A Mobile Manipulator

Matthew T. Mason
Carnegie Mellon
Pittsburgh, PA

Dinesh K. Pai
U. British Columbia
Vancouver, BC

Daniela Rus
Dartmouth
Hanover, NH

Lee R. Taylor
Carnegie Mellon
Pittsburgh, PA

Michael A. Erdmann
Carnegie Mellon
Pittsburgh, PA

Abstract

This paper describes a mobile manipulator that uses its wheels for manipulation as well as locomotion. This robot, named the mobipulator, *looks like a small car with four independently powered wheels, none of them steered. It is designed to manipulate paper and other objects on the surface of a desk. The wheels are used for locomotion or for manipulation, switching functions dynamically as the task demands. So far we have preliminary demonstrations of a variety of motions, and performance data for the task of moving a sheet of paper in a square while maintaining constant orientation.*

1 Introduction

This paper describes a mobile manipulator, which we call the *mobipulator*. It combines the two functions of locomotion and manipulation in a uniform way: it uses wheels for both. So far the robot has demonstrated the ability to manipulate a piece of paper on a desktop, and also shows potential for manipulating other common desktop objects. This paper introduces the robot, describes several different modes of manipulation and locomotion, outlines the kinematics of the device, and gives some experimental results.

Our goal is to create small robots that can be plugged into a desktop computer with the same ease as a camera or a CD ROM. We believe that motors, sensors, and electronic components have become small enough and inexpensive enough to build robotic systems that would be practical and useful in an office environment. A particular task might be to keep a desktop organized: to store, retrieve, or discard items on demand; to stack papers neatly; to scan papers or books; and so on. This is a complex problem—for the present we will focus on how the mobipulator handles single sheets of paper. Closely related work [15] describes a system that uses a camera to capture electronically and index the contents of the papers contained on a desktop. In the future, we envision combining this previous work with a robot (perhaps some variation on the mobipulator design) to address the larger problems of desktop robotics.

The mobipulator looks like a small car, with four independently powered wheels. None of the wheels are steered. We envision several different modes of manipulation/locomotion, of which four have been demonstrated so far:

- **Translation mode.** This mode is pure locomotion. To translate forward or backward, all four wheels are driven at equal rates.

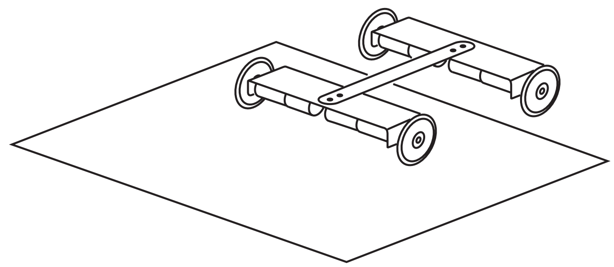

Figure 1: Mobipulator in dual diff drive mode.

- **Dual diff drive mode.** This mode (Figure 1) combines locomotion and manipulation. To maneuver a piece of paper on a desktop, one pair of wheels drives the paper relative to the robot, while the other pair drives the robot relative to the desktop.
- **Inchworm mode.** This mode (Figure 2) is pure manipulation. With all four wheels on the paper, by quick alternation of the front wheels and rear wheels, the paper can be advanced incrementally beneath the robot.
- **Cylinder rolling mode.** This mode (Figure 3) is inspired by dung beetles. The robot's front wheels are placed on a cylinder. It uses its rear wheels to propel itself forward while its front wheels turn in the opposite direction to roll the cylinder forward.
- **Scoot mode.** For this mode (Figure 4) the robot first accelerates hard and then decelerates hard. During acceleration the paper slips backward, and during deceleration the paper slips forward. There is a net forward motion of the paper.

While we have seen inchworm, translation, cylinder rolling, and scoot modes work, so far we have quantitative results only for dual diff drive mode.

Although dual diff drive mode is fundamental to the design of the robot, no single mode is sufficient by itself. For planar motions of the paper, we expect to rely primarily on dual diff drive mode, but also to use translation mode and others not yet demonstrated, such as skid-steering and dragging a paper with just one wheel. To move the paper out of the plane, we expect to use inchworm mode to create a hump in the paper, which might then be maneuvered on top of a corner of a stack of papers. To manipulate other objects, we expect to use a variety of pushing and other modes, perhaps including the cylinder-rolling mode for pencils.

This paper describes the robot in some detail, and documents the performance of the robot in dual diff drive mode. It also outlines some of the many technical challenges and

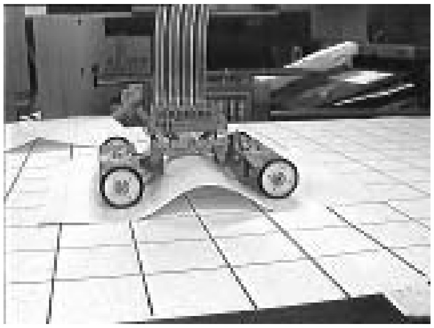

Figure 2: Inchworm mode.

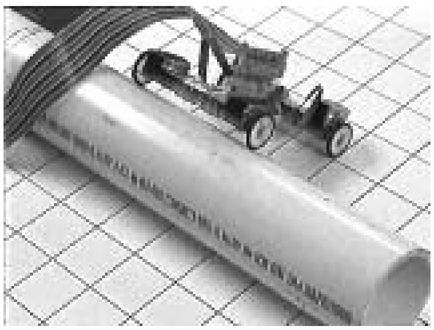

Figure 3: Cylinder rolling mode.

describes some variations of the robot. Finally we discuss some of the motivating ideas and possible applications.

2 Related Work

This section reviews some of the history of mobile manipulation and its relation to the present work. Then we describe the present paper's relation to recent work on minimalism and nonprehensile manipulation.

Several preceding systems have explored the connection between manipulation and locomotion. One of the earliest influential robots, Shakey [11], was a mobile manipulator. It pushed boxes and other objects. Many subsequent mobile robots have incorporated manipulation in some way. Even the recent activities of Sojourner on Mars included elements of manipulation. These and other examples provide a context of broad insights and experience with mobile manipulation.

A direct approach is to attach a manipulator to a mobile platform. The JPL Cart [17] provides an early example, while Romeo and Juliet [8] provide a current example. These projects have demonstrated effective coordination of wheels and arm joints in manipulation tasks.

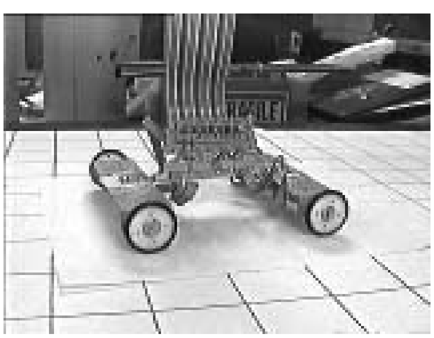

Figure 4: Scoot mode.

One goal of the work described here is to explore the relation of manipulation and locomotion, a goal shared with the distributed manipulation work of Donald et al [4]. This work included a set of mobile robots pushing objects, as if each robot were a finger in a multi-fingered grasp. The OSU Hexapod [16] used dynamic stability analysis and algorithms quite similar to those sometimes used to coordinate or analyze dexterous manipulation with multi-fingered hands.

The Platonic Beast [12] is probably closest in spirit to the present work. It had several limbs, each of which could be used for locomotion or manipulation.

All of these works illustrate that there is a deep connection between locomotion and manipulation, and the present work seeks to take the connection even further: the robot is just one of several movable objects in the task. The job of each actuator is resolved according to the task, be it manipulation, locomotion, or something not clearly classifiable as either.

The present work can also be viewed in relation to other manipulation systems. In fact it fits naturally with work on nonprehensile manipulation. A few examples are manipulation of objects in a tilting tray [6] or on a vibrating plate [14], manipulation of parts on a conveyor belt with a single-joint robot [1], manipulation of planar objects dynamically with simple robots [10], use of a passive joint on the end of a manufacturing arm to reorient parts [13], control of an object by rolling it between two plates [3], and manipulation of planar shapes using two palms [5]. These examples are simpler than a conventional general purpose manipulator, they can manipulate a variety of parts without grasping, and they exploit elements of the task mechanics to achieve goals. In the case of the present work, the robot manages six freedoms with just four motors. Manipulation is accomplished by friction, gravity, and dynamics, without grasping. Perhaps the most relevant previous work is the business card manipulation work of Kao and Cutkosky [7], which addressed manipulation of laminar objects by fingers pressing down from above.

We are just beginning work on the analysis, planning, and control of the mobipulator, which will depend heavily on well-established techniques for non-holonomic robots. [2, 9]

3 Dual diff drive mode

This section examines dual diff drive mode in detail. We begin with some remarks on the configuration space and the quality of paths. Then we develop the kinematic equations describing the paper's motion as a function of wheel velocities. The section concludes with some experimental results.

3.1 Paths

We assume that one pair of wheels, either the front or the back, is on the paper, and that the other pair is on the desk. Each pair defines a differential drive. For small motions we can consider the system to be decomposed into a chain of the

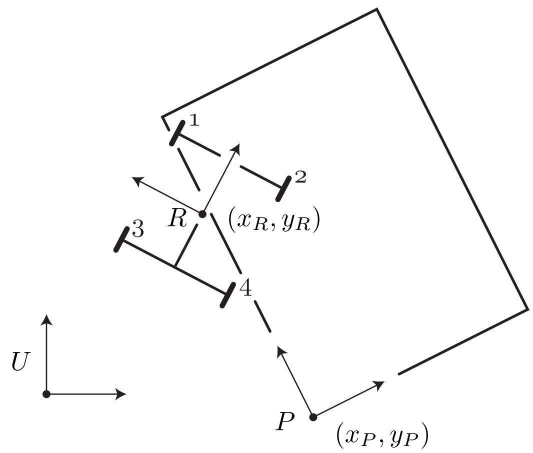

Figure 5: Notation for dual diff drive

two diff drives acting independently. The robot relative to the desk is a nonholonomic three degree-of-freedom system, and so is the paper relative to the robot. The entire system has six degrees of freedom but only four controls. It follows directly from well known results of differential drives that the entire system is small-time controllable, and that paths in the configuration space can be followed as closely as desired.

One important complication is an obstacle that is hard to ignore: the edge of the paper. To remain in dual diff drive mode, the paper's edge must pass between the front and rear wheels. Consider the problem of driving an automobile down the highway, while keeping the front wheels in the left lane and the rear wheels in the right lane. A long path would require a long sequence of tedious parallel parking maneuvers.

The configuration space comprises two components, what we might call *front-wheels-on* C-space, and *rear-wheels-on* C-space. Each of these components is very short in the car's preferred direction of travel, and long in the tedious direction.

For that reason, dual diff drive by itself is limited. However, if we also allow the robot to use translation mode, then it can drive directly across the paper, switching from front-wheels-on to rear-wheels-on, for example. The resulting plans are much more efficient.

3.2 Kinematic analysis

This section derives kinematic equations and shows that the motion of the paper relative to the desktop is unconstrained. We assume a coordinate frame U fixed to the desktop, a coordinate frame R fixed to the robot, and a coordinate frame P fixed to the paper. (See figure 5.) Let the origins of the R and P frames be (x_R, y_R) and (x_P, y_P) respectively, measured in the U frame. Let the angles be given by θ_R and θ_P. Then the configurations of the robot and the paper are given by $q_R = (x_R, y_R, \theta_R)$ and $q_P = (x_P, y_P, \theta_P)$. The velocities of the two frames will be taken relative to the fixed U frame,

$$\begin{pmatrix} u_R \\ v_R \\ \omega_R \end{pmatrix} = \frac{d}{dt} \begin{pmatrix} x_R \\ y_R \\ \omega_R \end{pmatrix}$$

$$\begin{pmatrix} u_P \\ v_P \\ \omega_P \end{pmatrix} = \frac{d}{dt} \begin{pmatrix} x_P \\ y_P \\ \omega_P \end{pmatrix}$$

We will follow the common convention of using pre-superscripts to indicate coordinate frames. Thus $({}^R x_P, {}^R y_P)$ gives the origin of the paper frame relative to the robot frame, in robot frame coordinates, while $({}^R u_P, {}^R v_P)$ gives the paper frame origin velocity in robot frame coordinates. But observe that velocities are relative to the U frame, then transformed to other frames.

Let ω_i be the angular velocity of wheel i, with the positive sense defined to correspond with a forward motion of the car. Let c be the wheel radius, and let a and b correspond to the chassis half-length and half-width respectively. If we assume no wheel slip, then with some calculation one can obtain for the robot's motion:

$$\begin{pmatrix} {}^R u_R \\ {}^R v_R \\ \omega_R \end{pmatrix} = \begin{pmatrix} \frac{c}{2}(\omega_3 + \omega_4) \\ \frac{ac}{2b}(\omega_4 - \omega_3) \\ \frac{c}{2b}(\omega_4 - \omega_3) \end{pmatrix}$$

from which the nonholonomic constraint may be derived:

$${}^R v_R - a\omega_R = 0$$

A similar set of equations may be derived for the paper's motion relative to the robot. Combining, we obtain the paper's motion relative to the fixed U frame, but for simplicity expressed in robot frame R coordinates:

$$\begin{pmatrix} {}^R u_P \\ {}^R v_P \\ \omega_P \end{pmatrix} = f_1\omega_1 + f_2\omega_2 + f_3\omega_3 + f_4\omega_4$$

where each f_i is the vector field arising from driving just one wheel:

$$f_1 = \begin{pmatrix} -\frac{c}{2} - \frac{c}{2b}{}^R y_P \\ \frac{c}{2b}{}^R x_P - \frac{ac}{2b} \\ \frac{c}{2b} \end{pmatrix}, \quad f_2 = \begin{pmatrix} -\frac{c}{2} + \frac{c}{2b}{}^R y_P \\ -\frac{c}{2b}{}^R x_P + \frac{ac}{2b} \\ -\frac{c}{2b} \end{pmatrix}$$

$$f_3 = \begin{pmatrix} \frac{c}{2} + \frac{c}{2b}{}^R y_P \\ -\frac{ac}{2b} - \frac{c}{2b}{}^R x_P \\ -\frac{c}{2b} \end{pmatrix}, \quad f_4 = \begin{pmatrix} \frac{c}{2} - \frac{c}{2b}{}^R y_P \\ \frac{ac}{2b} + \frac{c}{2b}{}^R x_P \\ \frac{c}{2b} \end{pmatrix}$$

Then we see that the paper's motion is unconstrained by finding linear combinations of the four vector fields which span the space of paper velocities.

$$f_1 + f_2 - f_3 - f_4 = \begin{pmatrix} -2c \\ 0 \\ 0 \end{pmatrix}$$

$$f_1 - f_2 + f_3 - f_4 = \begin{pmatrix} 0 \\ -\frac{2ac}{b} \\ 0 \end{pmatrix}$$

$$f_1 - f_2 - f_3 + f_4 = \begin{pmatrix} -\frac{2c}{b} {}^R y_P \\ \frac{2c}{b} {}^R x_P \\ \frac{2c}{b} \end{pmatrix}$$

It is readily observed that the above fields span $\Re^3$ for all configurations of the system. The system also has one "redundant" degree of freedom, corresponding to translation mode, where the robot rolls forward but the paper doesn't move:

$$f_1 + f_2 + f_3 + f_4 = 0$$

Thus, although the robot and paper system is nonholonomic, if we consider just the paper relative to the desktop, we can achieve velocities of any direction, as if it were a holonomic system. In principle we can close a feedback loop about the paper's configuration, although we have yet to do so.

3.3 Experiments

The goal of our experiments was to test the basic concept of dual diff drive, especially three key assumptions: that the paper will slip on the desktop, that the wheels will not slip on the desktop, and that the wheels will not slip on the paper. We were also concerned with the possibility that the paper would gather between two wheels, creating an unwanted hump. (Our chassis design allows some control of sideways humps, as in inchworm mode, but provides no control of humps parallel to the wheels.) Our conclusions are that it works fairly well, and that the slip assumptions are satisfied well enough for effective locomotion and manipulation. It came as no surprise that unwanted slip is significant enough to defeat accurate odometry. The uncontrollable hump problem was never observed.

Experimental Setup The car has a square wheelbase of approximately 125mm along each edge. All four wheels are independently powered by DC servomotors with gearheads and incremental encoders. The front wheels form one module and the rear wheels a second. The front and rear modules are joined by a steel strap that is flexible in torsion to allow the load to be evenly distributed between the four wheels.

The mobipulator was operated in dual diff drive mode, and was programmed to move a piece of square mylar about a desk in a rounded square, while maintaining constant orientation of the mylar. The pattern was composed of 8 separate segments: 4 arcs and 4 straight edges. Each straight edge moved the rear wheels (those not on the mylar) 100mm while keeping the front still. The corner moves consisted of moving the inner rear wheel along an arc of 89.5mm, while the outer rear wheel traversed an arc of 303mm. The front wheels traveled 106.8mm in opposite directions, rotating the mylar 90 degrees in reference to the mobipulator, while maintaining constant orientation in reference to the desk.

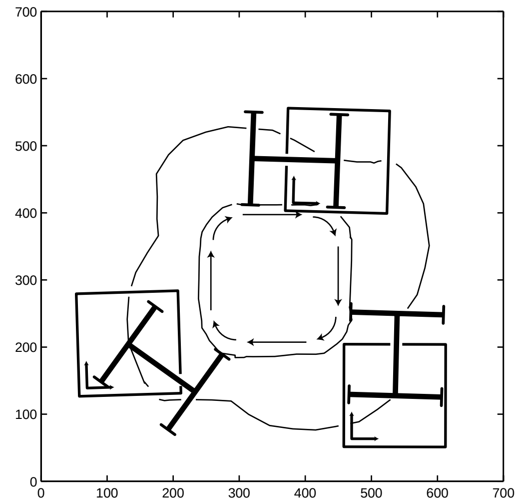

Figure 6: The mobipulator driving a rounded square, while holding the orientation of a mylar page constant. Dimensions are mm.

We acquired data using a video camera positioned above the desk. The surface of the desk was covered with paper containing a 50mm square reference grid. We recorded the mobipulator motions, captured individual frames, and digitized points by hand to generate plots of both the mylar and the mobipulator motion.

The control hardware generates trapezoidal velocity profiles. The parameters of the commanded motions were chosen to ensure that accelerations commence and cease simultaneously for all wheels. In joint space each motion segment is a straight line, even during acceleration and deceleration. The turning motions took 5 seconds, while the edges took 2 seconds. The system operated open-loop, except for the PID control of the individual wheels.

Results Figure 6 shows typical paths of the mobipulator and the mylar sheet it is maneuvering. One curve is a trace of the mylar sheet's center point; the other curve is a trace of the right back wheel.

Figure 7 is a plot of the mylar sheet's orientation. It was commanded to maintain a constant zero angle, so any motion is an error. After rotating 2π relative to the robot, the max error is around 0.1 radians, or about $6°$.

There seem to be two main sources of error. The first problem is slip of the wheels against either the paper it is manipulating, or more commonly the desk surface. This problem is worst when moving the paper about the corners of the square, which requires more effort from the rear wheels due to an increased lever arm. The second problem is tracking error during the motions, which affects synchronization of the wheels. The worst case occurs during a turn when the loads on the rear wheels differ, and so of course the servo error differs, leading to a departure from the programmed path.

Lessons learned We learned four main things from these experiments. First, we feel encouraged that the basic idea

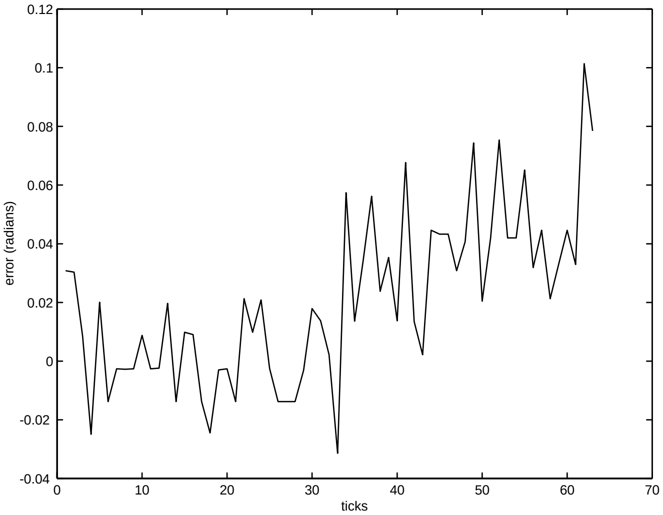

Figure 7: Control error for page orientation, which was commanded to be a constant zero, during the motion of figure 6. "Ticks" is a nonuniform sampling of time. Duration of the motion was one minute, with about half that time spent at rest between motion commands, due to communication delays.

worked so well on our first attempt. Second, not surprisingly, we have to provide means to deal with error due to wheel slip. Third, materials and wheel loads are important. We tried various combinations until we found one that worked best. We even cheated a little, using ballast to put a higher wheel load on the rear wheels than on the front, and using mylar to reduce slip. Fourth, independent control of the wheels is a problem. Although the servos seem to settle accurately enough, because of the nonholonomic nature of the system errors during the trajectory lead to significant errors in the final robot position.

4 Other modes

We envision a number of different motion modes. Besides the four modes described above, we plan attempts to drag a page with one wheel, to translate the page with two wheels on the same side of the robot, to use microslip to turn the device without a paper, to turn the pages of a book, and a variety of other techniques. However, we will focus on the four modes we have already tried.

Inchworm mode. If we first advance the front wheels of the car, and then the rear wheels, and repeat, what happens to the paper? The answer depends on the normal loads applied at each wheel, the frictional characteristics of each contact, and on the nature of the wheel motions. If these factors can be managed appropriately, the paper will incrementally advance beneath the car. Our first attempt with inchworm mode worked. It appears that a hard acceleration with one set of wheels will advance the paper without moving the robot. Possibly the reaction force is filtered by the mass of the car in combination with the servo stiffness, so that the other set of wheels never sees the peak force and hence does not reach the limit of static friction. We are presently working with simulations to explore this hypothesis.

Cylinder rolling mode. The design of the mobipulator was partly inspired by the noble dung beetle, which manipulates its dung ball with its rear legs as it walks with its front legs. On the desktop this seems a natural way to manipulate pencils. We had no difficulty rolling a small brass cylinder and a large PVC pipe several inches, but our first attempts to repeat the feat with a pencil were not successful. We suspect that slip of the pencil is a problem, especially with the hexagonal cross section.

5 Discussion

What is the difference between locomotion and manipulation? The answer seems obvious: it is the difference between moving onself relative to the earth, versus moving some collection of objects relative to oneself. The distinction is straightforward when comparing most mobile robots with most robotic manipulators. However, the distinction becomes awkward for some tasks, as illustrated by the mobipulator.

The design of hands, arms, legs, feet, and wheels really depends on our relation to the objects we interact with, not on whether it is locomotion or manipulation. The mobipulator is designed for a manipulation task that closely resembles indoor locomotion. A differential drive is a good way to manipulate a large flat object, just as it is a good way to locomote on a large flat floor.

The similarity between the mobipulator's locomotion and manipulation tasks means that its motors can be dynamically reassigned to either task. This is a characteristic of animal behavior that we might profitably emulate. It allows a more opportunistic approach to planning and control, and should increase the capabilities of a robot.

There are many other ways of driving the mobipulator. We hope it will be able to pick one sheet from a stack of paper, to place a sheet back on a stack, to turn a page of a book, and so on. There are a number of interesting technical problems to pursue: coordinated control of the wheels, automatic planning, analysis of different motion modes, and variations on the mechanical design to name just a few.

6 Conclusion

The mobipulator is a simple system for manipulating laminar objects on a desktop, and perhaps for a variety of other small-scale mobile manipulation tasks. Early experiments are encouraging, although there are a number of unresolved issues including control, planning, mechanical design, perception, design, and applications.

Acknowledgments

Many of these ideas arose during discussions with Illah Nourbakhsh. Stephen Moore and Jason Powell helped in the design and construction of the mobipulator. This work is supported by the National Science Foundation under Grant IRI-9318496. DKP would like to thank E. P. Krotkov, MAE, and MTM for hosting his sabbatical at CMU, and NSERC for financial support.

References

[1] 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.

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

[3] A. Bicchi and R. Sorrentino. Dexterous manipulation through rolling. In *IEEE International Conference on Robotics and Automation*, pages 452–457, 1995.

[4] B. R. Donald, J. Jennings, and D. Rus. Analyzing teams of cooperating mobile robots. In *Proceedings 1994 IEEE International Conference on Robotics and Automation*, 1994.

[5] M. A. Erdmann. An exploration of nonprehensile two-palm manipulation. *International Journal of Robotics Research*, 17(5), 1998.

[6] M. A. Erdmann and M. T. Mason. An exploration of sensorless manipulation. *IEEE Transactions on Robotics and Automation*, 4(4):369–379, Aug. 1988.

[7] I. Kao and M. R. Cutkosky. Dextrous manipulation with compliance and sliding. In H. Miura and S. Arimoto, editors, *Robotics Research: the Fifth International Symposium*. Cambridge, Mass: MIT Press, 1990.

[8] O. Khatib, K. Yokoi, K. Chang, D. Ruspini, R. Holmberg, A. Casal, and A. Baader. Force strategies for cooperative tasks in multiple mobile manipulation systems. In G. Giralt and G. Hirzinger, editors, *Robotics Research: The Seventh International Symposium*, pages 333–342, 1996.

[9] J. P. Laumond. Nonholonomic motion planning for mobile robots. Technical report, LAAS, 1998.

[10] K. M. Lynch, N. Shiroma, H. Arai, and K. Tanie. The roles of shape and motion in dynamic manipulation: The butterfly example. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pages 1958–1963, 1998.

[11] N. J. Nilsson. Shakey the robot. Technical Report 323, SRI International, 1984.

[12] D. K. Pai, R. A. Barman, and S. K. Ralph. Platonic beasts: a new family of multilimbed robots. In *Proceedings 1994 IEEE International Conference on Robotics and Automation*, pages 1019–1025, 1994.

[13] A. Rao, D. Kriegman, and K. Y. Goldberg. Complete algorithms for reorienting polyhedral parts using a pivoting gripper. In *IEEE International Conference on Robotics and Automation*, pages 2242–2248, 1995.

[14] D. Reznik and J. Canny. A flat rigid plate is a universal planar manipulator. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pages 1471–1477, 1998.

[15] D. Rus and P. deSantis. The self-organizing desk. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1997.

[16] S.-M. Song and K. J. Waldron. *Machines that Walk: The Adaptive Suspension Vehicle*. MIT Press, 1988.

[17] A. M. Thompson. The navigation system of the jpl robot. In *Fifth International Joint Conference on Artificial Intelligence*, pages 749–757, 1977.