Nonholonomic Path Planning for Pushing a Disk Among Obstacles

Pankaj K. Agarwal* Jean-Claude Latombe † Rajeev Motwani ‡
Prabhakar Raghavan §

Abstract

We consider the path-planning problem for a robot *pushing* an object in an environment containing obstacles. This new variant of the classical robot path-planning problem has several interesting geometric aspects, which we explore. We focus on the setting where the robot makes a point contact with the object which is assumed to be a unit disk, while the obstacles are assumed to be polygonal.

1 Introduction

Nonprehensile manipulation exploits task mechanics to achieve a goal without grasping [5]. Pushing is one important type of nonprehensile manipulation which can allow complex tasks to be performed by simple mechanisms; it may also enable a robot to move parts that are too heavy to be lifted. The mechanics of pushing is studied in [7], where a simple rule is established to qualitatively determine the motion of a pushed object. This analysis has inspired several nonprehensile manipulation algorithms, in particular the design of part feeders to re-orient parts arriving on a conveyor belt [1, 8, 9].

* Computer Science Department, Duke University, Box 90129, Durham, NC 27708-0129. Supported by National Science Foundation Grant CCR-93–01259, by an Army Research Office MURI grant DAAH04-96-1-0013, by a Sloan fellowship, by an NYI award, and by matching funds from Xerox Corporation, and by a grant from the U.S.-Israeli Binational Science Foundation.

† Department of Computer Science, Stanford University, Stanford, CA 95305. Supported by an ARO MURI Grant DAAH04-96-1-0007.

‡ Department of Computer Science, Stanford University, Stanford, CA 94305. Supported by an Alfred P. Sloan Research Fellowship, an IBM Faculty Partnership Award, an ARO MURI Grant DAAH04-96-1-0007, and NSF Young Investigator Award CCR-9357849, with matching funds from IBM, Mitsubishi, Schlumberger Foundation, Shell Foundation, and Xerox Corporation.

§ IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120.

In this paper we consider the path-planning problem for a robot pushing an object on a plane cluttered with obstacles, under quasi-static assumptions. This problem has been previously studied in [3, 4, 5, 6], where it was shown that conditions for stable pushing induce *nonholonomic constraints* in the motion of the pushed object. The controllability of the pushed object was analyzed in depth and an effective path planner was implemented using a variant of the algorithm in [2]. This algorithm discretizes the search space by integrating, at each step, a small number of feasible motions over a short duration δt (for a given constant velocity). It is asymptotically complete, i.e., if a solution path exists, the planners finds it, provided that δt is selected small enough. Here, our goal is to build the foundations of a complete (or exact) planning algorithm for the pushing problem, by exploring in detail the geometric properties of the possible pushing paths.

We have initiated our investigation by considering a simple setting where the robot makes a point contact with the object which is assumed to be a unit disk, while the obstacles are assumed to be polygonal. The first part of our study (Section 2) concerns the nonholonomic trajectory followed by the disk-center as the robot pushes the disk. We study the geometric properties of this trajectory and establish a variety of properties, some of which are counter-intuitive. For example, the hooked shape of the trajectory allows the pushing robot to move the disk around corners of the obstacles, with only a single push of fixed direction. In Section 3, we explicate the structure of the reachable region in the presence of a point obstacle, as a stepping stone to devising a complete planning algorithm. Based on this, in Section 4 we provide a scheme for the path-planning problem in the presence of obstacles. In Section 5, we conclude by arguing that the properties of the nonholonomic trajectories of a pushed disk, which underlie our planning scheme, should hold in more general settings.

2 The Model

Let D be a disk of unit radius, and let P be a point pusher that chooses a point p on the boundary of D and pushes D at unit velocity by maintaining the contact at p. Once p is chosen, it remains fixed during the course of a single push (stable pushing [6]), in the sense that P always maintains the contact at p and that it never slips around the periphery. The contact between P and D breaks when P's motion is tangent to D. As in [6, 7], we assume the *quasi-static* model, where friction between the disk and its supporting plane is large enough so that there is no motion after pushing ceases. Let O denote the center of D. We study the trajectory followed by O as P pushes D.

A configuration of the pair (P, D) can be represented as a point (x, y, θ) in $\mathbb{R}^2 \times \mathcal{S}^1$, where (x, y) denote the x- and y-coordinates of the center of D and θ denotes the angle between the ray emanating from p in direction $\vec{pO}$ and the x-axis; see Figure 1. Let $(0, 0, \varphi)$ denote the initial position of D, i.e., D is originally centered at the origin and the initial direction of $\vec{pO}$ is φ.

We first derive the equation for the trajectory of D, assuming that P is pushing D to right at p in the horizontal direction. Later we will consider the case where P pushes D at p in an arbitrary direction.

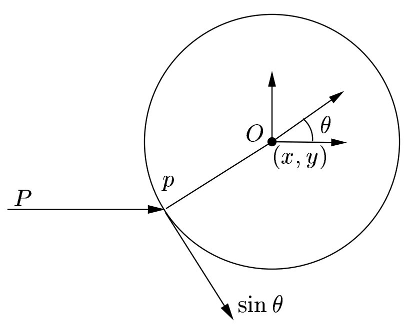

Figure 1: *Resolution of the unit-velocity push at p.*

If $\varphi = 0$, i.e., p is the leftmost point of D, then the center of D translates along the x-axis; and if $\alpha = \pi/2$, then D does not move at all. Suppose the position of D at time t is $(x(t), y(t), \theta(t))$. Then the unit velocity motion at p may be resolved into a radial component (towards the center of the disk) of magnitude $\cos\theta$, and a tangential component of magnitude $\sin\theta$. Thus we have

$$d\theta/dt = \sin\theta,$$

which implies

$$t = \int \frac{1}{\sin\theta} d\theta = \ln\left(\tan\frac{\theta}{2}\right) + c,$$

where c is a constant. Since $\theta = \varphi$ at $t = 0$, we obtain

$$\theta(t) = 2\tan^{-1}\left(\tan\left(\frac{\varphi}{2}\right) \cdot e^t\right). \qquad (1)$$

D stops moving once $\theta = \pi/2$, therefore

$$0 \le t \le \ln(\cot(\varphi/2)).$$

The radial component of the unit velocity determines the x- and y-motion of the center of D as follows.

$$\begin{aligned} dx/dt &= \cos^2\theta \\ dy/dt &= \sin\theta\cos\theta. \end{aligned}$$

Solving these equations and using the fact that the initial position of D is $(0, 0, \varphi)$, we obtain the following:

$$\begin{aligned} x(\theta) &= \ln\left(\frac{\tan\theta/2}{\tan\varphi/2}\right) + \cos\theta - \cos\varphi \qquad (2) \\ y(\theta) &= \sin\theta - \sin\varphi. \qquad (3) \end{aligned}$$

Substituting (1) in these equations, we obtain

$$\begin{aligned} x(t) &= t + \frac{1 - \tan^2(\varphi/2) \cdot e^{2t}}{1 + \tan^2(\varphi/2) \cdot e^{2t}} - \cos\varphi \qquad (4) \\ y(t) &= \frac{2\tan(\varphi/2) \cdot e^t}{1 + \tan^2(\varphi/2) \cdot e^{2t}} - \sin\varphi \qquad (5) \end{aligned}$$

for $0 \le t \le \ln\cot(\varphi/2)$. Figure 2 (i) plots θ, x, and y, as a function of time; and Figure 2 (ii) shows how x- and y-positions vary with θ. (Note that for $\varphi = 0$, the trajectory will be a straight line; this is the infinite prefix of the curve "before" it enters the regime starting from the small φ plotted in Figure 2.)

The trajectory of O, as D is being pushed, resembles a "hockey stick" (see Figure 3). We will use Π to denote this trajectory.

The vertical displacement from initial position $\theta = \varphi$ to $\theta = \pi/2$ is, as expected, $1 - \sin\varphi$ (this is exactly the initial height of p below the center of the disk). The curve has the following "memoryless" property: the trajectory subsequent to any particular value of θ is independent of the initial angle $\varphi \le \theta$ at which the pushing motion was begun. We now state some basic properties of Π; the proofs are omitted here.

Theorem 1 *(i) Π is a convex curve.*
(ii) Let ℓ be a line intersecting Π at two points, then Π is monotone with respect to ℓ.
(iii) The curvature of Π at a point with orientation θ is $\tan\theta$. If we denote by $\kappa(t)$ the curvature of Π as a function of time, then

$$\kappa(t) = \frac{2\tan(\varphi/2) \cdot e^t}{1 - \tan^2(\varphi/2) \cdot e^{2t}}$$

for $0 \le t \le \ln\cot(\varphi/2)$.

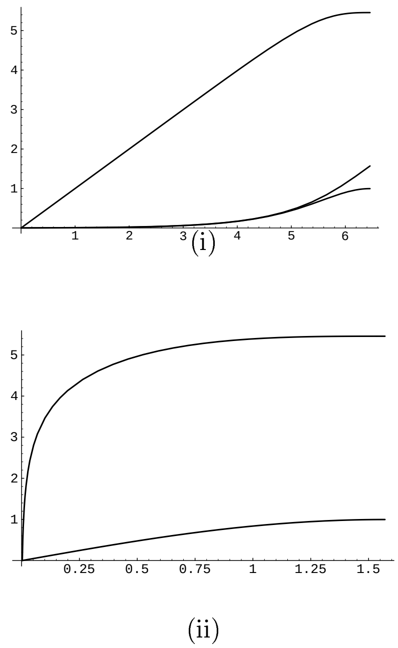

Figure 2: (i)$\theta(t)$, $x(t)$, and $y(t)$, with $\varphi = \pi/1000$; (ii) $x(\theta)$ and $y(\theta)$, with $\varphi = \pi/1000$.

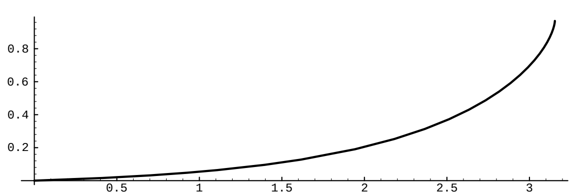

Figure 3: The hockey stick curve, with $\varphi = \pi/100$.

Instead of pushing D at p in the horizontal direction, P can push D in any direction while keeping the contact at p. Suppose P pushes D at p (with unit speed) in a direction ψ, then $\theta(t), x(t)$, and $y(t)$ can be rewritten as follows. Let $\alpha = \varphi - \psi$.

$$\theta(t) = 2\tan^{-1}\left(\tan(\alpha/2)\cdot e^t\right) + \psi \quad (6)$$

$$x(t) = t + \frac{1-\tan^2(\alpha/2)\cdot e^{2t}}{1+\tan^2(\alpha/2)e^{2t}} - \cos\alpha \quad (7)$$

$$y(t) = \frac{2\tan(\alpha/2)e^t}{1+\tan^2(\alpha/2)\cdot e^{2t}} - \sin\alpha \quad (8)$$

If we want to push D to the right, then $\varphi - \pi/2 \le \psi \le \varphi$; otherwise $\varphi - \pi/2 \le \psi \le \varphi + \pi/2$. If we fix a pair φ, ψ, the trajectory of O is uniquely defined. We denote it by $\Pi_{\varphi,\psi}$, and refer to such a push as the *(φ,ψ)-push*. Let

$$\mathbf{\Pi} = \{\Pi_{\varphi,\psi} \mid -\pi/2 \le \varphi \le \pi/2,\ \varphi - \pi/2 \le \psi \le \varphi\}$$

denote the set of all possible trajectories that O can follow, starting from the origin. For a trajectory $\Pi \in \mathbf{\Pi}$, we will refer to the final position on Π as the *tip* of Π. We say that a point q is *reachable* from the origin by a *(φ,ψ)-push* if $q \in \Pi_{\varphi,\psi}$. Suppose $q = \Pi_{\varphi,\psi}(t)$, then $\theta(t) - \psi$ is called the *arriving angle*. If we allow to vary the pushing direction, we represent a *configuration* of (P, D) by $(x, y, \theta - \psi)$, where (x, y) denote the x- and y-coordinates of O and $\theta - \psi$ denotes the arriving angle of disk at (x, y).

Using Theorem 1, one can prove the following properties of curves in $\mathbf{\Pi}$.

Lemma 2 *Any two curves Π_1, Π_2 in $\mathbf{\Pi}$, such that Π_1 is not contained in Π_2, intersect in at most one point other than the origin.*

Lemma 3 *For any point q and for any $-\pi/2 \le \theta \le \pi/2$, there is a unique (φ,ψ) pair such that D reaches q by the (φ,ψ)-push with arriving angle θ.*

Corollary 4 *Given two points q_1, q_2 in the plane, different from the origin, there is at most one curve $\Pi \in \mathbf{\Pi}$ that passes through q_1 and q_2. (Recall that all curves in $\mathbf{\Pi}$ pass through the origin.)*

3 The Reachable Region

Let $\mathcal{O}$ be a set of polygonal obstacles, with a total of n vertices. We study the set R of points to which the center of D can be moved from the origin *using a single push* without intersecting any of the obstacles. In the process we develop the shape of the reachable region for motion via pushing, which plays a fundamental role analogous to the visibility region in conventional path planning. In Section 4 below we will build on this notion to devise a path-planning scheme.

Let p and q be any two points in the plane, and let θ be an arbitrary arriving angle in $[-\pi/2, \pi/2]$. In the absence of obstacles it is possible to push a disk from p to q, or vice versa, such that the arriving angle is θ. The situation is very different in the presence of obstacles: one may have trajectories from p to q for several arriving angles, and yet have no trajectory from q to p at all. (Consider a single trajectory from p to q starting from a small value ϕ, say $\pi/100$ as in Figure 3; mark out a swath of width 2 about this trajectory. Imagine the rest of the plane being covered by small obstacles that allow no gaps between them. Now the asymmetric nature of the hockey stick curve prevents motion from q to p.)

Let (x, y, θ) be a configuration of (P, D). By Lemma 3, there is a unique (φ, ψ)-push by which D reaches the point (x, y) at the angle θ. The configuration (x, y, θ) is called *reachable* if D does not intersect any obstacle as it is moved from the origin to the point (x, y) along $\Pi_{\varphi,\psi}$. A point $q = (x_q, y_q)$ is called *reachable* if there exists an arriving angle α $(-\pi/2 \leq \alpha \leq \pi/2)$ so that the configuration (x_q, y_q, α) is reachable. We will refer to the set of nonreachable points from the origin as the *shadow region* of $\mathcal{O}$, and denote it by R.

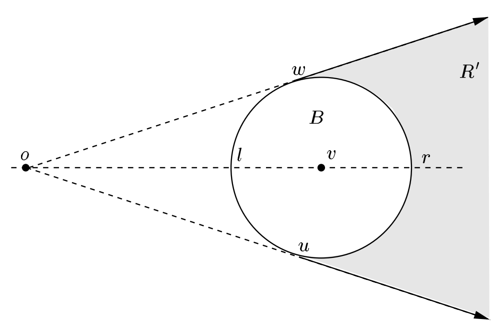

Figure 4: Disk B and the region R'.

In order to have a better understanding of the shadow region, we first study the simplest case in which $\mathcal{O}$ is just a point v on the x-axis, with $x(v) > 1$. Let B be the unit-radius disk centered at v. Then a point q is reachable if there is a curve $\Pi_{\varphi,\psi} \in \mathbf{\Pi}$ that contains q and $\Pi_{\varphi,\psi}[o, q]$ does not intersect B. Obviously $B \subseteq R$. For any two points $a, b \in \partial B$, let $\widehat{ab} \subseteq \partial B$ be the clockwise-oriented arc from a to b. Let l and r be the leftmost and the rightmost points of B, respectively; l and r partition ∂B into two semicircles — lower and upper semicircles. Let u (resp. v) be the point on the lower (resp. upper) semicircle so that the line tangent to B at u (resp. w) passes through the origin. Let ρ_u be the rightward-directed ray emanating from u in direction $\vec{ou}$; similarly define ρ_w. Let R' be the region bounded by the rays ρ_u, ρ_w, and the arc $\widehat{wu}$ (shaded region in Figure 4). For any point $q \notin D \cup R'$, the segment oq does not intersect B, therefore q is reachable, thereby implying that $R \subseteq D \cup R'$. We next describe the portion of R' that forms the shadow region.

Lemma 5 *For any point $q \in \partial B$, there is at most one curve in $\mathbf{\Pi}$ that is tangent to B at q.*

Lemma 6 *There is a unique point ξ^- on the arc $\widehat{ur}$ so that the curve $\Pi \in \mathbf{\Pi}$ with ξ^- as its tip is tangent to the arc $\widehat{u\xi^-}$ at another point χ^-. Moreover, for any point $s \in \widehat{u\chi^-}$, the curve in Π tangent to B at s does not intersect the interior of D. There are corresponding points ξ^+, χ^+ on the upper semicircle.*

We now define two arcs γ^-, γ^+ as follows. For a point $s \in \widehat{u\chi^+}$, let $\tau(s)$ denote the tip of the hockey-stick curve that is tangent to B at s. Define $\gamma^- = \{\tau(s) \mid s \in \widehat{u\chi^-}\}$. Similarly define $\gamma^+ = \{\tau(s) \mid s \in \widehat{\chi^+ w}\}$. We claim that γ^- is a x-monotone curve whose leftmost endpoint if ξ^- and whose asymptote is the line parallel to ρ_u and passing through v, the center of B; γ^+ also satisfies similar properties.

Lemma 7 *A point $q \in R$ lies in the shadow region if and only if q lies above γ^- and below γ^+.*

The above lemma implies that R is the region bounded by the circular arc $\widehat{\xi^+\xi^-}$ and the arcs γ^- and γ^+; see Figure 5. The shadow region of a segment can also be analyzed in a similar way. It is also bounded by $O(1)$ curves, and has a similar structure.

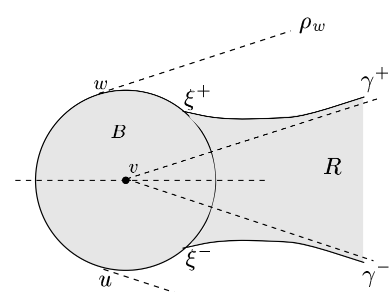

Figure 5: Reachable region; the shaded region is not reachable; $\varphi = \pi/100$

The shadow region induced by multiple obstacles is quite complex. If R_e denotes the shadow region of an obstacle edge e, then one is tempted to say that the overall shadow region R is equal to $\bigcup_{e \in \mathcal{O}} R_e$. Although it is true that $\bigcup_e R_e \subseteq R$, some other points may also be nonreachable. Let $\overline{R}$ be the set of non-reachable configurations of (P, D), then by definition, the shadow region is $R = \{(x, y) \mid (x, y, \alpha) \in \overline{R}, \forall \alpha \in [-\pi/2, \pi/2]\}$. In a companion paper we will prove that R has $O(n^2)$ vertices and edges, and describe an efficient algorithm for computing R in an appropriate model of computation. One can construct a scene with n obstacle edges so that the shadow region has $\Omega(n^2)$ vertices.

4 Path Planning

We consider the problem of moving a disk through a scene with obstacles by a sequence of pushes. Formally, given a set $\mathcal{O}$ of obstacles and a pair of points I and F, we want to compute a path Γ for the center of D from I to F that satisfies the following properties: (i) D does not intersect any obstacle as its center moves from I to F along Γ, and (ii) $\Gamma = \gamma_1 \| \gamma_2 \| \cdots \| \gamma_k$ for some $k \geq 1$ such that γ_i is a (φ_i, ψ_i)-push from a point q_{i-1} to q_i ($q_0 = I$ and $q_k = F$). We propose an efficient scheme for computing such a path Γ. The goal is to plan a path that minimizes the number of distinct pushes needed to move O from I to F. Here we are assuming a model of computation in which various *basic operations* on curves in $\mathbf{\Pi}$ (such as intersection points between two curves, intersection detection with a line, etc.) can be performed in constant time.

The fact that there is an infinite family of trajectories (rather than just a straight line segment) that the disk can use to reach a point from a fixed point could be used to go "around" obstacles with a single push, giving rise to plans that require fewer pushes than if only straight-line motion were used. We now explore the power of such pushes and how they can be used to plan motion. A push, in general, will consist of a pair of angles φ and φ' – the push is begun as before at φ; when the position $\theta = \varphi'$ is reached, we cease pushing.

We assume that the obstacles are disjoint polygons in the plane; let n be the total number of obstacle vertices. We expand each obstacle $\mathcal{O}_i \in \mathcal{O}$ by taking its Minkwoski sum with a unit-radius disk. Let $\mathcal{C}$ denote the set of expanded obstacles. For a point q in the plane, let $\mathbf{\Pi}[q]$ denote the set of trajectories of O induced by (φ, ψ)-pushes, assuming that the O was initially at q.

We propose the following preprocessing algorithm. The idea is to construct a *roadmap* in the plane by constructing a *directed reachability graph* G on a set $\mathcal{M}$ of m points selected in the plane; I and F are also added to $\mathcal{M}$. A milestone M_2 is said to be *reachable* from a milestone M_1 when it is possible to move the disk center from M_1 to M_2 using only one push, i.e., if there exists a trajectory $\Pi \in \mathbf{\Pi}[M_1]$ such that $M_2 \in \Pi$ and $\Pi[M_1, M_2]$ does not intersect any expanded obstacle in $\mathcal{C}$. The reachability graph G has the set of points in $\mathcal{M}$ as its vertex set. There is a directed edge from a vertex u to a vertex v if and only if v is reachable from u; in addition, we label the edge (u, v) with a representation of the trajectory Π that reaches from u to v. The path-planning algorithm is summarized below.

Algorithm Path planning
Input: Obstacles $\mathcal{C}$ and points I and F.
Output: A path from I to F.

1. Choose an integer m.
2. Select a set $\mathcal{M}$ of m milestones on the boundaries of the obstacles in $\mathcal{C}$, and add I, F to $\mathcal{M}$.
3. **for** all $M_1, M_2 \in \mathcal{M}$ **do**
   - check whether there exists $\Pi \in \mathbf{\Pi}[M_1]$ such that $M_2 \in \Pi$ and $\Pi[M_1, M_2]$ does not intersect any expanded obstacle in $\mathcal{C}$;
   - if so, add the edge from M_1 to M_2 into G and label it with Π.
4. Compute a shortest path in G from I to F.

Several details need to be filled in before we can analyze this algorithm. The milestones are selected as uniformly spaced points along the boundaries of the polygons in $\mathcal{C}$, although various other selection strategies are possible such as selecting the milestones at random. To decide whether milestone M_2 is reachable from the milestone M_1, we discretize the family of curves in $\mathbf{\Pi}[M_1]$ by suitably discretizing the angle φ that parametrize the curves $\Pi \in \mathbf{\Pi}[M_1]$. For each φ, there is at most one $\Pi_{\varphi,\psi} \in \mathbf{\Pi}[M_1]$ that contains M_2. Assuming that we can compute this trajectory in constant time, we can easily check in $O(n)$ time whether $\Pi_{\varphi,\psi}[M_1, M_2]$ intersects any expanded obstacle. If we find such a $\Pi_{\varphi,\psi}$, we add (M_1, M_2) as an edge in G. The granularity of the discretization determines the trade-off between the efficiency and the robustness of the algorithm. If we choose $1/\epsilon$ different values of φ, then a naive implementation of this algorithm requires $\Theta(m^2 n/\epsilon)$ time. For each φ, we can however compute in $O(n \log n)$ time the set of points that are reachable from M_1 by a single (φ, ψ)-push, for some ψ. After having computed this set, all the milestones in $M_2 \in \mathcal{M}$ that are reachable from M_1 by a (φ, ψ)-push can be computed in $O(m \log n)$ time. Hence, the overall running time of the algorithm is $O((1/\epsilon) \cdot m(m + n) \log n)$.

Theorem 8 *Algorithm Path-planning can be implemented to run using $O((1/\epsilon) \cdot m(m+n)\log n)$ basic operations.*

5 Conclusions

In this paper we have laid a foundation for the design of algorithms for planning nonholonomic motion via pushing. We now remark that our work extends beyond the case when the object being pushed is a disk, to other geometries. In general, all we require is (a) a trajectory for the object being pushed that is defined by one parameter (in our case θ); (b) a finite termination point for the trajectory (in our case the point $\theta = \pi/2$); (c) the properties in the statements of Lemmas 2 and 3, as well as properties (i)–(iii) of Theorem 1. These properties suffice to allow us to build up machinery similar to that in Section 3. Our work opens up the study of such more general settings, and the descriptive complexities of shadow regions.

References

[1] Akella, S., Huang, W., Lynch, K., and Mason, M. 1996. Planar manipulation on a conveyor with a one joint robot. In *Robotics Research*, G. Giralt and G. Hirzinger (eds.), Springer, New York, NY, pp. 265–276.

[2] J. Barraquand and J.C. Latombe. 1993. Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles, *Algorithmica*, 10(2-3-4):121–155.

[3] Lynch, K. 1992. The mechanics of fine manipulation by pushing. In *Proc. Int. Conf. on Robotics and Automation*, pp. 2269–2276.

[4] Lynch, K. 1993. Estimating the friction parameters of pushed objects. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, pp. 186–193.

[5] Lynch, K. 1996. *Nonprehensile Robotic Manipulation*. PhD Thesis, The Robotics Institute, Carnegie Mellon University.

[6] Lynch, K. and Mason, M.T. 1995. Stable pushing: Mechanics, controllability, and planing. In *Algorithmic Foundations of Robotics*, Goldberg, K. et al. (eds.), A K Peters, Wellesley, MA, pp. 239–262.

[7] M.T. Mason. 1986. Mechanics and planning of manipulator pushing operations, *Int. J. of Robotics Research*, 5(3):53–71.

[8] M.A. Peshkin and A.C. Sanderson. 1988. Planning robotic manipulation strategies for workpieces that slide. *IEEE Tr. on Robotics and Automation*, 4(5):524–531.

[9] Wiegley, J., Goldberg, K., Peshkin, M., and Brokowski, M. 1996. A complete algorithm for designing and passive fences to orient parts. In *Proc. Int. Conf. on Robotics and Automation*, pp. 1133–1139.