PROGRAMMABLE FORCE FIELDS FOR DISTRIBUTED MANIPULATION, AND THEIR IMPLEMENTATION USING MICRO-FABRICATED ACTUATOR ARRAYS

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Karl-Friedrich Böhringer
August 1997

PROGRAMMABLE FORCE FIELDS FOR DISTRIBUTED MANIPULATION, AND THEIR IMPLEMENTATION USING MICRO-FABRICATED ACTUATOR ARRAYS

Karl-Friedrich Böhringer, Ph.D.
Cornell University 1997

Programmable force vector fields can be used to control a variety of flexible planar parts feeders such as massively-parallel micro actuator arrays or transversely vibrating (macroscopic) plates. These new automation designs promise great flexibility, speed, and dexterity—they may be employed to position, orient, singulate, sort, feed, and assemble parts.

A wealth of geometric and algorithmic problems arise in the control and programming of manipulation systems with many independent actuators. The theory of programmable force fields represents the first systematic attack on massively-parallel distributed manipulation based on geometric and physical reasoning. We show how to develop combinatorially precise planning algorithms that synthesize force field strategies for controlling a very large number of distributed actuators in a principled, geometric, task-level fashion.

When a part is placed on our devices, the programmed force field induces a force and moment upon it. Over time, the part may come to rest in a dynamic equilibrium state. By chaining together sequences of force fields, the equilibrium states of a part in the field may be cascaded to obtain a desired final state. The resulting strategies require no sensing and enjoy efficient planning algorithms.

This thesis introduces new experimental devices that can implement programmable force fields. In particular, we describe the M-CHIP (Manipulation Chip), a massively-parallel array of programmable micro-motion pixels. Both the M-CHIP, as well as macroscopic devices such as transversely vibrating plates, may be programmed with force fields, and their behavior predicted and controlled using our *equilibrium analysis*. We demonstrate lower bounds (i.e., impossibility results) on what the devices *cannot* do, and results on a *classification of control strategies* yielding design criteria by which well-behaved manipulation strategies may be developed. We define *composition operators* to build complex strategies from simple ones, and show the resulting fields are also well-behaved.

Finally, we consider parts feeders that can only implement a very limited "vocabulary" of force fields. We show how to plan and execute parts-posing and orienting strategies for these devices, but with a significant increase in planning complexity and some sacrifice in completeness guarantees. We discuss the tradeoff between mechanical complexity and planning complexity.

Biographical Sketch

Karl-Friedrich Böhringer was born in Freudenstadt, Germany and grew up in the Black Forest town of Baiersbronn. He attended the Technical University of Karlsruhe as a student in computer science, and graduated with the degree of Diplom-Informatiker. He then joined Cornell University in Ithaca, New York to pursue a doctorate degree. During his graduate studies, he spend a year as a Visiting Scholar at the Robotics and the Transducers Laboratories at Stanford University. After receiving his Ph.D. in computer science, he departed Cornell for a postdoctoral researcher position at the University of California at Berkeley.

Acknowledgements

This thesis, and the part of my life that I spent working on it, has been largely influenced and supported by many people who were of immeasurable help and inspiration to me. My thesis advisor and committee chair, Bruce Randall Donald, has been a never-ending source of new ideas. He has provided me with support and guidance throughout my time as a graduate student. In many exciting conversations, my advisor Noel C. MacDonald has given me vision, determination, and optimism. Thanks also to my committee members Daniel Huttenlocher and Pradeep Chintagunta whose continuing help and advice I greatly appreciated.

The following is hardly a complete list of people who have all contributed to this thesis — I am very grateful for their ideas and help: Tamara Lynn Abell, Eric Babson, Vivek Bhatt, John Canny, Bernard Chazelle, Paul Chew, Perry Cook, Subhas Desa, Mike Erdmann, Steve Glander, Ken Goldberg, Danny Halperin, Srikanth Kannapan, Lydia Kavraki, Greg Kovacs, Jean-Claude Latombe, Matt Mason, Al Rizzi, Ivelisse Rubio, Andy Ruina, Ken Steiglitz, John Suh, and Andy Yao. Thanks especially to Jean-Claude Latombe and Greg Kovacs for their hospitality during my stay at their labs at Stanford University.

MEMS fabrication would have been impossible for a “theorist” without the help of the staff and students at the Cornell Nanofabrication Facility, in particular the MacDonald group members Scott Adams, Arturo Ayon, Liang Chen, Jon Das, Dan Haronian, Wolfgang Hofmann, Trent Huang, Ali Jazairy, Rob Mihailovich, Scott Miller, Ikuo Ogo, Rama Prasad, Bryan Reed, Taher Saif, Kevin Shaw, Russ Webb, Peter Kenji Yamasaki, and Yang Xu.

My fellow roboticists at the Cornell Robotics & Vision Lab, Amy Briggs, Russell Brown, Jim Jennings, Dinesh Pai, Jonathan Rees, Daniela Rus, Pat Xavier, as well as Vivek Bhatt and Ken Goldberg have all contributed to a creative and productive environment for my research.

Greg Kovacs and John Suh at the Transducers Lab, Stanford University were great colleagues who generously made their micro cilia chips available to us.

Thanks to Bruce Land, Chris Pelkie, and Martin Berggren at the Cornell Theory Center - Scientific Visualization for their help in preparing videos and animations, and for giving us access to the latest developments and tools for digital image and video processing.

Support for our robotics research was provided in part by the National Science Foundation (NSF) under grants IRI-8802390, IRI-9000532, IRI-9201699, and by a Presidential Young Investigator award to Bruce Donald, in part by an NSF/ARPA Small Grant for Exploratory Research IRI-9403903, and in part by the AFOSR, the Mathematical Sciences Institute, Intel Corporation, and AT&T Bell laboratories. This work was supported by

ARPA under contract DABT 63-69-C-0019. All fabrication of single-crystal silicon actuator arrays was performed at the Cornell Nanofabrication Facility at Cornell University which is supported by ARPA, NSF, and Industrial Affiliates.

*

During my stay at Cornell, I was lucky to meet truly special people, who have been with me at all times. Hoai Huong Tran supported me in any possible way. My friend Vivek Bhatt has shared and given so much to me — thank you! The same is true for Errol Montes-Pizarro and Hideki "Trip" Suzuki.

I am infinitely grateful to our great teacher Errol Montes-Pizarro, as well as to Kjartan Stefansson, Hoai Huong Tran, and all those who spent many late nights with me playing the literally divine game of Dominoes.

I thank Vivek Bhatt, Hideki "Trip" Suzuki, Andrew Davidson, and all the other rock climbers for low-gravity days at the Lindseth climbing wall, the Gunks, Joshua Tree, the still unconquered Grumman Wall, and the Chariot's.

Thanks to my Tae Kwon Do instructors Sandy Glattner, Raphael Gilbert, Mary Wimsatt, Grand Master Duk Sun Son, and to all the members of our do-jo.

Thanks also to my house mates in the Cornell golf course, in the famous Albany Street House, on Slaterville Road, in Cascadilla Park, on Maple Hill, and on University Avenue.

Daniel Scharstein provided good ideas and conversation in various German dialects – fia'di! My office mates Michael Slifker, Richard Chapman, Takako Hickey, and Rick Aaron have kept me in good company while at work. Thanks to my office neighbor Aswin van den Berg for his excellent selection of music, bringing to life the Cornell computer science department on lonely weekend nights.

Jan Batzer, Jessie Dimick, Linda Mardel, and Debbie Smith have all helped me a lot in dealing with Cornell bureaucracy and administration.

Finally, I am indebted to my roots back home in Baiersbronn, Germany, which include Bent Beilharz, Roman Husz, Markus Kieninger, Mirco Rizutto, Dirk Schmid, the Daunoras family and all other friends, and my family (includng my new nephew Manuel).

Thanks to "KK" Krishnamurty Kambhampati, Gisela Leclere and Mostafa from Morocco, Filippo Tampieri, Jennifer Thambayah, Lien Tran, and Mona Matar for a good start, and to Vivek Bhatt and Angelina Kalianda for a good finish at Cornell.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Programmable force fields offer a fundamentally new approach to automated parts manipulation. Instead of handling a part directly (e.g. with a robot gripper), a force field surrounding the part causes it to move. Programmable force fields promise great flexibility, speed, and dexterity for a wide variety of tasks such as parts orienting, positioning, singulating, sorting, feeding, and assembly. Recently, several devices have been invented that can implement programmable force fields: in particular, actuator arrays fabricated with micro electro mechanical system (MEMS) technology, as well as macroscopic vibrating plates. These new automation designs permit distributed, parallel, non-prehensile, sensorless manipulation tasks that make them particularly attractive for handling batch microfabricated parts, whose small dimensions and large numbers would prohibit conventional pick-and-place operations.

A wealth of geometric and algorithmic problems arise in the control and programming of manipulation systems with many independent actuators. The theory of programmable force fields represents the first systematic, computational attack on massively-parallel distributed manipulation based on geometric and physical reasoning. The goal of this thesis is to develop a science base for manipulation using programmable force fields, and to demonstrate experiments with prototype devices that support this theory. We present combinatorially precise planning algorithms that synthesize strategies for controlling and coordinating a very large number of distributed actuators in a principled, task-level fashion.

When a part is placed on such a device, the programmed vector field induces a force and moment upon it. Over time, the part may come to rest in a dynamic equilibrium state. In principle, we have tremendous flexibility in choosing the vector field, since using e.g. MEMS array technologies, the force field may be programmed pixel-wise. Hence, we have a lot of control over the resulting equilibrium states. By chaining together sequences of vector fields, the equilibria may be cascaded to obtain a desired final state — for example, this state may represent a unique orientation or pose of the part. A system with such a behavior exhibits the *feeding property* [AHLM95]:

> A system has the *feeding property* over a set of parts $\mathcal{P}$ and a set of initial configurations $\mathcal{I}$ if, given any part $P \in \mathcal{P}$, there is some output configuration $\mathbf{q}$ such that the system can move P to $\mathbf{q}$ from any location in $\mathcal{I}$.

Our work on programmable vector fields is related to nonprehensile manipulation [DJR95, ZE96,EM96,Erd96]: in both cases, parts are manipulated without form or force closure.

This thesis describes our experimental devices, a technique for analyzing them called *equilibrium analysis*, lower bounds (i.e., impossibility results) on what the devices *cannot*

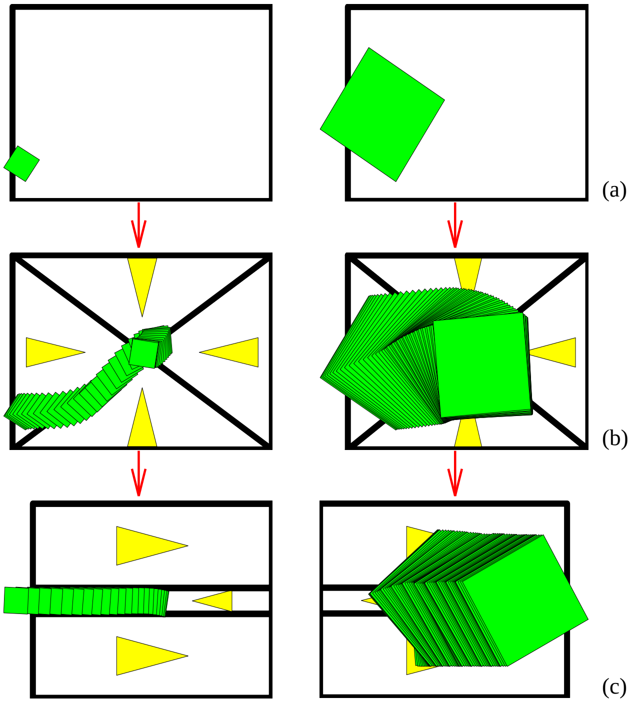

Figure 1.1: Sensorless sorting using force vector fields: parts of different sizes are first centered and subsequently separated depending on their size.

do, and results on a classification of control strategies yielding design criteria for useful manipulation strategies. Then we describe new manipulation algorithms using these tools. In particular, we improve earlier planning algorithms by a quadratic factor, show how to simultaneously orient and pose a part, and we relax dynamic and mechanical assumptions to obtain more robust and flexible strategies.

One corollary of our results is a method for coordinating the actions of a large distributed actuation system. The method is applicable to any controllable array capable of generating force vector fields. Such systems comprise arrays with up to tens of thousands of independently-servoable actuator cells, which we call *motion pixels*. We show how these systems can be programmed in a fine-grained, SIMD (Single Instruction Multiple Data) fashion to exert force fields on the manipulated object, thereby accomplishing massively-parallel distributed manipulation. Moreover, the theory of programmable force fields gives a method for controlling a very large number of distributed actuators in a principled, geometric, task-level fashion. Whereas many control theories for multiple independent actuators break down as the number of actuators becomes very large, our systems should only become more robust as the actuators become denser and more numerous.

We pose the question *Which force fields are suitable for manipulation strategies?* In particular, we ask whether the fields may be *classified*. That is: can we characterize all those force fields in which every part has stable equilibria? While this question has been well-studied for a point mass in a field, the issue is more subtle when lifted to a body with finite area, due to the moment covector. To answer, we first demonstrate impossibility results, in the form of "lower bounds:" there exist perfectly plausible fields which induce *no* stable equilibrium in very simple parts.

Fortunately, there is also good news. We present conditions for fields to induce well-behaved equilibria when lifted, by exploiting the theory of potential fields. While potential fields have been widely used in robot control [Kha86,KR88,RK92,RW95], micro-actuator arrays present us with the ability to *explicitly* program the applied force *at every point* in a vector field. Whereas previous work has developed control strategies with *artificial* potential fields, our fields are non-artificial (i.e., *physical*). Artificial potential fields require a tight feedback loop, in which, at each clock tick, the robot senses its state and looks up a control (i.e., a vector) using a state-indexed navigation function (i.e., a vector field). In contrast, physical potential fields employ no sensing, and the motion of the manipulated object evolves open-loop (for example, like a particle in a gravity field). This alone makes our application of potential field theory to micro-devices unique and novel. Moreover, such fields can be composed using addition, sequential composition, "parallel" composition by superposition of controls, or by a new kind of "morphing" of control signals which we will define.

Previous results on array manipulation strategies may be formalized using *equilibrium analysis*. In [BDMM94a] we proposed a family of control strategies called *squeeze patterns* and a planning algorithm for parts-orientation. This first result proved an $O(n^2)$ upper bound on the number E of orientation equilibria of a non-pathological (see Section 2.2) planar part with n vertices. This yields an $O(E^2) = O(n^4)$ planning algorithm to uniquely orient a part, under certain geometric, dynamic and mechanical assumptions. In this thesis, we argue that this bound on equilibria appears tight. This results in a high planning and execution complexity.

Using our equilibrium analysis, we introduce *radial* fields, which satisfy our stability property. Radial fields can then be combined with squeeze fields. We show this has several benefits:

1. The number of equilibria drops to $E = O(n)$.
2. The planning complexity drops to $O(E^2) = O(n^2)$.
3. Throughout the strategy execution, every part rotates about one fixed, unique point (after the first step).
4. This means that we can dispense with one critical assumption (called 2PHASE in [BDMM94a]): We no longer need to assume that the transitional and rotational motions induced by the array interact in a "quasi-static" and "sequential" manner.

We motivate our results by beginning with a description of the experimental devices we are interested in programming. In particular, we describe our progress in building the M-CHIP (Manipulation Chip), a massively parallel array of programmable micro-motion pixels. As proof of concept, we demonstrate a prototype M-CHIP containing up to 15,000 silicon actuators in one square inch. Our strategies are also applicable to macroscopic parts-feeders. We describe a planar, vibratory orienting and manipulation device which also uses our novel strategies.

Both of these devices foreground several key practical issues. First, the strategies employed by our improved algorithms and analysis require significant mechanical and control complexity — even though they require no sensing. While we believe such mechanisms are feasible to build using the silicon MEMS (Micro Electro Mechanical System) technologies we advocate, it is undeniable that no such device exists yet (the M-CHIPS will have pixel-wise programmability, but the first generation does not have sufficient directional resolution to implement highly accurate radial strategies). For this reason, we introduce and analyze strategies composed of field sequences that we know are implementable using current (microscopic or macroscopic) technology. Each strategy is a sequence of pairs of squeezes satisfying certain "orthogonality" properties. Under these assumptions, we can ensure

(a) equilibrium stability,
(b) relaxed mechanical and dynamical assumptions (the same as (4), above), and
(c) complexity and completeness guarantees.

The framework is quite general, and applies to any set of primitive operations satisfying certain "finite equilibrium" properties (which we define) — hence it has broad applicability to a wide range of devices. In particular, we view the restricted class of fields as a *vocabulary* and their rules of composition as a *grammar*, resulting in a "language" of manipulation strategies. Under our grammar, the resulting strategies are guaranteed to be well-behaved.

Finally, both our radial strategies and our finite manipulation grammar have the following advantage over previous manipulation algorithms for programmable vector fields: previous algorithms such as those described in [BDMM94a,BDM96b] guarantee to uniquely orient a part, but the translational position of the part is unknown at the strategy's termination. Both of our new algorithms guarantee to position the part uniquely (up to part symmetry) in translation *as well as* orientation space. Like the algorithms in [BDMM94a, BDM96b], the new algorithms require no sensing, and work from *any* initial configuration to uniquely pose the part. In particular, the initial configuration is never known to the (sensorless) execution system, which functions open-loop.

The complexity and completeness guarantees we obtain for manipulation grammars are considerably weaker than for the ideal radial strategies. For radial strategies, we show that *any* non-pathological planar part with finite area contact can be placed in a unique pose in $O(E) = O(n)$ steps. Under the simplified manipulation grammar, our planner is guaranteed to find a strategy if one exists (if one does not exist, the planner will signal this). However, it is not known whether there exists a strategy for every part. This lack of completeness of manipulation grammar strategies stands in contrast to the *complete* general squeeze and radial algorithms for which a *guaranteed* strategy exists for *all* parts. Moreover, the planning algorithm is worst-case exponential instead of merely quadratic.

Table 1.1 summarizes the various force fields discussed in this thesis, and lists their corresponding manipulation tasks as well as planning and execution complexities. These results illustrate a tradeoff between mechanical complexity (the density and force resolution of field elements) and planning complexity (the computational difficulty of synthesizing a strategy). If one is willing to build a device capable of radial fields, then one reaps great benefits in planning and execution speed. On the other hand, we can still plan for simpler devices, but the plan synthesis is more expensive (worst-case exponential in the number of equilibria), and we lose some completeness properties.

Finally, the desire to implement complicated fields raises the question of control uncertainty. We close by describing how families of potential functions can be used to represent control uncertainty, and analyzed for their impact on equilibria, and we will give an outlook on still open problems and future work.

1.1 Parts Feeders

It is often extremely costly to maintain part order throughout the manufacture cycle. For example, instead of keeping parts in pallets, they are often delivered in bags or boxes, whence they must be picked out and sorted. A parts feeder is a machine that orients such parts before they are fed to an assembly station. Currently, the design of parts

Table 1.1: Summary of programmable force fields, and their corresponding manipulation tasks for polygonal parts (with n vertices and k combinatorially distinct bisectors).

Task	**Field(s)**	**Complexity**		
		Fields	Planning	Plans
Translate	Constant	Constant magnitude and direction	-	1
Center	Radial	Constant magnitude, continuous directions	-	1
	Orthogonal squeezes	Piecewise constant magnitudes and directions	$O(1)$	$O(1)$
Orient uniquely (up to symmetry)	Sequence of squeezes	Piecewise constant magnitudes and directions	$O(k^2n^4)$	$O(kn^2)$
	Inertial	Smooth magnitude, piecewise constant direction	$O(1)$	$O(1)$
Pose uniquely (up to symmetry)	Manipulation grammar	m simple, arbitrary fields with at most E stable equilibria	$O(m^2 2^E)$	$O(m 2^E)$ (not complete)
	Sequence of radial + squeeze	Piecewise continuous magnitude and direction	$O(k^2n^2)$	$O(kn)$
	Elliptic	Smooth magnitude and direction	$O(1)$	$O(1)$
	Universal (conjecture)	Continuous magnitude and direction	-	1

feeders is a black art that is responsible for up to 30% of the cost and 50% of workcell failures [NW78,BPM82,FD86,Sch87,SS87]. *"The real problem is not part transfer but part orientation."*, Frank Riley, Bodine Corporation [Ril83, p.316, his italics]. Thus although part feeding accounts for a large portion of assembly cost, there is not much scientific basis for automating the process.

The most common type of parts feeder is the *vibratory bowl feeder*, where parts in a bowl are vibrated using a rotary motion, so that they climb a helical track. As they climb, a sequence of baffles and cutouts in the track create a mechanical "filter" that causes parts in all but one orientation to fall back into the bowl for another attempt at running the gauntlet [BPM82,Ril83,San91]. To improve feed rate, it is sometimes possible to design the track so as to mechanically rotate parts into a desired orientation (this is called *conversion*). Related methods use centrifugal forces [FD86], reciprocating forks, or belts to move parts through the filter [RL86].

Sony's APOS parts feeder [Hit88] uses an array of nests (silhouette traps) cut into a vibrating plate. The nests and the vibratory motion are designed so that the part will remain in the nest only in one particular orientation. By tilting the plate and letting parts flow across it, the nests eventually fill up with parts in the desired orientation. Although the vibratory motion is under software control, specialized mechanical nests must be designed for each part [MJU91].

The reason for the success of vibratory bowl feeders and the Sony APOS system is the underlying principle of *sensorless manipulation* [EM88] that allows parts positioning and orienting without sensor feedback. This principle is even more important at small scales, because sensor data will be less accurate and more difficult to obtain. The APOS system or bowl feeders are unlikely to work in the micro domain: instead novel device designs for micro manipulation tasks are required. The theory of sensorless manipulation is the science base for developing and controlling such devices.

Reducing the amount of required sensing is an example of *minimalism* [CG94,BBD$^+$95], which pursues the following agenda: For a given robot task, find the minimal configuration of resources required to solve the task. Minimalism is interesting because doing task A without resource B proves that B is somehow inessential to the information structure of the task. In robotics, minimalism has become increasingly influential. Raibert [RHPR93] showed that walking and running machines could be built without static stability. Erdmann and Mason [EM88] showed how to do dexterous manipulation without sensing. McGeer [McG90] built a biped, kneed walker without sensors, computers, or actuators. Canny and Goldberg [CG94] argue that minimalism has a long tradition in industrial manufacturing, and developed geometric algorithms for orienting parts using simple grippers and accurate, low cost light beams. Brooks [Bro86] has developed online algorithms that rely less extensively on planning and world models. Donald et al. [DJR95,BBD$^+$95] have built distributed teams of mobile robots that cooperate in manipulation without explicit communication. We intend to use these results for our experiments in micro manipulation, and to examine how they relate to our theoretical proofs of minimalist systems.

1.1.1 Microfabricated Actuator Arrays

A wide variety of micromechanical structures (devices with features in the μm range) has been built recently by using processing techniques known from VLSI industry (see for example [Gab95,MAA$^+$95,Mac96b,Mac96a]). Various microsensors and microactuators

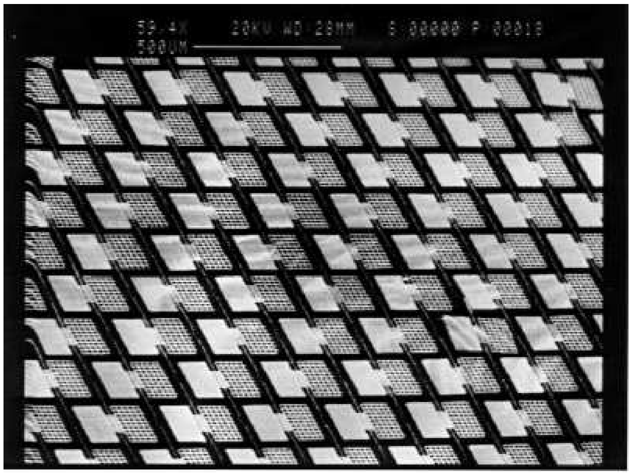

Figure 1.2: A large unidirectional actuator array (scanning electron microscopy). Each actuator is $180 \times 240\,\mu m^2$ in size. Detail from a $1\,in^2$ array with more than 15,000 actuators.

have been shown to perform successfully. E.g. a single-chip air-bag sensor is commercially available [Ana91]; video projections using an integrated, monolithic mirror array have been demonstrated recently [Sam93]. A fully integrated scanning tunneling microscope (STM) has been developed in our group [XMM95,MAA$^+$95]. However, the fabrication, control, and programming of micro-devices that can interact and actively change their environment remains challenging. Problems arise from

1. unknown material properties and the lack of adequate models for mechanisms at very small scales,
2. the limited range of motion and force that can be generated with microactuators,
3. the lack of sufficient sensor information with regard to manipulation tasks, and
4. design limitations and geometric tolerances due to the fabrication process.

MEMS manipulator arrays have been proposed by several researchers, among others Pister et al. [PFH90], Fujita et al. [Fuj93], Storment et al. [SBW$^+$94], Will et al. [LW95], Jacobson et al. [JGJB$^+$95], or Suh et al. [SGD$^+$96]. For an overview see Table 1.2, or [LW95,BDMM94b,BDMM94a]. Our arrays (Figure 1.2) are fabricated using a SCREAM (<u>S</u>ingle-<u>C</u>rystal Silicon <u>R</u>eactive <u>E</u>tching and <u>M</u>etallization) process developed at the Cornell Nanofabrication Facility [ZM92,SZM93]. The SCREAM process is low-temperature, and does not interfere with traditional VLSI [SM96]. Hence it opens the door to building monolithic micro electro mechanical systems with integrated microactuators and control circuitry on the same wafer.

Our design is based on microfabricated torsional resonators [MZSM93,MM96]. Each unit device consists of a rectangular grid etched out of single-crystal silicon suspended by two rods that act as torsional springs (Figure 6.1). The grid is about $200\,\mu m$ long and extends $120\,\mu m$ on each side of the rod. The rods are $150\,\mu m$ long. The current asymmetric design has $5\,\mu m$ high protruding tips on one side of the grid that make contact with an object lying on top of the actuator (Figure 1.3). The other side of the actuator consists of a denser grid above an aluminum electrode. If a voltage is applied between silicon substrate and electrode, the dense grid above the electrode is pulled downward by the resulting electrostatic force. Simultaneously the other side of the device (with the tips) is deflected out of the plane by several μm. Hence an object can be lifted and pushed sideways by the

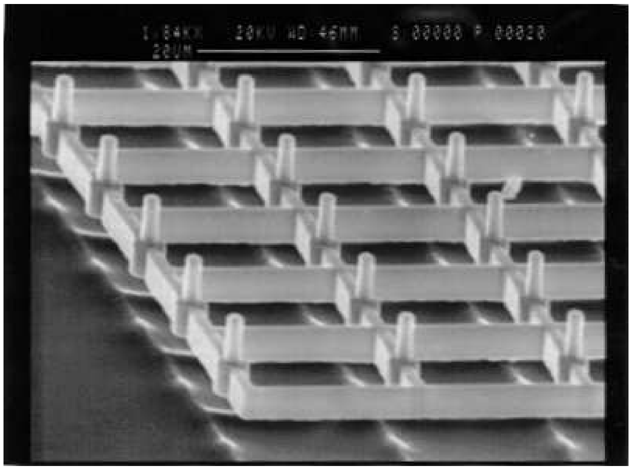

Figure 1.3: M-Chip microactuator consisting of a single-crystal silicon grid with $5\,\mu m$ high tips. There is a $5 - 10\,\mu m$ clearance between the released actuator and the silicon substrate.

Table 1.2: Devices that can implement programmable force fields.

Devices in this thesis	**Principle of Actuation**
SCREAM M-Chip (Section 6.1)	Electrostatic torsional micro oscillators
Polyimide cilia array (Section 6.2)	Thermobimoph and electrostatic micro actuators with CMOS circuits (built at Stanford [SGD$^+$96])
Transversely vibrating plate (Section 6.3)	Attraction to vibratory nodal lines
Other devices (selection)	
Planar MEMS actuator system [PFH90]	Air levitation with micro nozzles, electrostatic actuation
Biomimetic micro motion system [AOF93]	Thermobinorph micro actuators
Distributed micro motion system [KF93b]	Air levitation and actuation with micro nozzles
Intelligent motion surface [LTW$^+$95]	Magnetic actuators on CMOS
Virtual vehicle [LM97,LMC97]	Array of servoed roller wheels
Longitudinally vibrating plate [RCG97]	Nonlinearity between velocity and friction

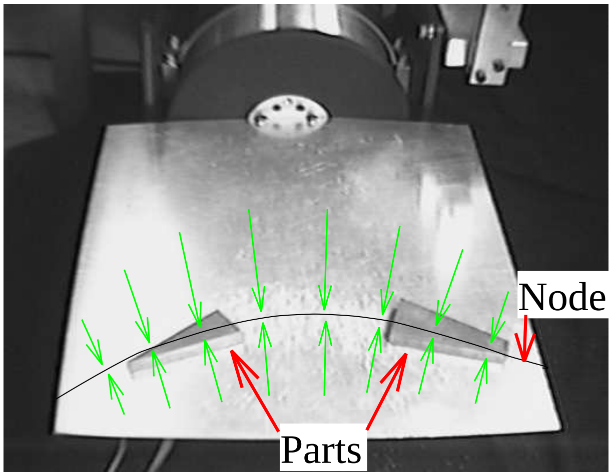

Figure 1.4: Vibratory plate parts feeder: an aluminum plate (size $50\,cm \times 40\,cm$) exhibits a vibratory minimum. Parts are attracted to this *nodal line* and reach stable equilibrium there.

actuator.

Because of its low inertia (resonance in the high kHz range) the device can be driven in a wide frequency range from DC to several $100\,kHz$ AC. Our actuators need not be operated at resonance: They can also be servoed to periodically "hit" an object on top, hence applying both lateral and vertical forces. Our calculations, simulations and experiments have shown that the force generated with a torsional actuator is approximately $10\,\mu N$, which corresponds to a force-per-area ratio of $100\,\mu N/mm^2$, large enough to levitate e.g. a piece of paper ($1\,\mu N/mm^2$) or a silicon wafer ($10\,\mu N/mm^2$).

Each actuator can generate motion in one specific direction if it is activated; otherwise it acts as a passive frictional contact. Figure 1.2 shows a small section of such a unidirectional actuator array, which consists of more than 15,000 individual actuators. The combination and selective activation of several actuators with different motion bias allows us to generate various motions in discrete directions, spanning the plane (Figure 6.3).

The fabrication process and mechanism analysis are described in more detail in Section 6.1, and in [BDMM94b,BDMM94a,BDM96a]. A different micro array design based on polyimide cilia with thermobimoph actuation is discussed in Section 6.2 and in [SGD$^+$96].

1.1.2 Macroscopic Vibratory Parts Feeders

Böhringer et al. [BBG95] have presented a device that uses the force field created by transverse vibrations of a plate to position and align parts. The device consists of an aluminum plate that is attached to a commercially available electrodynamic vibration generator,[1] with a linear travel of $0.02\,m$, and capable of producing a force of up to $500\,N$ (Figure 1.4). The input signal, specifying the waveform corresponding to the desired oscillations, is fed to a single coil armature, which moves in a constant field produced by a ceramic permanent magnet in a center gap configuration.

For low amplitudes and frequencies, the plate moves longitudinally with no perceptible transverse vibrations. However, as the frequency of oscillations is increased, transverse

---

[1] Model VT-100G, Vibration Test Systems, Akron, OH, USA.

vibrations of the plate become more pronounced. The resulting motion is similar to the forced transverse vibration of a rectangular plate, clamped on one edge and free along the other three sides. This vibratory motion creates a force field in which particles are attracted to locations with minimal vibration, called the *nodal lines*. This field can be programmed by changing the frequency, or by employing clamps as programmable fixtures that create various vibratory nodes.

Figure 1.4 shows two parts, shaped like a triangle and a trapezoid, after they have reached their stable poses. To better illustrate the orienting effect, the curve showing the nodal line has been drawn by hand. *Nota bene*: This device can only use the finite manipulation grammar described in Section 5.2 since it can only generate a constrained set of vibratory patterns, and cannot implement general squeeze and radial strategies.

Section 6.3 gives more details on our manipulation experiments with transversely vibrating plates.

Chapter 2

Equilibrium Analysis For Programmable Vector Fields

For the generation of manipulation strategies with programmable vector fields it is essential to be able to predict the motion of a part in the field. Particularly important is determining the stable equilibrium poses a part can reach in which all forces and moments are balanced. This *equilibrium analysis* was introduced in our short conference paper [BDMM94a], where we presented a theory of manipulation for programmable vector fields, and an algorithm that generates manipulation strategies to orient polygonal parts without sensor feedback using a sequence of *squeeze fields*. We now review the algorithm in [BDMM94a] and give a detailed proof of its complexity bounds. The tools developed here are essential to understanding our new and improved results, and will be used throughout this thesis to develop complexity bounds for our distributed manipulation algorithms.

We will in general assume that the dynamics of a part moving in the force field is governed by first-order dynamics. This assumption is based on extensive experimentation with the devices presented in Section 6. In a first-order system, the velocity of a part is directly proportional to the force acting on it. Basically, it is a rigid body dynamical system that is heavily damped.

2.1 Squeeze Fields and Equilibria

In [BDMM94a] we proposed a family of control strategies called *squeeze fields* and a planning algorithm for parts-orientation.

Definition 1 [BDM96b] *Assume l is a straight line through the origin. A* squeeze field *f is a two-dimensional force vector field defined as follows:*

1. *If $z \in \mathbb{R}^2$ lies on l then $f(z) = 0$.*
2. *If z does not lie on l then $f(z)$ is the unit vector normal to l and pointing towards l.*

We refer to the line l as the *squeeze line*, because l lies in the center of the squeeze field. See Figure 2.1 for examples of squeeze fields.

Assuming quasi-static motion, an object will move perpendicularly towards the line l and come to rest there. We are interested in the motion of an arbitrarily shaped (not necessarily small) part P. Let us call P_1, P_2 the regions of P that lie to the left and to the

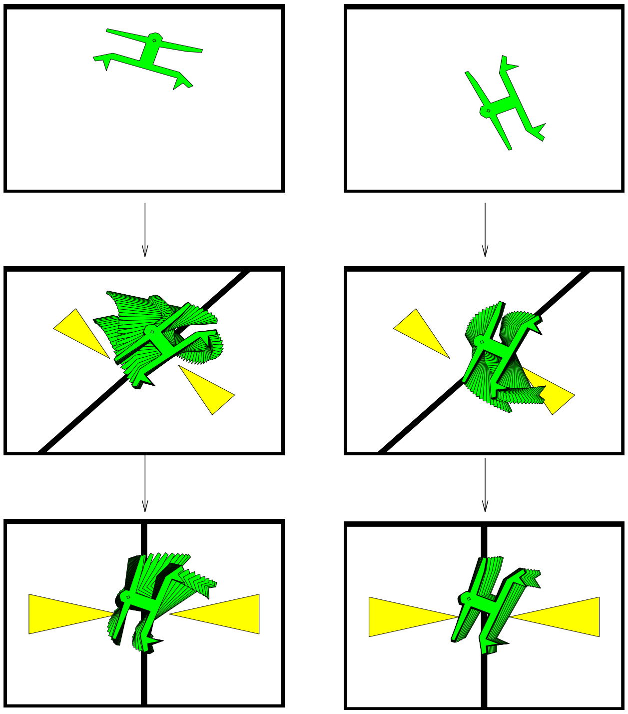

Figure 2.1: Sensorless parts orienting using force vector fields: The part reaches unique orientation after two subsequent squeezes. There exist such orientating strategies for all polygonal parts. See URL `http://www.cs.cornell.edu/home/karl/MicroManipulation` for an animated simulation.

right of l, respectively, and c_1, c_2 their centers of area. In a rest position both translational and rotational forces must be in equilibrium. We obtain the following two conditions:

I: The areas P_1 and P_2 must be equal.

II: The vector $c_2 - c_1$ must be normal to l.

P has a translational motion component normal to l if **I** does not hold. P has a rotational motion component if **II** does not hold (see Figures 2.2, 2.3, and 2.4)This assumes a uniform force distribution over the surface of P, which is a reasonable assumption for a flat part that is in contact with a large number of elastic actuators.

Definition 2 *A part P is in* translational equilibrium *if the forces acting on P are balanced. P is in* orientational equilibrium *if the moments acting on P are balanced.* Total equilibrium *is simultaneous translational and orientational equilibrium.*

Let (x_0, y_0, θ_0) be an equilibrium pose of P. (x_0, y_0) is the corresponding translation equilibrium*, and θ_0 is the corresponding* orientation equilibrium.

Note that conditions **I** and **II** do *not* imply that in equilibrium, the center of area of P has to coincide with the squeeze line l. For example, consider a large and a small square connected by a long rod of negligible width (Figure 2.5). If the rod is long enough, the center of area will lie outside of the large square. However, in equilibrium the squeeze line l will always intersect the large square.

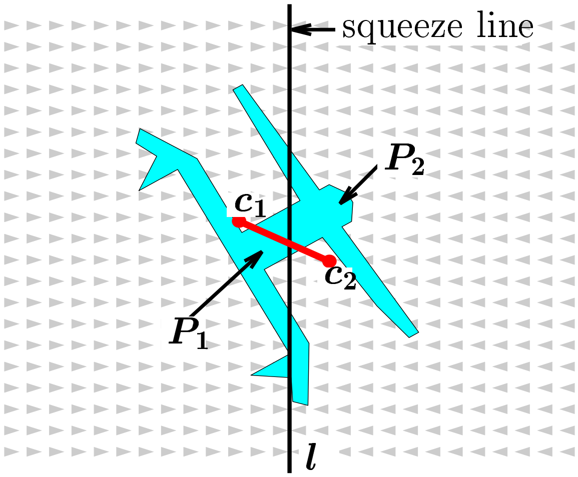

Figure 2.2: Equilibrium condition: To balance force and moment acting on P in a unit squeeze field, the two areas P_1 and P_2 must be equal (i.e., l must be a bisector), and the line connecting the centers of area c_1 and c_2 must be perpendicular to the node line.

2.2 Polygon Bisectors and Complexity

Consider a polygonal part P in a unit squeeze field as described in Section 2.1. In this section we describe how to determine the orientations θ_i in which P achieves equilibrium. This construction will show that equilibria *always exist* as long as the contact areas have finite size, and that for connected parts the orientation equilibria are *discrete.* More precisely, if a connected part is in equilibrium in a squeeze field, there are discrete values for its orientation, and its offset from the center of the squeeze line. The equilibrium is of course independent of its position along the squeeze line. Hence, in the remainder of Section 2, when using the term *discrete equilibria*, we mean that the orientation and offset of the part is discrete. We will derive upper bounds on the number of these discrete equilibria.

Definition 3 *A* bisector *of a polygon P is a line that cuts P into two regions of equal area.*

Proposition 4 [**BDM96b**] *Let P be a polygon whose interior is connected. There exist $O(k\,n^2)$ bisectors such that P is in equilibrium when placed in a squeeze field such that the bisector coincides with the squeeze line. n is the part complexity measured as the number of polygon vertices. k denotes the maximum number of polygon edges that a bisector can cross.*

If P is convex, then the number of bisectors is bounded by $O(n)$.

For most part geometries, k is a small constant.[1] However in the worst-case, pathological parts can reach $k = O(n)$. A (e.g. rectilinear) spiral-shaped part would be an example for such a pathological case, because every bisector intersects $O(n)$ polygon edges.

Lemma 5 *Given a polygon P and a line $l : y = mx + c$. Let n be the number of vertices of P.*

---

[1] In particular, in [BDMM94a] we assumed that $k = O(1)$.

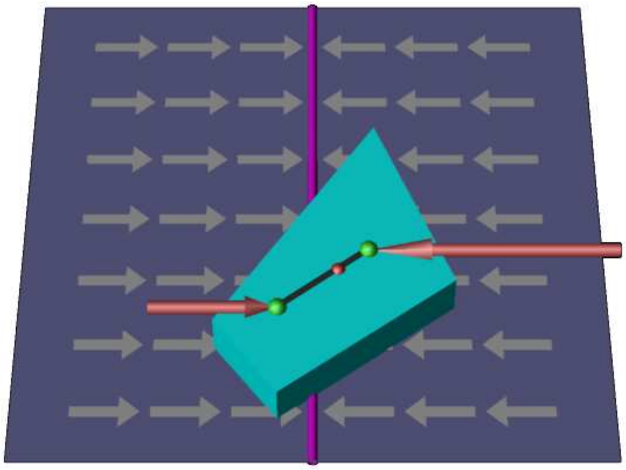

Figure 2.3: Part in a unit squeeze field. The resulting forces for the left and the right region of the part are shown acting at the respective centers of mass: the part experiences a translational force and a moment.

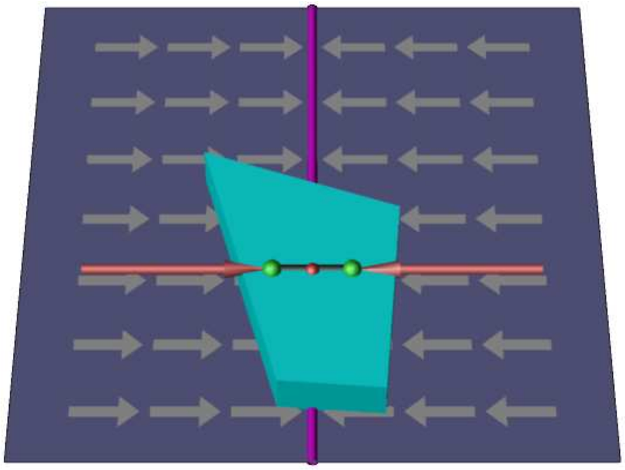

Figure 2.4: Part in equilibrium: The resulting forces for the left and the right region of the part are of equal magnitude and opposite direction, and the resulting moment is zero.

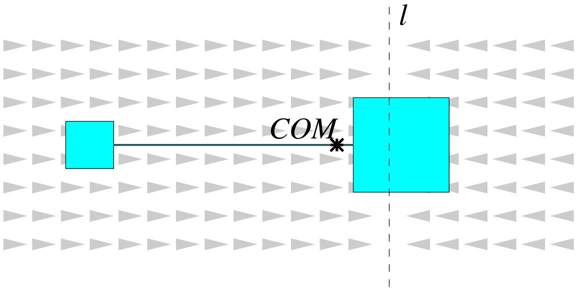

Figure 2.5: Part consisting of two squares connected by a long, thin rod. The part is in total equilibrium, but its COM does not coincide with the squeeze line l.

1. *There exist $O(n^2)$ combinatorially different ways how a line l can intersect P.*
2. *Let a and b be the intersections of bisector l with the convex hull of P. As m varies from $-\infty$ to $+\infty$, a and b progress monotonically counterclockwise about the convex hull of P.*
3. *If the interior of P is connected, then there exists a unique bisector of P for every $m \in \mathbb{R}$.*

Combinatorially equivalent intersections of polygon P are all those placements of the intersecting line l such that the sets of left and right polygon vertices are fixed. A necessary condition for combinatorial equivalence is that l intersects the same ordered set of polygon edges.

Proof:

1. There are $O(n^2)$ different placements for l such that it coincides with more than one vertex of P. Hence all placements of l fall into one of $O(n^2)$ combinatorially equivalent classes.
2. See [DO90, Lemma 3.1].
3. Assume l is a bisector of P with a fixed slope m. Since P is connected, the intersection between l and P must be a line segment of non-zero length. Hence a translation of l e.g. towards the left will cause a strictly monotonous decrease of the left area segment of P, and vice versa. Therefore the bisector placement of l for a given slope m is unique.

□

Consider the bisector l of polygon P for changing m values, as described in Lemma 5. The intersections of l with the convex hull of P, a and b, progress monotonically about the convex hull. In general, this progression corresponds to a rotation and a translation of l.

In the following proof for Proposition 4, we investigate the relationship between the location of the bisector, and the corresponding left and right areas of P and its respective centers of mass. This will allow us to show that for *combinatorially equivalent* bisector placements there are only a finite number of possible equilibria, and that this number is bounded by $O(k)$.

Proof: [**Proposition 4**] Consider two combinatorially equivalent placements of bisector l on polygon P. We will show that the number of equilibria for this bisector placement is bounded by $O(k)$. Since there are $O(n^2)$ such placements for P (see Lemma 5), the total number of equilibria will be $O(k\,n^2)$.

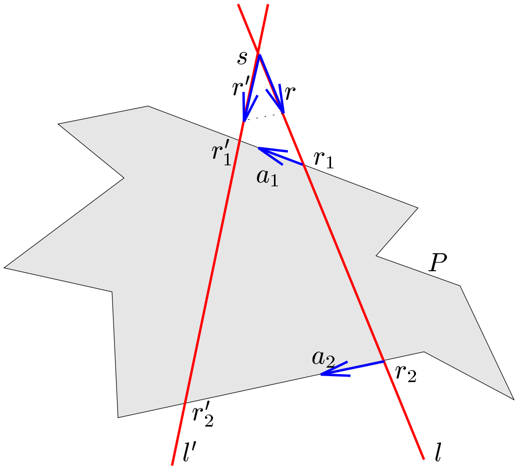

Figure 2.6: Two non-parallel lines l and l' in combinatorially equivalent intersection with polygon P.

Rotating the Bisector. Consider the line l and a point s that lies on l (Figure 2.6). The direction of l is given by a vector r. Assume for now that the line l intersects two edges of the polygon P in the points r_1 and r_2. Also assume that these edges have directions a_1 and a_2. Now consider another line l' with direction r' that intersects l in s. Assume that l and l' have combinatorially equivalent intersections with polygon P, and that l' intersects the polygon edges in r'_1 and r'_2. Let us write $r_i = s + \rho_i r$ and $r'_i = s + \rho'_i r'$. Then the polygon area between l and l' is:

$$\begin{aligned} A &= \frac{1}{2}(\rho'_2 r' \times \rho_2 r - \rho'_1 r' \times \rho_1 r) \\ &= \frac{1}{2}(\rho'_2 \rho_2 - \rho'_1 \rho_1)(r' \times r) \end{aligned}$$

In the general case where l and l' intersect multiple edges of some arbitrary polygon P at points $r_1, r_2, \ldots, r_k$ and $r'_1, r'_2, \ldots, r'_k$ (k even), the polygon area between l and l' is:

$$\begin{aligned} A &= \frac{1}{2}\sum_{i=1}^{k}(-1)^i A_i \\ &= \frac{1}{2}(r' \times r)\sum_{i=1}^{k}(-1)^i \rho'_i \rho_i \end{aligned}$$

W.l.o.g. let $\rho_k \neq 0$. Then r' can be written as $r' = r + \alpha a_k$ for some $\alpha \in \mathbb{R}$, and the above equation becomes:

$$\begin{aligned} &= \frac{1}{2}((r + \alpha a_k) \times r)\sum_{i=1}^{k}(-1)^i \rho'_i \rho_i \\ &= \frac{\alpha}{2}(a_k \times r)\sum_{i=1}^{k}(-1)^i \rho'_i \rho_i \qquad (2.1) \end{aligned}$$

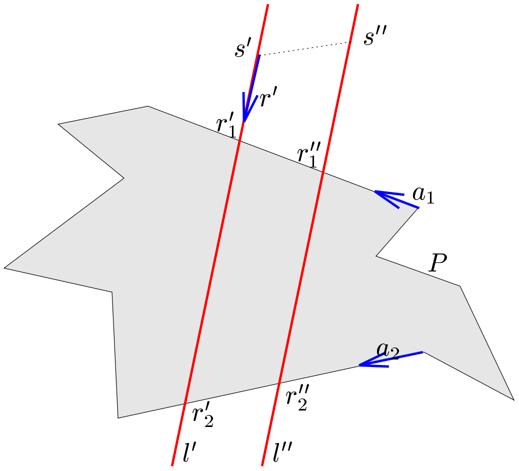

Figure 2.7: Two parallel lines l' and l'' in combinatorially equivalent intersection with polygon P.

From the two vector equations $r'_i = s + \rho'_i r'$ and $r'_i = s + \rho_i r + \lambda a_i$, $\lambda \in \mathbb{R}$, we can determine ρ'_i as:

$$\rho'_i = \frac{\rho_i(a_i \times r)}{(a_i \times r) + \alpha(a_i \times a_k)} \tag{2.2}$$

If we also choose the edge direction vectors a_i such that $(a_i \times r) = 1$, then Equations (2.2) and (2.1) simplify to the following rational functions in α:

$$\rho'_i = \frac{\rho_i}{1 + \alpha(a_i \times a_k)} \tag{2.3}$$

$$A = \frac{\alpha}{2}\sum_{i=1}^{k}(-1)^i \frac{\rho_i^2}{1 + \alpha(a_i \times a_k)} \tag{2.4}$$

Let us look at the denominator $d_i(\alpha) = 1 + \alpha(a_i \times a_k)$ in more detail. This is important because we shall see that in all formulas we will obtain, the denominators consist only of $d_i(\alpha)$. For an arbitrary polygon, d_i is a linear function of α. If all a_i are parallel then $d_i = 1$. If the polygon is rectilinear, i.e. all a_i are either parallel or perpendicular, then $d_i(\alpha) = 1$ if $a_i \parallel a_k$, and $d_i(\alpha) = 1 + \alpha a_\perp$ if $a_i \perp a_k$, where $a_\perp$ is constant. So in this case there are only two different denominators, one of which is 1.

Translating the Bisector. We now consider the case where l' shifts parallel (Figure 2.7). Analogously to the previous paragraph, let $r'_i = s' + \rho'_i r'$, and $r''_i = s'' + \rho''_i r'$. Also let the vector between s' and s'' be $s'' - s' = \beta a_2$. Then the polygon area between l' and l'' is:

$$\begin{aligned} B &= \beta a_2 \times \frac{1}{2}((r'_2 + r''_2) - (r'_1 + r''_1)) \\ &= \frac{\beta}{2}(\rho'_2 + \rho''_2 - \rho'_1 - \rho''_1)(a_2 \times (r + \alpha a_2)) \\ &= \frac{\beta}{2}(\rho'_2 + \rho''_2 - \rho'_1 - \rho''_1) \end{aligned} \tag{2.5}$$

In the general case l' and l'' intersect multiple edges of some arbitrary polygon P at points $r'_1, r'_2, \ldots, r'_k$ and $r''_1, r''_2, \ldots, r''_k$. Now the ρ''_i can be determined from the two vector equations $r''_i = r'_i + \lambda a_i$, $\lambda \in \mathbb{R}$, and $r''_i = s'' + \rho''_i r'$:

$$\begin{aligned}\rho''_i &= \rho'_i - \beta\frac{a_i \times a_k}{a_i \times r'} \\ &= \rho'_i - \beta\frac{a_i \times a_k}{1+\alpha(a_i \times a_k)} \\ &= \frac{\rho_i - \beta(a_i \times a_k)}{1+\alpha(a_i \times a_k)} \qquad (2.6)\end{aligned}$$

Then the polygon area between l' and l'' is:

$$\begin{aligned}B &= \frac{\beta}{2}\sum_{i=1}^{k}(-1)^i(\rho'_i + \rho''_i) \\ &= \frac{\beta}{2}\sum_{i=1}^{k}(-1)^i\frac{2\rho_i - \beta(a_i \times a_k)}{1+\alpha(a_i \times a_k)} \qquad (2.7)\end{aligned}$$

This is a quadratic polynomial in β (unless all a_i are parallel, in which case it simplifies to the linear equation $B = \beta\sum_{i=1}^{k}(-1)^i\rho_i$).

Maintaining the Bisector Property. From the above two paragraphs we see that if the bisector l is rotated to l', then the left and right areas are changed by a value A ($\neq 0$ in general) as described in Equation (2.4). Hence a subsequent shift of l' is necessary to restore the bisector property, by changing the areas by a value B, as described in Equation (2.7).

This implies the condition $A + B = 0$, with A and B given by Equations (2.4) and (2.7):

$$\begin{aligned}A + B &= \frac{1}{2}\sum_{i=1}^{k}(-1)^i\frac{\alpha\rho_i^2 + 2\beta\rho_i - \beta^2(a_i \times a_k)}{1+\alpha(a_i \times a_k)} \\ &= 0 \qquad (2.8)\end{aligned}$$

This equation ensures that l is a bisector of P. It is a necessary and sufficient condition for translational equilibrium in a unit squeeze field. Equation (2.8) is a rational function in α, and a quadratic polynomial in β. Hence for all combinatorially equivalent bisectors, we can obtain an explicit formula to describe β as a function of α. In case that all a_i are parallel, Equation (2.8) simplifies to a linear equation: $\sum_{i=1}^{k}(-1)^i\left(\alpha\frac{\rho_i}{2} + \beta\right)\rho_i = 0$.

In general, Equation (2.8) is equivalent to a polynomial in α and β whose degree depends on the number k of polygon edges intersected by the bisectors l, l', or l''. The degree of this polynomial is limited by k for α, and by 2 for β. In the rectilinear case the degrees for α and β are limited by 2.

Zero Moment. After rotating (parameter α, obtain l') and translating (parameter β, obtain l'') the bisector l, its intersections with the polygon edges move from r_i to

$$\begin{aligned}r''_i &= s + \rho''_i r' + \beta a_k \\ &= s + \frac{\rho_i - \beta(a_i \times a_k)}{1+\alpha(a_i \times a_k)}(r + \alpha a_k) + \beta a_k \qquad (2.9)\end{aligned}$$

If all a_i are parallel, this simplifies to $r_i'' = s + \rho_i r + (\alpha\rho_i + \beta)\, a_k$.

Suppose that c_l and c_r are the left and the right centers of area of P, and A_l and A_r are the respective area sections, so $A_l + A_r = A$. We are interested in how these points change when the bisector changes. Note that always $c = \frac{1}{A}(A_l c_l + A_r c_r)$, and if P is bisected (i.e. $A_l = A_r = \frac{1}{2}A$) then $c = \frac{1}{2}(c_l + c_r)$.

We consider the area between l and l'', which can be written as a sum of quadrangles (r_i, r_k, r_k'', r_i''). The weighted center of mass of this area can be determined as:

$$\begin{aligned} C = \sum_{i=1}^{k} (-1)^i \frac{1}{6}((r_i + r_k)(r_i \times r_k) + (r_k + r_k'')(r_k \times r_k'') \\ +(r_k'' + r_i'')(r_k'' \times r_i'') + (r_i'' + r_i)(r_i'' \times r_i)) \end{aligned} \tag{2.10}$$

For the left areas the following relationship holds:

$$\begin{aligned} A_l'' c_l'' &= A_l c_l + C \\ c_l'' &= \frac{A_l}{A_l''} c_l + \frac{1}{A_l''} C \end{aligned}$$

and similarly, for the right areas:

$$c_r'' = \frac{A_r}{A_r''} c_r - \frac{2}{A_l''} C$$

Hence

$$c_l'' - c_r'' = \frac{A_l}{A_l''} c_l - \frac{A_r}{A_r''} c_r + \left(\frac{1}{A_l''} + \frac{1}{A_r''} \right) C$$

Both l and l'' are bisectors, so $A_l = A_r = A_l'' = A_r'' = \frac{A}{2}$, and

$$c_l'' - c_r'' = c_l - c_r + \frac{4}{A} C$$

For orientational equilibrium we require that the line connecting the centers of mass, $c_r'' - c_l''$, and the direction of the bisector r' are perpendicular:

$$\begin{aligned} (c_l'' - c_r'') \cdot r' &= (c_l - c_r + \frac{4}{A} C) \cdot r' \\ &= 0 \end{aligned} \tag{2.11}$$

The value of C can be determined by using Equations (2.10) and (2.9), and $r' = r + \alpha a_k$. Equation (2.11) is a necessary and sufficient condition for orientational equilibrium.

By using the expressions derived in Equations 2.1 to 2.10, both Equations (2.8, for translational equilibrium) and (2.11, for orientational equilibrium) can be expressed as rational functions in α and β whose numerator / denominator degrees are $O(k)$ / $O(1)$ for α and $O(1)$ for β. Hence we can obtain two polynomial functions of degree $O(k)$ for α and $O(1)$ for β. There are at most $O(k)$ solutions to these two equations. □

2.3 Planning of Manipulation Strategies

In this section we present an algorithm for sensorless parts alignment with squeeze fields [BDMM94a,BDM96b]. Recall from Section 2.2 that in squeeze fields, the equilibria for connected polygons are discrete (modulo a neutrally stable translation parallel to the squeeze line which we will disregard for the remainder of Section 2).

To model our actuator arrays and vibratory devices, we made the following assumptions:

Density: The generated forces can be described by a vector field, i.e., the individual microactuators are dense compared to the size of the moving part.

2Phase: The motion of a part has two phases: (1) Pure translation towards l until the part is in translational equilibrium. (2) Motion in translational equilibrium until orientational equilibrium is reached.

Note that due to the elasticity and oscillation of the actuator surfaces, we can assume continuous area contact, and not just contact in three or a few points. If a part moves while in translational equilibrium, in general the motion is not a pure rotation, but also has a translational component.

Definition 6 **[BDM96b]** *Let θ be the orientation of a connected polygon P in a squeeze field, and let us assume that condition* **I** *holds. The* turn function $t : \theta \rightarrow \{-1, 0, 1\}$ *describes the instantaneous rotational motion of P:*

$$t(\theta) = \begin{cases} 1 & \textit{if } P \textit{ will turn counterclockwise} \\ -1 & \textit{if } P \textit{ will turn clockwise} \\ 0 & \textit{if } P \textit{ is in total equilibrium (Fig. 2.8).} \end{cases}$$

See Figure 2.8 for an illustration. The turn function $t(\theta)$ can be obtained e.g. by taking the sign of the lifted moment $M_P(\mathbf{z})$ for poses $\mathbf{z} = (x, y, \theta)$ in which the lifted force $f_P(\mathbf{z})$ is zero.

Definition 6 immediately implies the following lemma:

Lemma 7 **[BDM96b]** *Let P be a polygon with orientation θ in a squeeze field such that condition* **I** *holds. P is stable if $t(\theta) = 0$, $t(\theta+) \leq 0$, and $t(\theta-) \geq 0$. Otherwise P is unstable.*

Proof: Assume the part P is in a pose (x, y, θ) such that condition **I** is satisfied. This implies that the translational forces acting on P balance out. If in addition $t(\theta) = 0$, then the effective moment is zero, and P is in total equilibrium. Now consider a small perturbation $\delta_\theta > 0$ of the orientation θ of P while condition **I** is still satisfied. For a stable equilibrium, the moment resulting from the perturbation δ_θ must not aggravate but rather counteract the perturbation. This is true if and only if $t(\theta + \delta_\theta) \leq 0$ and $t(\theta - \delta_\theta) \geq 0$. □

Using this lemma we can identify all stable orientations, which allows us to construct the squeeze function [Gol93] of P (see Figure 2.8c), i.e. the mapping from an initial orientation of P to the stable equilibrium orientation that it will reach in the squeeze field:

Lemma 8 **[BDM96b]** *Let P be a polygonal part on an actuator array $\mathcal{A}$ such that assumptions* Density *and* 2Phase *hold. Given the turn function t of P, its corresponding* squeeze function $s : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ *is constructed as follows:*

1. *All stable equilibrium orientations θ map identically to θ.*

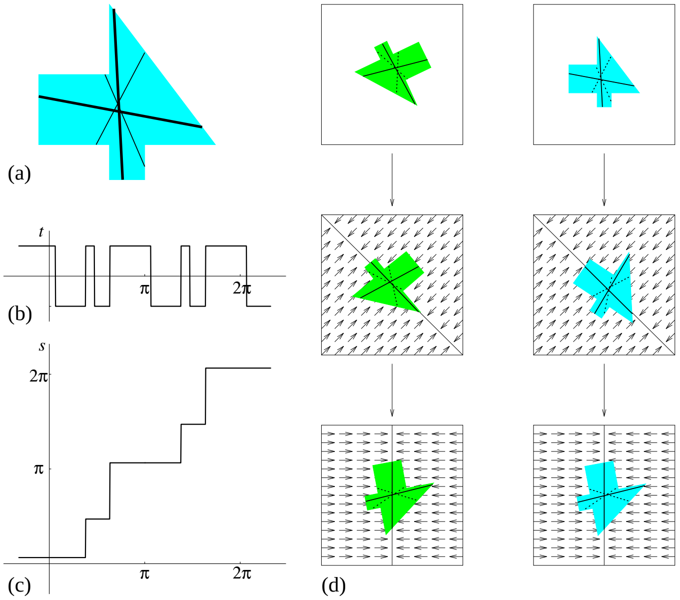

Figure 2.8: (a) Polygonal part. Stable (thick line) and unstable (thin line) bisectors are also shown. (b) Turn function, which predicts the orientations of the stable and unstable bisectors. (c) Squeeze function, constructed from the turn function. (d) Alignment strategy for two arbitrary initial configurations. See URL `http://www.cs.cornell.edu/home/karl/Cinema` for an animated simulation.

2. All unstable equilibrium orientations map (by convention) to the nearest counterclockwise stable orientation.

3. All orientations θ with $t(\theta) = 1$ (-1) map to the nearest counterclockwise (clockwise) stable orientation.

Then s describes the orientation transition of P induced by $\mathcal{A}$.

Proof: Assume that part P initially is in pose (x, y, θ) in array $\mathcal{A}$. Because of 2PHASE, we can assume that P translates towards the center line l until condition **I** is satisfied without changing its orientation θ. P will change its orientation until the moment is zero, i.e. $t = 0$: A positive moment ($t > 0$) causes counterclockwise motion, and a negative moment ($t < 0$) causes clockwise motion until the next root of t is reached. □

We conclude that any connected polygonal part, when put in a squeeze field, reaches one of a *finite* number of possible orientation equilibria [BDMM94a,BDM96b]. The motion of the part and, in particular, the mapping between initial orientation and equilibrium orientation is described by the squeeze function, which is derived from the turn function as described in Lemma 8. Note that all squeeze functions derived from turn functions are monotone step-shaped functions.

Goldberg [Gol93] has given an algorithm that automatically synthesizes a manipulation strategy to uniquely orient a part, given its squeeze function. While Goldberg's algorithm was designed for squeezes with a robotic parallel-jaw gripper, in fact, it is more general, and can be used for arbitrary monotone step-shaped squeeze functions. The output of Goldberg's algorithm is a sequence of angles that specify the required directions of the squeezes. Hence these angles specify the direction of the squeeze line in our force vector

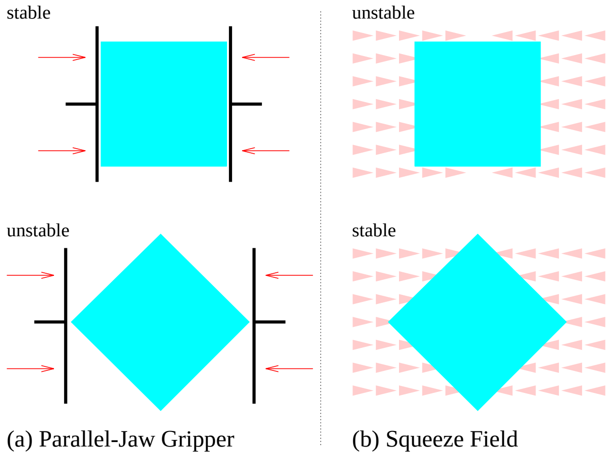

(a) Parallel-Jaw Gripper (b) Squeeze Field

Figure 2.9: Equilibrium configurations for a square-shaped part using (a) a frictionless parallel-jaw gripper and (b) a MEMS squeeze field. In this example, stable and unstable equilibria are reversed.

fields (for example the two-step strategies in Figures 2.1 and 2.8d).

It is important to note that the equilibria obtained by a MEMS squeeze field and by a parallel-jaw gripper will typically be different, even when the squeeze directions are identical. For example, to see this, consider squeezing a square-shaped part (Figure 2.9). Stable and unstable equilibria are reversed. This shows that our mechanical analysis of equilibrium is different from that of the parallel-jaw gripper. Let us summarize these results:

Theorem 9 [**BDM96b**] *Let P be a polygon whose interior is connected. There exists an alignment strategy consisting of a sequence of squeeze fields that uniquely orients P up to symmetries.*

Since the strategies of Theorem 9 consist of fields with squeeze lines at arbitrary angles through the origin, we call them *general* $\mathbb{S}^1$ *squeeze strategies*, or henceforth *general squeeze strategies*.

Corollary 10 [**BDM96b**] *The alignment strategies of Theorem 9 have $O(k\,n^2)$ steps, and they may be computed in time $O(k^2\,n^4)$, where k is the maximum number of edges that a bisector of P can cross. In the case where P is convex, the alignment strategy has $O(n)$ steps and can be computed in time $O(n^2)$.*

Proof: Proposition 4 states that a polygon with n vertices has $E = O(k\,n^2)$ stable orientation equilibria in a squeeze field ($O(n)$ if P is convex). This means that the image of its corresponding squeeze function is a set of E discrete values. Given such a squeeze function, Goldberg's algorithm [Gol93] constructs alignment strategies with $O(E)$ steps. Planning complexity is $O(E^2)$. □

Goldberg's strategies [Gol93] have the same complexity bounds for convex and non-convex parts, because when using squeeze grasps with a parallel-jaw gripper, only the convex hull of the part need be considered. This is not the case for programmable vector fields, where manipulation strategies for non-convex parts are more expensive. As described

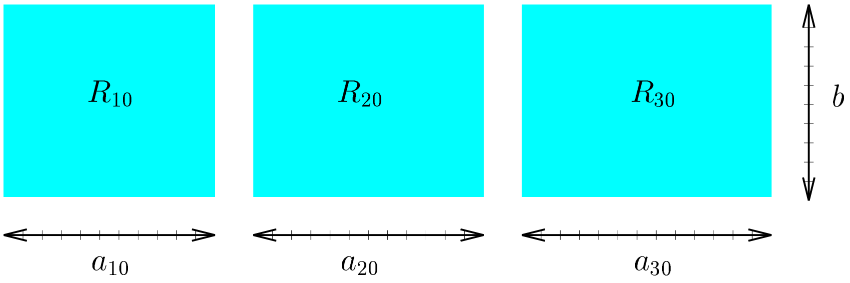

Figure 2.10: Sample rectangles R_{10}, R_{20}, and R_{30}. Edge a is 10, 20, and 30% longer than edge b, respectively.

in Section 2.2, there could exist parts that have $E = \Omega(k\,n^2)$ orientation equilibria in a squeeze field, which would imply alignment strategies of length $\Omega(k\,n^2)$ and planning complexity $\Omega(k^2\,n^4)$.

Note that the turn and squeeze functions have a period of π due to the symmetry of the squeeze field; rotating the field by an angle of π produces an identical vector field. Rotational symmetry in the part also introduces periodicity into these functions. Hence, general squeeze strategies (see Theorem 9) orient a part *up to symmetry*, that is, up to symmetry in the part *and* in the squeeze field. Similarly, the grasp plans based on squeeze functions in [Gol93] can orient a part with a macroscopic gripper only modulo symmetry in the part and in the gripper.[2] Since we reduce to the squeeze function algorithm in [Gol93], it is not surprising that this phenomenon is also manifested for squeeze vector fields as well. For a detailed discussion of parts orientation modulo symmetry see [Gol93].

In Section 5.1 we will present new and improved manipulation algorithms that reduce the number of equilibria to $E = O(k\,n)$.

This scheme may be generalized to the case where l is slightly curved, as in the "node" of the vibrating plate in Figure 1.4. See [BBG95] for details. The remaining sections of this paper investigate using more exotic fields (not simple squeeze patterns) to

1. allow disconnected polygons,
2. relax assumption 2PHASE,
3. reduce the planning complexity,
4. reduce the number of equilibria,
5. reduce the execution complexity (strategy length), and
6. determine feasibility results and limitations for manipulation with general force fields.

2.4 Example: Uniquely Orienting Rectangular Parts

To demonstrate the equilibrium analysis from Section 2.1 and the alignment algorithm from Section 2.3, we will generate plans for uniquely orienting several planar polygonal parts (up to part symmetry). In particular, here we will consider the simple case of three rectangles R_{10}, R_{20}, and R_{30}, which have sides a and b such that a is 10, 20, and 30 percent longer than b, respectively (Figure 2.10).

Our algorithm first determines stable and unstable equilibria of the parts, which correspond to the negative and positive steps in the turn function, respectively (see Lemma 7).

---

[2] Parallel-jaw gripper symmetry is also modulo π.

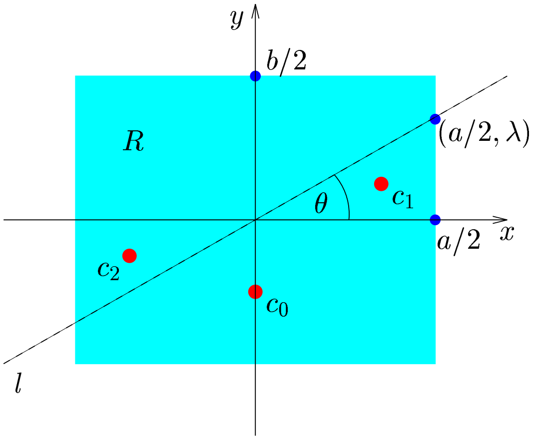

Figure 2.11: Analytically determining the moment function for a rectangular part R with sides of length a and b. c_0 is the center of mass of the segment below the x-axis. c_1 and c_2 are the centers of the triangular segments between x-axis and line l.

Table 2.1: Equilibria of rectangular parts R_{10}, R_{20}, and R_{30} in a unit squeeze field with vertical squeeze line.

Part	Equilibrium orientations θ	
	stable	unstable
R_{10}	$0.97, 2.18, 4.11, 5.32$	$0, \pi/2, \pi, 3\pi/2$
R_{20}	$1.29, 1.85, 4.43, 4.99$	$0, \pi/2, \pi, 3/pi/2$
R_{30}	$\pi/2, 3\pi/2$	$0, \pi$

The turn function can be obtained as the sign of the moment function, which, for polygonal parts, is a piecewise rational function, and can be derived automatically from the part geometry. For example, consider the rectangle R in Figure 2.11: A line l through the origin bisects R. If l is placed such that it intersects the right edge of R at $(a/2, \lambda)$ with $-b/2 \leq \lambda \leq b/2$, then the COM of the segment below l is

$$\begin{aligned} c_\lambda &= \left(\frac{ab}{2}c_0 + \frac{a\lambda}{4}(c_1 - c_2)\right)\frac{2}{ab} \\ &= c_0 + \frac{\lambda}{2b}2c_1 \\ &= (\frac{a\lambda}{3b}, -\frac{b}{4} + \frac{\lambda^2}{3b}) \end{aligned}$$

The moment function is the inner product between the vector c_λ, and the direction of the line l. For balanced moment, this product must be zero, which gives us the following condition for equilibrium:

$$\begin{aligned} 0 &= (\frac{a\lambda}{3b}, -\frac{b}{4} + \frac{\lambda^2}{3b}) \cdot (\frac{a}{2}, \lambda) \\ &= \frac{a^2\lambda}{6b} - \frac{b\lambda}{4} + \frac{\lambda^3}{3b} \end{aligned}$$

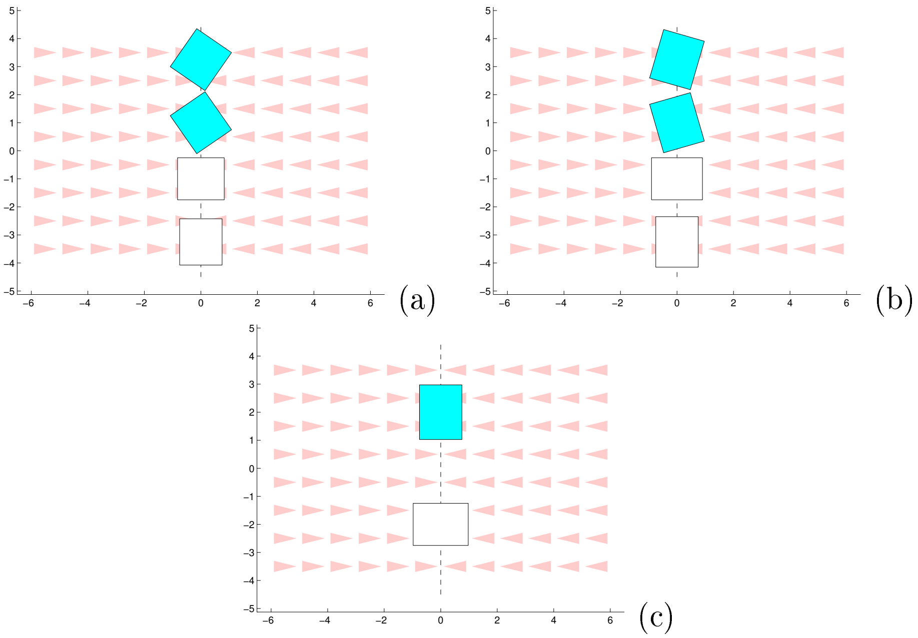

Figure 2.12: Stable (dark) and unstable (white) equilibria of three rectangular parts in a unit squeeze field with vertical squeeze line: (a) R_{10}, edge ratio 1.1; (b) R_{20}, edge ratio 1.2; (c) R_{30}, edge ratio 1.3. R_{10} and R_{20} exhibit two stable equilibria, R_{30} exhibits only one.

$$\begin{aligned} &= \frac{\lambda}{12b}(2a^2 - 3b^2 + 4\lambda^2) \\ \text{So } \lambda &= 0 \\ \text{or } \lambda &= \pm\frac{1}{2}\sqrt{3b^2 - 2a^2} \\ &= \pm\frac{b}{2}\sqrt{3 - 2c^2} \quad \text{for } a = cb \end{aligned}$$

This means that for rectangles with edge ratio $c \leq \sqrt{3/2} \approx 1.22$ (such as R_{10} and R_{20}), there exist equilibrium orientations at angles $\theta = \arctan(\pm\sqrt{3/c^2 - 2})$. For rectangles with larger edge ratio c (such as R_{30}), an equilibrium exists only at $\theta = 0$. A similar analysis can be performed for all other placements of the line l, see [BDM96b] for more details. Equilibrium orientations as determined by our planner are shown in Figure 2.12 and Table 2.1. Since all of our parts are symmetric with respect to rotation by π, for the remainder of this example we will consider all angles modulo π.

From the equilibrium orientations in Table 2.1 the algorithm generates the squeeze function, according to Lemma 8. Note that steps in the squeeze function occur at angles corresponding to unstable equilibria, while the image of the squeeze function is the set of all stable equilibrium orientations (see Figure 2.13).

Finally, the squeeze function is used as input for Goldberg's planning algorithm [Gol93], which returns as output a sequence of squeeze angles. A sequence of two squeeze fields, with a relative angle of $\pi/2$, is sufficient to uniquely orient both R_{10} and R_{20}. See Figure 2.14 for a sample execution of this plan for two arbitrary initial poses. R_{30} requires only one squeeze field at an arbitrary angle.

It was shown in [BDM96b] that this algorithm can uniquely orient arbitrary polygons

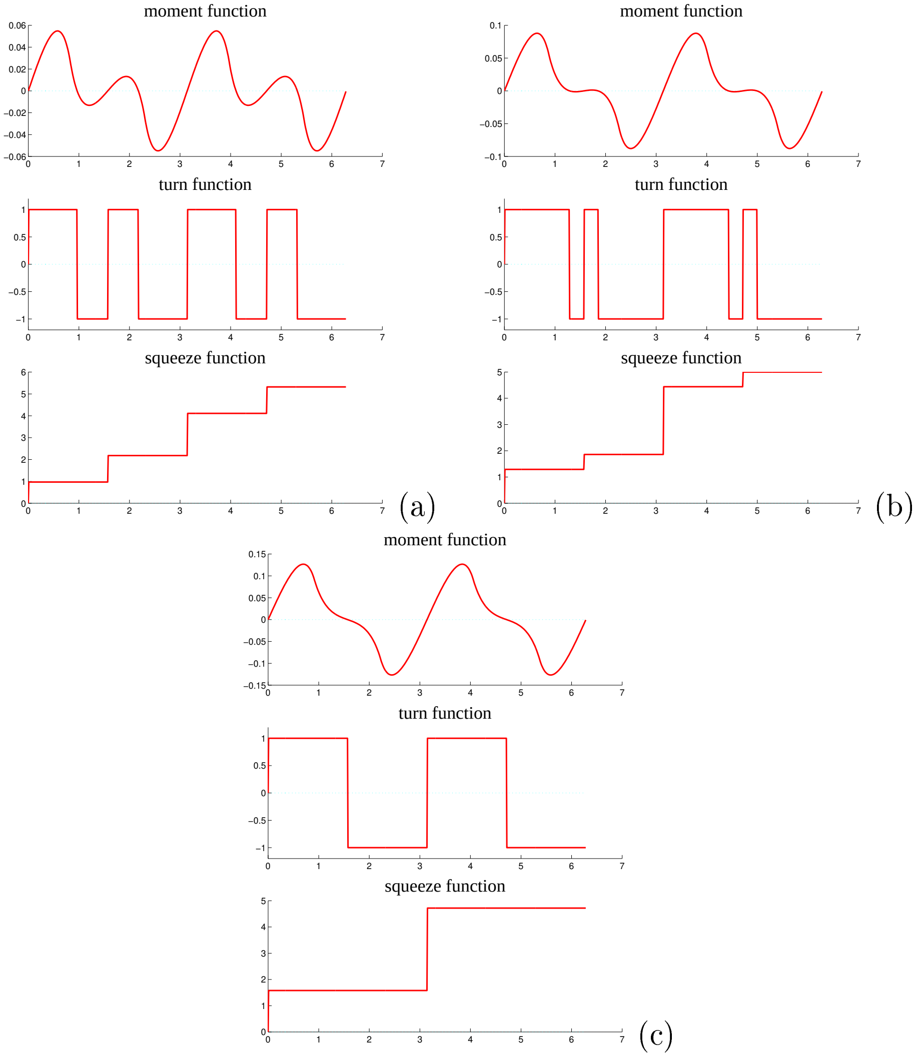

Figure 2.13: Moment function, turn function, and squeeze function for three rectangular parts: (a) R_{10}, edge ratio 1.1; (b) R_{20}, edge ratio 1.2; (c) R_{30}, edge ratio 1.3. R_{10} and R_{20} exhibit two stable equilibria for θ in the range $[0 \ldots \pi]$, R_{30} exhibits only one.

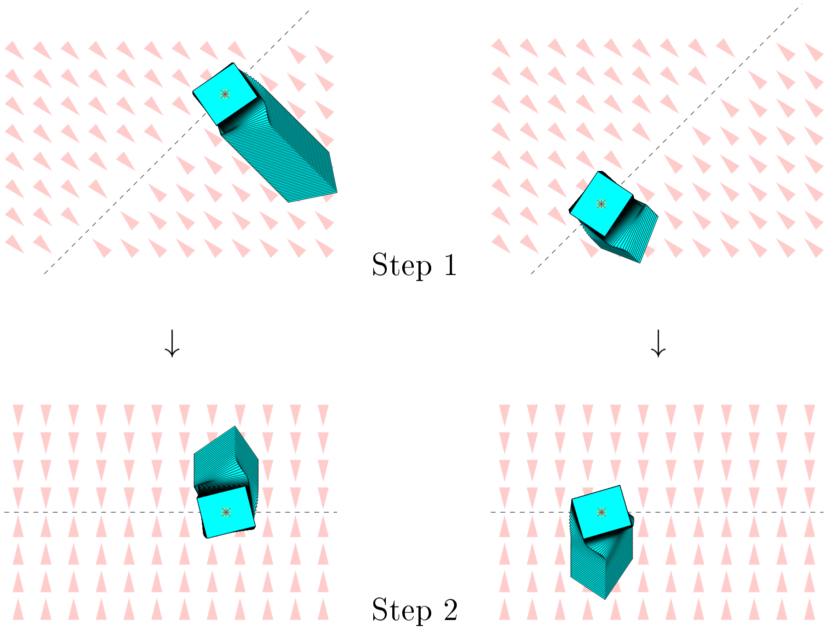

Figure 2.14: Two-step alignment plan for rectangle R_{20}. After two steps, R_{20} reaches a unique *orientation* θ independent of its initial pose. However, the *position* (x, y) is not unique.

from any initial configuration (up to part symmetry). However, recall that for this algorithm to work we have made several important assumptions that idealize the practical vibratory feeding devices presented in Section 6.3.

1. 2Phase assumption, which states that translational and rotational motion of the part is decoupled, implying that the turn function is independent of the initial offset of the part from the squeeze line; see also Section 2.5.
2. Depending on the part shape, the algorithm may generate alignment plans with unit squeeze fields at arbitrary angles. Due to mechanical design limitations, usually not all of these fields will be feasible to implement on most vibratory device setups.
3. The resulting plans uniquely *orient* a part, but the final *translational position* can not be predicted.

In the remainder of this thesis, we will investigate new manipulation strategies that address these key issues. In particular, in Section 5.2 we will develop algorithms for devices with a limited "vocabulary" of available force fields, which will result in a "manipulation grammar" for unique, sensorless posing strategies for arbitrary planar, polygonal parts.

2.5 Relaxing the 2Phase Assumption

In Section 2.3, assumption 2Phase allowed us to determine successive equilibrium positions in a sequence of squeezes, by a quasi-static analysis that decouples translational and rotational motion of the moving part. For any part, this obtains a *unique* orientation equilibrium (after several steps). If 2Phase is relaxed, we obtain a dynamic manipulation problem, in which we must determine the equilibria (x, θ) given by the part orientation θ and the offset x of its center of mass from the squeeze line. A stable equilibrium is a (x_i, θ_i)

pair in $\mathbb{R} \times \mathbb{S}^1$ that acts as an *attractor* (the x offset in an equilibrium is, surprisingly, usually not 0, see Figure 2.5). Again, we can compute these (x_i, θ_i) equilibrium pairs *exactly*, as outlined in Section 2.2.

Considering (x_i, θ_i) equilibrium pairs has another advantage. We can show that, even without 2PHASE, after two successive, orthogonal squeezes, the set of stable poses of any part can be reduced from $\mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$ to a *finite* subset of $\mathcal{C}$ (the configuration space of part P); see Claim 39 (Section 5.2.1). Subsequent squeezes will preserve the finiteness of the state space. This will significantly reduce the complexity of a task-level motion planner. Hence if assumption 2PHASE is relaxed, this idea still enables us to simplify the general motion planning problem (as formulated e.g. by Lozano-Pérez, Mason, and Taylor in [LPMT84]) to that of Erdmann and Mason [EM88]. Conversely, relaxing assumption 2PHASE raises the complexity from the "linear" planning scheme of Goldberg [Gol93] to the forward-chaining searches of Erdmann and Mason [EM88], or Donald [Don90].

Chapter 3

Lower Bounds: What Programmable Vector Fields Cannot Do

We now present "lower bounds" — constituting vector fields and parts with pathological behavior, making them unusable for manipulation. These counterexamples show that we must be careful in choosing programmable vector fields, and that, *a priori*, it is not obvious when a field is well-behaved.

3.1 Unstable Fields

In Section 2 we saw that in a vector field with a simple squeeze pattern (see again Figure 2.1), polygonal parts reach certain equilibrium poses. This raises the question of a *general classification of all those vector fields in which every part has stable equilibria.* There exist vector fields that do not have this property even though they are very similar to a simple squeeze.

Definition 11 *A* skewed field f_S *is a vector field given by* $f_S(x, y) = -\text{sign}(x)(1, \epsilon)$*, where* $0 \neq \epsilon \in \mathbb{R}$.

Proposition 12 *A skewed field induces* no *stable equilibrium on a disk-shaped part.*

Proof: Consider Figure 3.1, which shows a skewed field with $\epsilon = -\frac{2}{3}$: Only when the center of the disk coincides with the center of the squeeze pattern do the translational forces acting on the disk balance. But it will still experience a positive moment that will cause rotation. $\square$

Proposition 13 *A diagonally skewed field* ($\epsilon = \pm 1$) *induces* no *stable equilibrium on a square-shaped part.*

Proof: To reach equilibrium in a diagonally skewed field, the squeeze line has to bisect the part such that the connector between the left and the right centers of area is diagonal (i.e., parallel to the force vectors). An analysis similar to the proof of Proposition 4 (Section 2.2) or Figure 2.11 (Section 2.4) shows that for a square, no bisector placement can achieve an angle with the connector of less than 83°. $\square$

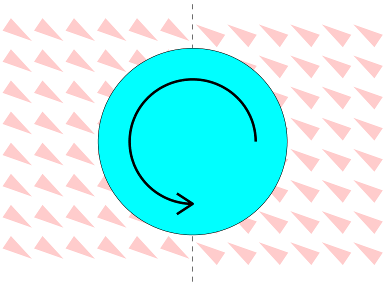

Figure 3.1: Unstable part in the skewed squeeze field ($\epsilon = -\frac{2}{3}$). The disk with center on the squeeze line will keep rotating. Moreover, it has *no* stable equilibrium in this field.

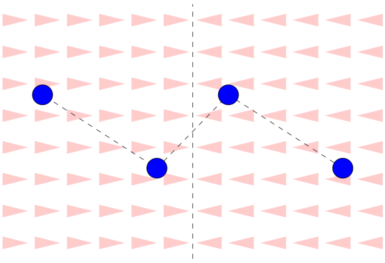

Figure 3.2: S-shaped part P_S with four rigidly connected point-contact "feet" in unstable total equilibrium (forces and moments balance). There exists *no* stable equilibrium position for this part in a vector field with a simple squeeze pattern. For an animation see URL `http://www.cs.cornell.edu/home/karl/MicroManipulation/Patho`.

3.2 Unstable Parts

Similarly we would like to identify the *class of all those parts that always reach stable equilibria* in particular vector fields. From Section 2 we know that connected polygons in simple squeeze fields satisfy this condition. This property relies on finite area contacts: it does not hold for point contacts. As a counterexample consider the part P_S in Figure 3.2.

Proposition 14 *There exist parts that do not have stable equilibria in a simple squeeze field.*

Proof: The S-shaped part in Figure 3.2 has four rigidly connected "feet" with small contact surfaces. As the area of each of these four feet approaches zero, the part has *no* stable equilibrium in a simple squeeze field. There is only one orientation for the part in which both force and moment balances out, and this orientation is unstable. $\square$

In Section 4.2 we will discuss this phenomenon in greater detail, after the tools necessary for a complete mathematical analysis have been developed.

Finally, the *number of stable equilibria* of a given part influences both the planning complexity and the plan length of an alignment strategy. It also affects the resolution of the vector field that is necessary to perform a strategy accurately. Even though all parts we have considered exhibit only one or two orientation equilibria, there exist no tight bounds on the maximum number of orientation equilibria in a unit squeeze field.

Proposition 15 *Let n be the number of vertices of a polygon P, and let k be the maximum number of edges that a bisector of P can cross:*

A. *Regular polygons have n stable orientation equilibria in a squeeze field.*

B. *Every connected polygon has $O(k\,n^2)$ stable orientation equilibria in a squeeze field.*

Proof:

A. Because of their part symmetry, regular polygons have $2n$ equilibria. Half of them are stable, the other n are unstable.

B. See Section 2.2.

□

As described in Section 2.2, there exist simple polygons with n vertices that can be bisected by a straight line in up to $O(k\,n^2)$ topologically different ways [BDH97]. This suggests that there could be parts that have $\Omega(k\,n^2)$ orientation equilibria in a squeeze field, which would imply alignment strategies of length $\Omega(k\,n^2)$ and planning complexity $\Omega(k^2\,n^4)$.

While the counterexample in Figure 3.2 may be plausibly avoided by prohibiting parts with "point contacts," the other examples (Figure 3.1 and Proposition 15B) are more problematic. In Section 4, we show how to choose programmable vector fields that exclude some of these pathological behaviors, by using the *theory of potential fields* to describe a class of force vector fields for which *all* polygonal parts have stable equilibria. In Section 5.1, we show how to combine these fields to obtain new fields in which all parts have only $O(k\,n)$ equilibria.

We believe parts with point contact (not having finite area contact) will behave badly in *all* vector fields. We can model a point contact with delta functions, such that e.g. for a point contact P_0 at (x_0, y_0): $\int_{P_0} f\;dA = \int f\;\delta(x_0, y_0)\;dA = f(x_0, y_0)$. This model is frequently used in mechanics (see e.g. [Erd94]). Point contact permits rapid, discontinuous changes in force and moment. Hence, bodies with point contact will tend to exhibit instabilities, as opposed to flat parts that are in contact with a large number of (elastic) actuators. Finally, we believe that as the area contact—the size of the "feet" of a part—approaches zero, the part may become unstable. This represents a design constraint on parts which are to be manipulated using programmable planar parts feeders.

The lower bounds we demonstrate are indications of the pathologies that can arise when fields without potential or parts with point contact are permitted. Each of our counterexamples (Figures 3.1 and 3.2) is "generic" in that it can be generalized to a very large class of similar examples. However, these lower bounds are just a first step, and one wishes for examples that delineate the capabilities of programmable vector fields for planar parts manipulation even more precisely.

The separating field shown in Figure 1.1c is not a potential field, and there exist parts that will spin forever, without equilibrium, in this field (this follows by generalizing the construction in Figure 3.1). However, for *specific parts*, such as those shown in Figure 1.1, this field is useful if we can pose the parts appropriately first (e.g., using the potential field shown in Figure 1.1b).

Finally, we may "surround" non-potential fields with potential fields to obtain reasonable behavior in some cases. Figure 1.1 shows how to "surround" a non-potential field in *time* by potential fields, to eliminate pathologies. Similarly, we can surround non-potential fields *spatially*. For example, if field 1.1c could be surrounded by a larger potential field, then after separation, parts can reach a stable equilibrium.

Non-potential fields can be used safely with the following methodology: Let $H \subset \mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$ be the undesirable limit set. For example, H could be a limit cycle where the part spins forever. Let $\widehat{P}_V(H)$ be the weak pre-image [LPMT84,Don89] of H under the field V. If we can ensure that the part starts in a configuration $z \notin \widehat{P}_V(H)$, it will not reach the unwanted limit cycle. For example, in Figure 1.1 the centering step (b) ensures that the part does not end up on the border between the two separating fields, where it would spin forever in step (c).

Chapter 4

Completeness: Classification Using Potential Fields

We are interested in a *general classification of all those vector fields in which every part has stable equilibria.* As motivation, recall that a skewed vector field, even though very similar to a regular squeeze field (see again Figure 2.1), induces *no* stable equilibrium in a disk-shaped part (Figure 3.1). In this section we give a family of vector fields that will be useful for manipulation tasks. These fields belong to a specific class of vector fields: the class of fields that have a *potential.*

We believe that fields without potential will often induce pathological behavior in many parts. Fields without potential admit paths along which a particle (point mass) will gain energy. Since mechanical parts are rigid aggregations of particles, this may induce unstable behavior in larger bodies. However, there are some cases where non-potential fields may be useful. For example, see Figure 1.1c, which is *not* a potential field. Such fields may be employed to *separate* but not to stabilize, pose, or orient parts. This strong statement devolves to our proof that fields like Figure 3.1 do not have well-behaved equilibria. Hence, they should only be employed when we want to induce an unstable system that will cast parts away from equilibrium, e.g. in order to sort or separate them.

Consider the class of vector fields on $\mathbb{R}^2$ that have a *potential*, i.e. fields f in which the work is independent of the path, or equivalently, the work on any closed path is zero, $\oint f \cdot ds = 0$. In a potential field each point (x, y) is assigned a real value $U(x, y)$ that can be interpreted as its potential energy. When U is smooth, then the vector field f associated with U is the gradient $-\nabla U$. In general, $U(x, y)$ is given, up to an additive constant, by the path integral $\int_\alpha f \cdot ds$ (when it exists and it is unique), where α is an arbitrary path from a fixed reference point (x_0, y_0) to (x, y). Assuming first-order dynamics, When U is smooth, an ideal point object is in stable equilibrium when it is at a local minimum of U.

Definition 16 *Let f be a force vector field on $\mathbb{R}^2$, and let $\mathbf{p}$ be a point that is offset from a fixed reference point $\mathbf{q}$ by a vector $\mathbf{r}(\mathbf{p}) = \mathbf{p} - \mathbf{q}$. We define the* generalized force *F as the force* and *moment induced by f at point $\mathbf{p}$:*

$$F(\mathbf{p}) \;=\; (f(\mathbf{p})\;,\;\mathbf{r}(\mathbf{p}) \times f(\mathbf{p})) \tag{4.1}$$

Let P be a part of arbitrary shape, and let $P_{\mathbf{z}}$ denote the part P in pose $\mathbf{z} = (x, y, \theta) \in \mathcal{C}$.

$$\begin{tikzcd}
F_P \arrow[r, "\int_\alpha"] & U_P \\
f \arrow[u, "\text{lifting: area integral}", "\int_P"'] \arrow[r, "\int_\alpha", "\text{potential: path integral}"'] & U \arrow[u, "\int_P"]
\end{tikzcd}$$

Figure 4.1: Determining the potential, and lifting are commutative operations on force vector fields.

We define the lifted force field f_P *as the area integral of the force induced by* f *over* $P_{\mathbf{z}}$*:*

$$f_P(\mathbf{z}) \;=\; \int_{P_{\mathbf{z}}} f \; dA \tag{4.2}$$

The lifted generalized force field F_P *is defined as the area integral of the force* and *moment induced by* f *over* P *in configuration* $\mathbf{z}$*:*

$$\begin{aligned} F_P(\mathbf{z}) \;&=\; \int_{P_{\mathbf{z}}} F \; dA \\ &=\; \left(\int_{P_{\mathbf{z}}} f \; dA \;,\; \int_{P_{\mathbf{z}}} \mathbf{r} \times f \; dA\right) \end{aligned} \tag{4.3}$$

Hence, F_P *is a vector field on* $\mathcal{C}$*. Finally, we define the* lifted potential $U_P : \mathcal{C} \to \mathbb{R}$*.* U_P *is the area integral of the potential* U *over* P *in configuration* $\mathbf{z}$*:*

$$U_P(\mathbf{z}) \;=\; \int_{P_{\mathbf{z}}} U \; dA \tag{4.4}$$

We now show that the category of potential fields is closed under the operation of lifting, and that U_P is the potential of F_P (see Figure 4.1). Note that U need not be smooth.

Let $g : X \to Y$ and $h : Y \to Z$. Let $k : X \to Z$ be the function which is the composition of g and h, defined by $k(x) = h(g(x))$. In the following proposition, we use the notation $h(g)$ to denote k, the function composition of g and h.

Proposition 17 *Let* f *be a force field on* $\mathbb{R}^2$ *with potential* U*, and let* P *be a part of arbitrary shape. For the lifted generalized force field* F_P *and the lifted potential* U_P *the following equality holds:* $U_P = \displaystyle\int_P U \; dA = \int_\alpha F_P \cdot d\mathbf{z} + c$*, where* α *is an arbitrary path in* $\mathcal{C}$ *from a fixed reference point, and* c *is a constant.*

Proof: Given a force field f with potential U, and a part P. We define P^* as the set $\{(r, \eta) \mid (r\cos\eta, r\sin\eta) \in P\} \subset \mathbb{R} \times \mathbb{S}^1$. P^* is a representation of P in polar coordinates: $\mathbf{p} = (r\cos\eta, r\sin\eta) \in P$ if and only if $(r, \eta) \in P^*$.

We write $P_{\mathbf{z}}$ to denote P in pose $\mathbf{z} = (x_{\mathbf{z}}, y_{\mathbf{z}}, \theta_{\mathbf{z}})$. If P is moved into pose $\mathbf{z}$, then the point $\mathbf{p}$ moves to $\mathbf{p}_{\mathbf{z}} = (x_{\mathbf{z}} + r\cos(\theta_{\mathbf{z}} + \eta)\;,\; y_{\mathbf{z}} + r\sin(\theta_{\mathbf{z}} + \eta)) = (x_{\mathbf{z}}, y_{\mathbf{z}}) + \mathbf{r}_{\mathbf{z}}$. Let us assume that for a given P, the COM of P is at $\mathbf{0}$; this implies that the COM of $P_{\mathbf{z}}$ lies at $(x_{\mathbf{z}}, y_{\mathbf{z}})$.

We define three families of functions ρ, ζ, and α as follows:

$$\rho_{r,\eta}: \quad [0,1] \quad \rightarrow \quad \mathbb{R}^2$$

such that $\rho_{r,\eta}$ is a smooth path in $\mathbb{R}^2$ with $\rho_{r,\eta}(0) = \mathbf{0}$ and $\rho_{r,\eta}(1) = \mathbf{p_0} = (r\cos\eta, r\sin\eta)$.

$$\zeta_{\mathbf{z}}: \quad [0,1] \quad \rightarrow \quad \mathbb{R}^2 \times \mathbb{S}^1$$

such that $\zeta_{\mathbf{z}}$ is a smooth path in $\mathbb{R}^2 \times \mathbb{S}^1$ with $\zeta_{\mathbf{z}}(\mathbf{0}) = \mathbf{0}$ and $\zeta_{\mathbf{z}}(1) = \mathbf{z} = (x_{\mathbf{z}}, y_{\mathbf{z}}, \theta_{\mathbf{z}})$.

$$\begin{aligned} \alpha_{r,\eta}: \quad \mathbb{R}^2 \times \mathbb{S}^1 \quad &\rightarrow \quad \mathbb{R}^2 \\ (x, y, \theta) \quad &\mapsto \quad (x + r\cos(\theta + \eta)\ , \\ &\qquad\quad y + r\sin(\theta + \eta)) \end{aligned}$$

So $\zeta_{\mathbf{z}}$ is an arbitrary smooth path from $\mathbf{0}$ to $\mathbf{z}$ in $\mathcal{C}$, and $\alpha_{r,\eta}(\zeta_{\mathbf{z}})$ is a smooth path in $\mathbb{R}^2$ from $\mathbf{p_0} = (r\cos\eta, r\sin\eta)$ to $\mathbf{p_z} = (x_{\mathbf{z}} + r\cos(\theta_{\mathbf{z}} + \eta)\ ,\ y_{\mathbf{z}} + r\sin(\theta_{\mathbf{z}} + \eta))$. Recall that $\alpha_{r,\eta}(\zeta_{\mathbf{z}})$ is the function composition of $\zeta_{\mathbf{z}}$ and $\alpha_{r,\eta}$.

We are interested in the potential of U at $\mathbf{p_z}$.

$$U(\mathbf{p_z}) \quad = \quad \int_{\beta} f \cdot ds$$

where β is some path from $\mathbf{0}$ to $\mathbf{p_z}$. The integral is path-independent because f has potential U. Since we can choose any path, we choose a path β that consists of two parts: (1) a path from $\mathbf{0}$ to the point $\mathbf{p_0}$, and (2) a path from $\mathbf{p_0}$ to $\mathbf{p_z}$.

$$= \quad \int_{\rho_{r,\eta}} f \cdot ds \ + \int_{\alpha_{r,\eta}(\zeta_{\mathbf{z}})} f \cdot ds$$

Note that the path $\alpha_{r,\eta}(\zeta_{\mathbf{z}})$, given by the composition of $\zeta_{\mathbf{z}}$ and $\alpha_{r,\eta}$, depends on $\mathbf{z}$ as well as on r and η, but $\rho_{r,\eta}$ is independent of $\mathbf{z}$. The left integral is the potential difference between $\mathbf{p_0}$ and $\mathbf{0}$. W.l.o.g. let us choose $U(\mathbf{0}) = 0$.

$$= \quad U(\mathbf{p_0}) \ + \int_{\zeta_{\mathbf{z}}} f(\alpha_{r,\eta}) \cdot (\mathbf{J}\ d\mathbf{z})$$

where $\mathbf{J}$ is the Jacobian

$$\mathbf{J} = \begin{pmatrix} \frac{\partial \alpha_x}{\partial x} & \frac{\partial \alpha_x}{\partial y} & \frac{\partial \alpha_x}{\partial \theta} \\ \frac{\partial \alpha_y}{\partial x} & \frac{\partial \alpha_y}{\partial y} & \frac{\partial \alpha_y}{\partial \theta} \end{pmatrix} = \begin{pmatrix} 1 & 0 & -r\sin(\theta+\eta) \\ 0 & 1 & r\cos(\theta+\eta) \end{pmatrix}$$

which is the derivative of $\alpha_{r,\eta}$. $f(\alpha_{r,\eta})$ is the function composition of $\alpha_{r,\eta}$ and f. Also note that $d\mathbf{z} = (dx, dy, d\theta)$.

$$\begin{aligned} = \quad U(\mathbf{p_0}) \ + \int_{\zeta_{\mathbf{z}}} \Big(f_x(\alpha_{r,\eta})\ ,\ f_y(\alpha_{r,\eta})\ , \\ r\cos(\theta + \eta) f_y(\alpha_{r,\eta}) - r\sin(\theta + \eta) f_x(\alpha_{r,\eta})\Big) \cdot d\mathbf{z} \end{aligned}$$

$$= \quad U(\mathbf{p_0}) \ + \int_{\zeta_{\mathbf{z}}} F(\alpha_{r,\eta}) \cdot d\mathbf{z} \tag{4.5}$$

Equation 4.5 states that the potential at a point $\mathbf{p_z} = (x_{\mathbf{z}} + r\cos(\theta_{\mathbf{z}} + \eta)\ ,\ y_{\mathbf{z}} + r\sin(\theta_{\mathbf{z}} + \eta))$ can be determined as the sum of two integrals: The first integrates the force f over a path from $\mathbf{0}$ to $\mathbf{p_0} = (r\cos\eta, r\sin\eta)$. If we choose $U(\mathbf{0}) = 0$, then the first integral is the potential at point $\mathbf{p_0}$. The right part of the expression can be understood as the path integral of the generalized force from $\mathbf{p_0}$ to $\mathbf{p_z}$.

With this result, we can now consider the *lifted* potential U_P at a point $\mathbf{z} =$

$(x_{\mathbf{z}}, y_{\mathbf{z}}, \theta_{\mathbf{z}}) \in \mathcal{C}$.

$$U_P(\mathbf{z}) = \int_{P_{\mathbf{z}}} U(\mathbf{p})\ dA$$

$$= \iint_{P^*} U(\mathbf{p_z})\ r\ dr\ d\eta$$

where $\mathbf{p_z} = (x_{\mathbf{z}} + r\cos(\theta_{\mathbf{z}} + \eta)\ ,\ y_{\mathbf{z}} + r\sin(\theta_{\mathbf{z}} + \eta))$ such that $(r, \eta) \in P^*$.

$$= \iint_{P^*} \left(U(\mathbf{p_0}) + \int_{\zeta_{\mathbf{z}}} F(\alpha_{r,\eta}) \cdot d\mathbf{z} \right) r\ dr\ d\eta$$

by using Equation 4.5. Again, $F(\alpha_{r,\eta})$ denotes the function composition of $\alpha_{r,\eta}$ and F.

$$= \iint_{P^*} U(\mathbf{p_0})\ r\ dr\ d\eta\ + \iint_{P^*} \left(\int_{\zeta_{\mathbf{z}}} F(\alpha_{r,\eta}) \cdot d\mathbf{z} \right) r\ dr\ r\eta$$

The first expression is the area integral of U over P. From Definition 16 it follows that this expression is equal to $U_P(\mathbf{0})$ (note that $U_P(\mathbf{0})$ is a constant which does not depend on $\mathbf{z}$).

$$= U_P(\mathbf{0}) + \iint_{P^*} \left(\int_0^1 F(\alpha_{r,\eta}(\zeta_{\mathbf{z}}(t))) \cdot \zeta'_{\mathbf{z}}(t)\ dt \right) r\ dr\ r\eta$$

where $\zeta'_{\mathbf{z}}$ is the derivative of $\zeta_{\mathbf{z}}$. The dot product yields a scalar value. We can now switch the integrals.

$$= \int_0^1 \left(\iint_{P^*} F(\alpha_{r,\eta}(\zeta_{\mathbf{z}}(t))) \cdot \zeta'_{\mathbf{z}}(t)\ r\ dr\ d\eta \right) dt + U_P(\mathbf{0})$$

$\zeta'_{\mathbf{z}}$ is constant with respect to the integration parameters r and η. Hence we can move $\zeta'_{\mathbf{z}}$ outside of the area integral.

$$= \int_0^1 \left(\iint_{P^*} F(\alpha_{r,\eta}(\zeta_{\mathbf{z}}(t)))\ r\ dr\ d\eta \right) \cdot \zeta'_{\mathbf{z}}(t)\ dt + U_P(\mathbf{0})$$

$$= \int_0^1 \left(\iint_{P^*} F(\zeta_{\mathbf{z},x}(t) + r\cos(\zeta_{\mathbf{z},\theta}(t) + \eta)\ , \right.$$

$$\left. \zeta_{\mathbf{z},y}(t) + r\sin(\zeta_{\mathbf{z},\theta}(t) + \eta))\ r\ dr\ d\eta \right) \cdot \zeta'_{\mathbf{z}}(t)\ dt + U_P(\mathbf{0})$$

where $\zeta_{\mathbf{z}} = (\zeta_{\mathbf{z},x}\ ,\ \zeta_{\mathbf{z},y}\ ,\ \zeta_{\mathbf{z},\theta})$.

$$= \int_0^1 \left(\iint_{P^*_{\zeta_{\mathbf{z}}(t)}} F(r\cos\eta, r\sin\eta)\ r\ dr\ d\eta \right) \cdot \zeta'_{\mathbf{z}}(t)\ dt + U_P(\mathbf{0})$$

where $P^*_{\zeta_{\mathbf{z}}(t)} = \{(r,\eta) \mid (r\cos\eta, r\sin\eta) \in P_{\zeta_{\mathbf{z}}(t)}\}$.

$$= \int_0^1 \left(\int_{P_{\zeta_{\mathbf{z}}(t)}} F\ dA \right) \cdot \zeta'_{\mathbf{z}}(t)\ dt + U_P(\mathbf{0})$$

$$= \int_0^1 F_P(\zeta_{\mathbf{z}}(t)) \cdot \zeta'_{\mathbf{z}}(t)\ dt + U_P(\mathbf{0})$$

by definition of the lifted force F_P.

$$= \int_{\zeta_{\mathbf{z}}} F_P \cdot d\mathbf{z} + U_P(\mathbf{0})$$

Hence U_P is indeed the potential of F_P. $\square$

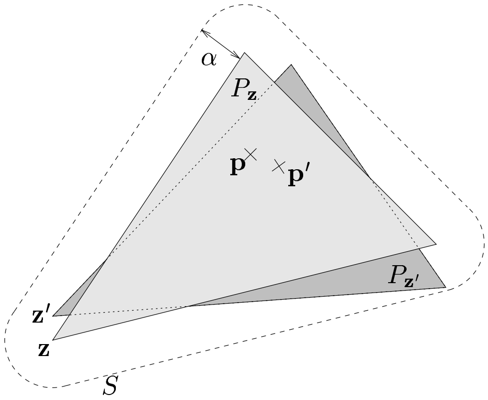

Figure 4.2: Two triangles $P_{\mathbf{z}}$ and $P_{\mathbf{z}'}$ with reference points $\mathbf{z}$ and $\mathbf{z}'$ whose symmetric difference is less than ϵ.

We believe that a shorter proof is possible by using differential forms for the case that both f and U are smooth. However, since the fields in consideration are usually not smooth (e.g. unit squeeze or radial fields) we give the longer general proof here.

Corollary 18 *Let f be a force field on $\mathbb{R}^2$ with potential U, and let P be a part of arbitrary shape. For the lifted generalized force F_P and the lifted potential U_P the following equality holds if U_P is differentiable: $\nabla U_P = -F_P$.*

Proof: Follows directly from Proposition 17. $\square$

So again, $U_P(x, y, \theta)$ can be interpreted as the potential energy of part P in configuration (x, y, θ). Therefore we obtain a lifted potential field U_P whose local minima are the stable equilibrium configurations in $\mathcal{C}$ for part P. Furthermore, potential fields are closed under addition and scaling. We can thus create and analyze more complex fields by looking at their components. In general, the theory of potential fields allows us to classify manipulation strategies with vector fields, offering new insights into equilibrium analysis and providing the means to determine strategies with stable equilibria. For example, it allows us to show that orientation equilibrium in a simple squeeze field is equivalent to the stability of a homogeneous boat floating in water, provided its density is $\rho = \frac{1}{2}\rho_{\text{water}}$ (for references on boat stability see e.g. [Gil56, pp. 42ff] or [New77, pp. 290ff]).

4.1 Properties of Lifted Force and Potential Fields

In this section we show that for a polygonal part P, the lifted force field is always continuous, and the lifted potential is always smooth.

Proposition 19 *Consider a polygon P at two configurations $\mathbf{z} = (x, y, \theta)$ and $\mathbf{z}' = (x', y', \theta')$, $\mathbf{z}, \mathbf{z}' \in \mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$. For all $\epsilon > 0$ there exists a $\delta > 0$ such that if $\mathbf{z}'$ lies within a δ-ball around $\mathbf{z}$, $\mathbf{z}' \in B_\delta(\mathbf{z})$, then $\mu(P_{\mathbf{z}} \triangle P_{\mathbf{z}'}) < \epsilon$ ($\mu(\cdot)$ denotes the size of an area, and $\triangle$ is the symmetric difference of two sets).*

Proof: First we will create a region S around $P_{\mathbf{z}}$ such that for any perturbed triangle $P_{\mathbf{z}'} \subset S$, the non-overlapping regions of $P_{\mathbf{z}}$ and $P_{\mathbf{z}'}$ are less than a given ϵ in size. Then we will show that there always exists a region $B_\delta(\mathbf{z})$ around $\mathbf{z}$ such that if $\mathbf{z}' \in B_\delta(\mathbf{z})$ then $P_{\mathbf{z}'}$ lies in S.

For now let us assume that $P_{\mathbf{z}}$ is a triangle in configuration $\mathbf{z}$, and let a be the length of its longest side. Consider the set $S = P_{\mathbf{z}} \oplus B_\alpha(\mathbf{0})$, for some $\alpha > 0$ (Figure 4.2). The area of $S - P_{\mathbf{z}}$ is $\mu(S - P_{\mathbf{z}}) < 3\alpha(a + 2\alpha)$. Let us choose $\alpha < \frac{1}{9}\min(1, \epsilon, \epsilon/a)$. Then if $\epsilon < 1$, $\mu(S - P_{\mathbf{z}}) < \frac{1}{3}(\epsilon + 2\epsilon^2) < \epsilon$. If $\epsilon \geq 1$, $\mu(S - P_{\mathbf{z}}) < \frac{1}{3}(\epsilon + 2) < \epsilon$. So in both cases, if $P_{\mathbf{z}'} \subset S$, then the area of the symmetric difference $P_{\mathbf{z}} \bigtriangleup P_{\mathbf{z}'}$ is at most ϵ.

We are interested in the distance between a point $\mathbf{p} \in P_{\mathbf{z}}$, and the corresponding perturbed point $\mathbf{p}' \in P_{\mathbf{z}'}$. We can describe the points $\mathbf{p}$ and $\mathbf{p}'$ as $\mathbf{p} = (x + r\cos(\phi+\theta), y + r\sin(\phi+\theta))$ and $\mathbf{p}' = (x' + r\cos(\phi+\theta'), y' + r\sin(\phi+\theta'))$, where r and ϕ are the length and the angle of a line from the reference point of P to the point $\mathbf{p}$. The distance between the x-coordinates is

$$\begin{aligned} |x_{\mathbf{p}} - x_{\mathbf{p}'}| &= |x + r\cos(\phi+\theta) - x' - r\cos(\phi+\theta')| \\ &\leq |x - x'| + |2r \sin\frac{2\phi+\theta+\theta'}{2} \sin\frac{\theta-\theta'}{2}| \\ &\leq |x - x'| + r|\theta - \theta'| \end{aligned}$$

If we choose $\delta = \frac{\alpha}{4\max(1,r)}$, we obtain

$$\begin{aligned} |x_{\mathbf{p}} - x_{\mathbf{p}'}| &\leq \alpha/4 + \alpha/4 \\ &= \alpha/2 \end{aligned}$$

Similarly, $|y_{\mathbf{p}} - y_{\mathbf{p}'}| < \alpha/2$, and hence $|\mathbf{p} - \mathbf{p}'| < \alpha$. We conclude that $\mathbf{p}' \in S$ whenever $\mathbf{z}' \in B_\delta(\mathbf{z})$. Hence we can always find a δ-ball around $\mathbf{z}$ such that the areas of $P_{\mathbf{z}}$ and $P_{\mathbf{z}'}$ differ by at most ϵ (by choosing $\delta < \frac{\alpha}{4\max(1,r)}$, i.e. $\delta < \frac{\min(1,\epsilon,\epsilon/a)}{36\max(1,r)}$).

This proof generalizes to arbitrary polygons (e.g. by using triangulations). □

Proposition 20 *Let P be a polygonal part in a force field f with potential U. The lifted force field $f_P(\mathbf{z})$, and the lifted generalized force field $F_P(\mathbf{z})$, with $\mathbf{z} = (x, y, \theta) \in \mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$, are continuous functions in x, y, and θ.*

Proof: For a given $\gamma > 0$, we want to determine an upper bound on the difference between $F(\mathbf{z})$ and $F(\mathbf{z}')$ for an arbitrary $\mathbf{z}' \in B_\gamma(\mathbf{z})$.

$$\begin{aligned} |F(\mathbf{z}) - F(\mathbf{z}')| &= \left| \int_{P_{\mathbf{z}}} f\, dA - \int_{P_{\mathbf{z}'}} f\, dA \right| \\ &\leq \left| \int_{P_{\mathbf{z}} - P_{\mathbf{z}'}} f\, dA \right| + \left| \int_{P_{\mathbf{z}'} - P_{\mathbf{z}}} f\, dA \right| \\ &\leq \widehat{f} \left(\int_{P_{\mathbf{z}} - P_{\mathbf{z}'}} dA + \int_{P_{\mathbf{z}'} - P_{\mathbf{z}}} dA \right) \\ &= \widehat{f}\, |P_{\mathbf{z}} \bigtriangleup P_{\mathbf{z}'}| \end{aligned}$$

where $\widehat{f} = \sup_S(|f(x,y)|)$ with $S = \{s \in P_{\mathbf{z}'} \mid \mathbf{z}' \in B_\gamma(\mathbf{z})\}$. This supremum exists whenever f is integrable, i.e. if f_P exists.

From Proposition 19 we know that we can make the area of $P_{\mathbf{z}} \bigtriangleup P_{\mathbf{z}'}$ arbitrarily small, by choosing an appropriate δ-ball around $\mathbf{z}$. In particular, we can force it to be less than

$1/\hat{f}$. Hence we can ensure that $|F(\mathbf{z}) - F(\mathbf{z}')| < \epsilon$ for any $\mathbf{z}' \in B_\delta(\mathbf{z})$, and any $\epsilon > 0$. This implies that F is continuous in $\mathbf{z} = (x, y, \theta)$.

An analogous argument holds for the lifted generalized force F_P. $\square$

Corollary 21 *For a polygonal part P, the lifted potential field $U_P(\mathbf{z}) = \int_{P_\mathbf{z}} U \; dA$ is C^1. Moreover, $\nabla U_P(\mathbf{z}) = -F_P(\mathbf{z})$, where F_P is the lifted generalized force acting on P.*

Proof: Because of Proposition 17, $U_P(\mathbf{z}) = \int_\alpha F_P \cdot d\mathbf{z} + c$ for some constant c. From Proposition 20 we know that the lifted generalized force F_P is continuous, hence the path integral of F_P must be C^1. $\nabla U_P(\mathbf{z}) = -F_P(\mathbf{z})$ because of Corollary 18. $\square$

4.2 Examples: Classification of Force Fields

Example: Radial fields. A *radial field* is a vector field whose forces are directed towards a specific center point. It can be used to center a part in the plane. The field in Figure 1.1b can be understood as a radial field with a rather coarse discretization using only four different force directions. Note that this field has a potential.

Definition 22 [**BDM96b**] *A unit radial field R is a two-dimensional force vector field such that $R(\mathbf{z}) = -\mathbf{z}/|\mathbf{z}|$ if $\mathbf{z} = 0$, and $R(\mathbf{0}) = \mathbf{0}$.*

Note that R has a discontinuity at the origin. A smooth radial field can be defined, for example, by $R'(\mathbf{z}) = -\mathbf{z}$.

Proposition 23 [**BDM96b**] *Given the radial fields R and R'. The corresponding potential fields are $U(\mathbf{z}) = ||\mathbf{z}||$, and $U'(\mathbf{z}) = \frac{1}{2}||\mathbf{z}||^2$, respectively.*

Note that U is continuous (but not smooth), while U' is smooth.

Counterexample: Skewed squeeze fields. Consider again the *skewed squeeze field* in Figure 3.1. Note that for example the integral on a cyclic path along the boundary of the disk is non-zero. This explains why the disk-shaped part has no equilibrium:

Proposition 24 *No skewed squeeze field has a potential.*

Counterexample: Parts with point contacts. Consider again the globally unstable S-shaped part P_S from Section 3 (Figure 3.2). At first glance, this example may seem counterintuitive. It can be shown (see Proposition 32) that there must exist a pose $\mathbf{z}_{\min}$ in which P_S achieves minimal potential, so why is P_S not stable in pose $\mathbf{z}_{\min}$? To better understand this problem, we investigate S-shaped parts with finite area contacts, and the transition as their contact areas are decreased towards 0.

Let us consider an S-shaped part with four square "feet." We choose the reference point at the COM, such that two of the feet are centered at $\pm(r_A, 0)$, and the other two feet are at $\pm(r_B \cos\phi, r_B \sin\phi)$, with ϕ constrained to $-\pi/2 < \phi < \pi/2$ (Figure 4.3). For the following discussion it will be sufficient to investigate the behavior in a squeeze field with its reference point fixed at $(0, 0)$.

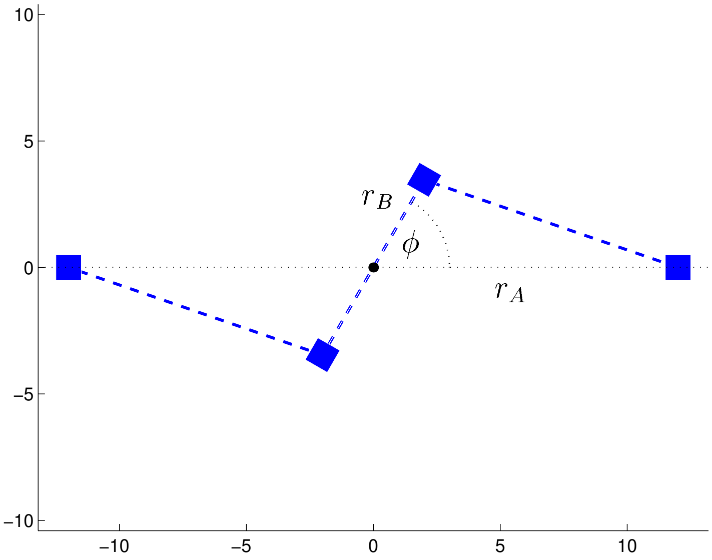

Figure 4.3: S-shaped part with four rigidly connected square "feet." $r_A = 12, r_B = 4, \phi = 60°$.

Figure 4.5a shows the moment function M_{P_S} and the potential U_{P_S} of an S-shaped part, where $r_A = 12$, $r_B = 4$, $\phi = 60°$, and the feet have area size 10. Notice that in poses with θ-angles corresponding to minima in the potential, the moment has a root with negative slope, which indicates a stable (orientational) equilibrium. Figures 4.5b and 4.5c show the (normalized) moments and potentials for parts with feet sizes 5 and 1, respectively. We observe that with decreasing contact areas, these functions become "less smooth," and the slope at the moment root increases. Figure 4.5d depicts moment and potential for a part with infinitesimally small feet. In this case, the moment function does not have a root at the minimum of the potential. Rather it exhibits a discontinuity at this orientation. This has the consequence that the part is not stable in this pose. In fact, for the moment function in Figure 4.5d there exist no roots with negative slope, hence there exists no stable equilibrium.

This observation can be made mathematically precise. The exact equations for the lifted potential and the moment of P_S are

$$
\begin{aligned}
U_{P_S} &= 2\ r_A\ |\cos\theta| + 2\ r_B\ |\cos(\theta+\phi)| && (4.6)\\
M_{P_S} &= 2\ r_A\ \mathrm{S}(\theta) + 2\ r_B\ \mathrm{S}(\theta+\phi) && (4.7)\\
\text{with } \mathrm{S}(\theta) &= \begin{cases} \sin\theta & \text{if } 0 \le \theta < \pi/2 \text{ or } 3/2\ \pi < \theta < 2\pi \\ -\sin\theta & \text{if } \pi/2 < \theta < 3/2\ \pi \\ 0 & \text{if } \theta = \pi/2 \text{ or } \theta = 3/2\ \pi \end{cases}
\end{aligned}
$$

The potential minimum is reached at $\theta = \pi/2$ or $\theta = 3/2\ \pi$. However, we see that for example $M_{P_S}(\pi/2) = -2\ r_B\ \mathrm{S}(\pi/2+\phi) = -2\ r_B\ \cos\phi \neq 0$. Furthermore, $M_{P_S}(\pi/2-) > 0$, and $M(\pi/2+) < 0$. This implies that the part P_S will oscillate about $\theta = \pi/2$. Under first-order dynamics, this oscillation will be infinitesimally small, because any infinitesimal

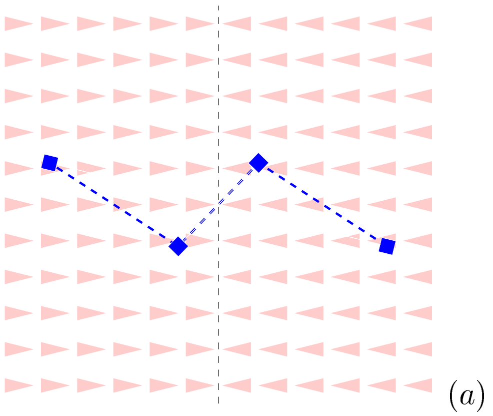

Figure 4.4: Total equilibria of an S-shaped part with area contacts in a squeeze field. (a) Maximum potential, $\mathbf{z}_{\max} = (0, 0, \theta_{\max})$ such that $r_A \sin\theta_{\max} = -r_B \sin(\theta_{\max} + \phi)$; $\theta_{\max} \approx -0.24$. (b) Minimum potential, $\mathbf{z}_{\min} = (0, 0, \theta_{\min})$; $\theta_{\min} \approx \pi/2$.

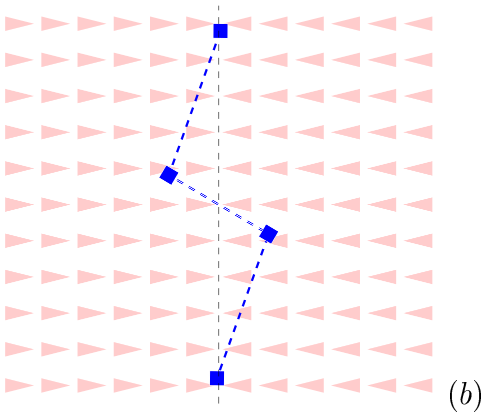

angular deflection of P_S results in a restoring moment with opposed orientation. Under second-order dynamics, the part may have a finite oscillation amplitude because of the inertia of the part. However, damping will reduce this amplitude over time.

We conclude that parts with point contacts can exhibit pathological behavior even in very simple and otherwise well-behaved potential fields: This example shows that, for such parts, it is possible that the generalized force is not zero in a pose that minimizes the potential of the part.

This pathology can not occur when only parts with finite area contact are allowed: From Corollary 21 we know that the (lifted) potential of a part with area contact is C^1, hence its gradient exists everywhere. In particular, the gradient is zero at the minimum of the potential. This means that in a pose with minimum potential the generalized force must be zero. Let us summarize these results:

Corollary 25 *Let P be a part with finite area contact in a force vector field f with potential U. In a configuration $\mathbf{z}_0$ that corresponds to a local extremum of the lifted potential U_P, the lifted generalized force $F_P(\mathbf{z}_0)$ is zero.*

In other words, for a first-order dynamical system and a part P with finite area contact, a local minimum (maximum) of U_P corresponds to a stable (unstable) equilibrium of P in f.

Example: Morphing and combining vector fields. Our strategies from Section 2 have *switch points* in time where the vector field changes discontinuously (Figure 2.1). This is because after one squeeze, for every part, the orientation equilibria form a finite set of possible configurations, but in general there exists no unique equilibrium (as shown in Section 2.3). Hence subsequent squeezes are needed to disambiguate the part orientation. Therefore these switches are necessary for strategies with squeeze patterns.

One may ask whether, using another class of potential field strategies, *unique* equilibria may be obtained without discrete switching. We believe that *continuously varying* vector fields of the form $(1-t)f + t\,g$, where $t \in [0, 1]$ represents time, and f and g are squeezes, may lead to vector fields that have this property. Here "+" denotes point-wise addition

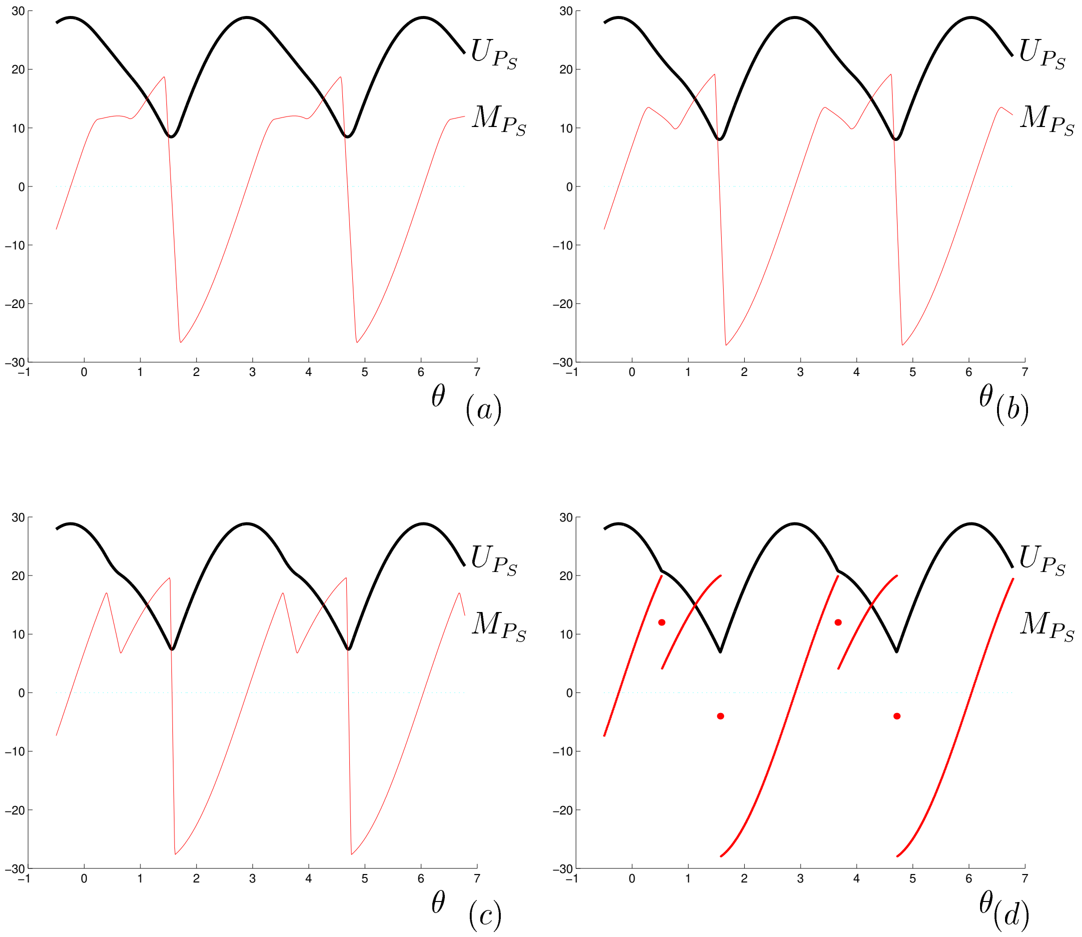

Figure 4.5: Moment function M_{P_S} (thin line) and potential U_{P_S} (thick line) for S-shaped parts: (a) feet have contacts of area size 10; (b) size = 5; (c) size = 1; (d) point contacts. Note how a discontinuity is created in the moment function when the contact area is decreased towards 0.

of vector fields, and we will write "$f \rightsquigarrow g$" for the resulting continuously varying field. By restricting f and g to be fields with potentials U and V, we know that $U + V$ and $(1 - t)U + tV$ are potential fields, and hence we can guarantee that $f + g$ and $f \rightsquigarrow g$ are well-behaved strategies. These form the basis of our new algorithms in Section 5.

Let us formalize the previous paragraphs. If f is a vector field (in this case a squeeze pattern) that is applied to move part P, we define the *equilibrium set* $E_P(f)$ as the subset of the configuration space $\mathcal{C}$ for which P is in equilibrium. Let us write $f * g$ for a strategy that first applies vector field f, and then vector field g to move part P. $f + g$ can be understood as applying f and g simultaneously. We have shown that in general $E_P(f)$ is not finite, but for two *orthogonal* squeezes f and g, the discrete switching strategy $f * g$ yields a finite equilibrium set $E_P(f * g)$ (see Section 5.2, Claim 39). Furthermore, for some parts the equilibrium is unique up to symmetry.

We wish to explore the relationship between equilibria in simple vector fields $E_P(f)$ or $E_P(g)$, combined fields $E_P(f + g)$, discretely-switched fields $E_P(f * g)$, and continuously varying fields $E_P(f \rightsquigarrow g)$. For example, one may ask whether there exists a strategy with combined vector fields, or continuously varying fields, that, in just one step, reaches the same equilibrium as a discretely switched strategy requiring multiple steps. Finally, let $f_1 * f_2 * \cdots * f_s$ be a sequence of squeeze fields guaranteed to uniquely orient a part P under assumption 2PHASE. We wish to investigate how continuously varying strategies such as $f_1 \rightsquigarrow f_2 \rightsquigarrow \cdots \rightsquigarrow f_s$ can be employed to dynamically achieve the same equilibria even when 2PHASE is relaxed. The distributed actuation strategy $F * G$ is distributed in space, but not in time. The strategy $F + G$ is parallel with respect to space and time, since F and G are simultaneously "run." Research in this area could lead to a *theory of parallel distributed manipulation* that describes *spatially distributed* manipulation tasks that can be *parallelized over time and space* by superposition of controls.

4.3 Upward-Shaped Potential Fields

So far we have presented specific force fields that *always* (e.g. squeeze and radial fields) or *never* (e.g. skewed squeeze fields) induce stable equilibria on certain classes of parts. We conclude this section with a criterion that provides a sufficient condition on force fields such that *all parts of a certain size reach a stable equilibrium.*

We have observed in Section 3 that *a priori* it is not obvious when a force field induces stable equilibria. Our Equilibrium Criterion will be based on two important properties:

1. The field has a potential. Potential fields do not allow closed paths (technically, limit cycles) along which the work is positive, which could induce infinite motion of a part.
2. The force field is "inward-directed," which implies that (assuming first-order dynamics) parts can never leave a certain region R. This useful property is a direct consequence of the definition of inward-directedness. An inward-directed force field corresponds to an "upward-shaped" potential, in which all paths that leave region R have an ascending slope.

We will require Property (1.) to hold for the entire force field, while Property (2.) devolves to a boundary condition.

4.3.1 Elementary Definitions

Definition 26 *Let* $\mathbf{z} \in \mathbb{R}^n$*. The* ϵ-ball around $\mathbf{z}$*, denoted* $B_\epsilon(\mathbf{z})$*, is the set* $\{\mathbf{r} \in \mathbb{R}^n \mid$

$|\mathbf{r} - \mathbf{z}| < \epsilon\}$ *of all points within a distance* ϵ *of* $\mathbf{z}$.

Definition 27 (Lozano-Pérez [LP83]) *Let* A, B *be sets in* $\mathbb{R}^n$. *The* Minkowski sum $A \oplus B$ *of two sets* A *and* B *is defined as the set* $\{\mathbf{a} + \mathbf{b} \mid \mathbf{a} \in A, \mathbf{b} \in B\}$.

From these definitions it follows that for a region R with boundary ∂R, the set $\partial R \oplus B_d(\mathbf{0}) = \{\mathbf{r} + \mathbf{z} \mid \mathbf{r} \in \partial R, \text{ and } |\mathbf{z}| \leq d\}$ comprises all points that are within a distance d from the boundary of R.

Definition 28 *Given a region* $R \subset \mathbb{R}^n$, *define the set* $CI(R, d) = R - (\partial R \oplus B_d(\mathbf{0}))$ *which is the region* R *shrunk by distance* d. *Note that* $CI(R, d)$ *is based upon the* configuration space interior [LP83] *of* R *for* $B_d(\mathbf{0})$. *Abusing terminology slightly, we call* $CI(R, d)$ *the* configuration space interior of R *in this thesis.*

Definition 29 *The* radius r_P of a part P *is the maximum distance between an arbitrary point of* P *and the center of mass (COM) of* P.

4.3.2 Equilibrium Criterion

We are now able to state a general criterion for a force field f to induce stable equilibria on all parts in a region S. As mentioned at the beginning of Section 4.3, this criterion is based on two main conditions: (1) if f has a potential, limit cycles with positive energy gain are avoided inside S. (2) if f is "inward-directed" (see the definition below), parts cannot leave the region S.

In the following we give a general definition of inward-directed vector fields on a manifold Z. We then specialize the definition to the special instances of $Z = \mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$ (the configuration space) and $Z = \mathbb{R}^2$, and give a sufficient, practical condition for inward-directed vector fields. We conclude with the presentation of the Equilibrium Criterion.

Definition 30 (Inward-Directed Force Fields) [1] *Let* Z *be an arbitrary smooth manifold, and let* $Y \subset Z$ *be a compact and smooth submanifold with boundary of* Z. *Assume that* ∂Y *has codimension 1 in* Z, *and that the boundary of* Y *is orientable. Let* $q \in \partial Y$ *be a point on the boundary of* Y, *and* $V_q \in T_q Z$ *be a tangent vector to* Z *at* q.

We say V_q is inward-directed to ∂Y at q *if there exists a sufficiently small* $\epsilon > 0$ *such that* $q + \epsilon V_q \in Y$.

Let V *be a vector field on* Z. *We say* V is inward-directed to ∂Y *if* $V(q)$ *is inward-directed to* ∂Y *at* q *for all* $q \in \partial Y$.

Assume the set $S \subset \mathbb{R}^2$ is compact and smooth. Consider the part P when it is placed into the force field f such that its COM lies in S. The set of all such poses is a subset of the configuration space $\mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$ which we call $\widetilde{S} = S \times \mathbb{S}^1$. The boundary of $\widetilde{S}$ is $\partial\widetilde{S} = \partial S \times \mathbb{S}^1$. Note that $\partial\widetilde{S}$ separates the interior $i\widetilde{S} = \widetilde{S} - \partial\widetilde{S}$ from the exterior $\mathcal{C} - \widetilde{S} = (\mathbb{R}^2 - S) \times \mathbb{S}^1$, and that $\partial\widetilde{S}$ is isomorphic to a torus $\mathbb{S}^1 \times \mathbb{S}^1$.

Now let $z = (x, y, \theta) \in \partial\widetilde{S}$, and let $F_z \in T_z\mathcal{C}$ represent the lifted generalized force acting on part P in pose z. F_z is *inward-directed* (with respect to $\partial\widetilde{S}$) if F_z points into the interior of $\widetilde{S}$. Note that this condition is equivalent to saying that the projection of F_z onto the tangent space at (x, y) to $\mathbb{R}^2$ points into S, because the rotational component of

---

[1] In this definition, for convenience we assume that Z is embedded in $\mathbb{R}^m$ for some m. This condition may be relaxed.

F_z is tangential to $\partial\widetilde{S}$. So for example, if $z = (x, y, \theta) \in \partial\widetilde{S}$, then $z' = (x, y, \theta') \in \partial\widetilde{S}$ for any θ'.

The following proposition gives a simple condition on a force field f that tells us if, for a given part P, its lifted generalized force field F_P is inward-directed:

Proposition 31 *Let P be a part with radius r whose COM is the reference point used to define its configuration space $\mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$. Let f be a force vector field defined on a region $R \subset \mathbb{R}^2$, with F_P the corresponding lifted generalized force field. Let $S \subset \mathbb{R}^2$ be a convex, compact, and smooth subset of the configuration space interior of R, and $S \subset CI(R, r)$.*

Consider a point $q \in \partial S$ with outward normal n_q, and a ball $B_r(q)$ with radius r about q. If for every point $q \in \partial S$, and for every point s in the corresponding ball $B_r(q)$, the dot product $g(s) = f(s) \cdot n_q$ is less than 0, then the lifted generalized force field F_P is inward-directed to $\partial\widetilde{S}$ (note: $(\,\cdot\,)$ is the standard inner product).

Proof: Consider the part P in pose $z = (x, y, \theta) \in \partial\widetilde{S}$ such that $q = (x, y)$. P has radius r, hence it lies completely inside the ball $B_r(q)$, independent of its orientation θ. As we know that $g(p) = f(p) \cdot n_q < 0$ for all $p \in B_r(q)$, we can conclude that the integral of $g(p)$ over P is also less than 0: $\int_P g(p)\, dA = \int_P f(p) \cdot n_q\, dA = f_P \cdot n_q < 0$. This implies that for f_P, which is the translational component of F_P (see Definition 16), the vector $q + \epsilon\, f_P(z)$ lies inside S, if ϵ is positive and sufficiently small. As mentioned on page 44 above, this suffices to ensure that the vector $z + \epsilon\, F_P(z)$ lies inside $\widetilde{S}$. □

Lemma 32 (Equilibrium Criterion) *Let P be a polygonal part with radius r, let f be a force field with potential U defined on a region $R \subset \mathbb{R}^2$, and let $S \subset R$ as specified in Proposition 31. Let us also assume that the motion of part P is governed by first-order dynamics.*

If the lifted force vector field F_P is inward-directed to $\partial\widetilde{S}$, then the part P will reach a stable equilibrium under f in $i\widetilde{S}$ whenever its COM is initially placed in S.

Proof: Assume that the COM of part P is placed at a point $(x, y) \in S$. This means that P is in some pose $z = (x, y, \theta) \in \widetilde{S}$. We now show that the COM of P cannot leave S when initially placed inside S. We know that $\partial\widetilde{S}$ separates $i\widetilde{S}$ from $\mathcal{C} - \widetilde{S}$. Hence every path from z to some $z^* \in \mathcal{C} - \widetilde{S}$ must intersect $\partial\widetilde{S}$ at some point $z' \in \partial\widetilde{S}$. Now consider part P in pose z'. Under first-order dynamics, its velocity must be in direction of $F_P(z')$. Because F_P is inward-directed, the velocity of P must be towards $i\widetilde{S}$. In particular, this means that the COM will move into iS, hence P cannot leave S, and that there is no equilibrium on ∂S.

f, and hence (because of Proposition 17) F_P have potential U and U_P, respectively. Therefore limit cycles with energy gain are not possible. Furthermore, $U_P(\widetilde{S})$ is the continuous image of a compact set, $\widetilde{S}$. Therefore the image $U_P(\widetilde{S})$ is a compact subset of $\mathbb{R}$, hence has a minimum value attained by some point $s \in \widetilde{S}$. Since f is inward-directed, s must lie in $i\widetilde{S}$. This minimum is a stable equilibrium of P in f, as shown in Corollary 25. □

Because of Lemma 32, the use of potential fields is invaluable for the analysis of effective and efficient manipulation strategies, as discussed in the following section. In particular, it is useful for proving the completeness of a manipulation planner.

Chapter 5

New and Improved Manipulation Algorithms

The part alignment strategies in Section 2.3 have *switch points* in time where the vector field changes discontinuously (Figure 2.8). We can denote such a *switched strategy* by $f_1 * f_2 * \cdots * f_s$, where the f_i are vector fields. In Section 2.3 we showed that a general squeeze strategy to align a (non-convex) polygonal part with n vertices may need up to $O(k\,n^2)$ switches, and require $O(k^2\,n^4)$ time in planning (k is the maximum number of polygon edges that a bisector can cross). To improve these bounds, we now consider a broader class of vector fields including simple squeeze patterns, radial, and combined fields as described in Section 4.

In Section 5.1 we show how, by using radial and combined vector fields, we can significantly reduce the complexity of the strategies from that of Section 2. In Section 5.2 we describe a general planning algorithm that works with a limited "grammar" of vector fields (and yields, correspondingly, less favorable complexity bounds).

5.1 Radial Strategies

Consider a part P in a force field f. Some force fields exhibit rotational symmetry properties that can be used to generate efficient manipulation strategies:

Property 33 *There exists a unique* pivot point *v of P such that P is in translational equilibrium if and only if v coincides with $\mathbf{0}$.*

Property 34 *There exists a unique* pivot point *v of P such that P is in (neutrally stable) orientational equilibrium if and only if v coincides with $\mathbf{0}$.*

We typically think of the pivot point v being a point of P; however, in generality, just like the center of mass of P, v does not need to lie within P, but instead is some fixed point relative to the reference frame of P. Now consider the part P in an ideal unit radial force vector field R as described in Section 4.

Proposition 35 *In a unit radial field R, Properties 33 and 34 hold.*

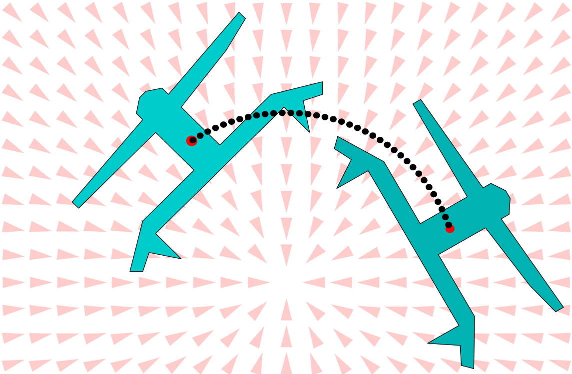

Figure 5.1: Rotating a part about the center of a unit radial field. The force and torque on the part remain constant with respect to its reference frame.

Proof: We fix the part P at an arbitrary orientation θ, and show that at this orientation P has a unique translational equilibrium $v(\theta)$. That is, placing $v(\theta)$ at the origin is necessary and sufficient for P to be in translational equilibrium at orientation θ. Second, we show that for any two distinct orientations θ and θ', $v(\theta) = v(\theta')$. We call this unique point v, dropping the orientation θ. Finally, we argue that whenever P is in translational equilibrium (i.e., v is at the origin), that P is neutrally stable with respect to orientation. This follows by the radial symmetry of R.

Consider the translational forces (but not the moments) acting on P in the radial field R. To do this, let us separate R into its x and y components, R_x and R_y, such that $R = (R_x, R_y)$. Assume for now that the orientation of P is fixed. If P is placed at a position $z_0 \in \mathbb{R}^2$, whose x-coordinate is sufficiently negative, the total force induced by R_x on P will point in the positive x direction. Symmetrically, placing P at a sufficiently large positive x coordinate will cause a force in the negative x direction. We claim that, by translating P rigidly with increasing x coordinate, this force decreases continuously and strictly monotonically, and hence has a unique root.

To verify this claim, consider a small area patch $\wp_0$ of P. A uniform translation t of $\wp_0$ in x direction can be described as $\wp(t) = \wp_0 \oplus (z_0 + t\hat{x})$ (with z_0 the initial position of the patch, $\hat{x}$ the unit vector in x direction, and $\oplus$ the Minkowski sum). The total force on $\wp(t)$ in x direction is $\int_{\wp(t)} R_x \, dA$. This force decreases continuously and strictly monotonically with t, because R_x is strictly monotone and continuous everywhere except on the x-axis, which has measure zero in $\mathbb{R}^2$. A similar argument applies for the y direction, and, because of the radial symmetry of R, for any direction.

If we choose the set S as a sufficiently large disk-shaped region around the origin and recall that R has a potential, we can apply Lemma 32 to conclude that there must exist at least one total equilibrium for P. Now assume that there exist two distinct equilibria $e_1 = (x_1, y_1, \theta_1)$ and $e_2 = (x_2, y_2, \theta_2)$ for P in R. We write "$P(e_i)$" to denote that P is in configuration e_i. Because of the radial symmetry of R, we can reorient $P(e_2)$ to $P(e_2')$ such that its orientation is equal to $P(e_1)$: $e_2' = (x_2', y_2', \theta_1)$, where $\binom{x_2'}{y_2'} = M \binom{x_2}{y_2}$, and M is a rotation matrix with angle $\theta_1 - \theta_2$ (Figure 5.1). This reorientation does not affect the equilibrium. Note that P can be moved from e_1 to e_2' by a pure translation. From above we know that such a translation of P corresponds to a strictly monotone change in the translational forces acting on P. Hence we conclude that $P(e_1)$ and $P(e_2')$ cannot both be in translational equilibrium unless e_1 and e_2' are equal. This implies that e_1 and e_2 cannot both be equilibria of P in R unless they both have the same pivot point v. $\square$

Surprisingly, v need not be the center of area of P. For example, consider again the part in Figure 2.5, which consists of a large and a small square connected by a long rod of negligible width. The pivot point of this part will lie inside the larger square. But if the rod is long enough, the center of area will lie outside of the larger square. However, the following corollary holds:

Corollary 36 *For a part P in a continuous radial force field R' given by $R'(\mathbf{z}) = -\mathbf{z}$, the pivot point of P coincides with the center of area of P.*

Proof: The force acting on P in R' is given by $F = \int_P -\,\mathbf{z}\; dA$, which is also the formula for the (negated) center of area. $\square$

Now suppose that R is combined with a unit squeeze pattern S, which is scaled by a factor $\delta > 0$, resulting in $R + \delta S$. The squeeze component δS of this field will cause the part to align with the squeeze, similarly to the strategies in Section 2.3. But note that the radial component R keeps the part centered in the force field. Hence, by keeping R sufficiently large (or δ small), we can assume that the pivot point of P remains within an ϵ-ball of the center of R. This implies that assumption 2Phase (see Section 2.3) is no longer necessary. Moreover, ϵ can be made arbitrarily small by an appropriate choice of δ.

Proposition 37 *Let P be a polygonal part with n vertices, and let k be the maximum number of edges that a bisector of P can cross. Let us assume that v, the pivot point of P, is in general position. There are at most $O(k\,n)$ stable equilibria in a field of the form $R + \delta S$ if δ is sufficiently small and positive.*

Proof: For a part in equilibrium in a pure radial field R (i.e., with $\delta = 0$), the pivot point v is essentially fixed at the origin. This is implied by Property 33. It is easy to see that Property 33 is not true in general for arbitrary fields of the form $R + \delta S$. Property 33 holds if $\delta = 0$, because then any orientation is an equilibrium when v is at the center of R. However, Property 34 does not hold if $\delta > 0$, because in general there does not exist a unique pivot point in squeeze fields (see Section 2.3).

We will conduct the combinatorial analysis of the orientation equilibria under the assumptions that (i) $\delta > 0$ and (ii) that v is fixed at the origin. Then we will relax the latter assumption (ii), and show that Property 33 holds, *approximately,* even in $R + \delta S$, for a sufficiently small $\delta > 0$. That is, we show that a sufficiently small δ can be chosen so that the combinatorial analysis is unaffected when assumption (ii) is relaxed.

First, we show that when δ is small but positive, and with v fixed at the center of R, there are only a linear number of orientation equilibria. (I.e., we constrain the pivot point v to remain fixed at the origin until further notice.) So let us assume that we are in a combined radial and small squeeze field $R + \delta S$.

Consider a ray $w(0)$ emanating from v. Assume w.l.o.g. that v is not a vertex of P, and that $w(0)$ intersects the edges $S(0) = \{e_1, \cdots, e_k\}$ of P in general position, $1 \le k \le n$. Parameterize the ray $w(\cdot)$ by its angle ϕ to obtain $w(\phi)$. As ϕ sweeps from 0 to 2π, each edge of P will enter and leave the *crossing structure* $S(\phi)$ exactly once. $S(\phi)$ is updated at *critical angles* where $w(\phi)$ intersects a vertex of P. Since there are n vertices, there are $O(n)$ critical angles, and hence $O(n)$ changes to $S(\phi)$ overall. Hence, since between critical angles $S(\phi)$ is constant, we see that $S(\phi)$ takes on $O(n)$ distinct values. Now place the squeeze line l to coincide with $w(\phi)$. For a given crossing structure $S(\phi) \cup S(\phi+\pi)$, satisfying conditions **I** and **II** as defined in Section 2.3 devolves to solving two equations. The first equation provides the condition for translational equilibrium, while the second equation implements the condition for orientational equilibrium. The latter equation is called the *moment function* $M(\phi)$, because it describes the moment acting on P as a function of ϕ. (But note that M is different from the moment function defined in Section 2.3, because here the part rotates about a fixed pivot point.) In analogy to Section 2.3 it can be shown that these equations are algebraic and of degree k, where k is the maximum number of edges intersected by the squeeze line as described in Section 2.2. This implies that between any two adjacent critical values there are only $O(k)$ orientations of l (given by $w(\phi)$) that satisfy conditions **I** and **II**. Hence, the overall number of orientations satisfying **I** and **II** is $O(k\,n)$.

If $\delta > 0$ the part P will be perturbed, so that Property 33 is only approximately satisfied. (That is, we now relax the assumption that v is constrained to be at the origin). However, we can ensure that v lies within an ϵ-ball around the origin (the center of the radial field). To see this, first consider P at some arbitrary configuration $\mathbf{z}$ in the squeeze field δS. The total squeeze force on $P_{\mathbf{z}}$ is given by the area integral $\delta S_P(\mathbf{z}) = \int_{P_{\mathbf{z}}} \delta S \; dA$. (Recall that S_P denotes the lifted force field of S; see Definition 16, Equation (4.2).) Now, δS_P is bounded above by $|\delta S_P| \leq \delta A$, where A is the area of P (note that S is a unit squeeze field).

P is in equilibrium with respect to the radial field R if v is at the origin. Now consider the lifted force R_P when the pivot point of P is not at the origin. More specifically, let $v_{\mathbf{z}}$ be the pivot point of $P_{\mathbf{z}}$, and let us define a set $\mathcal{R}_P(d) = \{|R_P(\mathbf{z})| \text{ such that } |v_{\mathbf{z}}| = d\}$. We also define a function $\widehat{R}_P(d) = \min\{\mathcal{R}_P(d)\}$. This function is well-defined, because $\mathcal{R}_P(d)$ is the continuous image of a compact set, hence the minimum exists. $\widehat{R}_P(d)$ is the minimum magnitude of the lifted force acting on $P_{\mathbf{z}}$ when its pivot point $v_{\mathbf{z}}$ is at distance d from the origin.

By decomposing R_P into its x- and y-components, we can write $|R_P|$ as $\sqrt{R_{P,x}^2 + R_{P,y}^2}$. Because of the radial symmetry of R let us assume w.l.o.g. that $v_{\mathbf{z}} = (d, 0)$. From the proof of Proposition 35 we know that, for any given orientation of $P_{\mathbf{z}}$, the magnitude of $R_{P,x}$ increases continuously and strictly monotonically with increasing $d \geq 0$. Furthermore, $R_{P,y}$ is continuous in d, and $R_{P,y}(\mathbf{0}) = 0$, so $R_{P,y}^2$ is continuous and monotonically increasing for all d less than some sufficiently small $d_0 > 0$. Hence for any fixed orientation of $P_{\mathbf{z}}$, R_P is a continuous and strictly monotonically increasing function for all $d \in [0, d_0]$.

Now suppose that $\widehat{R}_P(d)$ is not strictly monotone, i.e. that there exist d_1, d_2 with $0 \leq d_1 < d_2$, but $\widehat{R}_P(d_1) \geq \widehat{R}_P(d_2)$. Then there must exist $\mathbf{z}_1, \mathbf{z}_2$ with $|v_{\mathbf{z}_1}| = d_1$ and $|v_{\mathbf{z}_2}| = d_2$, and $|R_P(\mathbf{z}_1)| = \widehat{R}_P(d_1) \geq \widehat{R}_P(d_2) = |R_P(\mathbf{z}_2)|$. Let us define $\mathbf{z}'_2$ such that $\mathbf{z}'_{2,\theta} = \mathbf{z}_{2,\theta}$ and $v_{\mathbf{z}'_2} = c\,v_{\mathbf{z}_2}$ for some $c \in \mathbb{R}$, i.e., $v_{\mathbf{z}_2}$ and $v_{\mathbf{z}'_2}$ lie on a line through the origin. If we choose $0 \leq c < 1$ then $|R_P(\mathbf{z}'_2)| < |R_P(\mathbf{z}_2)|$ because $|R_P|$ is monotone, as shown in the previous paragraph. In particular, if we choose $c = d_1/d_2$, then $|R_P(\mathbf{z}'_2)| < |R_P(\mathbf{z}_2)| \leq |R_P(\mathbf{z}_1)| = \widehat{R}_P(d_1)$, and $|v_{\mathbf{z}'_2}| = |v_{\mathbf{z}_1}|$. This is a contradiction to the definition of $\widehat{R}_P(d_1) = \min\{\mathcal{R}_P(d)\}$. We conclude that $\widehat{R}_P$ is continuous and strictly monotone for sufficiently small $d \geq 0$.

Now consider $P_{\mathbf{z}}$ in the combined field $R + \delta S$, and again let d denote the distance between pivot point $v_{\mathbf{z}}$ and the origin. In equilibrium the lifted forces $R_{P_{\mathbf{z}}}$ and $\delta S_{P_{\mathbf{z}}}$

balance out, hence $\widehat{R}_P(d) \leq |R_{P_{\mathbf{z}}}| = |\delta S_{P_{\mathbf{z}}}| \leq \delta A$, with A the area of P. Since $\widehat{R}_P$ is continuous and strictly monotone in d for sufficiently small d, we can ensure that d is less than a given ϵ, by chosing an appropriately small δ. This implies that $v_{\mathbf{z}}$ must lie within an ϵ-ball of the center of the radial field. In particular, we can make this ϵ-ball small enough so that the crossing structure $S(\phi)$ is not affected.

Finally we have to ensure that the stable equilibria, as predicted by the moment function M, are approximated arbitrarily closely. This means that the disturbance in the moment function, caused by pivot point $v_{\mathbf{z}}$ not exactly coinciding with the center of the radial field, can be made arbitrarily small. To see this, first consider the original (unperturbed) moment function M which describes the moment acting on the part P if its pivot point coincides with the origin. In this case, the moment is caused solely by the squeeze field δS, while P is in equilibrium with respect to the radial field R.

Now consider the disturbance in M if the pivot point $v_{\mathbf{z}}$ is not exactly at the origin, but somewhere in an ϵ-ball around it. Let us call this disturbance ΔM, and note that ΔM has two components: $\Delta M_{\delta S}$, which is the change in moment caused by the squeeze field δS, and ΔM_R, the disturbance caused by the radial field R.

For a part P at a given orientation, any ϵ-displacement of $v_{\mathbf{z}}$ can change the force in a squeeze field δS by at most $|\Delta F_{\delta S}| \leq \delta\, \epsilon\, d_P$, where d_P is the maximum diameter of P. Hence $\Delta M_{\delta S} = r \times \Delta F_{\delta S}$ is proportional to the product of the disturbance in location ϵ and the magnitude of the squeeze field δ (r is the fixed distance between pivot point and COM of P), i.e. $|\Delta M_{\delta S}| = O(\delta\, \epsilon)$.

Since the force caused by the radial field R balances the force generated by the squeeze field δS, we obtain the same bounds for ΔM_R. We see that $\Delta M = \Delta M_{\delta S} + \Delta M_R = O(\delta\, \epsilon)$. Recall that ϵ decreases strictly monotonically with δ, hence ΔM decreases asymptotically faster than δ. This ensures that we can find a sufficiently small δ such that the moment function M is approximated arbitrarily closely, and the equilibria of the squeeze field δS are not affected.

We conclude that the number of equilibria in a field $R + \delta S$ is bounded by $O(k\, n)$, for sufficiently small δ. □

In analogy to Section 2.3 we define the turn function $t : \mathbb{S}^1 \to \mathbb{S}^1$, which describes how the part will turn under a squeeze pattern, and hence yields the stable equilibrium configurations. Given the turn function t we can construct the corresponding squeeze function s as described in Section 2.3. With s as the input for Goldberg's alignment planner [Gol93], we obtain strategies for unique part alignment (and positioning) of length $O(k\, n)$. They can be computed in time $O(k^2\, n^2)$.

The result is a strategy for parts positioning of the form $(R + \delta S_1) * \cdots * (R + \delta S_{O(kn)})$. Compared to the general squeeze algorithm in Section 2.3, it improves the plan length by a factor of n, and the planning complexity is reduced by a factor of n^2. The planner is complete: For any polygonal part, there exists a strategy of the form $*_i(R + \delta S_i)$. Moreover, the algorithm is guaranteed to find a strategy for any input part. By appending a step which is merely the radial field R without a squeeze component, we are guaranteed that the part P will be uniquely posed (v is at the origin) as well as uniquely oriented. We can also show that the continuously varying "morphing" strategy $(R + \delta S_1) \rightsquigarrow \cdots \rightsquigarrow (R + \delta S_{O(kn)}) \rightsquigarrow R$ works in the same fashion to achieve the same unique equilibrium.

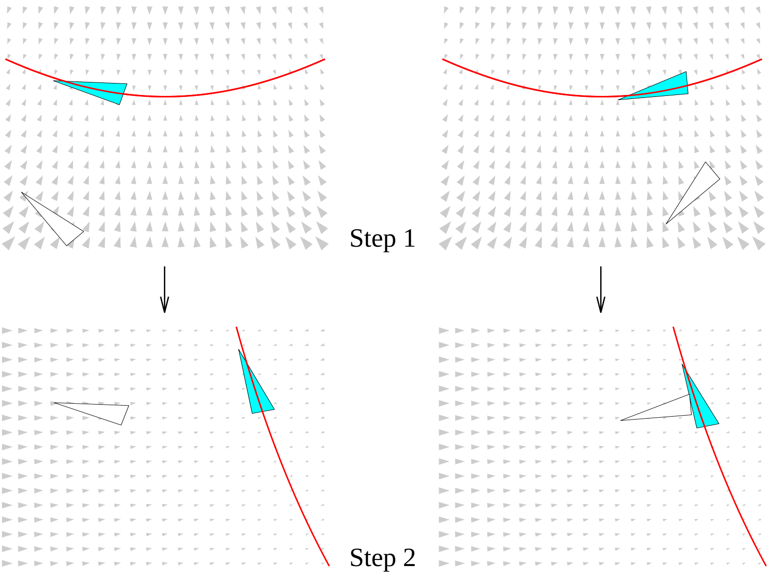

Figure 5.2: Manipulation vocabulary for a triangular part on a vibrating plate, consisting of two consecutive force fields with slightly curved nodal lines (attractors) which bring the part into (approximately) the same equilibria.

5.2 Manipulation Grammars

The development of devices that generate programmable vector fields is still in its infancy. The existing prototype devices exhibit only a limited range of programmability. For example, the prototype MEMS arrays described in Sections 6.1 and 6.2 currently have actuators in only four different directions, and the actuators are only row-wise controllable. Arrays with individually addressable actuators at various orientations are possible (see [BDMM94b,BDMM94a,LW95,BDM96a,SGD$^+$96]) but require significant development effort. There are also limitations on the resolution of the devices given by fabrication constraints. For the vibrating plate device from Section 6.3 the fields are even more constrained by the vibrational modes of the plate.

We are interested in the capabilities of such constrained systems. In this section we give an algorithm that decides whether a part can be uniquely positioned using a given set of vector fields, and it synthesizes an optimal-length strategy if one exists. Furthermore, in Chapter 5.2, the vector fields we consider may be arbitrary, and in particular can vary in magnitude (as opposed to unit squeeze fields). If we think of these vector fields as a vocabulary, we obtain a language of manipulation strategies. We are interested in those expressions in the language that correspond to a strategy for uniquely posing the part.

5.2.1 Finite Field Operators

We define two basic operations on vector fields. Consider two vector fields f and g. $f + g$ denotes point-wise addition, $f * g$ denotes sequential execution of f, and then g.

Definition 38 *Let P be an arbitrary planar part. A* finite field operator *is a sequence of vector fields that brings P from an arbitrary initial pose into a* finite set *of equilibrium poses.*

A field operator comes with the following guarantee: No matter where in $\mathbb{R}^2 \times \mathbb{S}^1$ the part starts off, it will always come to rest in one of E different total equilibria (Figure 5.2). That is: for any polygonal part P, either of these field operators is *always* guaranteed to reduce P to a *finite* set of equilibria in its configuration space $\mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$.

From Section 5.1 we know that combined radial-squeeze patterns $R + \delta S$ have this property. However, there are other simple field operators that also have this finiteness property:

Claim 39 [**BDM96b**] *Let f and $f_\perp$ be unit squeeze fields such that $f_\perp$ is orthogonal to f. Then the fields $f * f_\perp$ and $f + f_\perp$ induce a finite number of equilibria on every connected polygon P, hence $f * f_\perp$ and $f + f_\perp$ are finite field operators.*

Proof: First consider the field $f * f_\perp$, and w.l.o.g. assume that $f(x, y) = (-\mathrm{sign}(x), 0)$. Also assume that the COM of P is the reference point used to define its configuration space $\mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$. As discussed in Sections 2.2 and 2.3, P will reach one of a finite number of orientation equilibria when placed in f or $f_\perp$. More specifically, when P is placed in f, there exists a finite set of equilibria $E_f = \{(x_i, \theta_i)\}$, where x_i is the offset from f's squeeze line, and θ_i is the orientation of P (see Section 2.5). Similarly for $f_\perp(x, y) = (0, -\mathrm{sign}(y))$, there exists a finite set of equilibria $E_{f_\perp} = \{(y_j, \theta_j)\}$. Since the x-component of $f_\perp$ is zero, the x-coordinate of the reference point of P (the COM) remains constant while P is in $f_\perp$. Hence P will finally come to rest in a pose (x_k, y_k, θ_k), where $x_k \in \pi_1(E_f)$, $(y_k, \theta_k) \in E_{f_\perp}$, and π_1 is the canonical projection such that $\pi_1(x, \theta) = x$. Since E_f is finite, so is $\pi_1(E_f)$. $E(f_\perp)$ is also finite, therefore there exists only a finite number of such total equilibrium poses for $f * f_\perp$.

If P is placed into the field $f + f_\perp$, there exists a unique translational equilibrium (x, y) for every given, fixed orientation θ. In each of these translational equilibria, the squeeze lines of f and $f_\perp$ are both bisectors of P. Now consider the moment acting on P when P is in translational equilibrium as a function of θ. Since there are $O(n^2)$ topological placements for a single bisector, therefore there exist also only $O(n^2)$ topological placements for two simultaneous, orthogonal bisectors. In analogy to Proposition 4 in Section 2.2 we can show that for any topological placement of the bisectors, this moment function has at most $O(k)$ roots, where k is the maximum number of edges a bisector of P can cross. This implies that there exist only $O(k\,n^2)$ distinct total equilibria for $f + f_\perp$. ☐

Corollary 40 *Let f be a finite field operator for a part P, and let g be an arbitrary vector field. Then the sequence $g * f$ is a finite field operator.*

Proof: By definition of a finite field operator, f brings the part P into a finite set of equilibrium poses from arbitrary initial poses, in particular from the poses that are the result of field g. ☐

Thus by pre-pending an arbitrary sequence of fields to a finite field operator, one can always create a new finite field operator (possibly with a smaller set of discrete equilibria). In the remainder of this section, however, we will only consider finite field operators of minimal length, i.e. field sequences from which no field can be removed without losing the finiteness property (Definition 38).

We have seen in Sections 2 and 4 that for simple force fields such as e.g. squeeze or radial fields, we can predict the motion and the equilibria of a part with exact analytical methods. However, for arbitrary fields (e.g. the force fields described in Section 6.3 which are induced by vibrating plates), such algorithms are not known. Instead we can employ approximate methods to predict the behavior of the part in the force field. These methods are typically numerical computations that involve simulating the part from a specific initial pose, until it reaches equilibrium.[1] We call the cost for such a computation the *simulation complexity* $s(n)$. We write $s(n)$ because the simulation complexity will usually depend on the complexity of the part, i.e., its number of vertices n (for more details also see [DX95]).

Proposition 41 [BDM96b] *Consider a polygonal part P, and m finite field operators $\{F_i\}$, $1 \leq i \leq m$, each with at most E distinct equilibria in the configuration space $\mathcal{C}$ for P. There is an algorithm that generates an optimal-length strategy of the form $F_1 * F_2 * \cdots * F_l$ to uniquely pose P up to symmetries, if such a strategy exists. This algorithm runs in $O(m^2E\ (s(n) + 2^E))$ time, where $s(n)$ is the simulation complexity of P in F_i. If no such strategy exists, the algorithm will signal failure.*

Proof: Construct a transition table T of size m^2E that describes how the part P moves from an equilibrium of F_i to an equilibrium of F_j. This table can be constructed either by a dynamic analysis similar to Section 2.1, or by dynamic simulation. The time to construct this table is $O(m^2E\ s(n))$, where $s(n)$ is the simulation complexity, which will typically depend on the complexity of the part.

Using the table T, we can search for a strategy as follows: Define the *state* of the system as the set of possible equilibria a part is in, for a particular finite field operator F_i. There are $O(E)$ equilibria for each finite field operator, hence there are $O(m\,2^E)$ distinct states. For each state there are m possible successor states as given by table T, and they can each be determined in $O(E)$ operations, which results in a graph with $O(m\,2^E)$ nodes, $O(m^2 2^E)$ edges, and $O(m^2E\,2^E)$ operations for its construction. Finding a strategy, or deciding that it exists, then devolves to finding a path whose goal node is a state with a unique equilibrium. The total running time of this algorithm is $O(m^2E\,(s(n) + 2^E))$. ☐

Hence, as in [EM88], for any part we can decide whether a part can be uniquely posed using the vocabulary of field operators $\{F_i\}$ but (a) the planning time is worst-case exponential and (b) we do not know how to characterize the class of parts that can be oriented by a specific family of operators $\{F_i\}$. However, the resulting strategies are optimal in length.

This result illustrates a tradeoff between mechanical complexity (the dexterity and controllability of field elements) and planning complexity (the computational difficulty of synthesizing a strategy). If one is willing to build a device capable of radial fields, then one reaps great benefits in planning and execution speed. On the other hand, we can still plan for simpler devices (see Figures 1.4 and 5.2), but the plan synthesis is more expensive, and we lose some completeness properties.

5.2.2 Example: Uniquely Posing Planar Parts

In this section we will show how to accomplish tasks with manipulation grammars as developed in Section 5.2.1. Recall from Section 2.3 that we say a manipulation strategy orients (respectively, poses) a part uniquely if from *any* initial configuration, the part can be brought into a *unique* final orientation (respectively, pose). We will show how the synthesized plans uniquely pose parts from any initial configuration. As an example,

---

[1] See for example URL `http://www.cs.cornell.edu/home/karl/Cinema`.

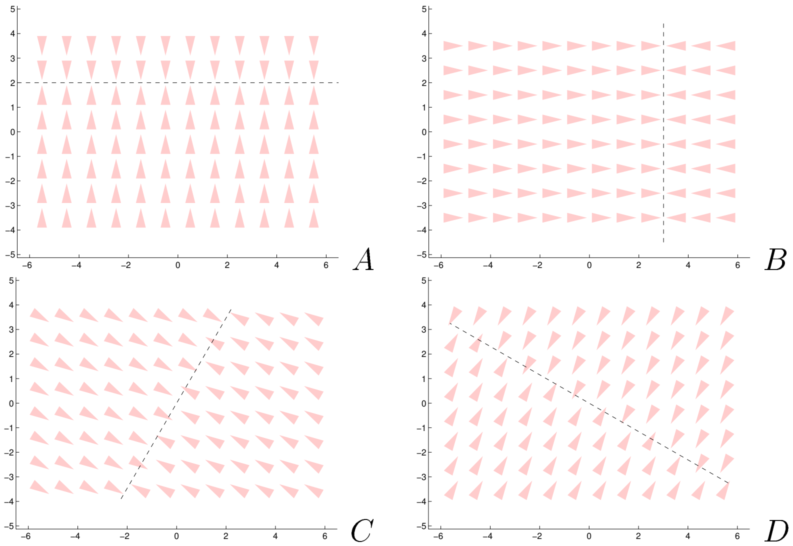

Figure 5.3: Manipulation vocabulary, consisting of 4 unit squeeze fields.

suppose our vibratory plate feeder can generate only a very limited vocabulary of four force vector fields, which are also not exactly centered on the plate. For simplicity we assume that the vocabulary consists of unit squeeze fields with squeeze lines at angles of 0, 90, 60 and 150 degrees. We call these fields A, B, C, and D, respectively. The squeeze line of field A is offset by 2 units from the origin, the squeeze line of B is offset by 3 units, and the squeeze lines of C and D intersect at the origin (see Figure 5.3).

The sequence $A * B$ is a finite field operator, since the squeeze lines of A and B are orthogonal (see Claim 39). In the remainder of this section, we will abbreviate "$A * B$" and simply write "AB." Other finite field operators besides AB are BA, CD, and DC, so that we obtain a vocabulary of $m = 4$ operators.

Note that using unit squeeze fields in this example is not essential; any fields that yield finite sets of equilibria could be used as well. However, for this "didactic" example it is advantageous to use unit squeeze fields because (a) it is easy to determine equilibria for unit squeeze fields, and (b) we can compare the result obtained here with the manipulation plans generated by the planner in Sections 2.3 and 2.4.

Uniquely Posing Rectangles

In this example we will attempt to generate plans for uniquely posing several rectangular parts with the manipulation vocabulary A, B, C, and D (up to part symmetry). As in Section 2.4, we consider three rectangles R_{10}, R_{20}, and R_{30} that have sides a and b such that a is 10, 20, and 30 percent longer than b, respectively (Figure 2.10). The stable equilibria of R_{10}, R_{20}, and R_{30} in a unit squeeze field were shown in Table 2.1. Modulo part symmetry, each squeeze field induces only two stable orientation equilibria for R_{10} and R_{20}, and only one stable orientation for R_{30}. Also note that in stable equilibrium, the COM of a rectangle lies on the squeeze line. This gives us a total of $mE = 4 \cdot 2 = 8$ discrete equilibria for R_{10} and R_{20}, when using the finite field operators AB, BA, CD, and DC. All equilibria are shown in Table 5.1 (compare with Table 2.1 and Figure 2.12). Finally, any one of the operators AB, BA, CD, and DC uniquely orients R_{30}, yielding

Table 5.1: Stable equilibria of rectangular parts R_{10} and R_{20} for the manipulation vocabulary AB, BA, CD, and DC.

Operator	Equilibrium	R_{10} (x, y, θ)	R_{20} (x, y, θ)
AB	1	(3,2,0.97)	(3,2,1.29)
	2	(3,2,2.18)	(3,2,1.85)
BA	3	(3,2,2.54)	(3,2,2.86)
	4	(3,2,0.61)	(3,2,0.28)
CD	5	(0,0,2.01)	(0,0,2.34)
	6	(0,0,0.08)	(0,0,2.90)
DC	7	(0,0,0.44)	(0,0,0.77)
	8	(0,0,1.65)	(0,0,1.33)

Table 5.2: Transition table for equilibria of the rectangles R_{10} and R_{20}, with finite field operators AB, BA, CD, and DC. For both rectangles, there exist a total of $E = 8$ equilibria, and $m = 4$ finite field operators.

			R_{10}				R_{20}			
			to				to			
			AB	BA	CD	DC	AB	BA	CD	DC
from	AB	1	1	4	6	7	1	4	5	8
		2	2	3	5	8	2	3	5	8
	BA	3	2	3	5	8	2	3	6	7
		4	1	4	6	7	1	4	6	7
	CD	5	2	3	5	8	2	3	5	8
		6	1	4	6	7	2	3	6	7
	DC	7	1	4	6	7	1	4	6	7
		8	2	3	5	8	1	4	5	8

trivial one-step plans to uniquely pose R_{30}. Hence we will omit R_{30} for the remainder of this example.

Given the discrete equilibria, the algorithm based on the constructive proof of Proposition 41 generates a transition table T that describes the mapping between initial equilibrium pose and final equilibrium pose of a part when one finite field operator is applied. This table has mE rows and m columns. Table 5.2 shows the transitions for parts R_{10} and R_{20}. Each entry in T can be determined either by dynamic analysis, or by simulation. The values in Table 5.2 were generated by our planner using simulation. Figure 5.4 shows a trace of such a simulation: The initial pose of part R_{20} is equilibrium $\mathbf{e}_3 = (3, 2, 2.86)$. In field C, R_{20} moves left and up until it reaches an equilibrium on the squeeze line of C. Subsequently, after field D is applied, R_{20} comes to rest in equilibrium $\mathbf{e}_6 = (0, 0, 2.90)$. In this case, using Claim 39, the equilibria (but not the transitions) can be calculated analytically.

Recall from Section 5.2.1 that this system has a state space of size $O(m\,2^E)$, because for each of the m finite field operators, there are $O(E)$ discrete equilibria in which the

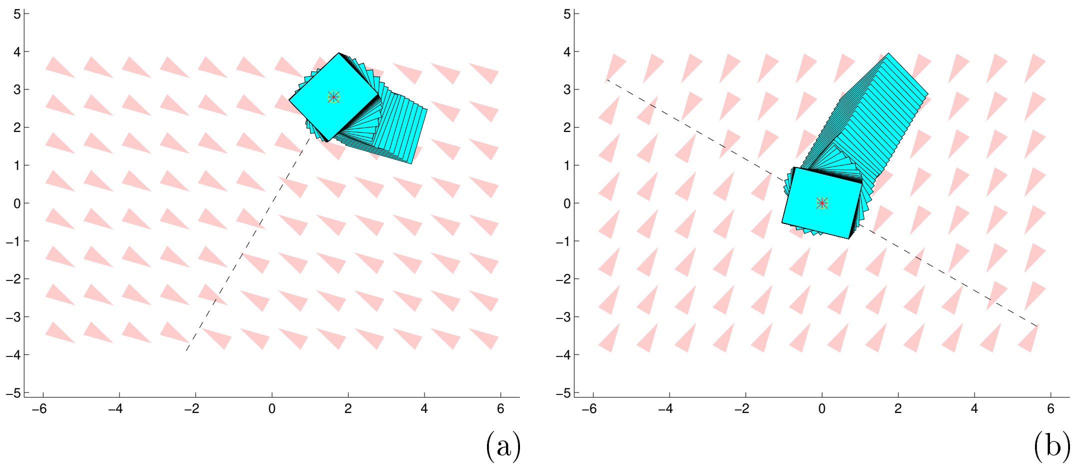

Figure 5.4: Simulation of part R_{20} from equilibrium 3 by using finite field operator CD, reaching equilibrium 6: (a) applying field C; (b) applying field D.

part could be. For example, a state could be "the part is in equilibrium 1, 2, or 4." We can represent such a state as a binary string, "11010000." Hence the transition table T can be used to define a transition graph whose nodes are the $O(m\,2^E)$ states, and whose $O(m^2 E\,2^E)$ edges are derived from the mE transitions in T. A simple breadth-first search of this graph, starting from the state in which all equilibria are possible, will yield optimal-length plans to reach any reachable state.[2] This algorithm will also decide which states are unreachable. Hence it can signal success when the shortest plan to reach a state with a unique equilibrium is found, or signal failure if no such plan exists. Figure 5.5 shows transition graphs for parts R_{10} and R_{20} with all reachable states, and the shortest paths to reach them from the initial "generic" state, in which the part has an arbitrary pose. Notice that there exists a two-step plan for uniquely posing R_{20}, but no such plan exists for R_{10}.

In summary, we observe that with our finite field operators AB, BA, CD, and DC, R_{30} can be uniquely posed in one step, R_{20} requires two steps, while there exists no strategy for R_{10}. Recall that the general squeeze algorithm in Section 2.4 found an alignment strategy for all three rectangles R_{10}, R_{20} as well as R_{30}. However, the algorithm required two squeeze fields at a relative angle of approximately 45°; for R_{10}, it would fail for squeeze lines at a relative angle of 60°. Apparently, parts that are closer to rotational symmetry (i.e., in this case, closer to square-shaped) are more difficult to pose uniquely than more asymmetric (i.e., long rectangular-shaped) parts.

Uniquely Posing and Feeding Arbitrary Parts

In this section we will demonstrate the manipulation grammar algorithm for a more realistic part (see Figure 5.6a), and for two different manipulation vocabularies. All strategies in this section (and Section 5.2.2) were computed using an automatic planner we implemented, using the techniques of Section 5.2.1. We will first extend our manipulation grammar by adding a field F that has a vertical squeeze line at $x = -3$ (Figure 5.7 left), which yields two new finite field operators, AF and FA. Analysis of the part shows that it has

---

[2] We could also imagine using A*-search to improve performance.

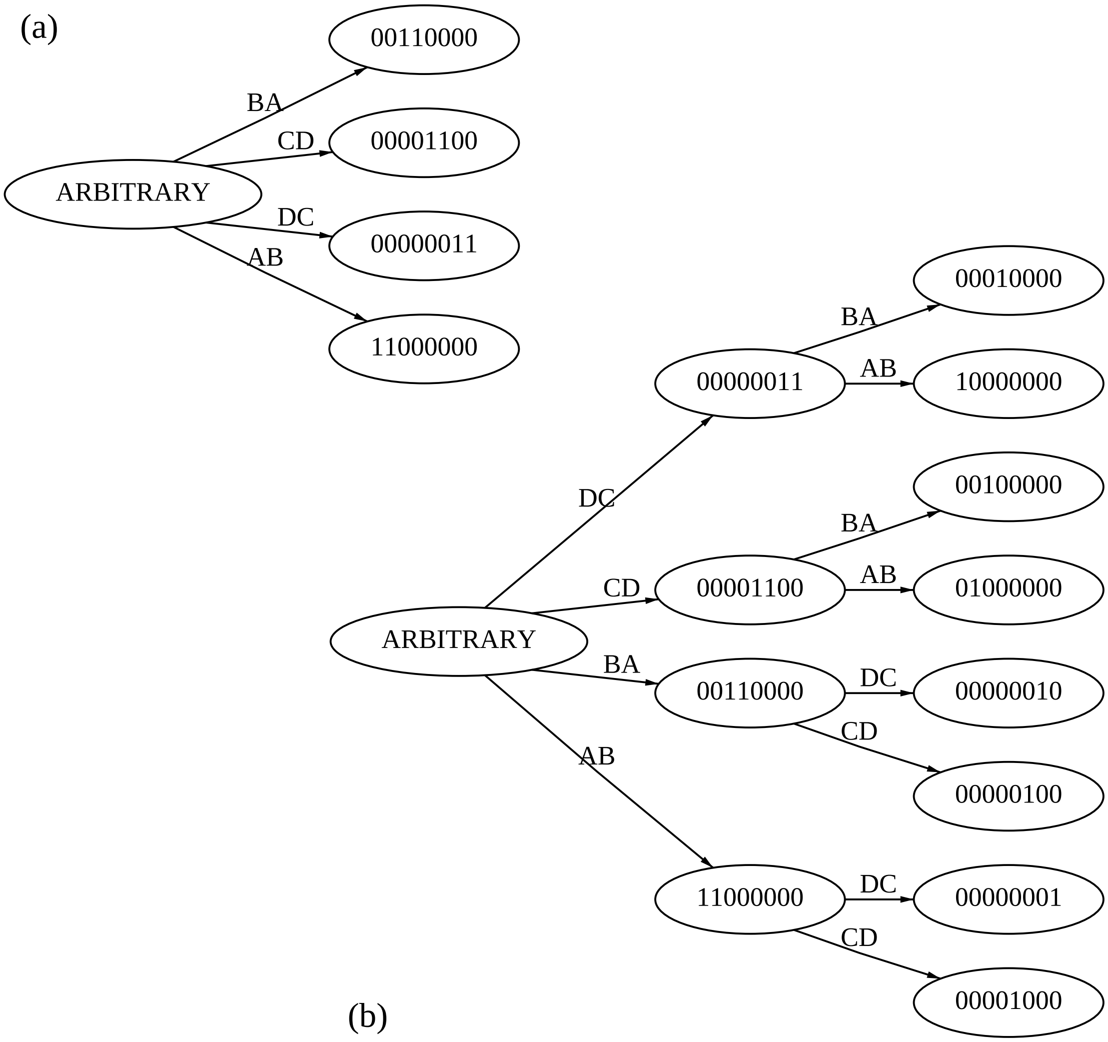

Figure 5.5: Minimum spanning trees of the state transition graphs for rectangles (a) R_{10}, and (b) R_{20}. All reachable states are shown, as well as the shortest paths to reach each of them. Non-spanning edges (e.g. an edge CD from 11000000 to 00001100) are omitted for simplicity.
(a) No state with unique equilibrium can be reached for R_{10}.
(b) There exist several two-step plans for R_{20} that reach states with unique equilibrium. (Graphs were generated automatically by our planner software.)

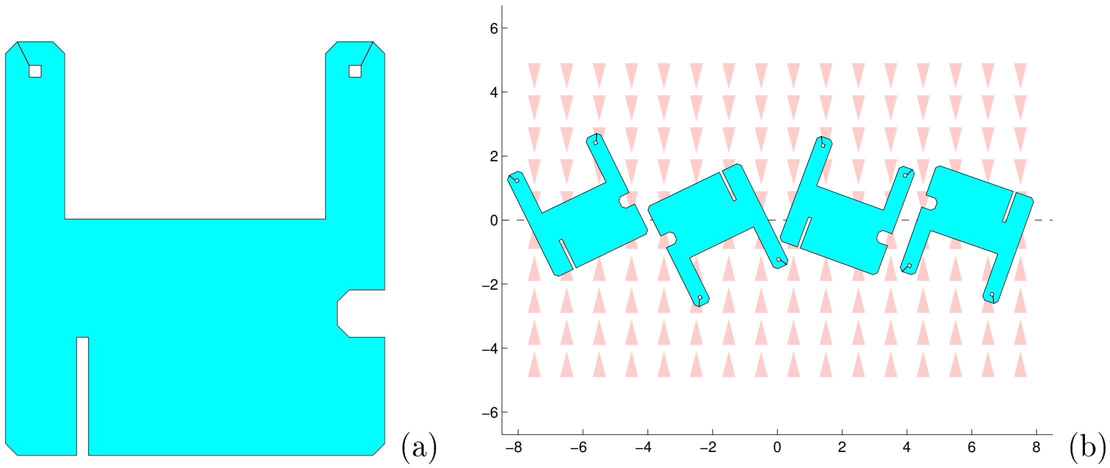

Figure 5.6: Sample part: (a) nonconvex shape with holes; (b) its four stable equilibria in a unit squeeze field.

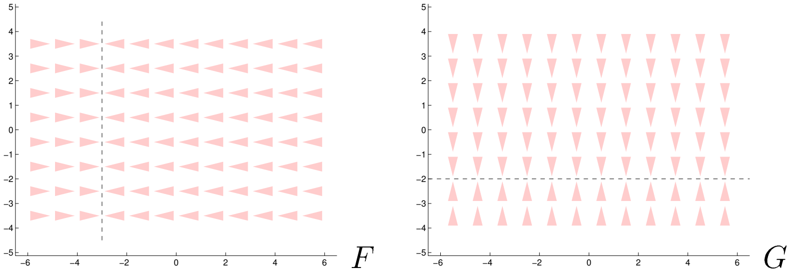

Figure 5.7: Extensions to the manipulation vocabulary, consisting of 2 unit squeeze fields.

4 stable orientation equilibria in a unit squeeze field (Figure 5.6b). It is not difficult to see that, after any two orthogonal squeezes, the part can be in $E = 8$ different poses. We obtain a transition table of size $m^2E = 228$, which results in a state transition graph with $m\,2^E = 1536$ nodes (states) and $m^2 2^E = 9216$ edges (transitions). The algorithm finds the following strategy: $CD\ BA\ AF\ FA$, which is equivalent to $CDBAFA$. Two sample executions of this strategy are shown in Figure 5.8, from different initial poses. A close look at the strategy reveals that operator CD approximately centers the part, such that B can move the part into one of four discrete orientation equilibria below the squeeze line of A. Then A reduces the number of orientation equilibria to two, and F to one (at a unique x-position). Finally, A brings the part into a unique pose: $\mathbf{e}_* \approx (-2.9, 1.9, 3.6)$.

It is important to note the following distinction between the general squeeze strategies for parts orienting of Section 2.3, and the manipulation grammar strategies: As mentioned in Section 2.3, turn and squeeze functions render planning algorithms based upon them susceptible to field symmetries, thereby introducing aliasing in orientation space and admitting completeness and uniqueness proofs of orientation only modulo field symmetry. Since manipulation grammars do not employ turn or squeeze functions, they are immune to this problem, and parts without rotational symmetry can be posed *uniquely*. In essence, turn and squeeze functions assume a global field symmetry. In manipulation grammars, such field symmetries may not exist, e.g. squeeze fields could have arbitrary angles and

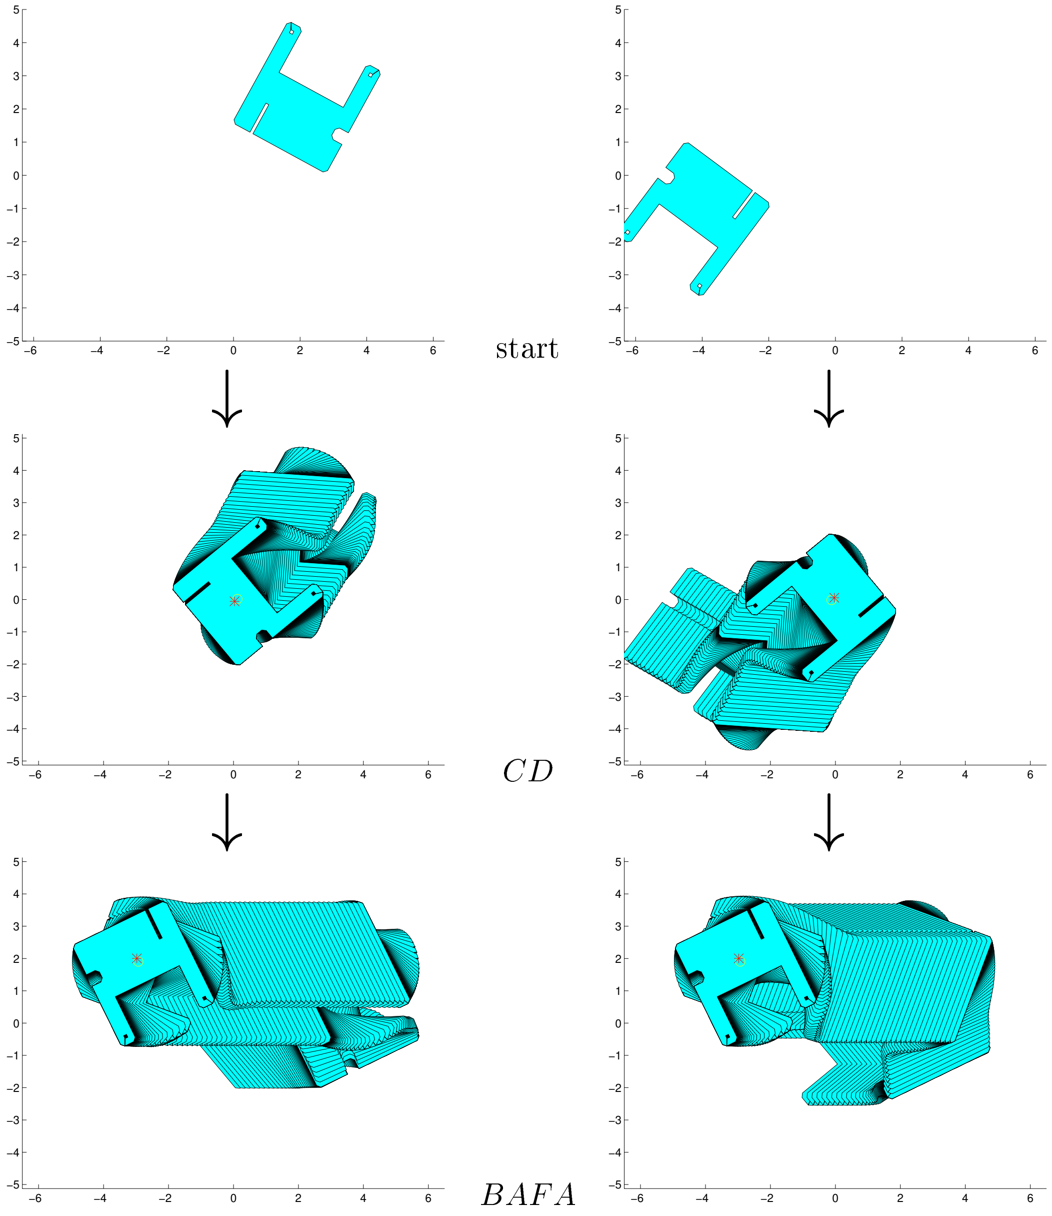

Figure 5.8: Two sample executions of the manipulation grammar strategy $CD\ BA\ AF\ FA$. For clarity, the simulation trace has been broken up into parts: initial pose (top), motion under CD (middle), and motion under $BA\ AF\ FA$ (bottom). Initial poses: (a) $\mathbf{z}_0 = (2, 2, -0.5)$, (b) $\mathbf{z}_0 = (-4, -1, 2.5)$. Final pose $\mathbf{e}_* \approx (-2.9, 1.9, 3.6)$.

offsets from the origin. In the first example of this section (Figure 5.8), the final pose is indeed unique.

As a second example, we add the field G, which has a horizontal squeeze line at $y = -2$ (Figure 5.7 right), and remove the fields C and D. This results in 8 finite field operators, hence we obtain $m^2 E = 512$ entries in the transition table, $m\,2^E = 2048$ states and $m^2 2^E = 16384$ transitions. We obtain the strategy $GB\ BA\ AF\ FG$, which is equivalent to $GBAFG$. During execution of this strategy, the COM of the part follows a counter-clockwise rectangular path, at each step reducing the number of possible equilibria, until, in the lower left corner, a unique pose is reached (Figure 5.9). This opens the possibility of *pipelining* the posing process, which could yield more efficient parts feeders: as long as we can ensure that the next part is initially placed sufficiently far to the right so not to interfere with its predecessor, the G field can be used simultaneously for two parts. Hence if the parts feeder periodically cycles through the fields $GBAF$, the next part can be introduced into the device each time *before* G is executed. A part is uniquely posed *after* each execution of G.

5.2.3 Summary

In this section we have defined manipulation grammars that consist of a vocabulary of planar force vector fields, and we presented an implemented planning algorithm that generates strategies to uniquely position and orient parts. In comparison with the general squeeze strategies of Section 2.3, manipulation grammars allow sets of arbitrary force vector fields, and are not limited to a 1-parameter family of squeeze fields. Consequently, depending on the available manipulation vocabulary, the resulting strategies can be more powerful or more restricted than the orienting strategies generated by the general squeeze algorithm of Section 2.3. In particular, parts can be uniquely posed even when only symmetric force fields are available. As a tradeoff, planning and execution complexity is worst-case exponential instead of merely quadratic in the number of equilibria of the part, and there exist no completeness guarantees that a strategy always exists for a given vocabulary or class of parts. Moreover, numerical simulation was employed to predict the transitions, whereas they may be exactly computed (Section 2.3) for simple squeezes.

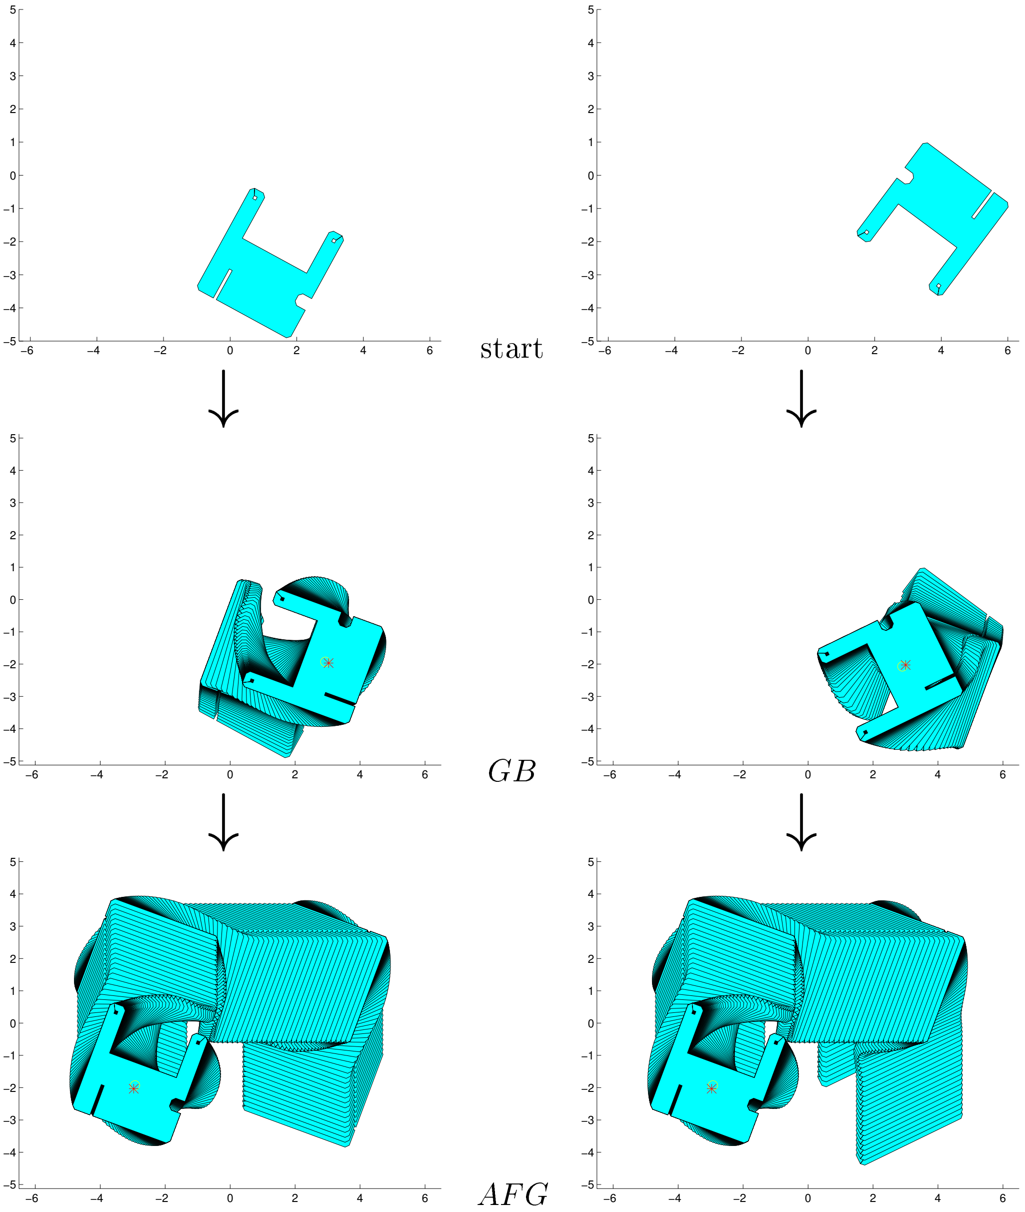

Figure 5.9: Two sample executions of strategy $GB\ BA\ AF\ FG = GBAFG$. For clarity, the simulation trace has been broken up into parts: initial pose (top), motion under GB (middle), and motion under AFG (bottom). Initial poses: (a) $\mathbf{z}_0 = (1, -3, -0.5)$, (b) $\mathbf{z}_0 = (4, -1, 2.5)$. Final pose $\mathbf{e}_* \approx (-2.9, -1.9, 5.9)$.

Chapter 6

Experimental Apparatus for Programmable Force Fields

The theory of programmable vector fields developed in this thesis arguably represents the first systematic attack on massively-parallel distributed manipulation based on geometric and physical reasoning. In this chapter we report on our experiments in implementing this theory using microfabricated actuator arrays and macroscopic vibrating plates.

6.1 Microfabricated Arrays of Single-Crystal Silicon Torsional Actuators

A wide variety of MEMS (Micro Electro Mechanical Structures) has been built recently by using processing techniques derived from VLSI fabrication. The first commercially successful MEMS structure was the Analog Devices, Inc. single-chip, bulk-fabricated air bag sensor [Ana91]. In this section, we investigate arrays of MEMS actuators that can perform manipulation tasks. There is a large number of possible applications from macroscopic to microscopic size scales, such as bulk-fabricated (cheap), ultra-thin transport mechanisms, e.g. for paper in copy machines or printers. At the other end of the scale spectrum, recent advances have brought within reach arrays equipped with tips that can probe and move objects consisting of only a few atoms [XMM95]. Such devices, employed in a massively parallel fashion, may yield tremendous data storage capacities.

Various actuation principles for micro actuator arrays and micro manipulation systems have been proposed in recent years. Pister et al. [PFH90] demonstrated a levitation system consisting of microfabricated nozzles, and electrostatic actuation. Takeshima and Fujita [TF90] introduce the concept of a distributed micro motion system that consists of an array of cooperating actuator modules. Furuhata et al. [FHF91] have built arrays of ultrasonic microactuators. An integrated, monolithic mirror array for video projections has been presented by Sampsell [Sam93]. Böhringer et al. [BDMM94b] have developed a theory of manipulation and control for micro actuator arrays. Konishi and Fujita [KF93a] use controlled, directed micro air valves to convey small objects. Ataka et al. [AOF93] simulate cilia with thermal bimorph micro structures. Liu and Will [LW95] propose the concept of intelligent motion surfaces for micro assembly tasks. Liu et al. [LTW$^+$95,LW95] present a micro assembly system based on magnetic actuator arrays. Electrostatic actuator arrays are described by Böhringer et al. [BDM96a] and Storment et al. [SBW$^+$94]. Suh et

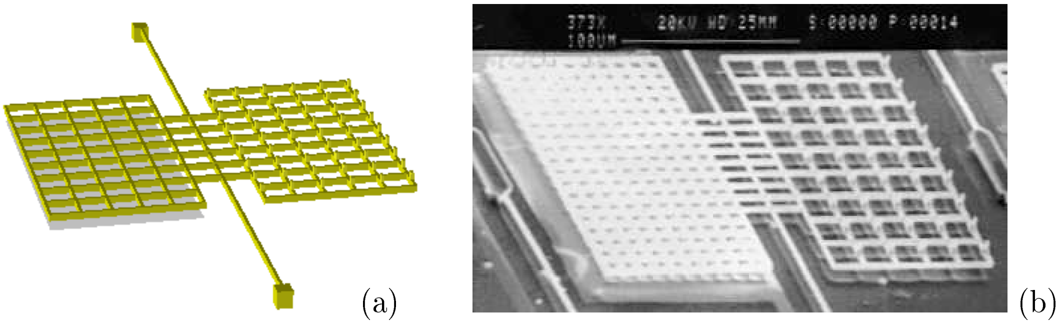

Figure 6.1: Torsional asymmetric actuator with electrode and tips, suspended on a torsional rod: (a) CAD model, (b) scanning electron microscopy. Left: Dense grid (10 μm spacing) with aluminum electrode underneath. Right: Grid with 5 μm high poles.

al. [SGD+96] present an array of thermobimorph micro cilia which are able to accurately move and position silicon chips. The MEMS array that we present here is designed for applications in which objects in the millimeter range are moved, e.g. for a programmable stage of a microscope, or for automated handling and assembly of small parts.

Reznik et al. [RBC97] performed thorough tests of the proposed manipulation strategies with a sophisticated dynamic simulator [MC95], verifying the predicted behavior of parts by our theory of force vector fields. Manipulation experiments that support our theory are described by Suh et al. [SGD+96].

Compatibility with standard VLSI processes allows integration of control circuitry on the same chip. This is of great importance if we want to employ powerful manipulation strategies that require individual control of the actuators in the array. The SCREAM process [ZM92,SZM93] (for Single-Crystal Silicon Reactive Etching and Metallization) is a low temperature process that can be performed *after* the fabrication of circuits has been completed [SM96]. In the following sections we will discuss design, fabrication, and testing of a SCREAM electrostatic actuator array (see Figure 1.2), and we discuss control strategies for massively-parallel actuator arrays.

Our goal is to generate motion plans at a high (task) level and automatically transform them into an actuator array control strategy, in analogy to a compiler that translates high level instructions into assembler code. We believe that such high level strategies of micro sensors and actuators are essential for efficient control of future complex MEMS systems. Conversely, the availability of high level descriptions for MEMS and micro robotic tasks may influence and improve the MEMS design, similarly to the concurrent design of omputer chips and compilers.

6.1.1 Actuator Design

A torsional actuator consists of a rectangular grid etched out of single-crystal silicon suspended by two rods that act as torsional springs (Figure 6.1). The design is based on torsional resonators [MZSM93,MM96]. In our current design, the grid is 180 μm long and extends $120 - 180\,\mu m$ on each side of the rod. The rods are 150 μm long. The current asymmetric design has $5 - 10\,\mu m$ high protruding tips (Figure 6.2) on one side of the grid that make contact with an object lying on top of the actuator. The other side of the grid consists of more densely spaced beams above an aluminum electrode. If a voltage is

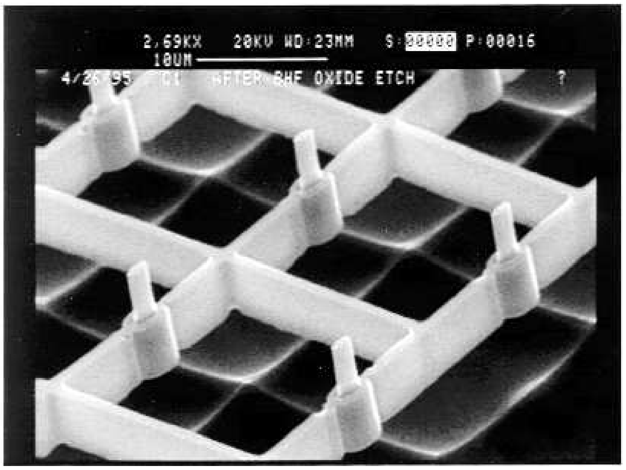

Figure 6.2: Released actuators consisting of single-crystal silicon with 5 μm high tips.

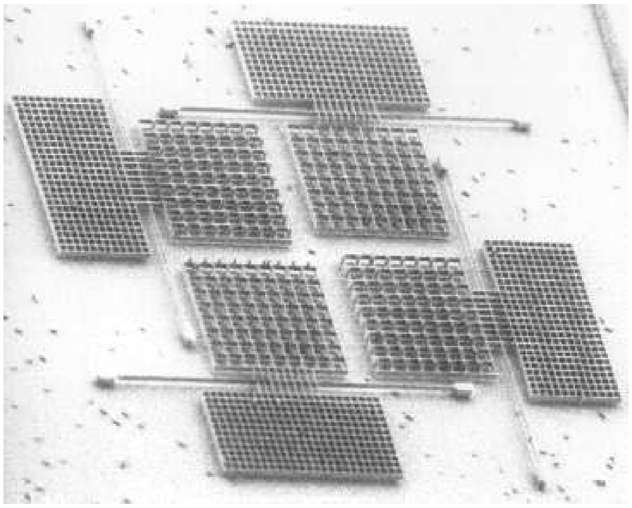

Figure 6.3: Micro motion pixel consisting of actuators oriented in four different directions.

applied between actuator grid and electrode, the half of the grid above the electrode is pulled downward by the resulting electrostatic force. Simultaneously the other side of the grid (with the tips) is deflected out of the plane by several μm. Hence an object can be lifted and pushed sideways by the actuator.

Because of its low inertia (resonance in the high kHz range) the device can be driven in a wide frequency range from DC to several 100 kHz AC. Due to the asymmetry in the actuator design, each actuator generates motion in one specific direction if it is activated; otherwise it acts as a passive frictional contact. Figure 1.2 shows a small portion of such a unidirectional actuator array, which consists of more than 15,000 individual actuators that densely cover the substrate surface.

The layout of the array can be changed such that the actuators point in various orientations. Then the combination and selective activation of several actuators in different orientations (a *motion pixel*, see Figure 6.3) allows us to generate various motions in discrete directions, spanning the plane.

Feasibility Calculations

To answer the question *Can microscopic actuators move macroscopic parts?* we investigate the (vertical) levitation forces generated by the micro actuators. Section 6.1.1 will then focus on lateral forces that induce motion on the part. Let us consider the specific weight (i.e., in this case, the weight per area ratio) of a variety of objects and materials. For example, the specific weight of paper is approximately $1\,N/m^2 = 1\,\mu N/mm^2$. A silicon wafer is about ten times heavier at $10\,\mu N/mm^2$. We compare these values with the electrostatic force generated by a parallel-plate capacitor. The force per area ratio can be computed as $F/A = \varepsilon_0 V^2/d^2$, where V is the applied voltage, d is the gap width between the capacitor plates, and ε_0 is the permittivity of free space. For a gap width of $d = 5\,\mu m$ and a voltage of 60 V (typical values for our actuators), we obtain a force per area ratio of approximately $1.2\,mN/mm^2$. This result is more than two orders of magnitude higher than e.g. the specific weight of a silicon wafer. Note that the force is proportional to the square of the applied voltage V as well as the gap width d. Therefore, higher voltages and smaller gap sizes will largely improve the effectiveness of our electrostatic devices. However, the force diminishes quickly with larger gap widths, which limits the range of motion of the current actuator design.

Due to fabrication constraints, the electrode of our torsional actuator does not consist of a continuous plate, but rather of a grid of approximately $1\,\mu m$ wide beams spaced at $10\,\mu m$ distance. This is necessary to release the structure from the substrate, as well as for the deposition of an Aluminum electrode underneath the actuator. To investigate the loss in electrostatic force incurred by these design limitations, we have analyzed torsional actuators using the finite element simulator COULOMB [Int89]. This analysis shows that for the given actuator geometry, the electrostatic force is reduced to about one-third of the force generated by an ideal parallel-plate capacitor.

Finally, due to layout constraints, not all of the wafer surface can be used for actuators. In a practical layout of a densely packed actuator array, approximately 30 % of the chip surface is covered with electrostatic capacitors. The remaining area is used for the actuator grids that support the tips, torsional beams, and electrical connections.

Summarizing the results of the feasibility calculations in this section, we conclude that even though the effective electrostatic force generated by our actuators is only about one tenth of an ideal parallel-plate capacitor, useful manipulation tasks can be performed with electrostatic micro actuators.

New designs with much higher performance are currently under development. Parallel-plate capacitors can be replaced by torsional comb-drive actuators, which combine very small gap widths with theoretically unlimited angular range of motion. Since these high aspect ratio comb drives employ *vertical* side walls as capacitors, the effective capacitor surface increases significantly. In particular, it can be larger than the (two-dimensional) chip area. Smaller gap sizes and larger surface areas will result in an increase in electrostatic force by several orders of magnitude.

Design of SCS Torsional Actuators

In this section we investigate the generation of lateral motion based on a kinematic analysis of torsional actuators, which will yield design rules to optimize actuator design. Calcula-

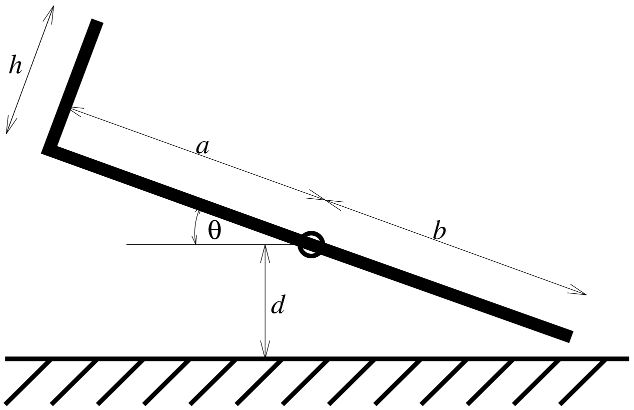

Figure 6.4: Schematic cross section of torsional actuator (not to scale): Actuator width (from torsional axis) a and b, height of tips h, gap width between actuator and electrode d, and torsional deflection θ.

tions show that the mass of an individual actuator is less than 100 ng; by comparison, the mass of a typical object that we want to move (paper, silicon wafer) is in the $10 - 100\,\mu g$ range, when considering an object of equal area size. Because of this difference of several orders in magnitude, impact dynamics are negligible, hence we will focus on actuator kinematics.

Figure 6.4 shows a schematic cross section of an actuator with tips on one edge. d is the clearance between actuator grid and trench, i.e. the gap width of the capacitor. a is the distance between tip and rotational axis, b the width of the other side of the oscillator, h the height of the tips, and θ the angular deflection of the actuator.

If the actuator is inactive, the object will rest on the tips. In case it is active and the tips are moving upward, they engage with the bottom surface of the object, pushing it up and sideways. During the downward motion of the actuator we assume that the object keeps its horizontal position (e.g. because of inertia, or because other actuators are holding it), and moves straight down (e.g. because of gravity or electrostatic forces). This results in a stepwise lateral motion with each oscillation of the actuator. We are interested in values for the design parameters that optimize this motion.

Let Δx be the horizontal component of the motion of a tip when it moves from $\theta = 0$ to maximum deflection $\theta_{\max}$. Thus

$$\Delta x = a - a\cos\theta_{\max} + h\sin\theta_{\max}$$

For small maximum deflections $\theta_{\max}$, this equation can be simplified to

$$\Delta x \approx \frac{1}{2}a\theta_{\max}^2 + h\theta_{\max}$$

Hence the step size Δx is linear in the size of the tips h and the actuator size a, and it is also proportional to the ratio of gap width and actuator size, $d/b = \sin\theta_{\max} \approx \theta_{\max}$. Also note that we want to choose $a \approx b$: otherwise, if $a \gg b$, the force of the actuator diminishes because the lever supporting the actuator tips is relatively longer, and the capacitor surface is reduced; if $a \ll b$, then the out-of-plane deflection $\theta_{\max}$, and subsequently Δx, becomes smaller.

Scaling Effects and Actuator Efficiency. To optimize the performance of the actuators, we investigate the lateral speed of an object as it is moved by the actuator array. Let us consider the ideal case where there is no slip between moving object and actuator during the out-of-plane motion. In this case the speed is the product of the step size Δx and the oscillation frequency f:

$$v = \Delta x \; f$$

Even though we can operate our devices at a wide range of frequencies, an upper bound is given by its natural frequency

$$f_n = \frac{1}{2\pi}\sqrt{K/I}$$

where K is the spring constant of the torsional rod, and I is the moment of inertia of the actuator. The spring constant can be computed as

$$K \approx x^3 y G / 3l$$

for a rectangular rod of cross-sectional width x and height y with $x < y$, rod length l, and shear modulus G [JM83]. The moment of inertia I of the actuator is proportional to its mass m and the square of its extension from the torsional axis r, $I \propto mr^2$.

Now consider scaling the area size of an individual actuator by a factor s, while leaving constant the vertical dimensions, and the torsional rod cross section. The spring constant scales as $K \propto 1/s$, while for the moment of inertia the dependency $I \propto s^4$ holds. Thus for the resonance frequency we obtain the relationship $f_n \propto s^{-5/2}$, which tells us that with decreasing scale s, the frequency increases much faster than the corresponding decrease in step size Δx. Also note that the force per area ratio remains constant during down-scaling (as long as the loss of area due to additional connections is not significant). We conclude that our actuators become more efficient the more they are scaled down. For a more thorough discussion of scaling effects in MEMS, see e.g. [Mac96a].

There are, however, limits to these advantageous effects of down-scaling.

1. Useful actuators require a certain minimum range of motion such that devices can still make contact with rougher or not perfectly flat objects.
2. The electrostatic force generated by an individual actuator capacitor must be sufficient to fully deflect the device, but down-scaling reduces the parallel-plate capacitor force and increases the torsional stiffness due to shortening of the torsional beam.
3. Finally, fabrication requires a certain minimum feature size (in our case about $0.8\,\mu m$).

Our current actuator design has a torsional rod length of $150\,\mu m$, grid sizes of $a = 180\,\mu m$ and $b = 140\,\mu m$, tip heights of $5 - 10\,\mu m$, and a capacitor gap size of $5\,\mu m$, which results in a kinematic constraint for the deflection angle at $\theta_{\max} \approx 2.4°$. Simple calculations show that with these dimensions, the electrostatic force is sufficient to overcome the torsional spring forces. We obtain a step size of $\Delta x \approx 500\,nm$. Resonance occurs in the $10 - 100\,kHz$ range, which yields a maximum lateral speed of $\Delta x \; f_n \approx 5 - 50\,mm/sec$.

6.1.2 Fabrication Process

The SCREAM process [ZM92,SZM93] (for Single-Crystal Silicon Reactive Etching and Metallization) is a reactive ion etching process for the fabrication of submicron, movable single-crystal silicon (SCS) electromechanical structures. Recently, several process variations have been developed to create a wide variety of SCS MEMS devices, e.g. an integrated scanning tunneling microscope (STM) [XMM95], opto-electronic devices [MJ95], loading devices for tribological measurements [PMT95] or micro assembly [PBM95] and actuators that generate forces in the milli-Newton range [SM95]. For an overview see [MAA+95, Mac96b, Mac96a]. Characteristics of the SCREAM process are high aspect ratio SCS structures, electrostatic actuation, high vertical stiffness, low in-plane and torsional stiffness, and compatibility with regular VLSI processes [SM96]. We have developed a modified, two-layer SCREAM process to fabricate devices with tips on movable micro structures [BDMM94b, BDM96a]. This process requires three photolithographic masks: the first defines the tips, the second the torsional actuators, and the third mask includes the patterns for electrodes and bonding pads. In the remainder of this section we give a concise summary of the fabrication process. A detailed description of all processing steps, with a complete list of all important process parameters, can be found in Appendix A.

Photolithographic Masks. The layout of the complete devices comprising torsional actuators with tips, electrodes, and contact pads is drawn using SYMBAD, a CAD package for integrated circuit (IC) and MEMS devices. For each of the three layers, the fractured CAD data is transferred to the pattern generator to create $5''$ Cr masks for use in a 10:1 optical stepper.

Tip Layer. On the substrate of an arsenic-doped, $0.005\,\Omega cm$ low resistivity, n-type (100) silicon wafer, a $3.2\,\mu m$ thick etch mask layer of SiO_2 is deposited using plasma enhanced chemical vapor deposition (PECVD) at $300^\circ\,C$, $450\,mT$, N_2O flow of $42\,sccm$ and SiH_4 flow of $12\,sccm$ for 90 min. Photolithography is used to transfer the pattern of the tips from the mask onto a layer of OCG 985i 50cs positive resist spun on the oxide. The minimum feature size in our devices is $0.8\,\mu m$ for tips and beams. The pattern is transferred from the resist layer to the SiO_2 layer using magnetron ion etching at a pressure of $2\,mT$ with a flow rate of $30\,sccm$ of CHF_3 at $1000\,W$. The photoresist is removed using an O_2 plasma etch. The pattern is transferred into the silicon substrate from the SiO_2 layer using Cl_2 reactive ion etching at $40\,mT$, $400\,V$, and at a flow rate of $50\,sccm$ for Cl_2 and $1.3\,sccm$ BCl_3 for 30 min to get $5.0\,\mu m$ deep trenches. This trench depth determines the height of the tips (see Figure 6.5a).

Actuator Grid Layer. A $2.0\,\mu m$ layer of mask SiO_2 is deposited, and the above steps are repeated for patterning of the actuator grids (Figure6.5b). Exposure of the actuator grid pattern requires careful characterization because lithography is performed on a highly non-planar surface. Accumulation of excessive photoresist around the tips has to be avoided as much as possible to prevent the supporting grid structure from becoming too wide, which would complicate the subsequent release process.

Following the second Cl_2 etch, an insulating layer of silicon dioxide is deposited for side wall passivation using 15 min of PECVD at the parameter values stated above, resulting in a $400\,nm$ thick silicon oxide layer (Figure 6.6c). Slow CHF_3 reactive ion etching is used to remove only the floor oxide for subsequent substrate etches.

A short Cl_2 reactive ion etch at $50\,sccm$ Cl_2 and $1.3\,sccm$ BCl_3, $40\,mT$, and $400\,V$ for 10 min generates $2\,\mu m$ deep trenches in the substrate to aid in the following release etch. All beams of thickness up to $2\,\mu m$ are released from the silicon substrate using an RIE process, in which an SF_6 plasma with a flow rate of $140\,sccm$ at $90\,mT$ and $150\,W$ releases the beams in 6 min. Finally, all remaining oxide is stripped in a 6:1 buffered HF etch.

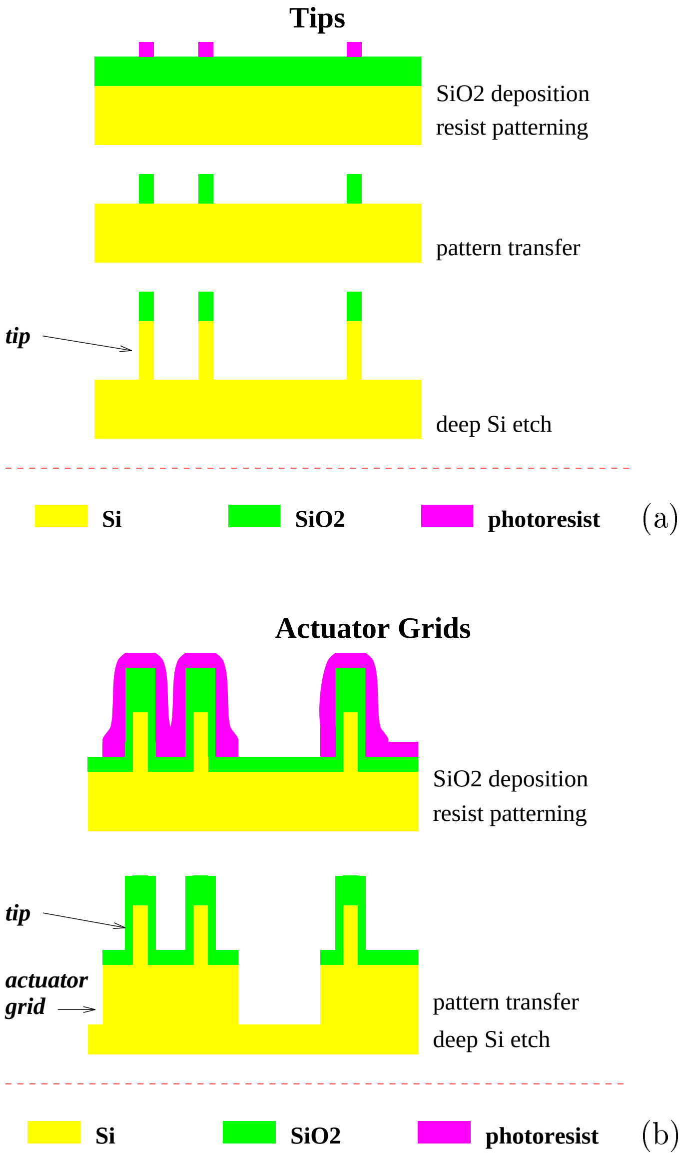

Figure 6.5: Device fabrication with a two-layer SCREAM process: (a) forming of tips using an RIE chlorine etch, (b) fabrication of actuator structures consisting of SCS beams of 1 μm width.

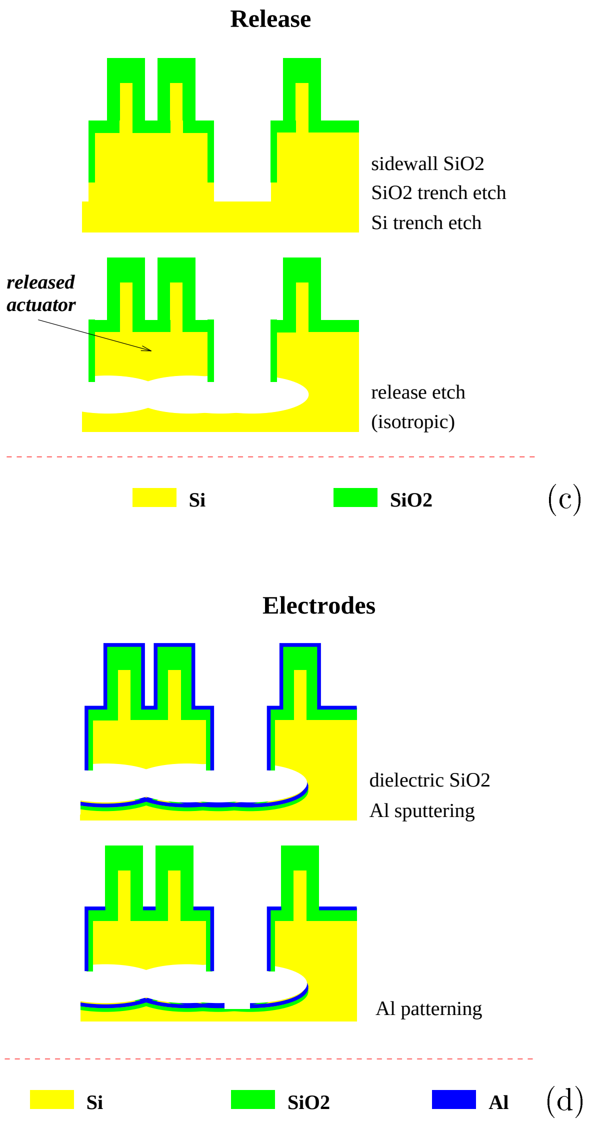

Figure 6.6: Device fabrication with a two-layer SCREAM process: (c) MIE/RIE trench bottom etch and RIE release etch, (d) Aluminum electrode deposition and patterning.

Electrode Layer. A 200 *nm* thin layer of dielectric PECVD silicon oxide is deposited. A 250 *nm* layer of Al is conformally deposited using sputtering performed at 30 mT pressure, with a beam current of 5 A and an Ar flow rate of 30 *sccm*. Due to overhanging sidewalls (see Figure 6.6d) trench and mesa are electrically isolated. Note however that, because of the rather high degree of isotropy of the Al sputtering, the electrodes under the capacitor grids consist of continuous Al. To pattern the Al in the trench we use lithography with OCG 985i 50cs positive resist. The pattern is transferred using Cl_2 reactive ion etching at 20 mT, 400 V, and flow rates of 20 *sccm* Cl_2, 40 *sccm* BCl_3, and 1.3 *sccm* CH_4 for 1.5 min.

Finally, a 100 *nm* thick layer of PECVD SiO_2 can be deposited on the device to avoid shorting when the actuators make contact with the trench electrodes during operation. The actuators consist of beams that are close to 1 μm wide and 5 μm high, with approximately 5 μm clearance underneath. Our actuator designs have grids of $120 - 240\,\mu m$ side length. The fabrication can be done in less than one week in the Cornell Nanofabrication Facility (CNF) at Cornell University.

Packaging and Testing. Chips of up to 6 cm × 2 cm with more than 3000 actuators have been packaged and wire-bonded. Larger arrays have been tested in the probe station. When an AC voltage (not necessarily at resonance frequency) is applied between the Al electrodes and the Si substrate, the actuators are brought into an oscillatory motion. The dielectric silicon oxide beween Al electrodes and Si substrate can withstand voltages of at least 60 V without breakdown.

6.1.3 Experimental Results

A wide variety of torsional resonators has been built and tested, yielding information on the optimal design of actuators and material properties such as stiffness, resonance properties, internal stresses, and energy dissipation. For a thorough discussion of torsional resonators see [MZSM93,MM96]. In preliminary manipulation experiments, we dispensed carbon particles of up to 10 μm diameter on the actuator and observed the devices in the SEM. At resonance, the particles were quickly tossed off the actuator.

Further manipulation experiments were performed on a prototype array with more than 15,000 individual actuators. While the first generation of torsional resonators [MM96] exhibited only small deflections even at resonance, the current devices were designed for larger displacements at DC voltages. This improvement was achieved by softer torsional springs and relatively larger capacitor grids (Figure 6.1). Out-of-plane motion of up to 5 μm (the maximum kinematically possible deflection) was measured under the SEM. The first set of manipulation experiments was performed with small glass pieces (microscope cover slips) of a few mm^2 size and close to 1 mg mass (note that the mass of an individual actuator measures less than 100 ng). We chose glass because it is flat, rigid, and transparent, which makes it a favorable material to test our devices. However, because of its smooth surface, the motion generated from the actuator array may be reduced because of slip between actuators and moving object. All manipulation experiments were performed in air. Therefore, damping and moisture on the hydrophilic surfaces of Aluminum and glass may also reduce the efficiency of the arrays.

We observed lifting of the glass pieces within the motion range of the actuators (several μm). The objects were also pushed sideways by several 100 μm, by actuating the array at about 50 Hz for a few seconds, so that the actuators repeatedly struck the glass

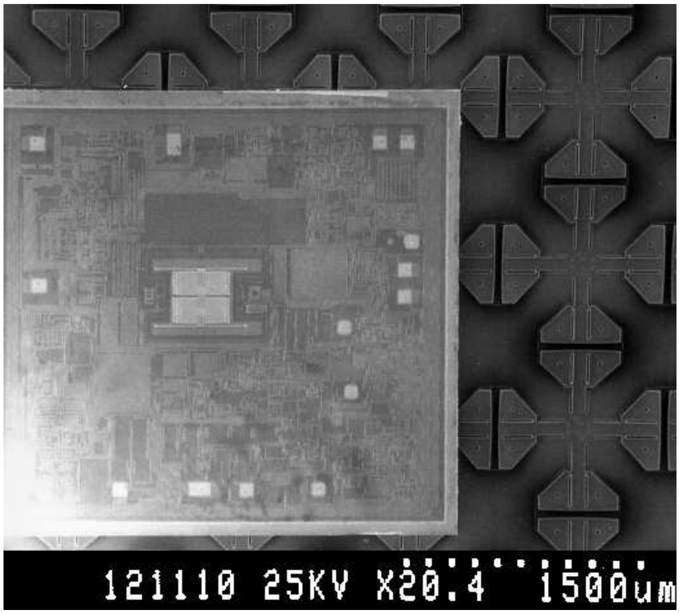

Figure 6.7: Micro cilia device manipulating an ADXL50 accelerometer chip [Ana91] (courtesy of Dr. R. Payne, Analog Devices, Inc.); its size is 2.7 mm squared. The SEM micrograph shows a portion of the cilia chip, which consists of an 8×8 array of motion pixels. The entire device includes four cilia chips that can be controlled independently to generate complex force vector fields. Picture reproduced from [SGD$^+$96].

piece. Larger distances were difficult to achieve with the current arrays because of low yield of our prototype devices. Depending on the shape of the object as well as its position on the actuator array, we also observed resonance between actuators and object at certain actuation frequencies.

A second round of experiments was performed with pieces of paper. Paper is considerably rougher and lighter than a glass piece of equal size. With the current actuator array no motion could be observed. We believe that the main reason for this result is the high surface roughness of paper which we measured as at least 10 - 100 μm, which is more than the height of our current actuator tips, and larger than their range of motion. Low yield in the current actuator arrays also reduces their effectiveness.

6.2 Polyimide Micro Cilia Arrays

We performed distributed manipulation experiments using an organic ciliary array of thin-film polyimide bimorph microactuators (Figure 6.7). These experiments constitute what may be the most convincing evidence to date in support of the theory of programmable force fields. The arrays were designed and built at the Center for Integrated Systems (CIS) at Stanford University by J. Suh, S. Glander, R. Darling, C. Storment, and G. Kovacs [SGD$^+$96], who generously made them available to us, and collaborated with us during our micro manipulation experiments. The actuators employ independent thermal actuation and electrostatic hold-down. Four orthogonally-oriented actuators are integrated into a unit cell called a “motion pixel” [BDMM94b]. The motion pixels are replicated to tile a square area, allowing for precise manipulation of small objects placed on top of the array. Our goal was to implement task-level, sensorless manipulation strategies for macroscopic objects. The tasks of parts-translation, -rotation, -orientation, and -centering were demonstrated using small IC dice. Strategies were programmed in a fine-grained SIMD (Single Instruction Multiple Data) fashion by specifying planar force vector fields. In our experiments, we employed vector fields with potential for parts-orientation and -posing tasks, and the theory was used to predict the equilibrium poses of specific parts (Figure 6.17). The poses predicted by the equilibrium analysis (see Section 2) were observed in our experiments (Figures 6.16-left).

Perhaps surprisingly, the theory has also predicted the existence of pathological fields which do not induce well-behaved equilibria. In particular, the "lower bounds" of Section 3 show that there exist perfectly plausible vector fields which induce *no* stable equilibrium in very simple parts. Although these fields are very simple, they result in limit cycles and quite complex behavior. We implemented such fields on the cilia array. Vector fields *without* potential were employed to cast parts into limit cycles, e.g. "infinite" rotation using a skew-symmetric squeeze field. The predicted behavior (Figure 6.18) for such "unstable" vector fields was also observed (Figure 6.16-right). This shows that rather complex—but potentially useful—behavior can be generated using very simple fields. We believe our experiments validate the theory and suggest the practicality of such minimalist approaches to distributed manipulation.

The programmable vector fields were implemented by actuating the organic cilia in a cyclic, gait-like fashion. Motion in non-principal (e.g. diagonal) directions was effected by a pairwise coupling of the cilia to implement *virtual cilia* and *virtual gaits,* (analogous to the virtual legs employed by Raibert's hopping and running robots [RHPR93]).

These experiments suggest that MEMS actuator arrays are useful for parts-orientation, -posing, -transfer, -singulation, and -sorting. Moreover, the theory of programmable vector fields and virtual gaits gives a method for controlling a very large number of distributed actuators in a principled, geometric, task-level fashion. Whereas many control theories for multiple independent actuators break down as the number of actuators becomes very large, our systems should only become more robust as the actuators become denser and more numerous. This section represents a step towards testing that theory.

6.2.1 Devices and Experimental Setup

Several groups have described efforts to apply MEMS (Micro Electro Mechanical System) actuators to component positioning, inspection, and assembly [PFH90, AOF93, Fuj93, BDMM94b, LW95, for example]. The Stanford cilia chip [SGD+96] employs independent thermal actuation and electrostatic hold-down of an array of 256 thin-film polyimide actuators (Figure 6.8) grouped together in 64 unit cells. Each cell consists of four orthogonally-oriented actuators called a "motion pixel" [BDMM94b], because it can generate motion in various directions, spanning the plane. The motion pixels are replicated to tile a square area, allowing for precise manipulation of small objects placed on top of the array.

Combined Thermobimorph and Electrostatic Actuators. Surface micromachining techniques were used to create the actuators using polyimide as the primary structural material and aluminum as a sacrificial layer. The fabrication process was designed to be compatible with CMOS or BiCMOS circuits which could be pre-fabricated on a silicon substrate.

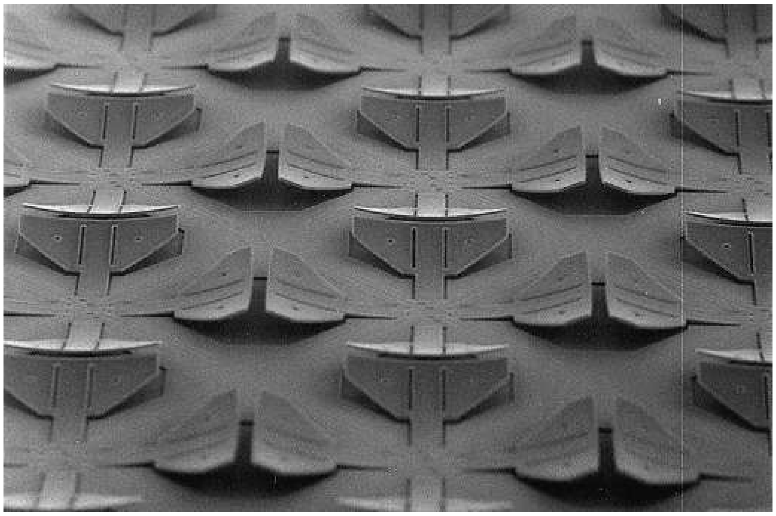

Figure 6.8: Portion of a polyimide cilia array (SEM micrograph). Four orthogonally-oriented actuators are integrated into a motion pixel, which covers a surface area of approximately $1.1 \times 1.1\,mm^2$. Picture reproduced from [SGD$^+$96].

Each actuator cilium consists of two layers of polyimide with different thermal expansion coefficients. The cilium also contains a Titanium-Tungsten (Ti:W) resistive heater loop for thermal actuation, Aluminum electrodes for electrostatic (low-power) hold-down, and a silicon nitride encapsulation/stiffening layer (Figure 6.9).

Since the polyimides were cured at high temperature and the upper polyimide layer has the larger thermal expansion coefficient, the actuators assume an upward, out-of-plane curvature when the process is complete (Figure 6.10). Heating by the Ti:W resistors partially flattens them back towards the substrate. Then, applying a voltage between cilia and substrate electrodes allows a low-power hold-down without heating current. For a detailed description of the fabrication process see [SGD$^+$96].

Vertical and horizontal displacements of the cilia tips are a function of the thermal mismatch in the actuator layers. For room temperature, these values can be calculated as $\delta_v \approx 120\,\mu m$ and $\delta_h \approx 20\,\mu m$ [SGD$^+$96]. Inspection under the scanning electron microscope (SEM) has verified these calculations.

The lifting capacity of an actuator can be estimated as the force required to deflect the actuator's tip to the substrate. The actuator load capacity has been calculated as $F_l = 76\,\mu N$, which gives a force-per-area ratio of $4 \times 76\,\mu N/(1.1\,mm)^2 \approx 250\,\mu N/mm^2$.

Chip Layout. The cilia array is composed of cells (*motion pixels*, each $1.1 \times 1.1\,mm^2$) which contain four orthogonally-oriented actuators (Figure 6.10). On the current cilia chip, the motion pixels are arranged in an 8×8 array which occupies approximately $0.77\,cm^2$ of a $1\,cm^2$ die. The four actuators of each pixel are independently activated by four thermal and four electrostatic control lines. Within each column of eight cells, the thermal elements for a given actuator are wired in series, whereas the electrostatic elements for a given actuator orientation are wired in parallel. The eight columns are then wired in parallel for both the thermal and electrostatic elements. Each electrostatic control line (E1 – E4) applies a potential to the electrodes of all 64 actuators of a given orientation. Similarly, each thermal control line (T1 – T4) runs a current through the Ti:W heaters of all 64 actuators of a

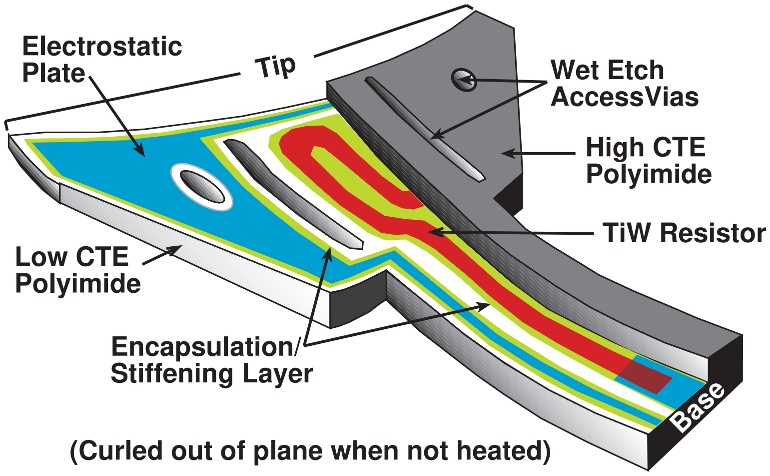

Figure 6.9: Organic thermal and electrostatic microactuator. Half of the upper polyimide and silicon nitride encapsulation/stiffening layer are shown removed along the cilium's axis of symmetry to show details. Picture reproduced from [SGD$^+$96].

given orientation, such that the load resistance for the overall array is the same as that of a single element ($\approx 1500\,\Omega$). A common ground bus provides a current return path for the thermal elements, and a substrate contact provides an opposing plate for the electrostatic elements (Figure 6.11).

Four 8×8 chips were diced and packaged together to make a quad-shaped 16×16 cilia array device, with a total of 1024 cilia. The device itself is attached to a hybrid package which is placed on a heat sink and thermo-electric cooler (Peltier effect module). The total input power to the chip can exceed $4\,W$, and without active cooling the package can become very hot. To observe the experiments, a long working-distance microscope is connected to a CCD camera, and a video cassette recorder is used to monitor and record both the movements of an individual cilium and the objects conveyed by the array.

Controller. The manipulation results described below were accomplished with the cilia array device interfaced to an IBM 486 personal computer. The PC provides speed control via the drive pulse frequency and directional control interactively via keyboard or mouse, or by actuator programs that can be specified using the MEMSA (MEMS Array) language[1] which we developed at Stanford. The control software including the MEMSA interpreter was written in PASCAL.

An 8-bit signal (4 electrostatic control lines E1 – E4, and 4 thermal control lines T1 – T4) is sent via the PC parallel port to D-type flip-flops which activate power transistors. Currently up to 4 cilia arrays can be controlled simultaneously by using a multiplexer with 2 address bits. (Figure 6.12).

6.2.2 Low-level Control: Actuator Gaits

To induce motion on a part that is placed on the array, the cilia are actuated in a cyclic, gait-like fashion. In each cycle, the part is moved in a certain direction by the motion

---

[1] MEMSA (named after MENSA) is the language for smart manipulation surfaces.

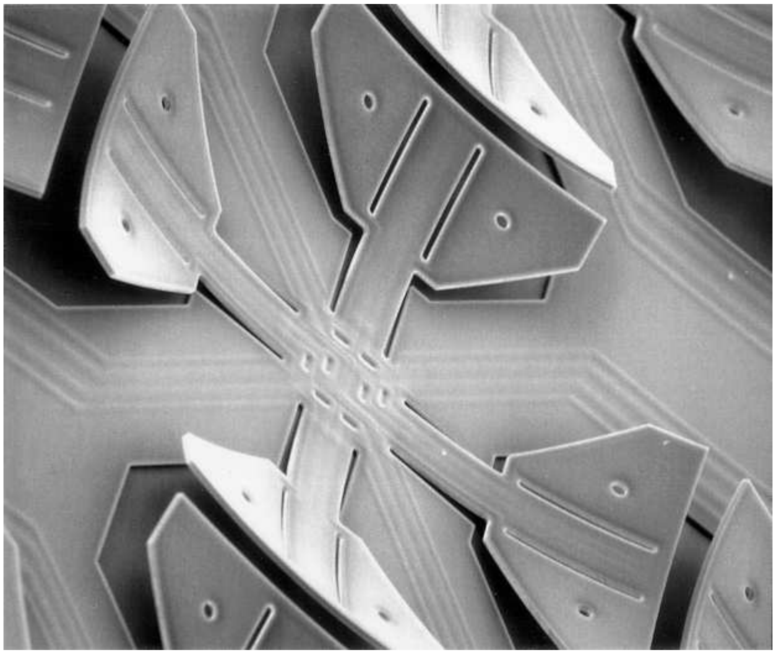

Figure 6.10: Polyimide cilia motion pixel (SEM micrograph). Four actuators in a common center configuration make up a motion pixel. Each cilium is $430\,\mu m$ long and bends up to $120\,\mu m$ out of the plane. Picture reproduced from [SGD$^+$96].

of the actuators that are in contact with it. The speed of the moving part depends on the (horizontal) displacement of the actuators per cycle as well as the frequency of cycle repetition. It also depends on the surface properties and weight of the moving part.

Task: Translation of Parts in Principal Directions. The simplest gait is the *two-phase gait*, in which all actuators of the same orientation repeatedly stroke the part while the remaining actuators are held down. Assuming that the orthogonal cilia within a motion pixel are oriented at the principal compass points, let us use capital letters `NEWS` to denote the North, East, West, and South actuators in the up position, and lower-case letters `news` to denote the actuators in the lowered position. Then the two-phase gait to effect motion in the East direction would be `news neWs` (see Figure 6.13).

The *four-phase gait* consists of four different actuation phases `news neWs nEWs nEws` such that motion is induced during upward as well as downward strokes of the cilia (Figure 6.14, see also [AOF93]). Note that the forces exerted on the moving part depend on the state of the motion pixel: e.g. in the transition from phase 1 to phase 2 the cilium `W` moves up while the opposing cilium `e` remains down. We denote the lateral force exerted on the part in this configuration $f_{\mathtt{W,e}}$. Analogously, during transitions 2–3, 3–4, and 4–1 we observe lateral forces $f_{\mathtt{E,W}}$, $f_{\mathtt{w,E}}$, and $f_{\mathtt{e,w}}$, respectively. $f_{\mathtt{W,e}}$ and $f_{\mathtt{e,w}}$ are in positive x-direction, while $f_{\mathtt{E,W}}$ and $f_{\mathtt{w,E}}$ are negative. Furthermore, from this analysis it follows that $|f_{\mathtt{W,e}}| \gg |f_{\mathtt{e,w}}| \gg |f_{\mathtt{E,W}}| \approx |f_{\mathtt{w,E}}| \approx 0$. Hence we expect a relatively large motion step $\Delta x_{\mathtt{W,e}}$ during transition 1–2, and a smaller step $\Delta x_{\mathtt{e,w}}$ during transition 4–1, while during the other transitions the part remains at its current location. This behavior has been observed in our experiments, where $\Delta x_{\mathtt{W,e}}$ was measured between $3\,\mu m$ and $10\,\mu m$ depending on input power, frequency, surface properties and weight of the part. $\Delta x_{\mathtt{W,e}}$ was usually about

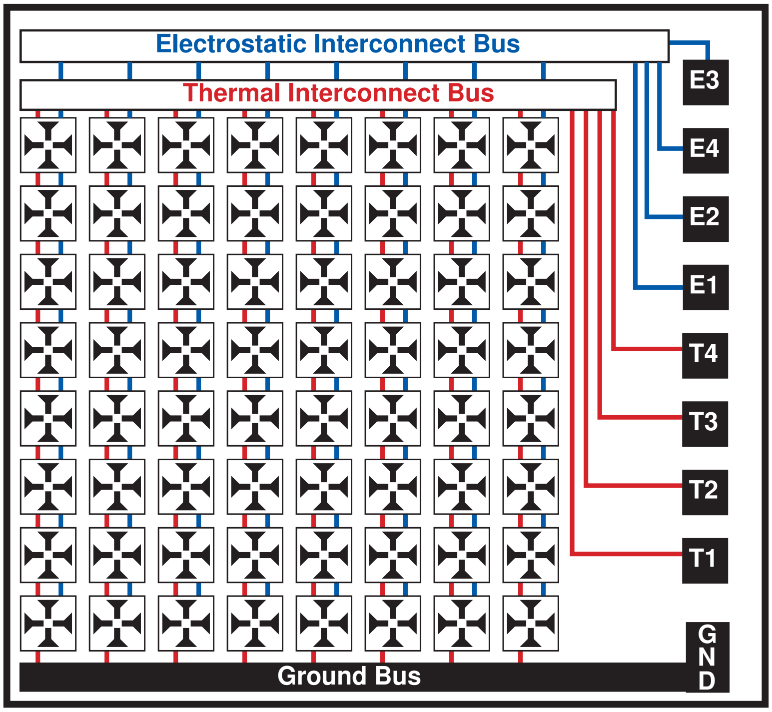

Figure 6.11: Polyimide cilia chip layout (not to scale), consisting of an 8×8 array of motion pixels with thermal and electrostatic control lines. Picture reproduced from [SGD$^+$96].

twice as large as $\Delta x_{\mathtt{e,w}}$.

Task: Translation of Parts in Arbitrary Directions. Motion in non-principal (e.g. diagonal) directions is effected by a pairwise coupling of two cilia of each pixel, implementing *virtual cilia* analogous to Raibert's concept of virtual legs for hopping and running robots [RHPR93]. Hence, several cilia can be coordinated to emulate a virtual cilium, which generates a force corresponding to the vector sum of its components. The diagonal gait to effect motion in the North-East direction would be `news neWS NEWS NEws` where the virtual cilia are `NE` and `WS`. Consequently, we obtain a *virtual gait* that moves the part in a diagonal direction. Note that in a section view through the array looking in the North-West direction, this gait looks virtually identical to the four-phase gait depicted in Figure 6.14.

Motion in arbitrary directions can be induced by alternating gaits that interleave principal (and virtual) gaits of different directions. For example, a translation at 25° from the x-axis requires motion in the y-direction and x-direction at a ratio of $\tan 25^\circ \approx 1:2$. Our control software determines the exact alternating sequence analogously to the Bresenham line scan algorithm [FVDFH96], which rasterizes lines at arbitrary angles, resulting in different fields that are interlaced in time.

Experiments and Results. A large number of translation experiments have been performed in which two-phase and four-phase gaits were used to implement principal and virtual gaits. These experiments show that a first-order dynamical system models the

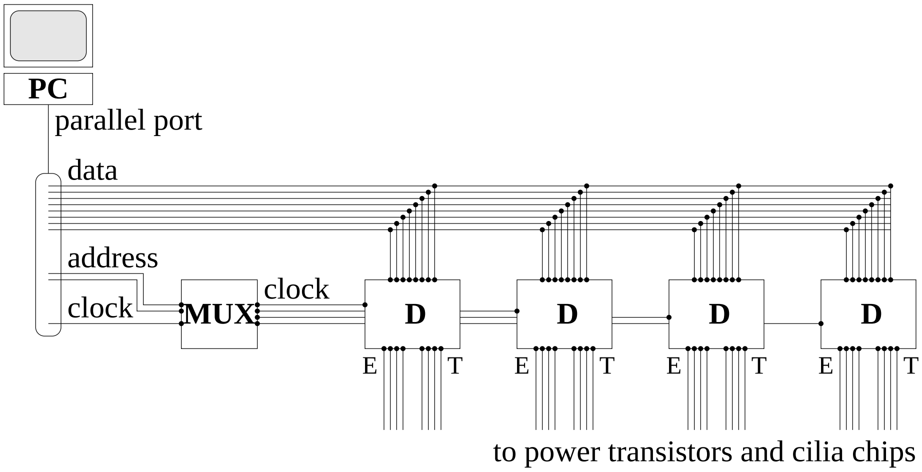

Figure 6.12: Schematic of cilia chip controller: eight data lines, two address lines, and one clock line are used from the PC parallel interface. The data lines provide the activation pattern for the four electrostatic and the four thermal control lines of the cilia chip. One out of four cilia arrays is selected by the address via the multiplexer. New data is stored in the D-type flip-flops at every clock signal.

device-part interaction well. Therefore, when describing and predicting the motion of parts in force vector fields, we have based our theory on a first-order system (see Section 6.2.3).

Silicon chips were moved with a motion resolution of a few μm and speeds up to $200\,\mu m/sec$. Four-phase gaits proved more effective than two-phase gaits, because during the downward motion in the two-phase gait, the part tends to slip backwards. The four-phase gait avoids this effect, because other cilia hold the part in place during the transition 3–4. In the subsequent downward motion in the transition 4–1, the part is also moved forward (Figure 6.14).

The diagonal gait also has the lowest power consumption (not considering electrostatic hold-down), due to the fact that its duty cycle for cilia hold-down is lowest (50 %), compared to 75 % for the principal four-phase gait, and 87.5 % for the two-phase gait.

As expected, diagonal (virtual) gaits induced the largest and fastest motion because all four cilia of each pixel were activated, whereas in principal gaits only two cilia are actively used, while the others have to be held down continuously (Figure 6.14).

6.2.3 High-level Control: Vector Fields

We believe that vector fields can be used as an *abstraction barrier* between applications requiring array micro-manipulation and MEMS devices implementing the requisite mechanical forces. That is, applications such as parts-feeding can be formulated in terms of the vector fields required. This then serves as a specification which the underlying MEMS device technology must deliver. Conversely, the capabilities of MEMS array technologies for actuation can be formulated in terms of the vector fields they can implement. For example, limitations in force magnitude are naturally expressed in vector field terms, as “small” vector fields. Restrictions in directional selectivity and magnitude control can also be manifested as restrictions on the vector field abstraction (resulting in discretization in orientation or modulus). This means that MEMS designers can potentially ignore certain

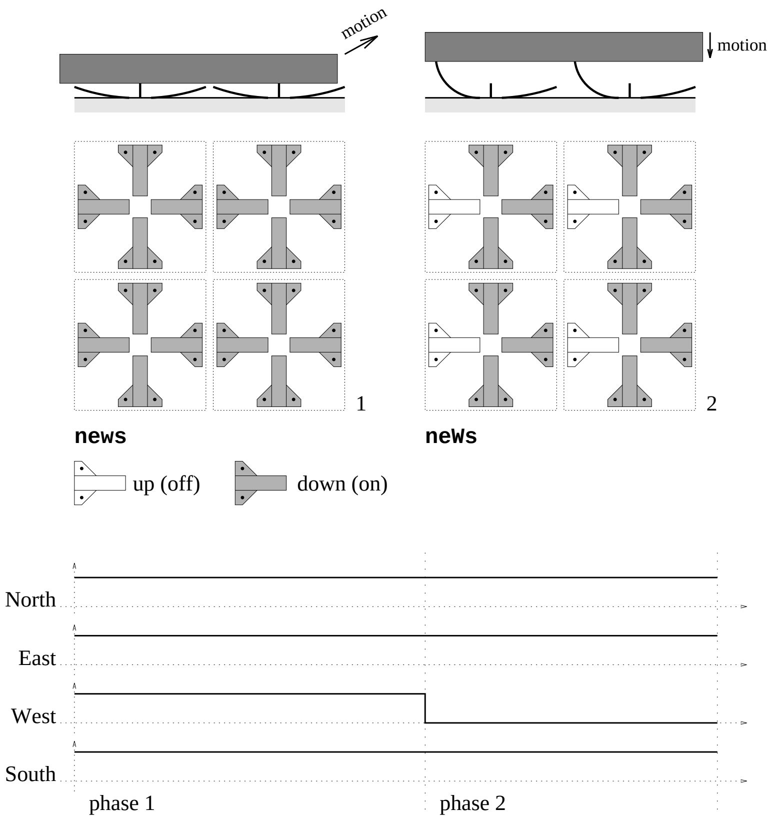

Figure 6.13: Two-phase gait. The `W` actuator is repeatedly switched on and off, while the other actuators always remain on, resulting in a `news neWs` sequence.

details of the application process, and instead focus on matching the required vector field specification. Then, once the capabilities of MEMS actuator arrays were published as vector fields and tolerances, an application designer could look in a catalog to choose a device technology based on the field specification it promises to implement. This would free application engineers from needing to know much about process engineering, in the same way that software and algorithm designers often abstract away from details of the hardware. Such an abstraction barrier could permit hierarchical design, and allow application designs with greater independence from the underlying device technology. At the same time, abstraction barriers could allow MEMS array technologies to be designed simultaneously with the (abstract) vector field control. This development pattern could be similar to the concurrent design of VLSI processors with their compilers, as is common in computer architecture.

In the remainder of this section, we give brief summaries of theoretical results (see Sections 2 – 5) and compare them to experiments that we performed with the micro cilia array device.

6.2.4 Squeeze Fields

In Section 2 we proposed a family of control strategies called *squeeze fields* and a planning algorithm for parts-orientation (see Figures 2.1 and 2.2).

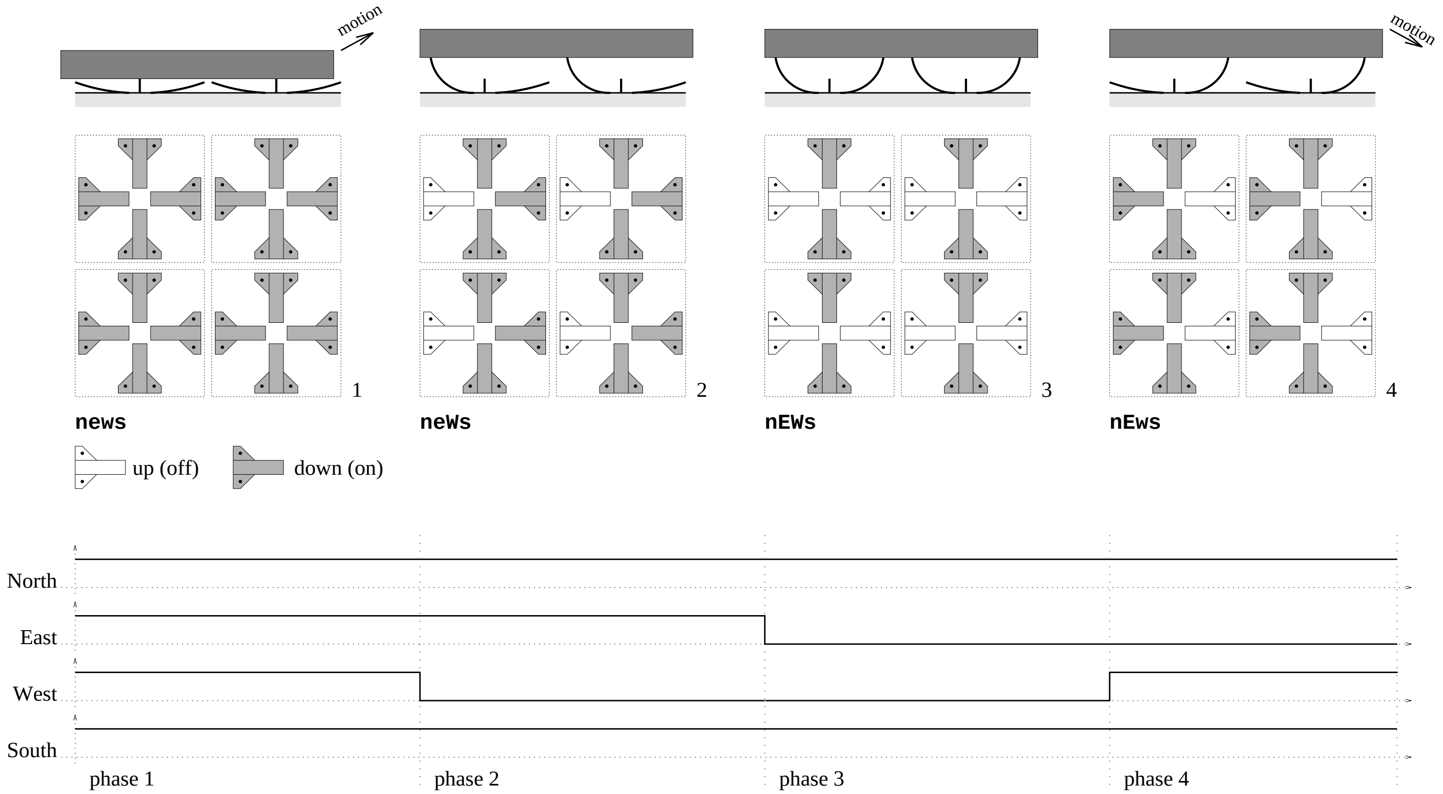

Figure 6.14: Four-phase gait consisting of the four-pattern sequence `news neWs nEWs nEws`.

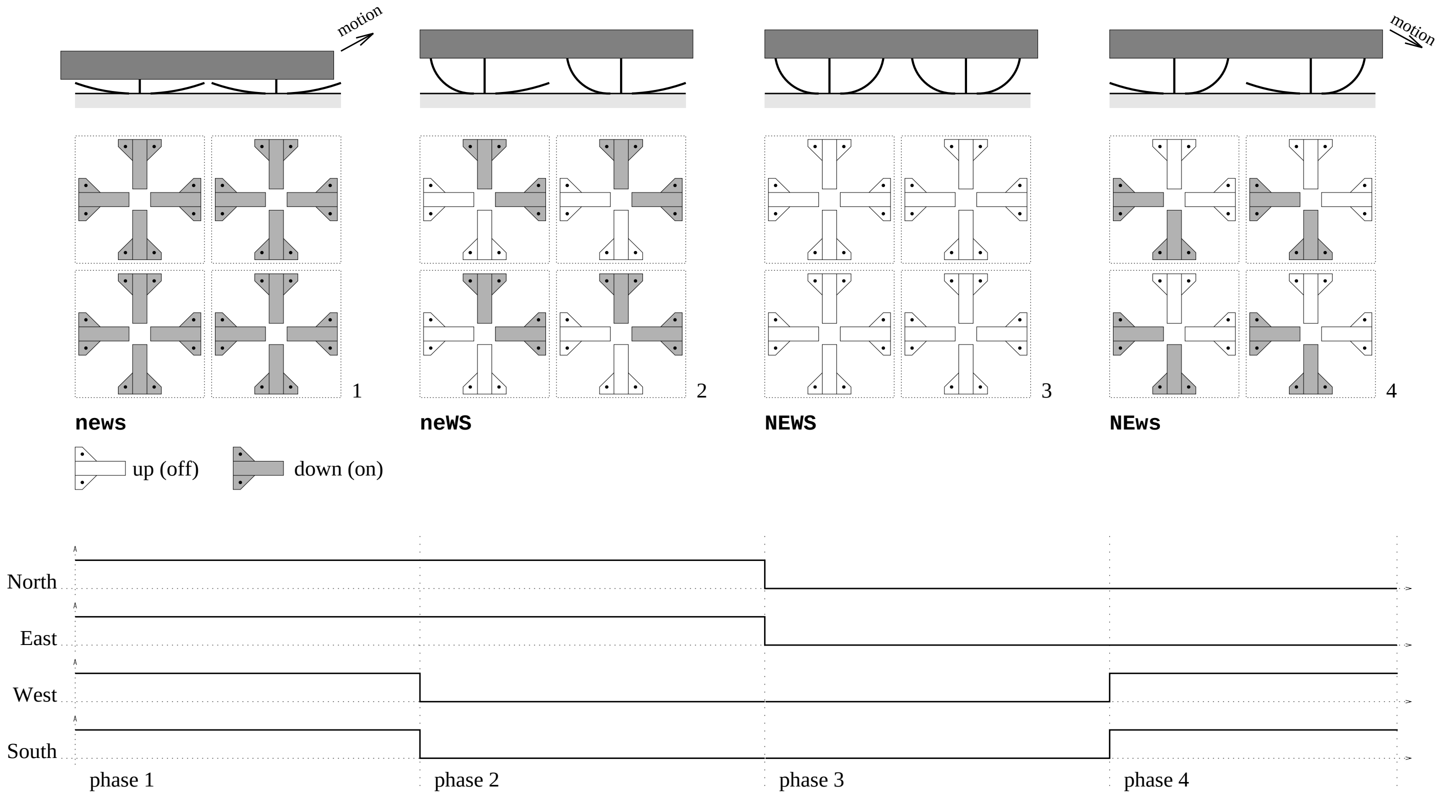

Figure 6.15: Diagonal (virtual) gait consisting of the four-pattern sequence `news neWS NEWS NEws`. The `N` and `E` cilia, and the `W` and `S` cilia, are coupled to form virtual cilia in North-East and South-West directions.

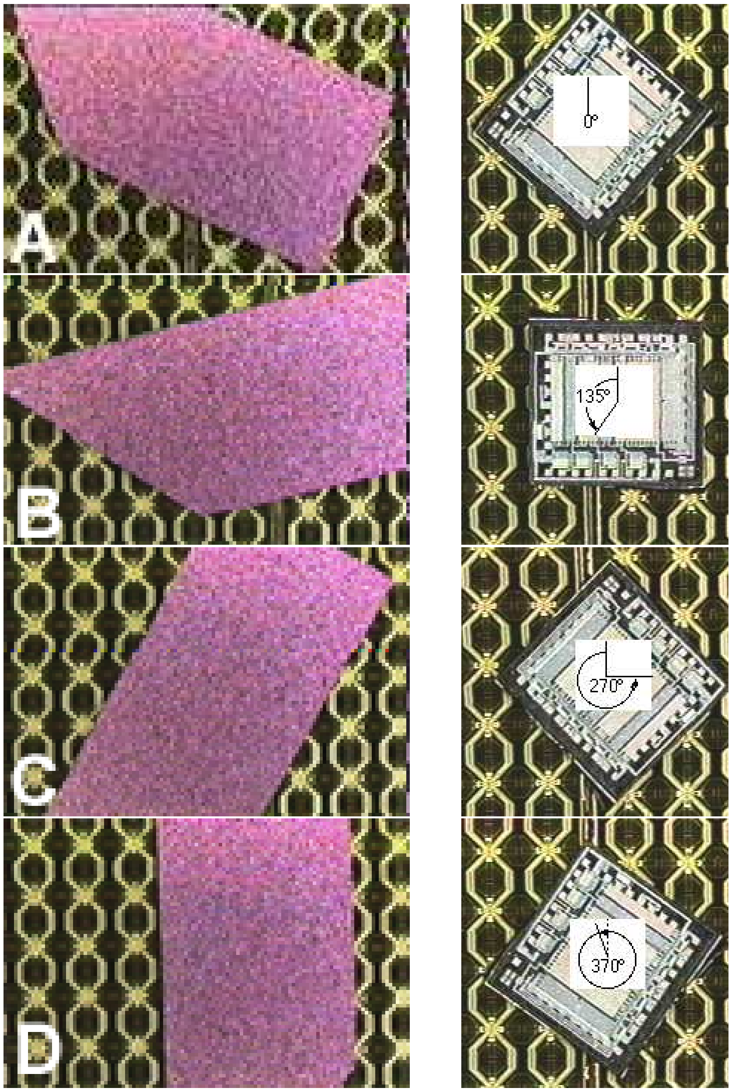

Figure 6.16: Manipulation of silicon chips in programmable vector fields induced by a micro cilia array (microscope video images).
Left: The chip is aligned with the vertical squeeze line of the field.
Right: Rotating a square-shaped chip counterclockwise in a skewed squeeze field. Pictures reproduced from [SGD+96].

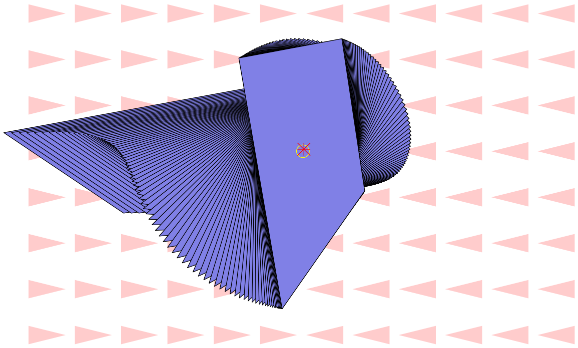

Figure 6.17: Simulation of alignment task with a squeeze field.

Definition 1 (Section 2.1) *Assume l is a straight line through the origin. A* squeeze field *f is a two-dimensional force vector field defined as follows:*

1. *If $z \in \mathbb{R}^2$ lies on l then $f(z) = 0$.*
2. *If z does not lie on l then $f(z)$ is the unit vector normal to l and pointing towards l.*

We refer to the line l as the *squeeze line*, because l lies in the center of the squeeze field.

Assuming quasi-static motion, an object will translate and rotate to an equilibrium configuration, as characterized in Figures 2.2, 2.3, and 2.4. This assumes a uniform force distribution over the surface of P, which is a reasonable assumption for a flat part that is in contact with a large number of elastic actuators.

Definition 2 (Section 2.1) *A part P is in* translational equilibrium *if the forces acting on P are balanced. P is in* orientational equilibrium *if the moments acting on P are balanced.* Total equilibrium *is simultaneous translational and orientational equilibrium.*

Claim 42 [BDM96b] *Every squeeze field f has potential, of the form $U(\mathbf{z}) = \int_\alpha f \cdot ds$, where α is an arbitrary path to $\mathbf{z}$ from a fixed reference point. If $d_{\mathbf{z}}$ denotes the perpendicular distance of $\mathbf{z}$ from the squeeze line, then $U(d_{\mathbf{z}}) = |d_{\mathbf{z}}|$.*

Proof: Follows from Definition 1 and Proposition 17, Section 4. □

Proposition 4 (Section 2.2) *Let P be a polygon whose interior is connected. There exist $O(k\,n^2)$ bisectors such that P is in equilibrium when placed in a squeeze field such that the bisector coincides with the squeeze line. n is the part complexity measured as the number of polygon vertices. k denotes the maximum number of polygon edges that a bisector can cross.*

If P is convex, then the number of bisectors is bounded by $O(n)$.

Equilibria can be calculated analytically or numerically, see Section 2.2 for details.

Theorem 9 (Section 2.3) *Let P be a polygon whose interior is connected. There exists an alignment strategy consisting of a sequence of squeeze fields that uniquely orients P up to symmetries.*

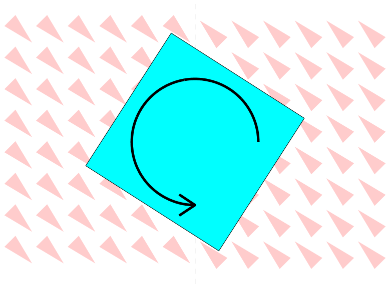

Figure 6.18: Unstable square-shaped part in a skewed squeeze field ($\epsilon = -1$). The square with center on the squeeze line will rotate indefinitely. Moreover, it has *no* stable equilibrium in this field.

Task: Orienting and Aligning Parts

If a part is placed in a squeeze field, it will translate and rotate until a stable equilibrium is reached (Proposition 4). Parts may exhibit several equilibria, hence after one squeeze the part orientation may be ambiguous. This ambiguity can be removed with the strategies of Theorem 9: by executing a sequence of squeezes at particular angles, the part is uniquely oriented (see Figure 2.1, for details see Section 2.3).

Experiments and Results. The long, thin part depicted in Figures 6.16-left and 6.17 exhibits a unique stable equilibrium (modulo 180° field symmetry). When placed in a squeeze field, its longitudinal axis aligns with the squeeze line. This dynamical process is predicted by simulation in Figure 6.17, and verified in experiment (see Figure 6.16-left). This part alignment experiment has also been performed with similar results for several other small pieces of glass and silicon of a few *mm* length and several *mg* of mass.

6.2.5 Skewed Squeeze Fields

Definition 11 (Section 3) *A* skewed field f_S *is a vector field given by* $f_S(x, y) = -\text{sign}(x)(1, \epsilon)$, *where* $0 \neq \epsilon \in \mathbb{R}$.

Proposition 24 (Section 4.2) *No skewed squeeze field has a potential.*

Proposition 12 (Section 3) *A skewed field induces* no *stable equilibrium on a disk-shaped part.*

Proposition 13 (Section 3) *A diagonally skewed field* ($\epsilon = \pm 1$) *induces* no *stable equilibrium on a square-shaped part.*

Task: Rotating Parts

According to Propositions 12 and 13, certain parts will rotate indefinitely in skewed squeeze fields (Figure 6.18). Note that even though our cilia device has more degrees of freedom, two areas of constant force are sufficient to implement a skewed field, resulting in a very simple task-level rotation strategy. In particular, the rotation algorithm resulting from the application of skew-symmetric squeeze fields is considerably simpler than rotation algorithms proposed in the MEMS literature (for example, the vortices suggested by Fujita [Fuj93] or by Liu and Will [LW95]). Vortices require at least four areas of the array to be pushing in different directions. That is, vortices can be implemented using four triangular or rectangular regions, upon each of which the vector field is constant. Skewed fields perform the same task with only two regions of constant force.

Experiments and Results. Figure 6.16-right shows video frames of a $3 \times 3\,mm^2$ IC chip rotating on the squeeze line of a skewed field. During the experiments, several full rotations of the part were performed.

6.2.6 Radial Fields

Definition 22 (Section 4.2) *A* unit radial field *R is a two-dimensional force vector field such that $R(\mathbf{z}) = -\mathbf{z}/|\mathbf{z}|$ if $\mathbf{z} = 0$, and $R(\mathbf{0}) = \mathbf{0}$.*

Proposition 23 (Section 4.2) *A unit radial field has a potential, $U(\mathbf{z}) = |\mathbf{z}|$.*

Proposition 35 (Section 5.1) *Given a polygonal part P in a radial field f, there exists a unique* pivot point *v of P such that P is in equilibrium if and only if v coincides with the center of the radial field.*

Task: Centering Parts

Radial fields can be used to center a part. With the current four-quadrant cilia device, we have implemented an approximation of an ideal radial field similar to the field in Figure 1.1-b. Note that this approximate radial field has a potential.

Experiments and Results. Small silicon and glass parts were centered using our cilia device. In this experiment, high positioning accuracy (in the μm range) was hard to achieve, because the center of the radial field coincides with the the location of the dice edges. Manual packaging of the four cilia chips leaves small gaps and non-planarities at these junctions. The next generation cilia device will overcome this problem, because it will allow us to implement the radial field with a single chip. Furthermore, because of its full pixel-wise programmability, the new chip will allow us to closely approximate ideal radial fields.

6.2.7 Conclusions

Small parts (sizes in the mm range) were successfully manipulated using an organic MEMS cilia array. The experiments demonstrated parts-translation, -rotation, -orientation, and -centering of IC dice, and other small pieces of silicon and glass, as predicted by the *theory of programmable vector fields*. With this theory, force vector fields can be cascaded into

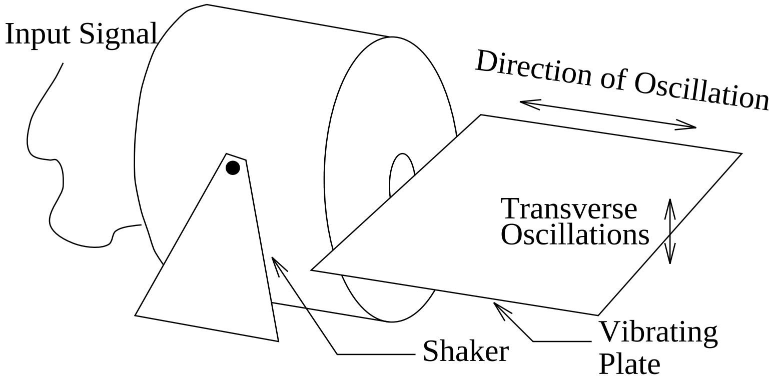

Figure 6.19: Schematic of experimental setup 1: A 50 $cm \times 40$ cm aluminum plate is forced to oscillate horizontally by the shaker armature. The forced oscillation causes a transverse vibration of the plate.

multi-step strategies. It has been shown that these open-loop strategies can sort as well as uniquely position and orient parts without sensor feedback.

Based on our experiments, a new, improved array is currently being fabricated. Benefits of the new, fully programmable cilia array will include among others the implementation of sequences of squeeze fields, nearly-ideal radial fields, and sensorless sorting strategies. The new design will also overcome current problems due to manual packaging.

6.3 Vibratory Plate Parts Feeders

In this section, we explore how controlled vibration can be used for a new device setup to systematically feed planar parts. The idea is to generate and change dynamic modes in a plate by varying applied frequencies. Depending on the frequency of vibration and the boundary conditions, nodes of different shapes are formed (this was experimentally studied by Chladni [Chl87]). If planar parts are put on this vibrating plate, they move to the node, and end up in a stable orientation [BBG95].

6.3.1 Setup and Calibration

Figures 6.19 and 6.20 are schematics of the experimental setup, which consists of an aluminum plate forced to oscillate in two different configurations. The shaker is a commercially available[2] electrodynamic vibration generator, with a linear travel of $0.02\,m$, and capable of producing a force of up to $500\,N$. The input signal, specifying the waveform corresponding to the desired oscillations, is fed to a single coil armature, which moves in a constant field produced by a ceramic permanent magnet in a center gap configuration.

In the first configuration (Figure 6.19), the plate is attached to the shaker armature such that it is forced to vibrate in the longitudinal direction. For low amplitudes and frequencies, the plate moves longitudinally with no perceptible transverse vibrations. However, as the frequency of oscillations is increased, transverse vibrations of the plate become more pronounced. The resulting motion is similar to the forced transverse vibration of a rectangular plate, clamped on one edge and free along the other three sides.

---

[2] Model VT-100G, Vibration Test Systems, Akron, OH, USA.

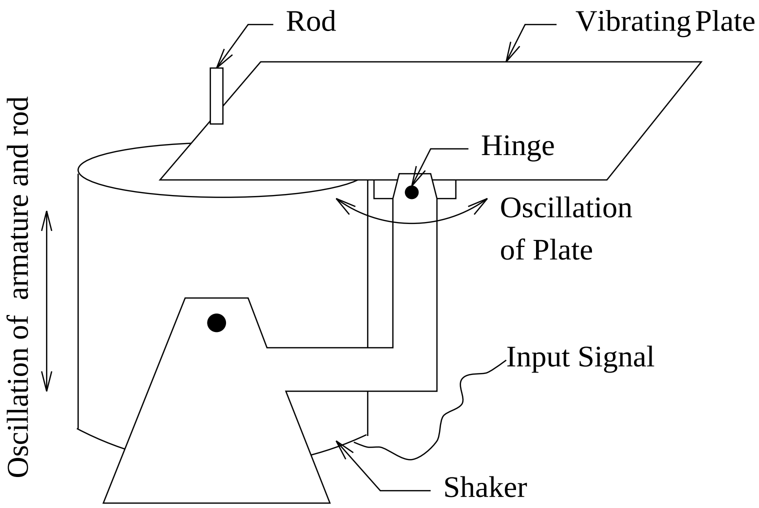

Figure 6.20: Schematic of experimental setup 2: The aluminum plate is hinged and can oscillate about an axis in its middle.

The nodes for these transverse oscillations can either be obtained theoretically (Rayleigh [Ray45], Timoshenko [Tim40]), or experimentally using the technique originally pioneered by Chladni [Chl87]. By sprinkling small sized particles[3] on a vibrating surface, the nodes can be experimentally identified as the regions where the particles tend to collect. The dynamics of "collecting" at the nodes is important in determining the effective force field that leads to the orienting and localization effect of our device, and is discussed in more detail in Section 6.3.3.

For the configuration in Figure 6.19, the location and shape of the node depends on the frequency of vibration. Figure 6.21 shows the experimental determination of the nodes for frequencies of $60\,Hz$ and $100\,Hz$.

The second configuration (Figure 6.20) forces the plate to undergo transverse vibrations such that the resulting shape of the node, and its location, are independent of the forcing frequency. The plate is hinged about an axis situated midway between, and parallel to, two of its sides. A rod connected to the armature of the shaker forces the plate to an oscillatory motion about the hinged axis. As expected, experimental determination shows that except for a slight distortion due to the effect of clamping at the rod, the node lines up with the hinge axis (Figure 6.22).

The second setup is run at lower frequencies, to ensure that only the mode where the plate oscillates about the hinge axis is excited. If we increase the operating frequency, modes corresponding to transverse vibration of a plate, clamped at the point of attachment to the rod and the hinged ends, become dominant, and the node shape gets complicated. This effect can be seen at $20\,Hz$ (Figure 6.22), where the node shows a tendency to get "pulled" towards the point where the plate is clamped to the rod.

6.3.2 Behavior of Planar Parts

---

[3] Chladni used sand, we use Urad lentils to get a better contrast on video.

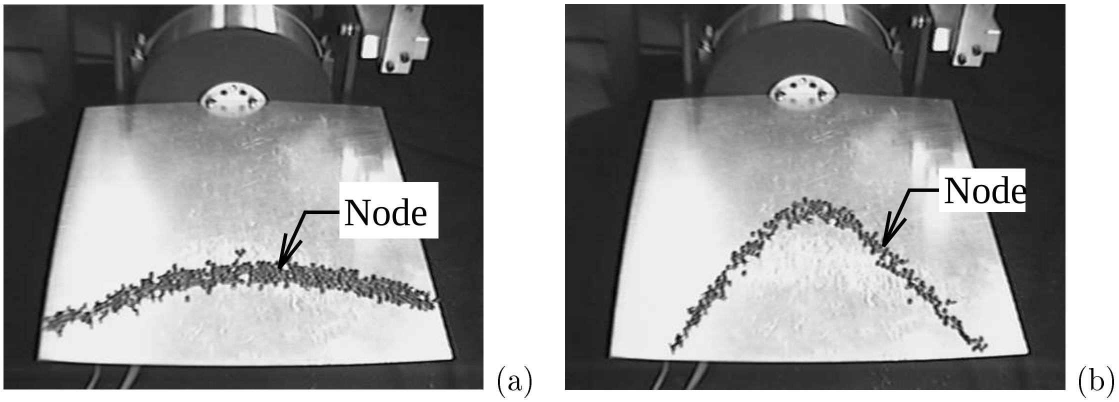

Figure 6.21: Experimentally determined nodes at (a) $60\,Hz$ and (b) $100\,Hz$, for experimental setup 1 (see Figure 6.19). After vibrating the plate for a short time, the particles form Chladni figures, which indicate the location of the vibrational nodes.

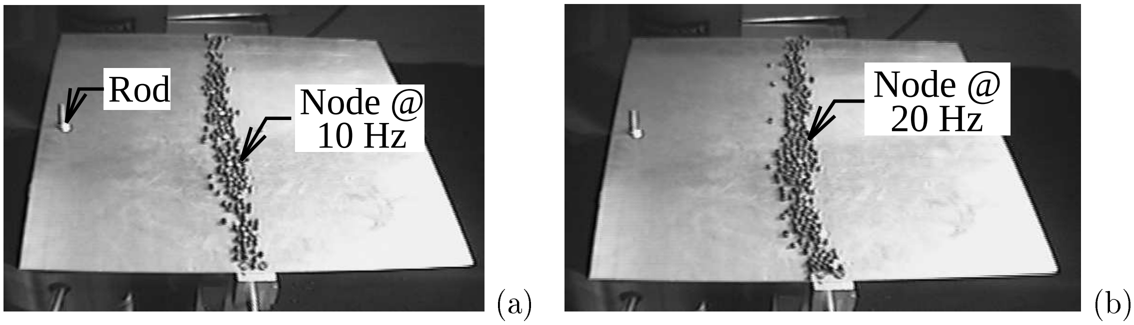

Figure 6.22: Experimentally determined vibrational nodes at (a) $10\,Hz$ and (b) $20\,Hz$, for experimental setup 2 (see Figure 6.20).

If we put planar parts on the vibrating surface, there is a marked tendency for them to move towards the node and end up in one of a finite number of stable orientations. We observe the following features over a wide range of frequencies in both the experimental setups:

- From all initial positions on the plate, the objects move towards the node. They end up in a stable position around some point on the node, which depends on the initial position of the object.
- As the object approaches the node (as we show later, after some portion of it crosses the node), there is a tendency for it to rotate until it reaches one of a finite number of stable orientations.

Figure 6.23 shows two planar shapes, a triangle and a trapezoid, after they have reached their stable position and orientation for the setup in Figure 6.19. To better illustrate the orienting effect, the curve showing the node has been drawn by hand. Figure 6.24, similarly shows the stable position of the planar parts for the second setup (Figure 6.24).

Over the large number of experimental runs performed, there are a couple of qualitative observations describing the ease and speed with which the parts get into a stable configuration:

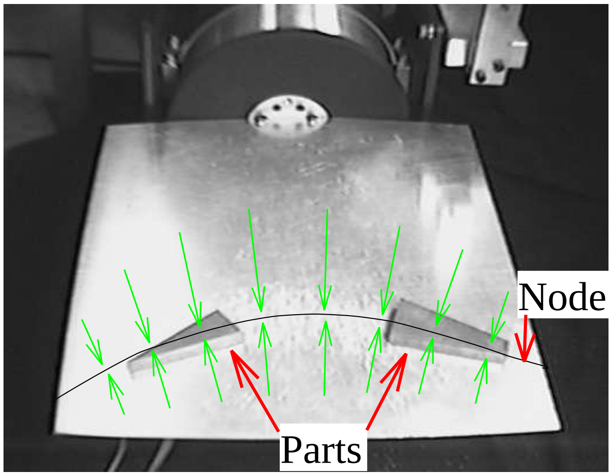

Figure 6.23: Stable position of planar parts in experimental setup 1, at a frequency of 60 Hz. The node is marked according to Figure 6.21.

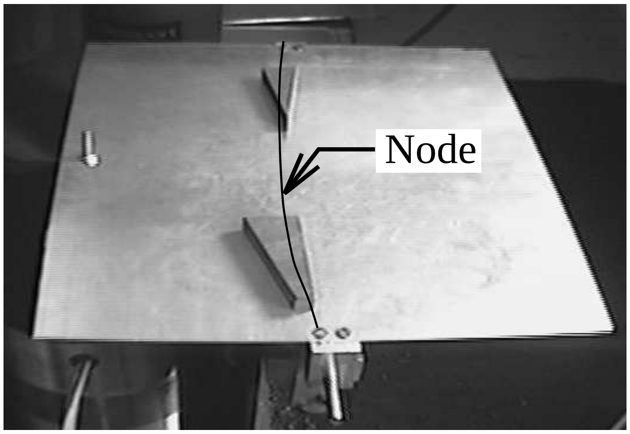

Figure 6.24: Stable position of planar parts in experimental setup 2, at a frequency of 20 Hz. The node is marked according to Figure 6.22.

- At higher frequencies of oscillation, both the velocity of the part towards the node, and the rate of orientation, are relatively faster.
- Objects with a higher degree of rotational asymmetry get into a stable orientation more easily.

Although the location of the node is better identified in the second setup, the lower operating frequencies make the localization of the part at the node, and the corresponding orienting behavior, much slower. These tradeoffs are design considerations that need to be investigated separately.

6.3.3 Dynamics of Particles and Planar Parts on a Vibrating Plate

The underlying dynamics that causes the objects placed on a vibrating surface to move towards the node give rise to an effective force field. In order to develop a theory for using our device as a viable method for sensorless manipulation, it is important to determine the genesis and variation of this force field over the vibrating plate.

Chladni Figures

When particles are spread on a vibrating surface, they collect at the nodes, resulting in patterns known as Chladni figures (after Chladni [Chl87], see Figures 6.21 and 6.22). Rayleigh [Ray45] describes the motion of the particles towards the nodes in the following words – "*the movement to the nodes is irregular in its character. If a grain be situated elsewhere than at a node, it is made to jump by a sufficiently vigorous transverse vibration. The result may be a movement either towards or from a node; but after a succession of such jumps the grain ultimately finds its way to a node.*"

The forces that cause the particles to move to the node act on any object placed on the vibrating surface, generating an effective force field. The underlying dynamics of this phenomenon are still poorly understood. In Appendix C we give an approach towards an analytical model for the more tractable case of the planar motion of a particle bouncing on a string in transverse vibration.

Results of our experiments indicate that the forces at each location of the plate can be thought of as proportional to the amplitude of vibration, and perpendicular to the sinusoidal "envelope" surface of the oscillating plate. They are also proportional to the oscillation frequency, and the coefficient of restitution. The first-order approximation near the nodes can be used to show that the force field is perpendicular to and points towards the node, and its magnitude varies linearly with the distance from the node. To simulate the field far from the node, we approximate its magnitude by a sine variation normal to the node curve.

Motion and Equilibria of Planar Parts

The case of general objects on the plate is more complicated than individual particles, because the determination of the point on the object that undergoes impact, and the resulting impulses, are both difficult problems to solve. For our analysis, we ignore effects such as rolling and tilting of the parts and assume that the contact geometry remains constant over the impacts.

Instead of dealing with general parts, for which a more complicated formulation would be required, we look at planar parts for which deriving the stability properties is relatively straightforward.

We can consider the planar parts as a conglomeration of "particles," each of which interacts with the plate and experiences the effective force field discussed in Section 6.3.3. The forces have to be summed up over the area of contact, giving a specific force (per unit area), f, that acts at every point of the planar object.

Let P be the planar part in contact with the vibrating plate, and c denote the center of mass of P. The total ("lifted") force f_P and moment M_P around c can be obtained by integrating the force field f over the contact surface of P:

$$\begin{aligned} f_P &= \int_P f \, dA && (6.1) \\ M_P &= \int_P (r - c) \times f \, dA && (6.2) \end{aligned}$$

Consider a part P on the vibrating plate. We assume that a first-order dynamical system describes the motion of P on the plate. In a first-order system, the velocity of a part is directly proportional to the force acting on it. Hence, an *equilibrium* is a placement of P such that P remains stationary. In an equilibrium, the force and moment acting on P are balanced. This *equilibrium condition* is met when the lifted force f_P and moment M_P (Equations (6.1) and (6.2)) are both zero.

We have made a series of assumptions to suggest that a force field exists for parts on a planar plate. Our experimental results indicate that they are good engineering assumptions when we observe the system over time, due to an averaging effect caused by the vibration of the plate. An "exact" modeling of the impact dynamics between part and plate, even though possible (see e.g. [MC95]), is not necessary for our purposes.

Chapter 7

Conclusions and Open Problems

7.1 Universal Feeder-Orienter (UFO) Devices

It was shown in Proposition 4 that every connected polygonal part P with n vertices has a finite number of stable orientation equilibria when P is placed into a squeeze field S. Based on this property we were able to generate manipulation strategies for unique part alignment. We showed in Section 5.1 that by using a combined radial and squeeze field $R + \delta S$, the number of equilibria can be reduced to $O(k\,n)$. Using elliptic force fields $f(x, y) = (\alpha x, \beta y)$ such that $\alpha \neq \beta$ and $\alpha, \beta \neq 0$, this bound can be reduced to two [Kav95,Kav97]. An "inertial" squeeze field $f(x, y) = (-\text{sign}(x)x^2, 0)$ uniquely orients a part modulo field symmetry π. In a stable equilibrium, the part's major principal axis of inertia lines up with the squeeze line to minimize the second moment of inertia.

Does there exist a *universal field* that, for every part P, has only one unique equilibrium (up to part symmetry)? Such a field could be used to build a *universal parts feeder* [AE96] that uniquely positions a part without the need of a clock, sensors, or programming.

We propose a combined radial and "gravitational" field $R + \delta G$ which might have this property. δ is a small positive constant, and G is defined as $G(x, y) = (0, -1)$. This device design is inspired by the "universal gripper" in [AE96]. Such a field could be obtained from a MEMS array that implements a unit radial force field. Instead of rectangular actuators in a regular grid, triangular actuators could be laid out in a polar-coordinate grid. The array could then be tilted slightly to obtain the gravity component. Hence such a device would be relatively easy to build. Alternatively, a resonating speaker, or a vibrating disk-shaped plate that is fixed at the center, might be used to create a radial force field. Extensive simulations show that for every part we have tried, one unique total equilibrium is always obtained. We are working toward a rigorous proof of this experimental observation.

7.2 Magnitude Control

Consider an array in which the *magnitude* of the actuator forces cannot be controlled. Does there exist an array with constant magnitude in which all parts reach one unique equilibrium? Or can one prove that, without magnitude control, the number of distinct equilibria is always greater than one?

7.3 Geometric Filters

This thesis focuses mainly on sensorless manipulation strategies for *unique positioning* of parts. Another important application of programmable vector fields are *geometric filters*, which would be useful for sorting and singulation of parts. Figure 1.1 shows a simple filter that separates smaller and larger parts. We are interested in the question *Given n parts, does there exist a vector field that will separate them into specific equivalence classes?* For example, does there exist a field that moves small and large rectangles to the left, and triangles to the right? In particular, it would be interesting to know whether for any two different parts there exists a sequence of force fields that will separate them.

7.4 Force Field Computers

In this thesis we have demonstrated that even with a rather limited vocabulary of simple force fields, useful and quite complex tasks such as sensorless posing or sorting of parts can be performed. It might be possible that force fields could be used to solve certain classes of problems, by encoding them in particular force fields, part shapes, and initial and goal poses, resulting in a "force field computer" that provides a physical implementation of the problem. Identifying the class of encodable problems might yield deeper insights into the complexity of parts manipulation with force vector fields.

7.5 Performance Measures

Are there performance measures for how fast (in real time) an array will orient a part? In some sense the actuators are fighting each other (as we have observed experimentally) when the part approaches equilibrium. For squeeze grasps, one measure of "efficiency," albeit crude, might be the integral of the magnitude of the moment function, i.e., $\int_0^{2\pi} |M(\theta)|\, d\theta$. The issue is that if, for many poses, $|M(\theta)|$ is very small, then the orientation process will be slow. Better measures are also desirable.

7.6 Uncertainty

In practice, neither the force vector field nor the part geometry will be exact, and both can only be characterized up to tolerances [Don89]. This is particularly important at micro scale. Within the framework of potential fields, we can express this uncertainty by considering not one single potential function U_P, but rather *families of potentials* that correspond to different values within the uncertainty range. Bounds on part and force tolerances will correspond to limits on the variation within these function families. An investigation of these limits will allow us to obtain upper error bounds for manipulation tasks under which a specific strategy will still achieve its goal.

A family of potential functions is a set $\{U_\alpha : \mathcal{C} \to \mathbb{R}\}_{\alpha \in J}$ where J is an index set. For example, we may start with a single potential function $U : \mathcal{C} \to \mathbb{R}$, and define a family of potential functions $\mathcal{F}(U, \epsilon, z)$ as $\{U_\alpha : \mathcal{C} \to \mathbb{R} \mid ||U_\alpha(p) - U(p)||_z < \epsilon\}$ for some ϵ and norm z. This is analogous to defining a neighborhood in function space, using e.g. the compact-open topology.

When we differentiate a family of potential fields (using the gradient) we obtain a differential inclusion instead of a differential equation. So if $\mathcal{F}(u) = \mathcal{F}(u, \epsilon, z)$, then $\nabla\mathcal{F}(u) = \{\nabla U_\alpha\}_{\alpha \in J}$.

When considering families of potentials, the equilibrium may be known to lie only within a set E_i, although we may know that it is always a point in E_i. If the sets E_i are of a small diameter less than some $\epsilon > 0$, our algorithms could be extended to handle ϵ-approximations.

As a more general approach, we propose an algorithm based on *back-projections*: For a given part, let $B_{F_i}(G) \subset \mathcal{C} = \mathbb{R}^2 \times \mathbb{S}^1$ be the back-projection [LPMT84] of the set G under F_i, where $G \subset \mathcal{C}$, and F_i is a family of fields on $\mathbb{R}^2$. Then we wish to calculate a sequence of fields $F_1, F_2, \cdots, F_k$ such that $B_{F_1}(B_{F_2}(\cdots B_{F_k}(G)\cdots)) = \mathcal{C}$, where G is a single point in $\mathcal{C}$ (cf. [LPMT84,EM88,Bro88,Don89,Bri92]).

7.7 Output Sensitivity

We have seen in Sections 2.1, 5.1, and 5.2 that the efficiency of planning and executing manipulation strategies critically depends on the number of equilibrium configurations. Expressing the planning and execution complexity as a function of the number of equilibria E, rather than the number of vertices n, is called *output sensitive analysis.* In practice, we have found that there are almost no parts with more than two distinct (orientation) equilibria, even in squeeze fields. This is far less than the $E = O(k\,n^2)$ upper bound derived in Section 2.2. If this observation can be supported by an exact or even statistical analysis of part shapes, it could lead to extremely good expected bounds on plan length and planning time, even for the less powerful strategies employing manipulation grammars (note that the complexity of the manipulation grammar algorithm in Proposition 41 is output-sensitive).

7.8 Discrete Force Fields

For the manipulation strategies described in this thesis we assume that the force fields are continuous, i.e. that the generated forces are dense compared to the moving part (assumption DENSITY in Section 2.3). When manipulating very small parts on microactuator arrays, this condition may be only approximately satisfied. We are interested in the limitations of the continuous model, and we would like to know the conditions under which it is necessary to employ a different, discrete model of the array that takes into account individual actuators, as well as the gaps between actuators. In [BDMM94b] we propose a model for the interaction between parts and arrays of individual actuators based on the theory of limit surfaces [GR88,GRP91].

7.9 Resonance Properties

Is it possible to exploit the dynamic resonance properties of parts to tune the control signal of the array or plate to perform efficient dynamic manipulation?

7.10 3D Force Fields

It may be possible to generate 3D force fields by using Lorentz electromagnetic forces. Tunable electric coils could be attached to various points of a 3D body, suspending the resulting object in a strong permanent magnetic field using magnetic levitation (the Lorentz effect) [HS93,SWH95]. The tuning (control) of the electric coils could be effected as follows: Integrated control circuitry could be fabricated and co-located with the coils, and conceivably a power supply. The control could be globally effected using wireless communication, or, the control of each coil evolves in time until the part is reoriented as desired. The Lorentz forces could then be deactivated to bring the object to rest on the ground. Planning for such a 3D device might reduce to [EMV93].

Appendix A

SCREAM Process

A.1 Processing Template

A.1.1 Photolithographic Masks

CAD Symbad polygon editor, Dracula polygon fracturing 1. generate design .... `ped` 2. fracture design .... `symbpg` 3. check pattern ..... `pgflash`, `pgcheck`, `jebcad` 4. transfer pattern to pattern generator . `pgnet`	⟹ file: .... ⟹ flashcount: ....

Pattern Generation GCA Mann 3600 Pattern Generator 1. load mask 2. load data file 3. set parameters 4. expose	⇒ file: .... ⇒ mask type: .... ⟹ mask name: ....
Chrome Mask Making 1. MF320 develop .... 2 *min* 2. DI rinse .... 1 *min* 3. DI soak .... 2 *min* 4. **resist descum** (O_2 plasma) 5. Chrome Etchant .... 1.5 *min* 6. DI rinse .... 1 *min* 7. DI soak .... 2 *min* 8. **resist strip** (1165 remover or O_2 plasma)	⇒ mask name: ....

A.1.2 Tip Layer

PECVD SiO_2 **Deposition** IPE System 1000 1. thorough chamber clean before deposition 2. deposition → N_2O ... 42 *sccm* (70 %) → SiH_4 (Silane) ... 12 *sccm* (17 %) → pressure ... 450 *mT* → power ... 50 *W* (4.5 %) → temperature ... 300° *C* → nominal deposition rate ≈ 380 *Å/min* 3. SiO_2 **thin film measurement (Leitz)**	⇒ wafer: ... ⇒ time: ... 90 *min* ⟹ thickness ... ⟹ deposition rate ...

Photolithography 1. **vapor prime (HMDS)** 2. spin on photoresist 3. hotplate bake ... 3 *min*, 90° *C* 4. **exposure on 10:1 stepper** (see below) 5. develop ... OCG 945 6. DI rinse ... 1 *min* 7. DI soak ... 3 *min* 8. dry (N_2 jet / spin)	⇒ wafer: ... ⇒ resist: .. OCG 895i 50cs ⇒ spin, time 4000 *rpm*, 30 *sec* ⇒ nominal resist thickness 2.4 *μm* ⇒ development time .. 2.5 *min*

Exposure on 10:1 Stepper GCA 6300 10:1 optical stepper 1. load mask 2. load wafer 3. expose	⇒ wafer: ... ⇒ mask: ... ⇒ file: ... ⇒ time: ... 5.25 *sec* ⇒ focus: ... 251

MIE SiO_2 Etch

MIE MRC-720

1. etch
   - → CHF_3 .... 30 *sccm*
   - → O_2 .... 0 *sccm*
   - → power .... 1000 *W*
   - → pressure .... 1–2 *mT*
   - → nominal etch rate .... 3500 *Å/min*
   - → nominal selectivity (SiO_2:resist) ... 2:1
2. **resist strip**
3. SiO_2 mesa **thin film measurement (Leitz)**

⇒ wafer ....
⇒ time .... 12 *min*
⟹ SiO_2 thickness ....

Resist Strip (O_2 Plasma)

O_2 plasma in Applied Materials RIE
- → O_2 flow .... 30 *sccm*
- → pressure .... 30 *mT*
- → power .... 90 *W*
- → time .... ≈ 10 *min*

Cl_2 RIE Deep Si Etch

RIE (old) Plasma Therm PK-1250

1. moisture removal .... 1 *min*
2. native oxide removal .... 1 *min*
3. Si etch

step	1	2	3	
Cl_2	0	2	50	*sccm*
BCl_3	14	14	1.3	*sccm*
H_2	7	7	0	*sccm*
pressure	20	20	40	*mT*
voltage	200	300	400	*V*

- → nominal etch rate .... 13 *μm/h*
- → nominal selectivity (Si:SiO_2) 18:1

4. **post-etch clean (Nanostrip)**
5. SiO_2 **thin film measurement (Leitz)**
6. **trench measurement (Alpha Step)**

⇒ wafer ....
⇒ time .... 40 *min*
⟹ SiO_2 thickness ....
⟹ trench depth (Si only) ..
⟹ etch rate ....
⟹ selectivity ....

A.1.3 Actuator Grid Layer

PECVD SiO_2 **Deposition** IPE System 1000 1. thorough chamber clean before deposition 2. deposition → N_2O .... 42 *sccm* (70 %) → SiH_4 (Silane) .... 12 *sccm* (17 %) → pressure .... 450 *mT* → power .... 50 *W* (4.5 %) → temperature .... 300° *C* → nominal deposition rate ≈ 380 *Å/min* 3. SiO_2 **thin film measurement (Leitz)**	⇒ wafer: .... ⇒ time: .... 40 *min* ⟹ thickness .... ⟹ deposition rate ....
Photolithography 1. **vapor prime (HMDS)** 2. spin on photoresist 3. hotplate bake .... 3 *min*, 90° *C* 4. **exposure on 10:1 stepper** (see below) 5. develop .... OCG 945 6. DI rinse .... 1 *min* 7. DI soak .... 3 *min* 8. dry (N_2 jet / spin)	⇒ wafer: .... ⇒ resist: OCG 895i 16.5cs ⇒ spin, time 3000 *rpm*, 30 *sec* ⇒ nominal resist thickness 1.5 *μm* ⇒ development time .. 2.5 *min*
Exposure on 10:1 Stepper GCA 6300 10:1 optical stepper 1. load mask 2. load wafer 3. expose	⇒ wafer: .... ⇒ mask: .... ⇒ file: .... ⇒ time: .... 7.25 *sec* (non-planar) / 4.75 *sec* (planar) ⇒ focus: . 284 (transparent mode)

MIE SiO_2 **Etch**

MIE MRC-720

1. etch
   - → CHF_3 .... 30 *sccm*
   - → O_2 .... 0 *sccm*
   - → power .... 1000 *W*
   - → pressure .... 1–2 *mT*
   - → nominal etch rate .... 3500 *Å/min*
   - → nominal selectivity (SiO_2:resist) ... 2:1
2. **resist strip**
3. SiO_2 mesa **thin film measurement (Leitz)**

⇒ wafer ....
⇒ time .... 5 *min*
⟹ SiO_2 thickness ....

Resist Strip (O_2 **Plasma)**

O_2 plasma in Applied Materials RIE

- → O_2 flow .... 30 *sccm*
- → pressure .... 30 *mT*
- → power .... 90 *W*
- → time .... ≈ 10 *min*

Cl_2 **RIE Deep** Si **Etch**

RIE (old) Plasma Therm PK-1250

1. moisture removal .... 1 *min*
2. native oxide removal .... 1 *min*
3. Si etch

step	1	2	3	
Cl_2	0	2	50	*sccm*
BCl_3	14	14	1.3	*sccm*
H_2	7	7	0	*sccm*
pressure	20	20	40	*mT*
voltage	200	300	400	*V*

   - → nominal etch rate .... 13 *μm/h*
   - → nominal selectivity (Si:SiO_2) 18:1
4. **post-etch clean (Nanostrip)**
5. SiO_2 **thin film measurement (Leitz)**
6. **trench measurement (Alpha Step)**

⇒ wafer ....
⇒ time .... 37.5 *min*
⟹ SiO_2 thickness ....
⟹ trench depth (Si only) ..
⟹ etch rate ....
⟹ selectivity ....

PECVD SiO_2 Deposition

IPE System 1000

1. thorough chamber clean before deposition
2. deposition
   - → N_2O .... 42 *sccm* (70 %)
   - → SiH_4 (Silane) .... 12 *sccm* (17 %)
   - → pressure .... 450 *mT*
   - → power .... 50 *W* (4.5 %)
   - → temperature .... 300° *C*
   - → nominal deposition rate ≈ 380 *Å/min*
3. SiO_2 **thin film measurement (Leitz)**

⇒ wafer: ....
⇒ time: .... 15 *min*
⟹ thickness ....
⟹ deposition rate ....

CHF_3 RIE Trench Bottom Etch

(New) Plasma Therm System 72

1. CHF_3 RIE .... process 2
   - → CHF_3 .... 100 %
   - → O_2 .... 2 %
   - → pressure .... 40 *mT*
   - → power (step 4) .... 30 %
   - → nominal etch rate .... 255 *Å/min*
2. SiO_2 (mesa) **thin film measurement (Leitz)**

⇒ wafer ....
⇒ time .... 18 *min*
⟹ SiO_2 thickness (mesa) ..
⟹ etch rate ....

Cl_2 RIE Deep Si Etch

RIE (old) Plasma Therm PK-1250

1. moisture removal .... 1 *min*
2. native oxide removal .... 1 *min*
3. Si etch

→

step	1	2	3	
Cl_2	0	2	50	*sccm*
BCl_3	14	14	1.3	*sccm*
H_2	7	7	0	*sccm*
pressure	20	20	40	*mT*
voltage	200	300	400	*V*

   - → nominal etch rate .... 13 *μm/h*
   - → nominal selectivity (Si:SiO_2) 18:1
4. **post-etch clean (Nanostrip)**
5. SiO_2 **thin film measurement (Leitz)**
6. **trench measurement (Alpha Step)**

⇒ wafer ....
⇒ time .... 10 *min*
⟹ SiO_2 thickness ....
⟹ trench depth (Si only) ..
⟹ etch rate ....
⟹ selectivity ....

Release Etch

(New) Plasma Therm System 72

1. SF_6 release etch .... process 5
   - → SF_6 .... 140 *sccm* (70 %)
   - → O_2 .... 0 %
   - → pressure .... 90 *mT*
   - → power (step 4) .... 30 %
   - → nominal etch rate (isotropic) .... 1300 *Å/min*
   - → nominal selectivity ($Si:SiO_2$) 300:1
2. check release in **SEM**

⇒ wafer ....
⇒ time .... 6 *min*
⟹ etch rate ....

Oxide Strip (BHF)

1. BHF 6:1 soak .... 15 – 30 *min*
2. thorough rinse

A.1.4 Electrode Layer

PECVD SiO_2 Deposition

IPE System 1000

1. thorough chamber clean before deposition
2. deposition
   - → N_2O .... 42 *sccm* (70 %)
   - → SiH_4 (Silane) .... 12 *sccm* (17 %)
   - → pressure .... 450 *mT*
   - → power .... 50 *W* (4.5 %)
   - → temperature .... 300° *C*
   - → nominal deposition rate ≈ 380 *Å/min*
3. SiO_2 **thin film measurement (Leitz)**

⇒ wafer: ....
⇒ time: .... 8.5 *min*
⟹ thickness ....
⟹ deposition rate ....

Sputter Deposition

CVC AST-601 Sputtering System

- → current .... DC 5 *A*
- → pressure .... 25 *mT*
- → Ar flow .... 40 *sccm*
- → nominal deposition rate .... 90 *Å/min*

⇒ wafer ....
⇒ ramp + presputter . 2 + 10 *min*
⇒ sputter time .... 28 *min*
⟹ thickness ....

Photolithography on Al

1. **vapor prime (HMDS)**
2. spin on photoresist
3. bake .... 3 *min*, 90° *C*
4. **exposure on 10:1 stepper**
5. develop .... MDC diluted 1:3
6. DI rinse .... 1 *min*
7. DI soak .... 3 *min*
8. dry (N_2 jet / spin)

⇒ wafer: ....

⇒ resist: .. OCG 895i 50cs

⇒ spin, time 3000 *rpm*, 60 *sec*

⇒ nominal resist thickness 2.8 *μm*

⇒ development time 3 *min*

Exposure on 10:1 Stepper

GCA 6300 10:1 optical stepper

1. load mask
2. load wafer
3. expose

⇒ wafer: ....

⇒ mask: ....

⇒ file: ....

⇒ time: .... 15.5 *sec* (non-planar) / 12.5 *sec* (planar)

⇒ focus: . 284 (transparent mode)

Cl_2 **RIE** Al **Etch**

RIE (old) Plasma Therm PK-1250

→

Cl_2	20	*sccm*
BCl_3	40	*sccm*
CH_4	1.3	*sccm*
pressure	20	*mT*
voltage	400	*V*

→ nominal etch rate .... 2000 *Å/min*

→ nominal selectivity (Al:resist) 2:1

⇒ wafer ....

⇒ time .... 1:30 *min*

Resist Strip (O_2 **Plasma)**

O_2 plasma in Applied Materials RIE

→ O_2 flow .... 30 *sccm*

→ pressure .... 30 *mT*

→ power .... 90 *W*

→ time .... ≈ 10 *min*

A.2 Processing Notes

Notes on CAD

- The pattern is fractured into rectangles. Therefore acute angles are not possible. However they are approximated to pattern generator resolution by a large number of small (fractal-like) rectangles. Thus acute angles should be avoided.
- Fracturing is an NP-hard problem. Several heuristics are used which sometimes give non-optimal or even surprising results.
- The order of the rectangles is optimized for pattern generator speed. Ordering precedence seems to be angle, size, location. `pgcheck` draws the rectangles in the order of exposure.
- It is advantageous to have the origin of the mask exactly at the center. `symbpg` puts the center at the exact geometric midpoint. Therefore the pattern should always be made symmetric w.r.t. the origin.
- Standard alignment marks for the 10:1 stepper are generated with `EXECUTE cadmacros:gca_key`. Only necessary on the first mask. Make alignment marks as trenches, otherwise they may get released and disappear.
- `ped` *library* — CAD polygon editor.
- `symbpg` (`symbpg_rev`) — fracturing (with tone reversal). `symbpg` starts a batch job. The batch queue can be checked with `show queue /batch`.
- `pgflash` — print flashcounts and other information after fracturing.
- `pgcheck /d=gpx` *file* or `jebcad` — draw generated rectangles.
- `pg_dump [/out=`*file*`]` — write rectangle file in ASCII format.
- `pgmerge` *file1,file2,...,fileN outfile* — merge rectangles of multiple files.
- (`pgtape` — write fractured pattern on tape for pattern generator.)
- `pgnet` — transfer pattern to Pattern Generator

Notes on Pattern Generation

- Maximum field size is 10 mm square with cut-off corners. Largest full inscribed square is 9.4 mm square.
- For chrome masks (shutter mode) approximately 7,000 to 8,000 flashes per hour.
- Closure test produces pattern that is partially made before and after exposure. This way alignment shifts can be detected. Center rectangle is supposed to be exactly centered. X:40 / Y:55 puts pattern in upper right corner. Avoid 55/55 because of non-uniform resist in corner, avoid 55/0 or 0/55 because of fiducials.
- Chrome masks (used here) have positive resist. Tone reversal is necessary to expose everything except pattern. Emulsion masks have negative resist.
- Mask needs fiducial marks (“C” for Cr masks) for alignment with the optical column of the stepper.
- Only first mask needs CAD alignment marks.

Notes on Chrome Mask Making

- Chrome etch over light until chrome in exposed areas is gone, i.e. mask is clear.

Notes on PECVD SiO_2 Deposition

- To avoid build-up of flakes as much as possible, the chamber has to be very clean. Clean Process is often not sufficient. Remove all depositions inside the chamber. Remove shower head and scrub it in sink. Blow and vacuum chamber.
- Flakes tend to fall during and after chamber venting. Remove wafers as quickly as possible and let them cool down outside chamber. Flakes falling down after processing can mostly be removed with the N_2 jet.
- Some users recommend short conditioning of chamber before deposition (a) to remove CF_4 and (b) to bond loose particles left after clean (debatable).
- Some users suggest letting the system sit for a while after pumpdown at high temperature to get rid of moisture.

Notes on Photolithography

- HMDS vapor prime is preferred. Alternatively, bake wafer on 115° C hot plate for 10 *min* to remove moisture. Cover whole wafer with HMDS 20 % and let sit for a few seconds. Then spin for 30 *sec* at 3000 *rpm*.
- Put a quarter-sized blob of resist on center of wafer before spinning, avoid bubbles as they may cause areas with no resist.
- Resist may be spread more evenly, and without the risk of leaving streaks of uncoated wafer, if the spin speed is ramped up.
- Use chuck slightly smaller than wafer so that excess resist can get off freely, avoiding edge beads.
- Resist (especially OCG 895i) does not stick well to Al (or rather Al_2O_3). Shipley System 8 sticks better. Very careful cleaning is necessary. It may be better to deposit about 1000 $\AA$ of SiO_2 on Al for better adhesion. However, Al does not take PECVD too well, it tends to form bumps and bubbles.
- Reflectivity of Al causes focusing problem. Stepper can be switched to "transparent" mode.
- Most developers (especially OCG 945) attack Al. For lithography on Al use Microposit Developer Concentrate (MDC) diluted (1:3) with DI.
- Nominal resist thicknesses for OCG 895i (μm):

rpm	3000	4000
5cs	.8	.6
10cs	1.1	.9
16.5cs	1.5	1.3
22cs	1.7	1.5
34cs	2.1	1.9
50cs	2.8	2.4

- Some users recommend contrast enhancing material (CEM). CEM is opaque but becomes transparent after a short exposure. It acts like an in-situ contact mask, improving resolution and sidewall slopes.
- Then, some users do not use BC5 but applies CEM365 directly to photoresist. Potential problem: CEM365 does not bond well to photoresist.
- AZ4903 thickness:

rpm	*μm*
750	30.0
1000	23.5
2000	14.6
3000	13.1
4000	10.7
5000	9.1
6000	8.0
7000	7.2

- Thick resist AZ4903 needs hold time (30 *min*) to absorb moisture before exposure.
- Strip AZ4903 with AZ Thinner, AZ EBR, acetone, or O_2 plasma.
- 10x stepper delivers 58.08 mW/cm^2.

Notes on Exposure on 10:1 Stepper

- Don't forget to get lamp into position after aligning mask.
- Check shutters, opaque/transparent switch.
- Exposure wavelength 365 *nm*.
- For optimal alignment after loading or switching wafers and masks, wait at least 5 *min* for temperature to settle inside the Stepper.
- Standard alignment marks are 63.5 *mm* apart.
- Positions and spacing can be controlled with 5 digits accuracy.
- To get standard alignment marks, spacing should be a divisor of 63.5 *mm*.
- If an exposure has multiple passes, masks can be switched after each pass without losing alignment.
- If parts of the exposure need different masks, use multiple passes and use `ARRAY-DROPOUT` to remove single dies from the exposure, or `PLUG-PLUG IN` to put single dies into the exposure. `PLUG IN` also allows an arbitrary offset, so the die can basically go anywhere.
- Syntax: `EXEC <file>, PASS1, PASS2`. After each pass the system waits and asks `CHANGE RETICLES?` Then replace masks.
- Masks have 4 fiducial marks, so they can be used in 4 orientations.
- Exposure of dies is done as shown by ASCII-graphics on screen, but each individual die is rotated by 180° due to the projection.

Notes on MIE SiO_2 etch

- Etch is uniform up to about 1/2 inch to edge of wafer, where it is slower and very nonuniform. Therefore patterns close to the edge should be avoided. They will produce a lot of grass.
- Watch the He pressure. If it drops, liquid He is leaking. This means the wafer will not be cooled properly, and the He may mess up the etch.
- After chamber has been opened, always run chamber clean (O_2 or CHF_3) until pressure reaches $\approx 5\ mT$.

Notes on Cl_2 RIE Deep Si Etch

- If previous user did not etch Si (especially Al) make sure chamber is clean.

- Steps 1 and 2 may not be absolutely necessary.
- Low Cl_2 flow in Step 2 is hard to achieve. Start at higher flow and lower it gradually. Still, 2 $sccm$ may not be stable.

Notes on CHF_3 RIE Trench Bottom Etch

- Slow SiO_2 etch for very thin films.
- Make sure gas selector switch is on correct setting.
- Watch DC bias voltage. Change in bias indicates build-up of polymers.

Notes on Release Etch

- Make sure gas selector switch is on correct setting.
- Etch is rather non-uniform.
- Some users recommend a short pre-release etch to get rid of oxide and other deposits: process 8, pressure 10 mT, power 15 %, CF_4 50 %, O_2 2 %.
- Some users use O_2 chamber clean before release. Don't! Plasma Therm staff recommends anything *except* O_2 (in particular SF_6) for chamber conditioning before release etch.

Notes on Sputter Deposition

- System takes long time to pump down.
- Use monitor wafer to measure thickness, e.g. with partial resist for acetone lift-off and alpha-step (usually provided by staff).
- To close shutter, *black* nob must be in target sector.
- Al with low percentage of Si supposedly has better properties for wire bonding, but is not very popular in the CNF.
- Breakdown voltage for PECVD SiO_2 is about 300 $V/\mu m$. However voltage may leak through defects in oxide. Devices with leaks may still function as the leaks are rather small.
- Leaks and breakdown decrease with thicker films, but film stress increases, which in turn can cause leakage and breakdown.

Appendix B

Microscopic Model For Actuator Contact

In this section we develop a model for the mechanics of microactuators. We make use of limit surfaces [GR88] that describe anisotropic frictional contact. We extend the model to active contacts and describe fast algorithms to compute the combined effect of many actuators.

Limit Surfaces. Assume we have a part P that moves on top of the actuator array. The *limit surface* $\mathcal{L}$ in load space (forces F_x and F_y, moment M) fully describes the relationship between generalized velocity v, effective applied load F_{eff}, and frictional load F_r of the moving part. It is based on the Maximum Work Inequality which is an engineering assumption commonly used when modeling friction or plasticity [GR88,PS88]. The Maximum Work Inequality generalizes Coulomb's friction law to anisotropic rate-independent friction. The following properties of limit surfaces are useful [GR88]:

1. $\mathcal{L}$ is a closed convex surface in load space.
2. $\mathcal{L}$ contains all possible frictional loads F_r on P.
3. If $v \neq 0$, $F_r \in \mathcal{L}$ and $-v$ is normal to $\mathcal{L}$ at F_r.
4. F_r is the vector inside or on $\mathcal{L}$ such that the length of $F_r - F_{\mathit{eff}}$ is minimal.
5. In the special case when the moment component M of F_r is 0, $\mathcal{L}$ is a *limit curve* in $(F_x - F_y)$ load space.

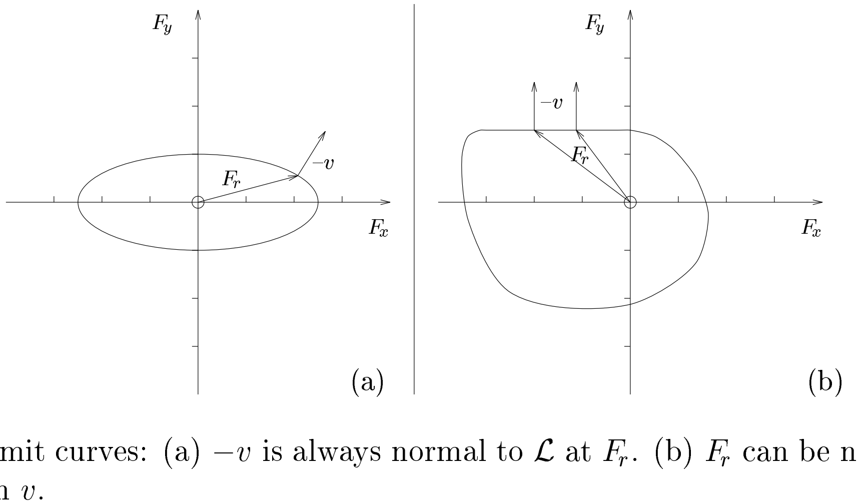

Figure B.1: Limit curves: (a) $-v$ is always normal to $\mathcal{L}$ at F_r. (b) F_r can be nonunique for a specific given v.

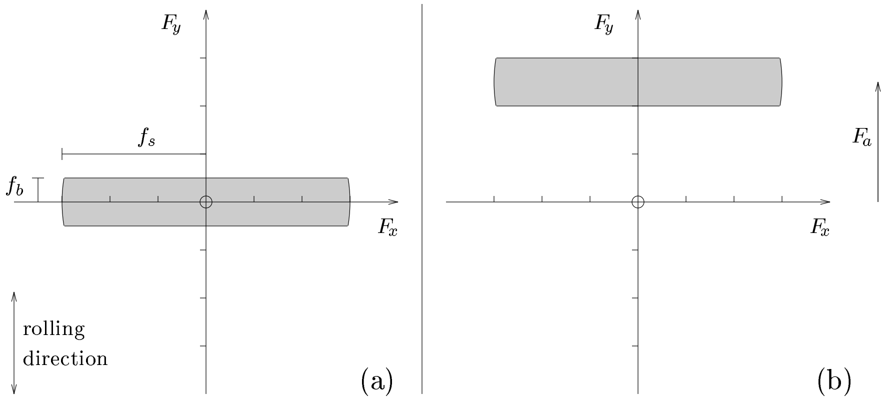

Figure B.2: Limit curve of (a) passive wheel, (b) driven wheel.

See Figure B.1 for examples. It follows that the inside of $\mathcal{L}$ contains all loads F_{eff} that can be applied to P without setting it in motion, and if $F_{\mathit{eff}} \neq 0$ and $v \neq 0$, F_r is determined uniquely. But note that for given v, F_r is not unique if $\mathcal{L}$ has a flat face with normal $-v$. Similarly for given F_r, v is not unique if $\mathcal{L}$ has a vertex or edge at F_r. This indeterminacy can be resolved by taking the inertia of P into account [GR88].

Consider as an example Figure B.2a. The anisotropic behavior of a wheel can be modeled with a long rectangular limit curve which gives low bearing friction f_b in the rolling direction and high sideways friction f_s.

Active Contacts. We now extend the limit surface model to "active" contacts that apply loads to P.

Definition 43 *The* active limit surface $\mathcal{L}$ *in load space is the set of loads that can be applied to P without resulting in motion of P.*

This definition includes limit surfaces for passive contacts, but it allows us to model for example a wheel driven by some torque τ_a. Figure B.2b shows that if no additional load is applied the wheel will move in y direction, accelerated by Fa minus the bearing friction f_b (where F_a is such that $\tau_a = r \times F_a$, and r is the radius of the wheel). In general we get motion if the origin of load space $\mathbf{0}$ lies outside of $\mathcal{L}$.

For the wheel accelerated with torque $\tau_a = r \times F_a$ the limit curve simply shifts in load space by F_a. For our actuators we expect the shape of active and passive limit surface to be different because of interactions between friction and oscillation. However, because the limit surface will represent the time average over frictional contacts, we believe that the theory of limit surfaces is a valid model.

Combining Limit Surfaces. We have already noted one advantage of limit surfaces: they offer a uniform, purely geometric representation of contact properties. In addition, multiple contacts can be described by a single joint limit surface. Goyal and Ruina [GR88] have shown that this joint limit surface is the convolution (Minkowski sum) of the individual limit surfaces. There exist fast algorithms from computational geometry [GS88] to compute this convolution. They are described in more detail in our technical report [BDMM93].

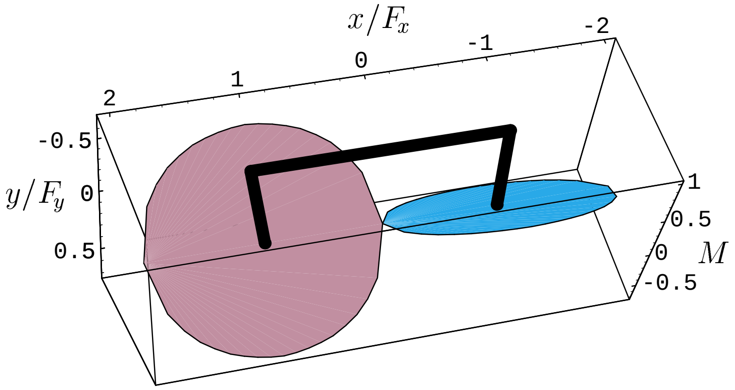

Figure B.3: Two rigidly connected point contacts at $(1, 0)$ and $(-1, 0)$, and their corresponding limit surfaces with respect to the center of mass $(0, 0)$.

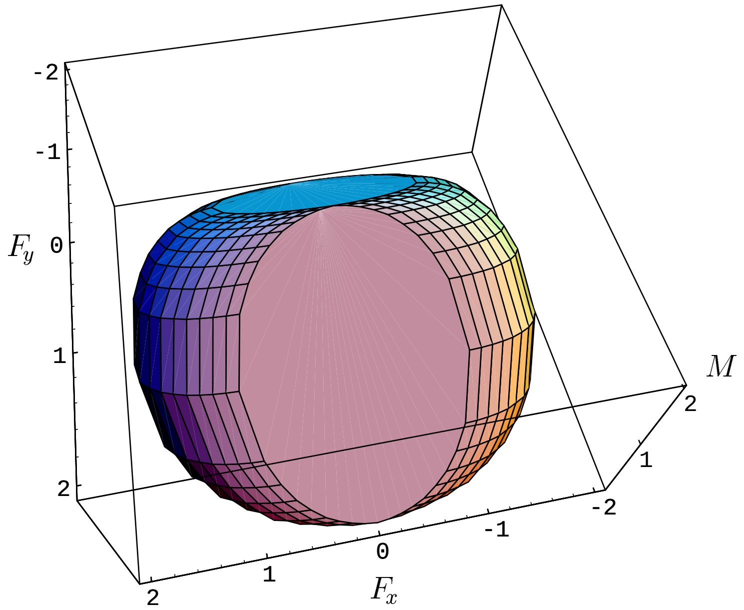

Figure B.4: Combined limit surface for two rigidly connected point contacts.

Thus limit surfaces offer a technique to efficiently determine and represent the behavior of an entire actuator array.

For illustration, Figure B.3 shows two rigidly connected point contacts and Coulomb friction, and their corresponding limit surfaces. The individual surfaces are flat because they can generate no moment about the contact point, but tilted because they can generate moments about the center of mass. The tilt angle θ can be determined by $\tan\theta = \frac{|M|}{|F|} = \frac{|r \times F|}{|F|} = |r|$, so for $|r| = 1$ we get $\theta = 45°$. Figure B.4 shows the combined limit surface.

Motion Prediction. For simplicity let us first consider the "upside-down" case where the actuator array "walks" on a homogeneous flat surface. The contacts and thus the limit surface $\mathcal{L}$ are fixed with respect to a coordinate system attached to the actuator array. If the origin of load space $\mathbf{0}$ is inside $\mathcal{L}$ there will be no motion. Otherwise the generated force F is the point on $\mathcal{L}$ closest to $\mathbf{0}$. At that point the surface normal is parallel to F, so velocity and accelerating force are parallel. The walker will move on a straight line or a circle. This is not unexpected for a fixed actuator strategy on a homogeneous surface.

Now consider the case where the object is on top of the actuator array. There are two major differences to the previous case: When the object moves (1) some actuators lose contact, others make contact, (2) the induced moment changes. (1) requires periodic updates of the limit surface. (2) results in stretching of the limit surface along the moment axis. Our technical report [BDMM93] gives details on how to handle these cases efficiently. There, we outline a simple yet efficient motion prediction algorithm: Numerically integrate the velocities computed as described above, and update the limit surface $\mathcal{L}$ accordingly. Each integration step can be done in constant time. Each update of $\mathcal{L}$ is linear in the complexity of $\mathcal{L}$.

Motion Planning. The shape of the limit surface is determined by the activation pattern of the actuator array, which depends on each individual actuator. The limit surface gives us a geometric representation of the forces and velocities generated with a specific actuator activation pattern. Though theoretically possible, there are practical limitations on using limit surfaces to *plan* microscopic manipulation strategies due to the combinatorial complexity and mechanical uncertainty. However, the microscopic model will prove important to *analyze and verify* strategies before fabrication. It forms the link between "actuator macros" (Section 2.3) that predict the global behavior of manipulation strategies, and individual microfabricated mechanisms (Sections 6.1 and 6.2).

Appendix C

Particle Bouncing on a Vibrating String

To understand the effective forces on particles on a vibrating surface, we look at the more tractable case of the planar motion of a particle bouncing on a string in transverse vibrations (Figure C.1).

The string vibrates in the first mode, and is not affected by its interaction with the particle. The shape of the string, at time t, for a given x location is:

$$y_s = A\ \sin x\ \sin 2\pi\nu t$$

where ν is the frequency of oscillation. The position of the particle is given by (x_p, y_p).

The interaction between the particle and the string is through a sequence of impacts. We use a model for particle impact with a finite friction coefficient μ, and a coefficient of restitution e. θ is the slope of the string at the point and instant of impact, and is small for small amplitudes of string vibration.

$$\tan\theta = A\ \cos x\ \sin 2\pi\nu t$$

The motion of the particle can be simulated as a series of impacts with the string, with the particle in free flight in-between. The change in the momentum of the particle during impact is calculated using a simple planar impact model. Figure C.2 shows the results of a numerical simulation of the model at two different values of e.

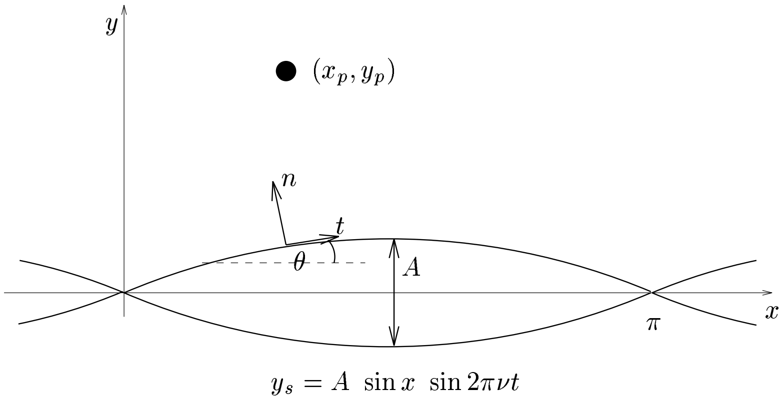

Figure C.1: Particle bouncing on a vibrating string.

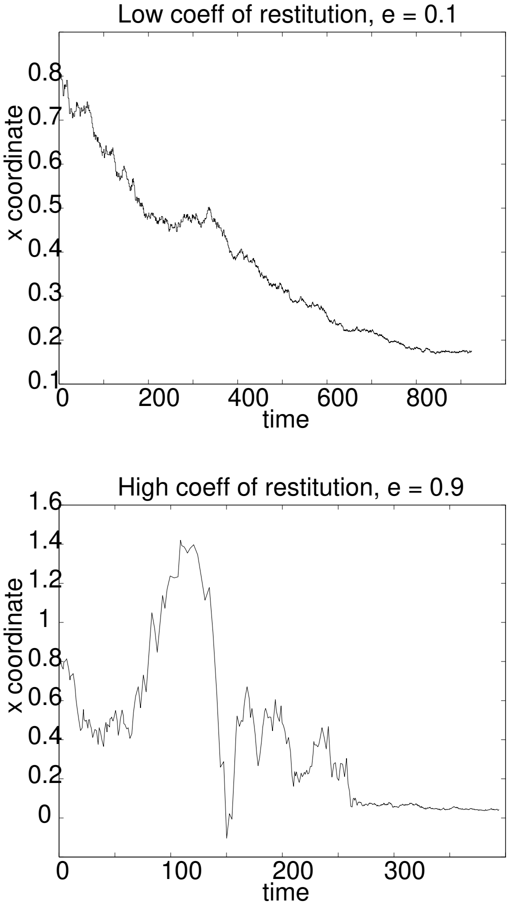

Figure C.2: Simulation results showing the position of a particle moving on a vibrating string.

For a particle starting at rest, at $t = 0$, we find that $\dot{y}_p \gg \dot{x}_p$. Using the assumption that the amplitude of oscillations is small, $\sin\theta \approx \tan\theta$, $\cos\theta \approx 1$. If $(\dot{x}_p^-, \dot{y}_p^-)$ represent the velocity just before impact, the velocity just after impact $(\dot{x}_p^+, \dot{y}_p^+)$, is:

$$
\begin{aligned}
\dot{x}_p^+ &= e\,(\dot{y}_p^- - \dot{y}_s^-)\sin\theta + \alpha\, v_{rel_t}^- && \text{(C.1)}\\
\dot{y}_p^+ &= \dot{y}_s^-(1+e) - e\,\dot{y}_p^- && \text{(C.2)}
\end{aligned}
$$

where $v_{rel_t}^- = (\dot{y}_p^- - \dot{y}_s^-)\sin\theta + \dot{x}_p^-$ is the relative velocity along the tangential direction before impact, and $\alpha \in [0,1]$ is the dissipation factor that depends on μ.

After the impact, $\dot{x}_p^+$ is a sum of the relative tangential velocity before impact, attenuated by friction; and a component from the impulse in the normal direction, which depends on e and the slope of the string at the point of impact. The portion of x impulse added purely due to the effect of the string can be approximated as $-e\,\dot{y}_s\sin\theta$, by setting $\dot{y}_p^- = 0$.

If this component of the impulse were spread uniformly over time, the effective force, F_{eff}, that the particle would experience is:

$$
F_{\mathit{eff}} \propto -\nu e A^2 \sin x\ \cos 2\pi\nu t\ \cos x\ \sin 2\pi\nu t \tag{C.3}
$$

We now use the argument that it is more probable for the particle to impact the string at times when the string is above the mean rest position, to show that over a large number of impacts, the time dependent terms in equation (C.3) average out to a positive quantity. Therefore, the time averaged effective force, F_{avg}, experienced by the particle is:

$$F_{avg} \propto -\nu e A^2 \sin 2x$$

This confirms the intuition and the observed behavior that the particle moves faster at higher amplitudes of string oscillation, coefficient of restitution, and oscillation frequency. The sine dependency of the force with x ensures that it points towards the corresponding nodes on either side of the anti-node at $x = \frac{\pi}{2}$.

Bibliography

[AE96] Tamara Lynn Abell and Mike Erdmann. A universal parts feeder, 1996. Personal communication / in preparation.

[AHLM95] Srinivas Akella, Wesley H. Huang, Kevin M. Lynch, and Matthew T. Mason. Planar manipulation on a conveyor by a one joint robot with and without sensing. In *International Symposium of Robotics Research (ISRR)*, 1995.

[Ana91] Analog Devices, Inc., Norwood, MA 02062. *Introducing the ADXL50 Micromachined Accelerometer Sensor*, 1991.

[AOF93] Manabu Ataka, Akito Omodaka, and Hiroyuki Fujita. A biomimetic micro motion system. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, pages 38–41, Pacifico, Yokohama, Japan, June 1993.

[BBD$^+$95] Karl-Friedrich Böhringer, Russell G. Brown, Bruce R. Donald, James S. Jennings, and Daniela Rus. Distributed robotic manipulation: Experiments in minimalism. In *Fourth International Symposium on Experimental Robotics (ISER)*, Stanford, California, June 1995. `http://www.cs.cornell.edu/home/brd`.

[BBG95] Karl-Friedrich Böhringer, Vivek Bhatt, and Kenneth Y. Goldberg. Sensorless manipulation using transverse vibrations of a plate. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 1989 – 1996, Nagoya, Japan, May 1995. `http://www.cs.cornell.edu/home/karl/VibratoryAlign`.

[BDH97] Karl-Friedrich Böhringer, Bruce R. Donald, and Dan Halperin. The area bisectors of a polygon and force equilibria in programmable vector fields. In *13th ACM Symposium on Computational Geometry*, Nice, France, June 1997.

[BDM96a] Karl-Friedrich Böhringer, Bruce R. Donald, and Noel C. MacDonald. Single-crystal silicon actuator arrays for micro manipulation tasks. In *Proc. IEEE Workshop on Micro Electro Mechanical Systems (MEMS)*, pages 7–12, San Diego, CA, February 1996. `http://www.cs.cornell.edu/home/karl/MicroActuators`.

[BDM96b] Karl-Friedrich Böhringer, Bruce R. Donald, and Noel C. MacDonald. Upper and lower bounds for programmable vector fields with applications to MEMS and vibratory plate parts feeders. In *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, Toulouse, France, July 1996. `http://www.cs.cornell.edu/home/karl/MicroManipulation`.

[BDMM93] Karl-Friedrich Böhringer, Bruce R. Donald, Robert Mihailovich, and Noel C. MacDonald. A geometric theory of manipulation and control for microfabricated actuator arrays. Technical Report 93–87, Cornell University, Mathematical Sciences Institute, Ithaca, NY 14853, November 1993. `http://www.cs.cornell.edu/home/karl/MicroActuators`.

[BDMM94a] Karl-Friedrich Böhringer, Bruce R. Donald, Robert Mihailovich, and Noel C. MacDonald. Sensorless manipulation using massively parallel microfabricated actuator arrays. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 826–833, San Diego, CA, May 1994. `http://www.cs.cornell.edu/home/karl/MicroManipulation`.

[BDMM94b] Karl-Friedrich Böhringer, Bruce R. Donald, Robert Mihailovich, and Noel C. MacDonald. A theory of manipulation and control for microfabricated actuator arrays. In *Proc. IEEE Workshop on Micro Electro Mechanical Systems (MEMS)*, pages 102–107, Oiso, Japan, January 1994. `http://www.cs.cornell.edu/home/karl/MicroActuators`.

[BPM82] Geoffrey Boothroyd, Corrado Poli, and Laurence E. Murch. *Automatic Assembly.* Marcel Dekker, Inc., 1982.

[Bri92] Amy J. Briggs. An efficient algorithm for one-step planar compliant motion planning with uncertainty. *Algorithmica*, 8(3), 1992.

[Bro86] Rod Brooks. A layered intelligent control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA(2), 1986.

[Bro88] Randy C. Brost. Automatic grasp planning in the presence of uncertainty. *Int. Journal of Robotics Research*, 7(1):3–17, 1988.

[CG94] John Canny and Ken Goldberg. “RISC” for industrial robotics: Recent results and open problems. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*. IEEE, May 1994.

[Chl87] Ernst Florens Chladni. *Entdeckungen über die Theorie des Klanges.* Weidmanns Erben und Reich, Leipzig, 1787.

[DJR95] Bruce Randall Donald, James Jennings, and Daniela Rus. Information invariants for distributed manipulation. In K. Goldberg, D. Halperin, J.-C. Latombe, and R. Wilson, editors, *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, pages 431–459, Wellesley, MA, 1995. A. K. Peters.

[DO90] M. Díaz and J. O’Rourke. Ham-sandwich sectioning of polygons. In *Proc. 2nd Canadian Conference on Computational Geometry*, pages 98–101, Ottawa, 1990.

[Don89] Bruce R. Donald. *Error Detection and Recovery in Robotics*, volume 336 of *Lecture Notes in Computer Science.* Springer Verlag, Berlin, 1989.

[Don90] Bruce R. Donald. The complexity of planar compliant motion planning with uncertainty. *Algorithmica*, 5(3):353–382, 1990.

[DX95] Bruce Randall Donald and Pat Xavier. Provably good approximation algorithms for optimal kinodynamic planning for cartesian robots and open chain manipulators. *Algorithmica*, 14(6):480–530, November 1995.

[EM88] Michael A. Erdmann and Matthew T. Mason. An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4(4), August 1988.

[EM96] Michael A. Erdmann and Matthew T. Mason. Nonprehensile manipulation. In *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, Toulouse, France, July 1996.

[EMV93] Michael A. Erdmann, Matthew T. Mason, and G. Vaneček, Jr. Mechanical parts orienting: The case of a polyhedron on a table. *Algorithmica*, 10, 1993.

[Erd94] M. A. Erdmann. On a representation of friction in configuration space. *Int. Journal of Robotics Research*, 13(3):240–271, 1994.

[Erd96] Michael A. Erdmann. An exploration of nonprehensile two-palm manipulation: Planning and execution. Technical report, Carnegie Mellon University, Pittsburgh, PA, 1996.

[FD86] Gregory T. Farnum and Bill Davis. Delivering the part. *Manufacturing Engineering*, March 1986.

[FHF91] T. Furuhata, T. Hirano, and H. Fujita. Array-driven ultrasonic microactuators. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, pages 1056–1059, Montreux, France, June 1991.

[Fuj93] Hiroyuki Fujita. Group work of microactuators. In *International Advanced Robot Program Workshop on Micromachine Technologies and Systems*, pages 24–31, Tokyo, Japan, October 1993.

[FVDFH96] J. D. Foley, A. Van Dam, Feiner, and Hughes. *Computer Graphics: Principles and Practice*. Addison Wesley, 2 edition, 1996.

[Gab95] Kaigham J. Gabriel. Engineering microscopic machines. *Scientific American*, 273(3), September 1995.

[Gil56] Thomas C. Gillmer. *Fundamentals of Construction and Stability of Naval Ships*. The United States Naval Institute, Annapolis, MD, 1956.

[Gol93] K. Y. Goldberg. Orienting polygonal parts without sensing. *Algorithmica*, 10(2/3/4):201–225, August/September/October 1993.

[GR88] Suresh Goyal and Andy Ruina. Relation between load and motion for a rigid body sliding on a planar surface with dry friction: Limit surfaces, incipient and asymptotic motion. *Wear*, 1988.

[GRP91] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. Planar sliding with dry friction. part 1: Limit surface and moment function. part 2: Dynamics of motion. *Wear*, 143:307–352, 1991.

[GS88] Leonidas J. Guibas and Raimund Seidel. Computing convolutions by reciprocal search. In *Proceedings of the ACM Symposium on Computational Geometry*, Urbana, IL, 1988.

[Hit88] Hajime Hitakawa. Advanced parts orientation system has wide application. *Assembly Automation*, 8(3), 1988.

[HS93] Ralph Hollis and S. E. Salcudean. Lorentz levitation technology: A new approach to fine motion robotics, teleoperation, haptic interfaces, and vibration isolation. In *International Symposium of Robotics Research (ISRR)*, Hidden Valley, PA., October 1993.

[Int89] Integrated Engineering Software Inc. *Coulomb*, 1989. Users and Technical Manual, Version 2.1.

[JGJB$^+$95] J. D. Jacobson, S. H. Goodwin-Johansson, S. M. Bobbio, C. A. Bartlet, and N. Yadon. Integrated force arrays: Theory and modeling of static operation. *Journal of Microelectromechanical Systems*, 4(3):139–150, September 1995.

[JM83] Robert C. Juvinall and Kurt M. Marshek. *Fundamentals of Machine Component Design.* Wiley, New York, 1983.

[Kav95] Lydia E. Kavraki. On the number of equilibrium placements of mass distributions in elliptic potential fields. Technical Report STAN-CS-TR-95-1559, Department of Computer Science, Stanford University, Stanford, CA 94305, 1995.

[Kav97] Lydia Kavraki. Part orientation with programmable vector fields: Two stable equilibria for most parts. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Albuquerque, New Mexico, April 1997.

[KF93a] Satoshi Konishi and Hiroyuki Fujita. A conveyance system using air flow based on the concept of distributed micro motion systems. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, pages 28–31, Pacifico, Yokohama, Japan, June 1993.

[KF93b] Satoshi Konishi and Hiroyuki Fujita. A proposal for a conveyance system with autonomous decentralized micro modules. In *IEEE International Symposium on Autonomous Decentralized Systems*, Kawasaki, Japan, March 1993.

[Kha86] Oussama Khatib. Real time obstacle avoidance for manipulators and mobile robots. *Int. Journal of Robotics Research*, 5(1):90–99, Spring 1986.

[KR88] Daniel E. Koditschek and Elon Rimon. Robot navigation functions on manifolds with boundary. *Advances in Applied Mathematics*, 1988.

[LM97] Jonathan E. Luntz and William Messner. A distributed control system for flexible materials handling. *IEEE Control Systems*, 17(1), February 1997.

[LMC97] Jonathan E. Luntz, William Messner, and Howie Choset. Parcel manipulation and dynamics with a distributed actuator array: The virtual vehicle. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 1541 – 1546, Albuquerque, New Mexico, April 1997.

[LP83] Tomás Lozano-Pérez. Spacial planning: A configuration space approach. *IEEE Transactions on Computers*, C-32(2):108–120, February 1983.

[LPMT84] Tomás Lozano-Pérez, Matt Mason, and Russell Taylor. Automatic synthesis of fine-motion strategies for robots. *Int. Journal of Robotics Research*, 3(1), 1984.

[LTW$^+$95] C. Liu, T. Tsao, P. Will, Y.C. Tai, and W.H. Liu. A micro-machined magnetic actuator array for micro-robotics assembly systems. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, Stockholm, Sweden, 1995.

[LW95] Wenheng Liu and Peter Will. Parts manipulation on an intelligent motion surface. In *Int. Conf. on Intell. Robots & Systems (IROS)*, Pittsburgh, PA, 1995.

[MAA$^+$95] N. C. MacDonald, S. G. Adams, A. A. Ayon, K.-F. Böhringer, L.-Y. Chen, J. H. Das, D. Haronian, W. Hofmann, X. T. Huang, A. Jazairy, R. E. Mihailovich, S. A. Miller, I. Ogo, R. Prasad, B. W. Reed, M. T. A. Saif, K. A. Shaw, R. Y. Webb, and Y. Xu. Micromachined microdevices and microinstruments. In *Micro- And Nano-Engineering (MNE)*, Aix-en-Provence, France, September 1995.

[Mac96a] Noel C. MacDonald. *Nanotechnology*, chapter Nanostructures in Motion: Micro-Instruments for Moving nm-Scale Objects. A&P Press, editor G. Timp, 1996. Forthcoming.

[Mac96b] Noel C. MacDonald. SCREAM microelectromechanical systems. *Journal of Microelectronic Engineering*, February/March 1996. Special Issue on Nanotechnology, invited paper.

[MC95] Brian Mirtich and John Canny. Impulse-based simulation of rigid bodies. In *Symposium on Interactive 3D Graphics*, Monterrey, CA, 1995.

[McG90] T. McGeer. Passive dynamic walking. *Int. Journal of Robotics Research*, 1990.

[MJ95] Noel C. MacDonald and Ali Jaizairy. Microelectromechanical Fabry-Perot interferometer. In *SPIE*, 1995.

[MJU91] Paul Moncevicz, Mark Jakiela, and Karl Ulrich. Orientation and insertion of randomly presented parts using vibratory agitation. In *ASME 3rd Conference on Flexible Assembly Systems*, September 1991.

[MM96] R. E. Mihailovich and N. C. MacDonald. Dissipation measurements of vacuum-operated single-crystal silicon resonators. *Sensors and Actuators*, 1996.

[MZSM93] R. E. Mihailovich, Z. L. Zhang, K. A. Shaw, and N. C. MacDonald. Single-crystal silicon torsional resonators. In *Proc. IEEE Workshop on Micro Electro Mechanical Systems (MEMS)*, pages 155–160, Fort Lauderdale, FL, February 1993.

[New77] John Nicholas Newman. *Marine Hydrodynamics.* MIT Press, Cambridge, MA, 1977.

[NW78] James L. Nevins and Daniel E. Whitney. Computer-controlled assembly. *Scientific American*, 1978.

[PBM95] Rama Prasad, Karl-Friedrich Böhringer, and Noel C. MacDonald. Design, fabrication, and characterization of SCS latching snap fasteners for micro assembly. In *Proceedings of the ASME International Mechanical Engineering Congress and Exposition (IMECE)*, San Francisco, California, November 1995. `http://www.cs.cornell.edu/home/karl/Snap`.

[PFH90] K. S. J. Pister, R. Fearing, and R. Howe. A planar air levitated electrostatic actuator system. In *Proc. IEEE Workshop on Micro Electro Mechanical Systems (MEMS)*, pages 67–71, Napa Valley, California, February 1990.

[PMT95] Rama Prasad, Noel C. MacDonald, and Dean Taylor. Micro-instrumentation for tribological measurements. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, Stockholm, Sweden, June 1995.

[PS88] Michael A. Peshkin and Arthur C. Sanderson. Minimization of energy in quasistatic manipulation. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 421–426, Philadelphia, PA, 1988.

[Ray45] John William Strutt Rayleigh. *The Theory of Sound.* Dover, New York, 2nd edition, 1945.

[RBC97] Dan Reznik, Stan Brown, and John F. Canny. Dynamic simulation as a design tool for a microactuator array. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Albuquerque, NM, April 1997. `http://www.cs.berkeley.edu/~dreznik/mems.html`.

[RCG97] Dan Reznik, John F. Canny, and Ken Y. Goldberg. Analysis of part motion on a longitudinally vibrating plate. In *Int. Conf. on Intell. Robots & Systems (IROS)*, Grenoble, France, September 1997.

[RHPR93] M. H. Raibert, J. K. Hodgins, R. R. Playter, and R. P. Ringrose. Animation of legged maneuvers: jumps, somersaults, and gait transitions. *Journal of the Robotics Society of Japan*, 11(3):333–341, 1993.

[Ril83] Frank J. Riley. *Assembly Automation, A Management Handbook.* Industrial Press, New York, 1983.

[RK92] Elon Rimon and Dan Koditschek. Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation*, 8(5), October 1992.

[RL86] A. H. Redford and E. Lo. *Robots in Assembly.* Halsted Press, New York, 1986.

[RW95] John Reif and H. Wang. Social potential fields: A distributed behavioral control for autonoomous robots. In K. Goldberg, D. Halperin, J.-C. Latombe, and R. Wilson, editors, *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, pages 431–459. A. K. Peters, Wellesley, MA, 1995.

[Sam93] J. B. Sampsell. The digital micromirror device and its application to projection displays. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, pages 24–27, Pacifico, Yokohama, Japan, June 1993.

[San91] Ben-Zion Sandler. *Robotics: Designing the Mechanisms for Automated Machinery.* Prentice Hall, 1991.

[SBW$^+$94] C. W. Storment, D. A. Borkholder, V. Westerlind, J. W. Suh, N. I. Maluf, and G. T. A. Kovacs. Flexible, dry-released process for aluminum electrostatic actuators. *Journal of Microelectromechanical Systems*, 3(3):90–96, September 94.

[Sch87] Berhard J. Schroer. Electronic parts presentation using vibratory bowl feeders. *Robotics*, 3, 1987.

[SGD$^+$96] John W. Suh, Steven F. Glander, Robert B. Darling, Christopher W. Storment, and Gregory T. A. Kovacs. Combined organic thermal and electrostatic omnidirectional ciliary microactuator array for object positioning and inspection. In *Proc. Solid State Sensor and Actuator Workshop*, Hilton Head, NC, June 1996.

[SM95] M. Taher A. Saif and Noel C. MacDonald. A milli newton micro loading device. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, Stockholm, Sweden, June 1995.

[SM96] Kevin A. Shaw and Noel C. MacDonald. Integrating SCREAM micromechanical devices with integrated circuits. In *Proc. IEEE Workshop on Micro Electro Mechanical Systems (MEMS)*, San Diego, CA, February 1996.

[SS87] N. C. Singer and W. P. Seering. Utilizing dynamic stability to orient parts. *Journal of Applied Mechanics*, 54:961–966, December 1987.

[SWH95] S. E. Salcudean, N. M. Wong, and R. L. Hollis. Design and control of a force-reflecting teleoperation system with magnetically levitated master and wrist. *IEEE Transactions on Robotics and Automation*, 11(6), December 1995.

[SZM93] Kevin A. Shaw, Z. Lisa Zhang, and Noel C. MacDonald. SCREAM I: A single mask, single-crystal silicon process for microelectromechanical structures. In *Transducers — Digest Int. Conf. on Solid-State Sensors and Actuators*, Pacifico, Yokohama, Japan, June 1993.

[TF90] N. Takeshima and H. Fujita. Design and control of systems with microactuator arrays. In *Proc. IEEE Workshop in Advanced Motion Control*, pages 219–232, Yokohama, Japan, March 1990.

[Tim40] Stephen Timoshenko. *Theory of Plates and Shells.* McGraw-Hill, New York, 1940.

[XMM95] Yang Xu, Scott A. Miller, and Noel C. MacDonald. Microelectromechanical scanning tunneling microscope. *Bulletin of the American Physical Society*, 40(1):63, 1995.

[ZE96] N. B. Zumel and M. A. Erdmann. Nonprehensile two palm manipulation with non-equilibrium transitions between stable states. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Minneapolis, MN, April 1996.

[ZM92] Z. Lisa Zhang and Noel C. MacDonald. An RIE process for submicron, silicon electromechanical structures. *Journal of Micromechanics and Microengineering*, 2(1):31–38, March 1992.