meaning of a "part" is intuitively obvious, we distinguish it from unactuated degrees-of-freedom of an underactuated robot by requiring an n-dimensional subset of the system state space where the control u has no influence on $\dot{z}_P$ (e.g., the robot can break contact with the part).

We call the system an *underactuated nonprehensile system* if $m_R < m_P$, and an *overactuated nonprehensile system* if $m_R > m_P$ and the evolution of z_P is affected by more than m_P actuators (i.e., the actuators are somehow connected to the part).

We will assume that a "part" is a rigid body. A frame $\mathcal{F}^P$ is attached to the center of mass of the rigid body part and aligned with the principal axes of inertia. The configuration of a spatial part is the matrix $g \in SE(3)$ representing the displacement of $\mathcal{F}^P$ relative to an inertial frame $\mathcal{F}^w$, where

$$g = \begin{bmatrix} R & p \\ 0 & 1 \end{bmatrix},$$

or simply $g = (p, R)$ for short, where $R \in SO(3)$ is a 3×3 rotation matrix and $p \in \mathbb{R}^3$ is the position of the origin of $\mathcal{F}^P$ in $\mathcal{F}^w$. The velocity of the part is an element ξ of $se(3) = \mathbb{R}^6$, the Lie algebra associated to $SE(3)$. We write $\xi = (v, \omega)$, where $\omega \in \mathbb{R}^3$ is the angular velocity and $v \in \mathbb{R}^3$ is the linear velocity of the body written in $\mathcal{F}^P$. The kinematic equations are

$$\dot{p} = Rv, \quad \dot{R} = R\widehat{\omega}, \quad \text{where } \widehat{\omega} = \widehat{\begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix}} = \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}. \tag{1}$$

Let f_P be the force applied by the robot to the part, written in the part frame $\mathcal{F}^P$. Let r_P be the vector to a point on the line of action of f_P, written in $\mathcal{F}^P$, so that the torque acting on the part is $\tau_P = r_P \times f_P$. Together, (f_P, τ_P) is called a generalized force or *wrench*. Let m be the part mass, $\mathbb{M} = \mathrm{diag}(m, m, m)$, and $\mathbb{J} = \mathrm{diag}(I_1, I_2, I_3)$ be the body-fixed inertia tensor.

The evolution of a nonprehensile system is of the form

$$h(q) \geq 0 \tag{2}$$

$$f_P = c(q, \dot{q}, u), \quad f_P h(q) = 0 \tag{3}$$

$$M(q_R)\ddot{q}_R + \dot{q}_R \Gamma(q_R)\dot{q}_R + G(q_R) = T(q_R)u - J^T(q) f_P \tag{4}$$

$$\begin{bmatrix} \dot{p} \\ \dot{R} \end{bmatrix} = \begin{bmatrix} Rv \\ R\widehat{\omega} \end{bmatrix}, \quad \begin{bmatrix} \mathbb{M} & 0 \\ 0 & \mathbb{J} \end{bmatrix} \begin{bmatrix} \dot{v} \\ \dot{\omega} \end{bmatrix} + \begin{bmatrix} \omega \times mv \\ \omega \times \mathbb{J}\omega \end{bmatrix} = \begin{bmatrix} f_P \\ \tau_P \end{bmatrix} + \begin{bmatrix} f_E(q_P, \dot{q}_P) \\ \tau_E(q_P, \dot{q}_P) \end{bmatrix} \tag{5}$$

$$\text{when } h(q) = 0, \dot{h}(q) < 0: \qquad (q^+, \dot{q}^+) = (q, \eta(q, \dot{q})), \tag{6}$$

where $h(q) \in \mathbb{R}$ is a "distance" function between the part and the robot, describing the unilateral configuration constraints that the part cannot penetrate the robot; $c(q, \dot{q}, u) \in \mathbb{R}^3$ gives the force applied to the part as a function of the system state (including the contact geometry), the control, and the friction law (the contact force is zero if the distance between the part and the robot is nonzero); (4) denotes the robot subsystem dynamics, and $J^T(q)$ describes how the force from the part affects the robot; (5) denotes the kinematics and dynamics of the part subsystem, where $(f_E, \tau_E) \in \mathbb{R}^6$ indicate any environmental forces such as gravity or frictional forces as a part slides over a stationary surface; and (6) denotes the impact law, which leaves the configuration of the system unchanged but changes the velocity discontinuously. Impacts between the part and the environment could also be modeled.

The coupling between the robot and part subsystems comes through the contact law c and the impact law η. Controllability and control of nonprehensile manipulation systems depends intimately on the friction and impact laws, and the vast majority of the complexity of controlling nonprehensile systems comes from them. In particular, the contact law c is nonsmooth due to geometry (changes in robot-part contact geometry as contacts are established and broken) and the friction law (switches from slipping to sticking contact). Each contact mode, which includes the set of contacts and their current state (slipping or sticking), leads to its own set of dynamic equations, and a switch in the contact mode results in a change in the equations. Some of these switches can be controlled, while others cannot. Thus nonprehensile systems inherit the difficulties of general hybrid control systems.

In Sections 3 and 4, we specialize the equations of motion to systems with much more structure.

2.1 Controllability Definitions

Let $R^{\mathcal{V}}(z)$ denote the set of reachable states from the initial state z by feasible trajectories remaining in $\mathcal{V} \subseteq \mathcal{Z}$. (In this definition, note that we are not concerned with time.) It is possible to define the standard controllability terms for the nonprehensile system on $\mathcal{Z}$, such as accessibility, controllability, etc. However, we are interested more in what can be done with the part, i.e., the subsystem (5), than the entire system. Let $\mathcal{V}_P \subset \mathcal{Z}_P$ be a neighborhood of z_P, and let $R_P^{\mathcal{V}_P}(z)$ be the reachable states of the part starting from the initial system state z by trajectories of the part confined to $\mathcal{V}_P$.

Definition 1. The part is *locally-locally accessible* from z if $\dim(R_P^{\mathcal{V}_P}(z)) = n_P$ for any $\mathcal{V}_P$ containing z_P in its interior.

Definition 2. The part is (globally) *controllable* from z if $R_P^{\mathcal{Z}_P}(z) = \mathcal{Z}_P$. For the rest of the paper, the term "controllable" refers to this global concept.

Definition 3. The part is *locally-locally controllable* from z if $z_P \in \text{int}(R_P^{\mathcal{V}_P}(z))$ for all neighborhoods $\mathcal{V}_P$ of z_P.

Definition 4. Given a set of initial states $\mathcal{W} \subset \mathcal{Z}$, we say the part is *controllable to* z_P from $\mathcal{W}$ if $z_P \in R_P^{\mathcal{Z}_P}(w)$ for all $w \in \mathcal{W}$. We will sometimes say the part is *feedable* to z_P from $\mathcal{W}$.

Since the definitions are not concerned with time, we borrow the "local-local" modifiers from Haynes and Hermes [16]. We could also define notions of configuration controllability and equilibrium controllability for a part [20].

3 Dynamic Underactuated Nonprehensile Manipulation

In this section we discuss the control of a rigid-body part by contact forces on its surface. The shape of the part, as well as the contact friction coefficient, determines the set of forces that can be applied to the part. Unlike grasping manipulation, where the robot makes multiple contacts with the part to establish a grasp which is capable of resisting all external wrenches, in this section we will focus on the case where the robot makes a single point of contact with the part at any given time, and study the resulting set of possible part trajectories. Let $d = 2, 3$ be the dimension of the part's configuration space. Since $m_P = \dim(SE(d)) = 6$ (3) for a spatial (planar) part and a point in space (the plane) has only three (two) degrees-of-freedom, the nonprehensile manipulation is necessarily underactuated.

We will first consider the part by itself, assuming a point robot that moves freely around the part. This is the subsystem (5) with (f_P, τ_P) as the control. We will then consider more realistic models of robot interaction with the part.

3.1 Part Geometry Only

Let ∂P be the surface (boundary) of the part P. Let r_P be the vector to a point on ∂P, and let $\hat{n}$ be the inward-pointing unit normal at this point, measured in $\mathcal{F}^P$. (We will assume that $\hat{n}$ is well-defined at each $r_P \in \partial P$, although the definition could be generalized to handle nondifferentiable points such as vertices.) Then, in the absence of friction, a contact at this point will give rise to a wrench on the part

$$f_P = k\hat{n}, \quad k \geq 0 \tag{7}$$

$$\tau_P = r_P \times f_P, \tag{8}$$

where $k \geq 0$ is the force magnitude (encoding the unilateral constraint) and f_P and τ_P are the force and torque measured in $\mathcal{F}^P$. In the presence of dry Coulomb friction with a friction coefficient μ, (7) becomes the constraint

$$f_P = k(\hat{n} + \gamma\hat{n}^{\perp}), \quad k \geq 0, \quad 0 \leq \gamma \leq \mu \tag{9}$$

where $\hat{n}^{\perp}$ is any unit vector satisfying $\hat{n}^T\hat{n}^{\perp} = 0$ (i.e., in the contact tangent plane). The constraint (9) defines a *friction cone*, where $\alpha = \tan^{-1}\mu$ is the friction cone half-angle (Fig. 2). If the contact is slipping, then

$$f_P = k(\hat{n} + \gamma\hat{n}^{\perp}), \quad k \geq 0, \quad \gamma = \mu \tag{10}$$

and $\hat{n}^{\perp}$ is opposite the relative slip direction in the tangent plane.

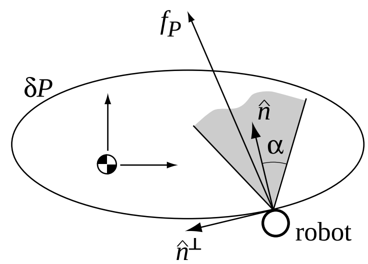

Fig. 2. The contact friction cone for a planar part. The contact force f_P must lie somewhere in the cone of half-angle $\alpha = \tan^{-1}\mu$.

For a point $r_P \in \partial P$, let $W(r_P)$ denote the set of wrenches that can be applied satisfying the friction cone constraint (9). Let $W(\partial P) = \bigcup_{r_P \in \partial P} W(r_P)$ be the set of all wrenches that can be applied to the part by point contact. Define the control system $(*)$ to be the part subsystem (5) with $(f_E, \tau_E) = 0$ and the control $(f_P(t), \tau_P(t)) : \mathbb{R} \to W(\partial P)$.

Theorem 1. *The control system* $(*)$ *is locally-locally controllable at zero velocity states* $(q_P, 0)$ *for* $d = 2$ *unless* $\mu = 0$ *and the part is a disk centered at its center of mass.*

In other words, the planar part is locally-locally controllable (Definition 3) if it can be pushed at all points along its boundary.

Now define the system $(\dagger)$ to be the part subsystem (5) with $(f_E, \tau_E) = 0$ and the control $(f_P(t), \tau_P(t)) : \mathbb{R} \to W(r_P)$ for a fixed point of contact r_P.

Theorem 2. *For any planar part* ($d = 2$) *and any* $\mu > 0$*, there exists a* $r_P \in \partial P$ *such that the control system* $(\dagger)$ *is controllable.*

In other words, any planar part is controllable (Definition 2), but not locally-locally controllable, by contact forces through a single point if the part is not frictionless. A sufficient condition for the choice of r_P to render $(\dagger)$ controllable for $d = 2$ is that the contact normal $\hat{n}$ pass through the center of mass.

Proofs for these theorems can be found in [29,25]. Both rely in part on the geometry of ∂P and the implied $W(\partial P)$. Theorem 1 is also satisfied by a finite set of contact points, i.e., $(f_P, \tau_P) \in W(r_{P_1}) \cup W(r_{P_2}) \cup \ldots \cup W(r_{P_k})$, and both theorems are satisfied by choosing a discrete set of wrenches belonging to each $W(r_{P_i})$ for the finite set of contacts $\{r_{P_1}, \ldots, r_{P_k}\}$. Therefore, the discrete set of contact forces can be considered unilateral jet thrusters arranged along the boundary of the part, with the added constraint that only one of the thrusters can be active at a time. Unilateral thruster control of a 2D or 3D spacecraft may be more intuitive than pushing a part on its boundary.

These results indicate the possibility of controlling a rigid-body part by point contact along its boundary, ignoring how these contacts may be established by a robot. Some open questions include:

Problem 1. Complete characterization of controllability and local-local controllability for the systems $(*)$ and $(\dagger)$ for spatial bodies ($d = 3$). A related problem is finding minimum unilateral thruster configurations for controlling spatial rigid bodies.

Problem 2. The proof of (global) controllability for the system $(\dagger)$ is a laborious near-constructive proof. On the other hand, it is not difficult to prove global controllability for systems (a) satisfying the Lie algebra rank condition (LARC), (b) subject to a phase-volume-conserving drift vector field, and (c) evolving on a compact configuration space (see, e.g., [21,18]). For nonprehensile (or jet thruster) manipulation of a rigid body, it is easy to establish the LARC, and the drift field is phase-volume-conserving. The configuration space is not compact, however; $SE(d)$ is a semi-direct product of a noncompact space $\mathbb{R}^d$ and a compact space $SO(d)$. We are not aware of general (global) controllability results for such systems, but we expect that the interaction between the compact and non-compact subsystems can be exploited in a systematic manner.

Problem 3. The trajectory generation problem is to find thrusts (wrenches) as a function of time yielding a feasible trajectory between the initial state and the goal state.

Problem 4. Imagine a controllable body which is not locally-locally controllable. What are the closed orbits achievable by this system? Although it is not possible to stabilize an equilibrium, is it possible to stabilize an orbit? Describe the control law. This question is relevant to trajectory generation and feedback control of severely underactuated space and underwater vehicles, with or without unilateral thrust constraints.

Another version of the control problem of the part subsystem (5) is pushing a planar body over a frictional support surface. In this case, the environmental forces (f_E, τ_E) are significant and uncertain due to the unknown distribution of support over the area of the body. A quasistatic version of this problem is studied in [28].

3.2 With Manipulator Constraints

Instead of manipulating a planar part at arbitrary points along its perimeter, imagine manipulating the part with a robot arm with a single revolute joint, as shown in Fig. 3. For the purposes of the example, let the part be a polygon. If the part remains rigidly attached to the robot, its accessible state space is two-dimensional and is parameterized by the robot state $(\theta, \dot{\theta})$. If the part can also roll relative to the robot with sticking at the contact point, it can achieve two more motion directions in its state space, parameterized by $(\phi, \dot{\phi})$. If the contact can also independently slip on the robot, the part can achieve two more motion directions, parameterized by the contact point on the robot arm and its rate of change. Thus it may be possible for the robot to transfer the part to an open subset of its six-dimensional state space $\mathcal{Z}_P$ without even breaking contact with the part. The robot can then throw the part, increasing the volume of the reachable state space. More interestingly, the robot could break and recontact the part to increase the reachable state space, or bat-juggle the part (Section 3.3).

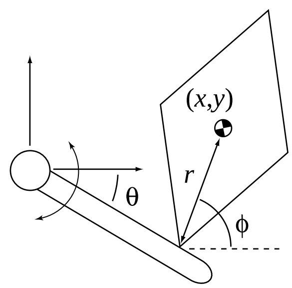

Fig. 3. A one-joint arm manipulating a polygon at a vertex.

Let's study how the system of Fig. 3 can be expressed in our canonical form (2 – 6). We can express four different systems depending on the contact mode: the robot either breaks contact or stays in contact with the part, and if in contact, the contact either sticks, slips left, or slips right. Which of these systems actually occurs can be determined by examining the current state and control and incorporating the friction law (except in cases of ambiguity and inconsistency, which may occur due to the Coulomb friction law [33]). Sticking contact implies a kinematic equality constraint, while slipping contact implies a force equality constraint (the force must lie on the appropriate edge of the friction cone).

We will assume that the robot moves so as to stay in contact with the vertex of the polygon, and we will assume zero friction at the contact. A slipping contact model with dry Coulomb friction is similar, except in this case the contact force lies along a friction cone edge, not the contact normal.

Let the mass of the part be m, its inertia about the center of mass be I_P, and the robot inertia about the pivot be I_R. We have $q_R = \theta$, $q_P = (x, y, \phi)$, and $q = (q_R, q_P)$. For simplicity in this example, we will break with the notation of (2 – 6) by expressing the contact wrench (f_P, τ_P) and the part velocity and acceleration in a frame fixed to the part center of mass but always aligned with the inertial frame fixed to the pivot of the robot. Except for this, Equations (11 – 14) correspond directly to Equations (2 – 5) in the canonical form:

$$h(q) = \cos\theta(y - r\sin\phi) + \sin\theta(r\cos\phi - x) \geq 0 \quad (11)$$

$$f_P = \begin{bmatrix} f_{Px} \\ f_{Py} \end{bmatrix} = k \begin{bmatrix} -\sin\theta \\ \cos\theta \end{bmatrix}, \quad \tau_{Pz} = -f_{Py} r\cos\phi + f_{Px} r\sin\phi \quad (12)$$

$$I_R\ddot{\theta} = u - J^T(q) f_P, \quad J^T(q) = (r\sin\phi - y, x - r\cos\phi) \quad (13)$$

$$\begin{bmatrix} m & 0 & 0 \\ 0 & m & 0 \\ 0 & 0 & I_P \end{bmatrix} \begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{\phi} \end{bmatrix} = \begin{bmatrix} f_{Px} \\ f_{Py} \\ \tau_{Pz} \end{bmatrix}. \quad (14)$$

Since we assume that contact is maintained, we have $h(q) = \dot{h}(q) = \ddot{h}(q) = 0$. It is the last of these that is used in solving the equations of motion. If in solving the equations of motion we find that the contact force magnitude k becomes negative, we know that the "maintain contact" assumption has been violated.

The exact equations of motion for even this simple system are complex, so we omit them (see [29]). The robot control torque u and the contact force magnitude k are related by

$$u = \alpha_1(z) + k\alpha_2(q),$$

where $\alpha_1(z)$ is the torque needed for the robot to stay in contact with the moving rod and $\alpha_2(q)$ is the extra torque needed to provide one unit of contact force to the rod. With a feedback transformation to use k as the control, the part dynamics can be written in the form

$$\dot{z}_P = g_0(z_P) + kg_1(z_P),$$

where $g_0(z_P) = (\dot{x}, \dot{y}, \dot{\phi}, 0, 0, 0)^T$, $g_1(z_P) = (0, 0, 0, -\frac{\sin\theta}{m}, \frac{\cos\theta}{m}, -\frac{r\cos(\phi-\theta)}{I_P})^T$, and $\theta = \mathrm{atan2}(y - r\sin\phi, x - r\cos\phi)$. An examination of the Lie algebra of $\{g_0, g_1\}$ readily shows that the part is locally-locally accessible (Definition 1) at generic states when the contact point is not at the robot's pivot. In other words, the robot can manipulate the part to a full six-dimensional subset of its state space. This is a desirable property, but a limited description of the capability of the system. Some relevant open problems are listed below.

Problem 5. Controllability of rolling manipulation has been heavily studied when the rolling velocities are directly controlled (a kinematic system). Controllability of dynamic nonprehensile rolling systems, where the body rolls on a single actuated surface, has received less attention (but see, e.g., [30,17,12]). If one body is free to roll on an actuated surface, what are the conditions for local and global controllability of the body? How much actuation does the surface need? How do we do motion planning? Can we derive feedback laws to stabilize equilibria or trajectories? Stabilizing rolling motion is inherently easier than stabilizing slipping motions, as slip depends critically on the friction coefficient, which may be unknown or varying.

Problem 6. In general, a part may go through a sequence of rolling, slipping, and free flight phases during dynamic nonprehensile manipulation. What are the reachable states for the part from a given initial state? Can we find a sequence of phases, and smooth motion plans within each phase satisfying the phase transition conditions at the switches, that will transfer the part to a desired state? Some early work in this direction is reported in [43,29].

Problem 7. The "dexterous kinematic workspace" of a robot arm is usually defined as the set of configurations a gripper can reach with any orientation. A "dynamic workspace" for a robot-part system might describe the set of reachable part states from a given initial system state. In particular, the robot can send the part to points outside the kinematic workspace by throwing.

3.3 Juggling

Bat-juggling is the process of manipulating a part by a series of impacts with the part moving in ballistic flight in between (e.g., Fig. 1(d)). A typical goal of a juggling system is to stabilize a desired limit cycle. Brogliato and Zavala-Rio [6] have defined the term "controllable through the impacts" to describe juggling controllability of a part, and Spong [41] considered the controllability of a batted air hockey puck. Control laws for juggling planar pucks and balls in space, and experimental implementations, are described in [7,8,39,40,11,27]. Brogliato and Zavala-Rio [6] outline a general framework for studying juggling systems. Here we simply mention some relevant problems.

Problem 8. Most work on juggling has focused on parts modeled as point masses. What about parts with non-trivial geometry? Are such parts controllable? Can we stabilize limit cycles? How much feedback is needed to do so? How do we plan a sequence of impacts to take the part to a desired state? What is the effect of different friction and impact laws?

Problem 9. For a given part of non-trivial geometry, is it possible to design an open-loop robot motion, as well as the robot shape, so that the impacts will cause the part to converge from a large basin of attraction to a desired resting state relative to the robot? This is essentially the APOS parts feeder (Fig. 1(e)) design problem for a single part.

4 Distributed Manipulation

Distributed manipulation is an area of robotics dedicated to the consideration of systems with many inputs affecting a part. One of the simplest examples, and the one we will consider here, is that of a planar array of actuators, like the one depicted in Fig. 1(c). If one places a part upon such an array, each of these actuators can exert a force on the part, together creating a net force and torque on the part. Therefore, using many actuators, one can control the location of the part in $SE(2)$. Hence, we can consider this system to be *over-actuated*, in the sense that there are many more inputs than outputs. It can be difficult to model these systems, partly because the underlying mechanics can be both nonsmooth and nondeterministic. There have been several attempts to address the issue of overactuation. The work of [4,13,23,24,42] presents a methodology where the forces produced by the individual actuators are idealized as a force field acting on the part. A different modeling methodology used in previous work by the second author in [34,36,37] has suggested both analytically and experimentally that a "quasistatic" approach may work well, and has the advantage of being first order. This is the approach we will describe here.

We should comment here that in some special cases, the feedback control problem for distributed manipulation is very straightforward. Consider the case where the distributed manipulation system is fully actuated, in the sense that every actuator can exert a force of an arbitrary magnitude in an arbitrary direction while all satisfying the kinematic constraints associated with contact between the actuators and the part. In the example in Fig. 1(c), these kinematic constraints are associated with the nonslip constraints between the wheels and the part. In this case, perhaps not surprisingly, one can effectively reduce such a system from an overactuated system to a system with three inputs and three outputs by enforcing kinematic compatibility constraints between the actuators. (This functions in precisely the same manner as the differential on an automobile, which ensures that both front wheels satisfy the same kinematic constraints to reduce wear on tires due to slipping.) By doing this, one gains direct control over the velocity of the part. This method additionally provides desirable robustness properties with respect to actuator uncertainties. For more details on this method, see [34]. This control law is smooth, which we will see cannot necessarily be expected for overactuated systems. Moreover, it can only be used for fully actuated systems, which requires that every actuator in the plane have two degrees of freedom. Many interesting examples do not satisfy this requirement, such as micro-electromechanical system (MEMS) arrays which typically only have a one degree of freedom actuator at every location. We will consider a similar system with only one degree of freedom actuators shortly as an example to illustrate the questions we are interested in.

Let us consider the dynamic representation (in 2-6) of distributed manipulation. As before, let q_R and q_P denote the configuration of the array/part system, consisting of the part's planar location, and the variables that describe the configuration of each actuator array element. We treat the part and the array element contact as a rigid body contact system. We assume point contacts between the part and the array elements, and that the part's contact with the manipulating surface is governed by the Coulomb friction law at each contact point. When the part is slipping against the i^{th} actuator we get a force f_P^i which is proportional to the slipping velocity between the actuator and the part. When it is not slipping we get a force which is just the actuator input u^i scaled to account for the actuator and part inertias. The robot dynamics (in (4)) are simply the inputs, again scaled to account for actuator and part inertias. This leaves the part dynamics, as in (5). The difficulty lies in the fact that as the actuators stick and slip against the part, the dynamics change discontinuously. Because the dynamics are different for every combination of stick and slip for each actuator, there are 2^n dynamics for n actuators. Rather than attempt to analyze such a complicated system, we instead choose a quasistatic approach, described next.

If $\omega_i(q)$ is the one form describing the kinematic constraint between the part and the i^{th} actuator, then $\omega_i(q)\dot{q}$ describes the relative motion of the contact between the part and the actuator. If $\omega_i(q)\dot{q} = 0$, the contact is not slipping, while if $\omega_i(q)\dot{q} \neq 0$, then $\omega_i(q)\dot{q}$ describes the slipping velocity. In general, the moving part will be in contact with the actuator array at many points. From kinematic considerations, one or more of the contact points must be in a slipping state, thereby dissipating energy. The *power dissipation function* measures the part's total energy dissipation due to contact slippage.

Definition 5. The *Dissipation* or *Friction Functional* for an n-contact state is defined to be

$$\mathcal{D} = \sum_{i=1}^{n} \mu_i N_i \mid \omega_i(q)\dot{q} \mid \tag{15}$$

where μ_i and N_i are the Coulomb friction coefficient and normal force at the i^{th} contact, which are assumed known.

Since there will generally not exist a motion where all of the contacts can be simultaneously slipless (although our example of a fully actuated system is a case when such a condition exists), we are led to the following concept for finding the governing equations.

Power Dissipation Principle: With $\dot{q}$ small, a part's motion at any given instant is the one that minimizes $\mathcal{D}$.

The *power dissipation method (PDM)* assumes that the part's motion at each instant is the one that instantaneously minimizes power dissipation due to contact slippage. This method is adapted from the work of [2] on wheeled vehicles. An interesting topic to explore in this context is finding a concrete correspondence between solutions to the PDM and solutions to the dynamics as in (2-6). In particular, one would like to know if for every choice of inputs for a control system on $\mathcal{Q}$ there exists a corresponding choice of inputs for a control system on $T\mathcal{Q}$ such that the trajectory on $\mathcal{Q}$ is merely the projection of the other trajectory on $T\mathcal{Q}$. Although we do not discuss them here, preliminary results have been made in [34] based on results in [19]. In [35], we showed that the power dissipation approach generically leads to multiple model systems defined next. The intuition behind this is based on the fact that both μ_i and N_i typically depend on the configuration, $q \in \mathcal{Q}$. Therefore, as the set of supports, coefficient of friction, and center of mass location all change, a different set of contacts are slipping. For each one of these choices of slipping or not slipping the model is different. Moreover, the transitions between these models can be modeled as discontinuous jumps from one model to the next. In the case of distributed manipulation, this switching corresponds to the switching among different contact states (i.e., different sets of slipping contacts) due to variations in contact geometry and surface friction properties. Other examples can be found in [38]. Now we proceed to our formal description of such systems.

Definition 6. A control system Σ is said to be a *multiple model driftless affine system (MMDA)* if it can be expressed in the form

$$\Sigma: \quad \dot{q} = g_1(q)u_1 + g_2(q)u_2 + \cdots + g_m(q)u_m \tag{16}$$

where for any q and t, $g_i \in \{\gamma_{\alpha_i} | \alpha_i \in I_i\}$, with I_i an index set, γ_{α_i} analytic in (q,t) for all α_i, and the controls $u_i \in \mathbb{R}$ piecewise constant and bounded for all i. Moreover, letting σ_i denote the "switching signals" associated with g_i (which will be referred to as "MMDA maps"),

$$\begin{aligned} \sigma_i : \mathcal{Q} \times \mathcal{R} &\longrightarrow \mathbb{N} \\ (q,t) &\longrightarrow \alpha_i \end{aligned}$$

then the σ_i are measurable in (q,t).

An MMDA system is a driftless affine nonlinear control system where each control vector fields may "switch" back and forth between different elements of a finite set. The σ_i which regulate this switching may not be known, so we have no guarantees about the nature of the switching except that it is, by assumption, measurable. In the context of control, we now have a synthesis problem. For instance, if the array is such that each array element can only exert a force in one direction (in $\mathbb{R}^2$), then the system generally *must* have some slipping, thereby leading to nonsmooth effects as the model describing the dynamics changes discontinuously.

Now we proceed to an example. Figure 4 shows on the left a photograph of an experiment at Caltech which has been used previously to test algorithms for distributed manipulation. In the photograph we see four wheels all oriented towards the origin. These wheels are not allowed to change direction–hence each actuator is a one degree of freedom actuator. We use a

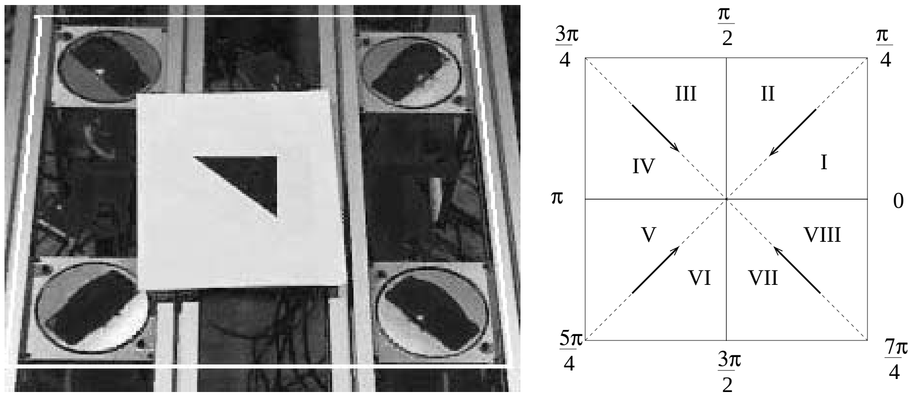

Fig. 4. Photograph and cartoon of 4 cell distributed manipulator.

piece of plexiglass (for purposes of visualization) on top of the four wheels. The white line seen in the photograph indicates the outline of the plexiglass. The goal, then, is to control the center of mass to the origin in $\mathbb{R}^2$ with a desired orientation of $\theta = 0$. To do this, we obtain feedback of the plexiglass' configuration by affixing a piece of paper with a black triangle (also seen in the photo) whose right angle corner coincides with the plexiglass' center of mass. Using this, we obtain the position and orientation of the plexiglass through visual feedback. The right hand side of Fig. 4 is a cartoon of the experiment, where the four arrows correspond to actuators and the regions denoted by **I-VIII** and **0-$\frac{7\pi}{4}$** will be important in our subsequent description of the kinematics described by the PDM.

Note that this system thus described is overactuated because there are four inputs and only three outputs. Now we describe the kinematics of this system. Assume the coefficient of friction is the same for all four actuators. Then the model describing the kinematics will only change as the center of mass moves across the array. The solution to minimizing $\mathcal{D}$ always consists of choosing the constraints $\omega_i(q)$ that have the most dissipation and finding the kinematics that satisfy those constraints (see [2,35]). Therefore, the actuator wheel nearest to the center of mass will have both its "rolling" constraint and its "sideways" slip constraint satisfied. The actuator wheel second closest to the center of mass will have one of its two constraints satisfied. In the case of the wheels shown in the figure, it will be the rolling constraint. For details on this analysis, see [36]. Denote the actuator input associated with the closest actuator by u_i and the actuator input associated with the second closest actuator by u_j. Then these considerations lead to kinematics of the form:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = g_1 u_i + g_2 u_j \tag{17}$$

where

$$g_1 \in \begin{bmatrix} \frac{-y_i}{(x_j-x_i)\sin(\theta_j)+(y_i-y_j)\cos(\theta_j)} \\ \frac{x_i}{(x_j-x_i)\sin(\theta_j)+(y_i-y_j)\cos(\theta_j)} \\ \frac{u_j}{(x_i-x_j)\sin(\theta_j)+(y_j-y_i)\cos(\theta_j)} \end{bmatrix} \tag{18}$$

$$g_2 \in \begin{bmatrix} \frac{\sin(\theta_j)((x_i-x_j)\cos(\theta_i)+y_i\sin(\theta_i))+\cos(\theta_i)\cos(\theta_j)y_j}{(x_j-x_i)\sin(\theta_j)+(y_i-y_j)\cos(\theta_j)} \\ \frac{-\cos(\theta_i)\cos(\theta_j)x_i-\sin(\theta_i)(x_j\sin(\theta_j)-(y_i-y_j)\cos(\theta_j))}{(x_j-x_i)\sin(\theta_j)+(y_i-y_j)\cos(\theta_j)} \\ \frac{-\cos(\theta_i-\theta_j)}{(x_i-x_j)\sin(\theta_j)+(y_j-y_i)\cos(\theta_j)} \end{bmatrix} \tag{19}$$

In these equations x_i, y_i, and θ_i refer to the planar coordinates and orientation of the i^{th} actuator. The set-valued notation of (18) and (19) refers to the fact that at a transition between actuators i and j being the two closest actuators to actuators k and l being the closest the kinematics are discontinuous. Therefore, at these points we must allow multi-valued differentials in order to guarantee existence of solutions to the differential equation in (17). See [34] for more details. It should be noted that here the index notation should be thought of as mapping (i, j) pairs to equations of motion in some neighborhood (not necessarily small) around the i^{th} and j^{th} actuator. In each region $I - VIII$ the kinematics are smooth, but when a trajectory crosses a boundary $\mathbf{0}$-$\frac{\mathbf{7\pi}}{\mathbf{4}}$, there is a discontinuity in the kinematics. It is possible to obtain point stabilization to $(x, y, \theta) = (0, 0, 0)$ from any initial condition using discontinuous control laws based on the kinematics and knowing the current model (see [34] for details of this control design). However, there are many questions relevant to this system which remain unanswered.

Problem 10. Is the system in (17) locally controllable near the origin, or anywhere else?

Problem 11. We know that point stabilization to the origin can be achieved using discontinuous control laws, but does there exist a continuous control law which will do so? More generally how do we test to see if a given nonsmooth system (such as Definition 6) has a smooth feedback law for the purposes of point stabilization?

Problem 12. How can we characterize the set of points in $\mathbb{R}^2$ to which the system in (17) can be stabilized?

Problem 13. What kind of trajectories can this system follow? Most importantly, if we want to move a part to an arbitrary point in $SE(2)$, how close can we get it to the desired point?

With the advent of more and more actuators being built cheaply at very small scales, massively overactuated systems will become prevalent in industry. It is therefore important to develop systematic tools for the analysis and design of control laws for these systems. Sometimes, by necessity, this will involve complicated control laws based on the full nonsmooth system. In such a case, we will need the tools to answer questions such as those posed in Problems 10-13 before we will be able to move parts efficiently and reliably at the MEMS scale, or in any other distributed manipulation system with limited actuator degrees of freedom.

References

1. S. Akella, W. Huang, K. M. Lynch, and M. T. Mason. Parts feeding on a conveyor with a one joint robot. *Algorithmica*, 26(3):313–344, March-April 2000.
2. J. C. Alexander and J. H. Maddocks. On the kinematics of wheeled vehicles. *The International Journal of Robotics Research*, 8(5):15–27, October 1989.
3. K. F. Böhringer and H. Choset, editors. *Distributed Manipulation*. Kluwer, 2000.
4. K. F. Böhringer, B. R. Donald, L. E. Kavraki, and F. Lamiraux. A distributed, universal device for planar parts feeding: unique part orientation in programmable force fields. In *Distributed Manipulation*, pages 1–28. Kluwer, 2000.
5. R. W. Brockett and L. Dai. Nonholonomic kinematics and the role of elliptic functions in constructive controllability. In Z. Li and J. Canny, editors, *Nonholonomic Motion Planning*. Kluwer Academic, 1993.
6. B. Brogliato and A. Zavala-Rio. On the control of complementary-slackness juggling mechanical systems. *IEEE Transactions on Automatic Control*, 45(2):235–246, February 2000.
7. 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.
8. M. Bühler, D. E. Koditschek, and P. J. Kindlmann. Planning and control of a juggling robot. *International Journal of Robotics Research*, 13(2):101–118, 1994.
9. F. Bullo, A. D. Lewis, and K. M. Lynch. Controllable kinematic reductions for mechanical systems: Concepts, computational tools, and examples. In *2002 International Symposium on the Mathematical Theory of Networks and Systems*, August 2002.
10. F. Bullo and K. M. Lynch. Kinematic controllability for decoupled trajectory planning of underactuated mechanical systems. *IEEE Transactions on Robotics and Automation*, 17(4):402–412, August 2001.
11. R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Toward a dynamical pick and place. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2: 292–297, 1995.
12. P. Choudhury and K. M. Lynch. Rolling manipulation with a single control. *International Journal of Robotics Research*, 2002. To appear.
13. B. R. Donald, J. Jennings, and D. Rus. *Algorithmic Foundations of Robotics (WAFR)*, chapter Information invariants for distributed manipulation, pages 431–459. A.K. Peters, Ltd, Wellesley, MA, 1995.

14. M. A. Erdmann. Understanding action and sensing by designing action-based sensors. *International Journal of Robotics Research*, 14(5):483–509, October 1995.
15. B. Goodwine and J. W. Burdick. Controllability of kinematic control systems on stratified configuration spaces. *IEEE Trans. on Automatic Control*, 46(3):358–368, 2000.
16. G. W. Haynes and H. Hermes. Nonlinear controllability via Lie theory. *SIAM Journal on Control*, 8(4):450–460, November 1970.
17. Y. Jia and M. Erdmann. Observing pose and motion through contact. In *IEEE International Conference on Robotics and Automation*, pages 723–729, 1998.
18. V. Jurdjevic. *Geometric Control Theory*. Cambridge University Press, 1997.
19. A. D. Lewis. When is a mechanical control system kinematic? In *IEEE Conference on Decision and Control*, pages 1162–1167, December 1999.
20. 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.
21. K. Lian, L. Wang, and L. Fu. Controllability of spacecraft systems in a central gravitational field. *IEEE Transactions on Automatic Control*, 39(12):2426–2440, December 1994.
22. J. Luntz and W. Messner. Closed-loop stability of distributed manipulation. In *Proc. American Control Conference (ACC)*, 2000.
23. J. Luntz, W. Messner, and H Choset. Velocity field design for parcel manipulation on the modular distributed manipulation system. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
24. J. Luntz, W. Messner, and H. Choset. Closed-loop distributed manipulation using discrete actuator arrays. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2000.
25. K. M. Lynch. Controllability of a planar body with unilateral thrusters. *IEEE Transactions on Automatic Control*, 44(6):1206–1211, June 1999.
26. K. M. Lynch. Locally controllable manipulation by stable pushing. *IEEE Transactions on Robotics and Automation*, 15(2):318–327, April 1999.
27. K. M. Lynch and C. K. Black. Recurrence, controllability, and stabilization of juggling. *IEEE Transactions on Robotics and Automation*, 17(2):113–124, April 2001.
28. K. M. Lynch and M. T. Mason. Stable pushing: Mechanics, controllability, and planning. *International Journal of Robotics Research*, 15(6):533–556, December 1996.
29. K. M. Lynch and M. T. Mason. Dynamic nonprehensile manipulation: Controllability, planning, and experiments. *International Journal of Robotics Research*, 18(1):64–92, January 1999.
30. Kevin M. Lynch, Naoji Shiroma, Hirohiko Arai, and Kazuo Tanie. The roles of shape an motion in dynamic manipulation: The butterfly example. In *IEEE International Conference on Robotics and Automation*, pages 927–932, 1998.
31. A. Marigo and A. Bicchi. Rolling bodies with regular surface: Controllability theory and applications. *IEEE Transactions on Automatic Control*, 45(9):1586–1599, September 2000.
32. A. Marigo, M. Ceccarelli, S. Piccinocchi, and A. Bicchi. Planning motions of polyhedral parts by rolling. *Algorithmica*, 26:560–576, 2000.
33. M. T. Mason. *Mechanics of Robotic Manipulation*. MIT Press, 2001.

34. T. D. Murphey. *Control of Multiple Model Systems.* PhD thesis, California Institute of Technology, May 2002.
35. T. D. Murphey and J. W. Burdick. Issues in controllability and motion planning for overconstrained wheeled vehicles. In *Proc. Int. Conf. Math. Theory of Networks and Systems (MTNS)*, Perpignan, France, 2000.
36. T. D. Murphey and J. W. Burdick. On the stability and design of distributed systems. In *Proc. IEEE Int. Conf. on Robotics and Automation*, Seoul, Korea, 2001.
37. T. D. Murphey and J. W. Burdick. Global exponential stabilizability for distributed manipulation. In *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington D.C., 2002.
38. F. Pfeiffer and C. Glocker. *Multibody Dynamics with Unilateral Controls.* Wiley, Chichester, 1996.
39. A. A. Rizzi and D. E. Koditschek. Progress in spatial robot juggling. In *IEEE International Conference on Robotics and Automation*, pages 775–780, Nice, France, 1992.
40. A. A. Rizzi and D. E. Koditschek. Further progress in robot juggling: The spatial two-juggle. In *IEEE International Conference on Robotics and Automation*, pages 3:919–924, Atlanta, GA, 1993.
41. M. W. Spong. Impact controllability of an air hockey puck. *Systems and Control Letters*, 42:333–345, 2001.
42. A. Sudsang and L. Kavraki. A geometric approach to designing a programmable force field with a unique stable equilibrium for parts in the plane. In *Proc. IEEE Int. Conf. Robotics and Automation*, Seoul, Korea, 2001.
43. M. Zefran and J. W. Burdick. Stabilization of systems with changing dynamics. In *Workshop on Hybrid systems: Computation and control*, Berkeley, CA, 1998.