# Experiments with Desktop Mobile Manipulators

Matthew T. Mason
Carnegie Mellon University
Pittsburgh PA, USA
matt.mason@cs.cmu.edu

Dinesh K. Pai
University of British Columbia
Vancouver, BC, Canada
pai@cs.ubc.ca

Daniela Rus and Jon Howell
Dartmouth
Hanover NH, USA
rus@cs.dartmouth.edu and jonh@cs.dartmouth.edu

Lee R. Taylor and Michael A. Erdmann
Carnegie Mellon University
Pittsburgh PA, USA
lrt+@andrew.cmu.edu and me@cs.cmu.edu

**Abstract:** This paper describes our work on Desktop Robotics. The main focus is two robots that locomote and manipulate paper on a desktop. One robot uses wheels for both manipulation and locomotion. The other robot uses wheels for locomotion and a sticky foot to lift and carry pieces of paper. We outline the goals of our work on desktop robotics, describe the design of the robots built so far, present experimental data, and outline some of the issues for future work.

## 1. Introduction

Desktop robotics addresses applications of robotics in office environments: storage and retrieval of documents and other desktop paraphernalia, integration of hardcopy documents with computer files, or perhaps routine maintenance and cleaning tasks. The goal is to create small robots that can be plugged into a desktop computer with the same ease as a camera, a CD ROM, or any other peripheral device. 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. Besides the possible practical value of such systems, we are interested in the fresh perspective on research issues. We believe that ideas arising from this work may be useful in a variety of applications, for example manipulating flexible objects, designing prosthetic devices, or human-robot interaction.

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 which may ultimately involve a large number of different sensor and motor systems. 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 robots described in the present paper) to address the larger problems of desktop robotics.

This paper focuses on two robots for manipulating paper, and perhaps other objects, on a desktop. The first is a mobile manipulator that we call the *mobipulator*. It looks like a small car, with four wheels independently driven, none of them steered (Figure 1). It uses its wheels both for locomotion and manipulation. The second robot, called *Fiat*, can manipulate paper on a desktop, even when the surface is crowded with several sheets of paper. Fiat has a sticky foot, which it uses to lift a sheet of paper, carry it to a designated location of the desk, and place it down on the desk. Section 2 describes the Mobipulator, and Section 3 describes Fiat. Section 4 address previous work, and Section 5 is a discussion and conclusion.

## 2. The Mobipulator

The Mobipulator (Figure 1) 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 five 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.
- **Dual diff drive mode.** This mode (Figure 2) 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 1) 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 2) 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.** With all four wheels on the paper (Figure 1) the robot accelerates hard and then decelerates hard. During the high acceleration the paper slips backward, and during the high deceleration the paper slips forward. There is a net forward motion of the paper.

While we have seen all of these techniques work in the lab, so far we have performance data only for dual diff drive mode.

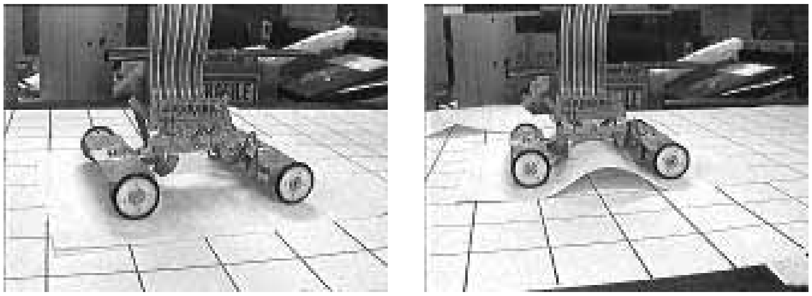

Figure 1. (Left) Mobipulator I, shown in scoot mode on a sheet of US standard letter paper. (Right) Inchworm mode.

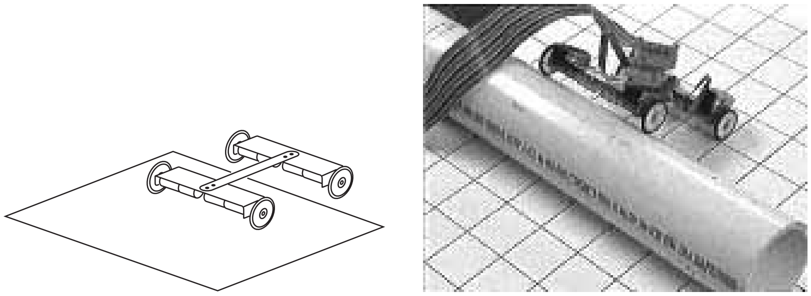

Figure 2. (Left) Dual diff drive mode. (Right) Cylinder rolling mode.

A typical scenario would switch modes freely, to avoid the limitations inherent to any single mode. The robot might begin with translation mode (and perhaps some slip-steering) to roll on to a piece of paper. Then would follow a sequence of dual diff-drive motions alternating with translations, to coarsely position the robot. Finally, the robot can use dual diff-drive mode to perform some closed-loop fine positioning of the paper. More exotic (and highly speculative) scenarios involve the robot using inchworm mode to maneuver a page onto a stack of paper, or using other modes to turn the pages of a book.

These scenarios will require research progress in a number of areas: the interaction between tire and paper, the behavior of a hump in the paper during inchworm mode, and planning techniques that integrate a variety of fundamentally different actions.

### 2.1. Experiments with the Mobipulator

The goal of our experiments was to test the basic concept of dual diff-drive mode, and specifically to see whether unwanted slip would be a problem.

The car has a square wheelbase of about 125mm along each edge. Each wheel is driven by a DC gearmotor with incremental encoder. The front half of the car is joined to the rear half by a thin piece of steel, providing a suspension to evenly distribute the load among the wheels. The tires are rubber O-rings. The "desktop" was cardboard, and the "paper" was mylar.

The mobipulator was operated in dual diff drive mode, moving the square piece of mylar along a path in the shape of a rounded square, while holding the mylar orientation constant. The motions had trapezoidal velocity profiles, with synchronized changes in acceleration. Each wheel was controlled independently by a PID servo. To record the motion, we hand-digitized data taken with a video camera positioned above the desk.

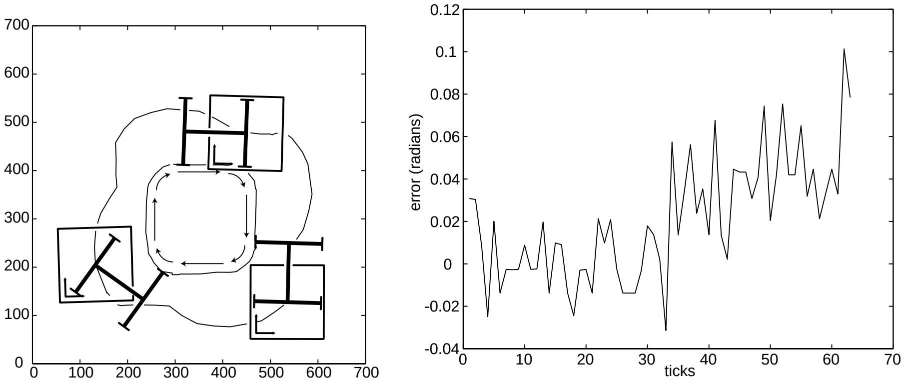

Figure 3. (Left) The mobipulator driving a rounded square, while holding the orientation of a mylar page constant. One curve is a trace of the mylar sheet center; the other curve is a trace of the right back wheel. Dimensions are mm. (Right) Control error for page orientation during the motion of the left figure. The page orientation was commanded to be a constant zero, so any motion is an error. "Ticks" is a nonuniform sampling of time. Duration of the motion was one minute, with about half that time spent at rest. After rotating $2\pi$ relative to the robot, the max error is around 0.1 radians, or about 6°.

Our conclusions were that unwanted slip is bad enough to require additional sensors, but not bad enough to prevent effective manipulation and locomotion. We also discovered that we need a better approach to control, coupling the wheel servos and attending to the nonholonomic nature of the device.

## 3. Fiat: the paper-lifting robot

The Mobipulator can move single sheets of paper in the $x$ and $y$ directions when the papers rest directly on a desktop. To cope with areas of desktops that are crowded with stacks of paper we designed the mobile robot called *Fiat* (see Figure 4). In the future, we plan to combine the functionality of Fiat and the Mobipulator into a single mobile robot.

Fiat can lift a single piece of paper from a stack and move it to a different location on the desktop. The mechanical design for Fiat was motivated by the goal of moving paper in the vertical direction, in order to stack sheets one above the other deterministically. The manipulator of this robot is a foot made sticky with a supply of removable tape. A spool of removable adhesive tape provides a supply of "fresh stickiness" for the foot (see Figure 4).

The manipulator needs a way to attach the foot to a piece of paper and to release the paper from the foot; this is accomplished by integrating the manipulator with the locomotive design. The foot travels on an eccentric path relative to the free rear wheels. In one direction, the foot emerges below the wheels, using the weight of the robot to stick the foot to the paper. As the eccentric motion continues, the foot moves behind the wheels, and lifts up the paper. The opposite motion detaches the foot from the paper, using the rear wheels to hold down the paper while the sticky foot is freed. Notice that the wheels have been placed over the edge of the sheet of paper, holding it in place and pulling it free from the sticky foot. The result is that we can reliably grab and release a flexible sheet of paper using only a single rotary motion.

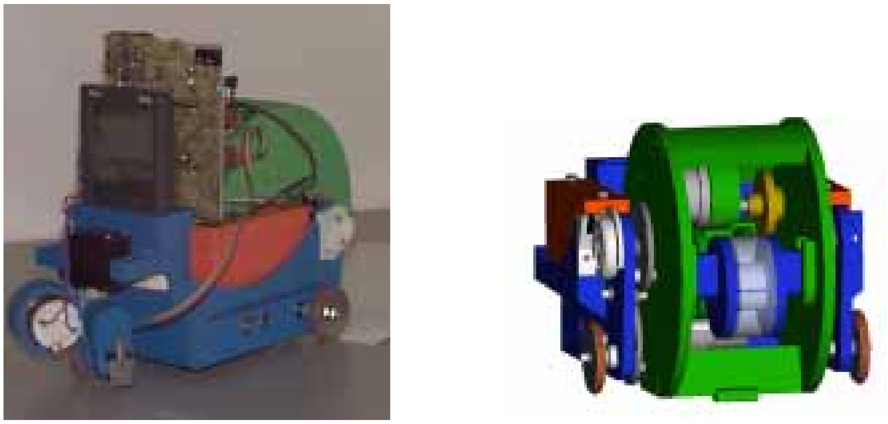

Figure 4. (Left) A picture of the paper lifting robot. (Right) Details of the sticky foot mechanism of the robot Fiat.

There are four total actuators (steering servo, drive motor, cam motor, tape winding motor). There are also two internal sensors: one shaft encoder to allow dead reckoning motion, and one optical transmissive sensor to detect the phase of the cam assembly.

The total weight of this robot is 2.7 kg (5.9 pounds), of which the batteries account for 0.6 kg (1.4 pounds). Exterior dimensions are approximately 29 cm long by 23 cm high by 17 cm wide (11.5 x 9 x 6.7 inches). Our fabrication method permits rapid redesign and prototyping. We expect that each linear dimension will be halved by the next design revision, to enable the robot to operate in much smaller spaces.

The Fiat robot contains a processor that is used for controlling the motors and future sensors for this robot. A 3Com Palm III personal organizer serves as the main processor on the robot. We chose this approach for several reasons: (1) A PalmOS/Motorola 68000 compiler is readily available for open systems; (2) Its built-in user interface makes "field tests" very convenient; (3)It has a built-in infrared data port, which allows the robot to interface to an external camera hosted by a computer without a tether; (4) Its generous 4MB of memory and 16MHz processor come in a small, inexpensive package that is hard to match with custom components; (5) It consumes very little power; and (6) Its "Hot Sync" function serves as a very convenient bootstrap mechanism. Figure 6 shows the control components for this robot.

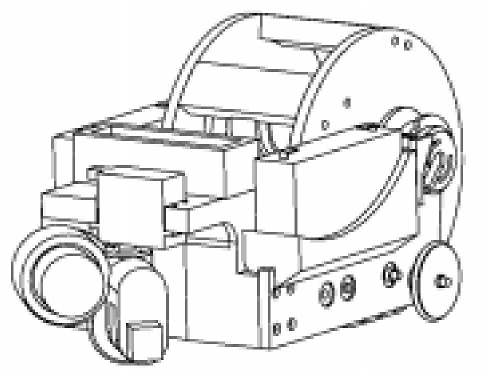

Figure 5. The CAD model of the paper-lifting robot Fiat used for fabrication.

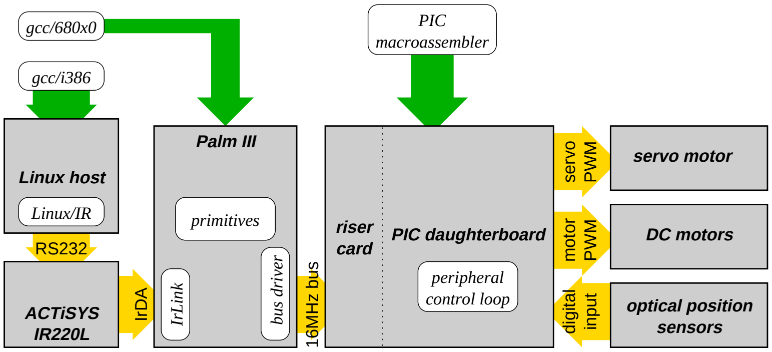

Figure 6. The schematic diagram of the paper-lifting robot Fiat. The Palm III interfaces with an expansion board that contains a 4MHz Microchip PIC 16C74A microcontroller and I/O pins through the Palm's memory bus.

### 3.1. Control and Capabilities

The basic capabilities of Fiat consist of planar movement on a desk top, lifting a paper, releasing a paper, and advancing the tape on the sticky foot. The following primitives were implemented on the Palm III to effect these capabilities:

1. `drive(float distance)` moves the robot straight for `distance` centimeters, dead reckoning with the drive wheel's shaft encoder. Speed is ramped along a trapezoidal envelope.

2. `rotate(float angle)` turns the steering wheel perpendicular to the rear wheels and rotates the robot `angle` degrees around the midpoint of the rear wheels, the location of the sticky foot when retracted.

3. `paperPickup()` winds the sticky tape spool a little, then rotates the cam out to pick up a sheet of paper. Each rear wheel should be approximately at the edge of the paper.

4. `paperDrop()` rotates the cam inward, detaching the paper from the foot.

Using these primitives, the robot can move a sheet of paper from one side of a desk to the other. Since the leading edge of the paper is at least a centimeter above the lowest point on the rear wheels, the motion of the paper is over other sheets of paper that are resting on the desk.

### 3.2. Experiments

We have run two sets of performance tests on the Fiat robot.

A first set of experiments was designed to quantify the reliability of shifting a single piece of paper on a desk top. Placement can be implemented as a sequential execution of `paperPickup()`, `rotate(float angle)`, and `paperDrop()`. A single sheet of paper was placed in front of the robot. The robot was commanded to pick up the paper, rotate it 180 degrees, and place it down. The location of the paper was then compared with the expected location to estimate the placement error. This experiment was repeated 54 times. We observed a 93% success rate for this cycle. Each experiment resulted in a reasonable location for the placement of the paper. The 7% error were due to problems in the cam alignment for the paper lifting mechanism. We believe that this problem can be solved by using a better shaft encoder. To measure the placement errors, we measured the location of one point (in terms of its $x$ and $y$ coordinates) and the orientation of the paper. We found that the standard deviations in the $(x, y)$ location of the point was $(0.47cm., 0.59cm.)$, and the standard deviation for the orientation was 3.29 degrees. Since the standard deviation numbers are low, we conclude that the experiment has a high degree of repeatibility.

We also conducted an experiment to test how well the robot can shift piles of paper. The basic experimental setup consists of a pile of three sheets of paper placed at a designated location on a desktop. The robot starts at a known location and follows a pre-coded trajectory to the location of the paper stack. The robot then uses the sticky foot to lift the top page and transport it to another pre-specified location on the desktop, following a given trajectory. After depositing the paper, the robot returns to its starting location and the process is ready to be repeated. We have repeated this loop many times. We observed failures in the system due to odometry. Because the robot uses no sensors to navigate, position errors accumulate and result in a misalignment between the robot and the paper stack. We plan to enhance the architecture of the robots with sensors that will provide guidance and paper detection capabilities.

## 4. Related Work

This section reviews some previous work and its relation to the present paper.

Several preceding systems have explored the connection between manipulation and locomotion. One of the earliest influential robots, Shakey [11], was a mobile manipulator. 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.

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.

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, each robot uses four motors to manage six freedoms (mobipulator) or seven freedoms (Fiat). Manipulation is accomplished by friction, gravity, dynamics, and adhesion, 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 still have a great deal of work to do on analysis, planning, and control, which will depend heavily on well-established techniques for non-holonomic robots. [2, 9]

## 5. Discussion and Conclusion

Our goal is to develop reliable desktop robots, capable of performing useful and interesting tasks on a desktop, while attached as peripherals to a desktop computer. In this paper we explore two designs for robot systems that allow the robots to move paper on a desktop in the $x$, $y$, and $z$ directions. Both robots are small in scale and architecturally minimalist.

More generally, our robots explore the deep connection between locomotion and manipulation. While several authors have noted the existence of this connection (see §4) 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.

Our experiments show that these robots are capable of manipulating individual pieces of paper in the $x$, $y$, and $z$ directions on a desktop[1]. Many components of these demonstrations were hard wired. The purpose of the experiments was to demonstrate the basic capabilities of the robots to manipulate paper and other common objects on the desktop. We are currently incorporating a planning module in the system.

Our systems are very preliminary. In the future, we plan to enhance these robots by combining their two functionalities into one mobile robot and to augment their mechanical designs with sensing. Another goal is to develop automated planners able to compute complex motions for desktop tidying tasks. We plan to incorporate (1) an external vision system to guide the robot; (2) reflective optical sensors for detecting edges of contrasting color in order to help the robot align to the edge of a sheet of paper on a contrasting background; (3) an inkjet head to be used to print barcodes on new papers arriving on the desk; and (4) optical sensors that, when combined with distance information from the shaft encoder, could be used to read barcodes on pages as the robot passes over them.

## 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. We thank Keith Kotay for many valuable discussions. 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. Daniela Rus has been supported under NSF Career award IRI-9624286.

## 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.

---

[1] Currently, neither robot can handle thick stacks of stapled papers, which is a challenge for the future.

[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.