ACM Doctoral Dissertation Award 1987

The Complexity of Robot Motion Planning

John F. Canny

The MIT Press

The Complexity of Robot Motion Planning

John Canny

The MIT Press
Cambridge, Massachusetts
London, England

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data
Canny, John.
The complexity of robot motion planning / John Canny.
p. cm. -- (ACM doctoral dissertation awards ; 1987)
Bibliography: p.
Includes index
ISBN 0-262-03136-1
1. Robots--Motion. I. Title. II. Series: ACM doctoral dissertation award: 1987.
T J211.4.C36 1988
629.8'92--dc 19

The Complexity of Robot Motion Planning

ACM Doctoral Dissertation Awards

1982

Area-Efficient VLSI Computation
by Charles Eric Leiserson

1983

Generating Language-Based Environments
by Thomas W. Reps

1984

Reduced Instruction Set Computer Architectures for VLSI
by Manolis G.H. Katevenis

1985

Bulldog: A Compiler for VLIW Architectures
by John R. Ellis

1986

Computational Limitations of Small Depth Circuits
by Johann Håstad

Full Abstraction and Semantic Equivalence
by Ketan Mulmuley

1987

The Complexity of Robot Motion Planning
by John F. Canny

In memory of Pat Canny

List of Figures

1.1 (a) A two link planar manipulator, and (b) the constraints on motion in its configuration space. 6

1.2 (a) Peg and hole environment, (b) Configuration space showing locus of reference point of peg during compliant motion. 11

2.1 Position of edge e_B outside the wedge $\hat{e}_A$, where inside/outside is determined by $C^+_{e_A,e_B}$ 32

2.2 The two types of contact between polyhedra 33

2.3 Collision detection example, (i) L-shaped object moves from configuration C_i to C_f, (ii) Collision points and contact types along the path. 44

4.1 The set S, and a slice through it by the plane P_c 79

4.2 Locus of extremal points of slices $x = c$ as c is varied. 80

4.3 The silhouette of S and the recursively computed roadmap of the slice through one of the critical points. 83

4.4 Critical points of the projection map from the sphere to $\Re^2$ 89

4.5 Example of a map f transversal to a submanifold V of the x-y-plane, the preimage of V is a circle in this case 91

4.6 Example where f fails to be transversal to V. The tangent to V and the image of the tangent plane of M are parallel at critical values of f. $f^{-1}(V)$ is not a manifold in this case 92

4.7 The Cartan Umbrella $y^2 = zx^2$ 96

4.8 Stratification of the Cartan Umbrella. (a) Stratification which is not Whitney regular, (b) Whitney regular stratification 96

4.9 Cylinder of example 4.6 and its stratification into sign-invariant sets. Surfaces have thin outline, curves are thick. The interior stratum (sign $< -, +, - >$) is not shown. 99
4.10 The silhouette of the torus 107
4.11 The roadmap of the torus, showing recursively computed roadmaps of slices 115

5.1 The simplified Voronoi diagram in the plane. The central triangle is the object, the circle is its reference point. Physical obstacles are shown in black, and the configuration space obstacles are grey. 130
5.2 True Voronoi diagram for the same obstacles. 131
5.3 Illustration of the homotopy of lemma 4.2 140
5.4 Homotopy to continuously link a deforming path to a point on a convex bisector 142

6.1 Path classes and virtual sources for a single slit. 150
6.2 Path splitter, showing the principle of path unfolding. 152
6.3 Path splitter showing doubling of the number of path classes. 152
6.4 Path shuffler. The first two plates are shown at the top, and plates two, three, and four at the bottom of the figure. 154
6.5 A literal filter with barrier to stretch paths having $b_i = 1$. 155
6.6 A clause filter. 157
6.7 Path splitter in an asteroid environment. 165
6.8 Path "reflection" for the asteroid problem. Path classes on object B move normal to B at the maximum velocity. 165
6.9 A filter for motion in the x-y plane. 171
6.10 A filter for horizontal motion in the left half-plane. 171
6.11 A one-way gate. 172
6.12 The legal move filter, consisting of three boxes. 174
6.13 Tape encoding. 175
6.14 Tape and state logic (one level only). 176
6.15 State logic (global view). 177
6.16 tape logic for shifting beyond the end of the tape 178

Series Foreword

The ACM Doctoral Dissertation Award is presented each year by the Association for Computing Machinery (ACM). This award recognizes the best English-language dissertation written in a computer-related field during the academic year. The winning author receives a cash award of $1,000 from ACM. In addition, the winning dissertation is published by the MIT Press, with the author receiving appropriate royalties.

This was the sixth competition to be held. Schools were asked to nominate their best dissertations among those accepted between July 1, 1986, and June 30, 1987. By a screening process in which each of the forty-one nominated dissertations was reviewed by several external referees, five dissertations were selected for final review. A committee with representatives from both industry and academia then read the five final selections and chose the winner. The members of the committee were Doug Degroot, David Johnson, Fred Maryanski, Jack Minker, Steve Muchnik, and Larry Snyder. The dissertation they chose was "The Complexity of Robot Motion Planning," by John F. Canny. Dr. Canny's thesis work was supervised by Professor Tomás Lozano-Pérez of the Massachusetts Institute of Technology.

Canny's thesis is a ground-breaking study of the algorithmic aspects of robotics, with material enough for two award winners. In particular he is concerned with the problem of planning 3-dimensional movement in the presence of obstacles and other constraints. He resolves long-open problems concerning the complexity of motion planning problems, and for the central problem of finding a collision free path for a jointed robot in the presence of obstacles, he obtains exponential speedups over existing algorithms by applying high-powered new mathematical techniques. Moreover, he has paid much attention to his representations and algorithms, so that both the high level mathematics and the low level implementation details are presented with unusual precision and clarity. The work presents a truly outstanding example of doctoral research in computer science.

David S. Johnson, Chairman
ACM Doctoral Dissertation Award Subcommittee

Preface

Robotics is a discipline that lies on the boundary between several fields. In this thesis we are interested in the computational complexity of planning collision-free motions for a robot in an environment filled with obstacles. The tools for analysis of algorithms and complexity come from computer science, but we also need ideas from differential topology and commutative algebra. The results we obtain depend crucially on the use of all of these ideas, and would be impossible without the same level of cross-fertilization.

The main problem we tackle is the "generalized movers' problem", which is the problem of moving a general robot, which could be an arm or a mobile robot, in a three dimensional environment filled with obstacles. The first general solution for this problem was given by Schwartz and Sharir [SS], who showed that it could be solved in time double-exponential in the number of degrees of freedom. Since then, there have been steady improvements in algorithms for special cases of the problem, and these often have exponents that are linear in the number of degrees of freedom. However, they involve enormous constants due to the cost of exact computation on algebraic numbers. The main contribution of this thesis is the roadmap algorithm, which solves the generalized movers' problem in single-exponential time, with exponent equal to the number of degrees of freedom. This algorithm equals or betters the asymptotic performance of most special purpose algorithms, and has dramatically lower constants.

The improvement derives from the following ideas: Firstly, we use a coarse partition or *stratification* of the configuration space of the robot. This stratification divides configuration space into a small number of simple geometric pieces. Secondly, we use *multivariate resultants* for solution of system of polynomials in many variables. The multivariate resultant allows elimination of several variables in single-exponential time. Finally, from bounds on the size of the multivariate resultant, we derive a very general gap theorem for systems of polynomials. The gap theorem gives lower bounds on the separation of distinct algebraic numbers defined by any system of polynomials of a certain size. The theorem makes it possible to use binary approximations to algebraic numbers with no loss of accuracy.

The roadmap algorithm can either be used directly on the set of collision-free configurations, or on a subset of maximal clearance configurations. This suggests the use of Voronoi diagrams, but the usual definition is very difficult to work with in a high-dimensional configuration space. Instead we introduce a new kind of Voronoi diagram which is based on a distance function derived directly from the boundary surfaces in configuration space. Its algebraic complexity is lower in many cases, and it is shown to be complete for motion planning.

We also give lower bounds for some extensions of the generalized movers' problem. While the movers' problem is solvable in polynomial time for any fixed number of degrees of freedom, when a shortest path is sought, or when the obstacles can move faster than the robot, the planning problem is provably hard even in three dimensions. In particular we show that the euclidean shortest path problem in three dimensions, a major open problem in computational geometry, is NP-hard. The "2-d asteroid avoidance problem" is also shown to be NP-hard. Finally, when compliant (sliding) motion of the robot is allowed, and uncertainty is taken into account, motion planning in three dimensions becomes hard for non-deterministic exponential time.

The thesis contains several chapters with fairly heavy mathematical content. For many of the geometric ideas, figures have been included. These should give the general reader a good flavor of what is going on. The dependencies between chapters have been kept to a minimum, and each can be read in isolation. The exception to this is chapter 1, which contains definitions and terminology used throughout the later chapters. Chapter 2 concerns computation of configuration space descriptions from physical descriptions of the robot and obstacles. It contains a short summary of the properties of quaternions, which are used to represent the configuration of a free polyhedron. Chapter 3 describes the multivariate resultant and the gap theorem. Chapter 4 is the heart of the thesis. Stratified sets are introduced, and the roadmap algorithm is described. Some familiarity with singularity theory is useful here. Chapter 5 describes the simplified Voronoi diagram, and uses simple homotopy arguments to show that the diagram is complete for motion planning. Chapter 6 is pure computational geometry. It gives lower bounds for problems of 3-d shortest path, 2-d asteroid avoidance, and 3-d complaint motion planning with uncertainty.

Acknowledgements

My sincere thanks to Prof. Tomás Lozano-Pérez, my thesis advisor at MIT for all his support, understanding and perspective throughout the work on this thesis. In addition to Tomás, I was lucky enough to be able to interact with some other truly exceptional individuals. John Reif taught me a geat deal about complexity theory in robotics and nature, and David Mumford gave me invaluable direction in the search for the right kind of mathematics to apply in this new field.

Thanks to Bruce Donald and Mike Erdmann for detailed comments on the thesis and for sharing the learning process. This version has greatly benefited from comments from Ken Goldberg, Barbara Moore and Doug Ierardi. Finally thanks to Stuart Russell, Robert Wilensky and the people in BAIR for providing, and keeping together, the text processing resources I needed.

The Complexity of Robot Motion Planning

Contents

List of Figures xi
Series Foreword xiii
Preface xiv

1 Introduction **1**
1.1 Robot Motion Planning Problems 4
1.1.1 The Generalized Movers' Problem 4
1.1.2 The Shortest Path Problem 8
1.1.3 Dynamic Motion Planning with Velocity Bounds 9
1.1.4 Compliant Motion Planning with Uncertainty 9
1.2 Algebraic Decision Procedures 12
1.2.1 The Theory of the Reals 13
1.2.2 Existential Theories, Complex Numbers and the Reals 15
1.2.3 Standard or Gröbner Bases 17
1.3 Overview of the Thesis 18

2 Motion Constraints **21**
2.1 Object-Obstacle Constraints 24
2.1.1 Disjunctive Form 24
2.1.2 Conjunctive Form 29
2.2 Contact Conditions 33
2.3 Constraints for Free Polyhedra 34
2.3.1 Review of Quaternions 34
2.3.2 Transformation of Object Features 37
2.3.3 Primitive Predicates 39

2.4 An Algorithm for Collision Detection 40
2.5 Collision Detection for Several Moving Objects 45
2.6 Constraints for Jointed Manipulators 47

3 Elimination Theory 52
3.1 The Multivariate Resultant 53
3.1.1 Partial Coefficient Specialization 56
3.1.2 Method 1 58
3.1.3 Method 2 59
3.2 The u-Resultant 66
3.3 Degree bounds and Gap Theorems 66
3.4 Solution of Polynomial Systems 71

4 The Roadmap Algorithm 76
4.1 An Example 78
4.2 Preliminaries 85
4.3 Stratifications 94
4.4 The Silhouette 102
4.4.1 Tubes Around Varieties 108
4.5 The Roadmap 112
4.5.1 The Basic Roadmap 113
4.5.2 Linking Curves 117
4.5.3 The Full Roadmap 118
4.6 The Roadmap Algorithm 120

5 Performance Improvements 128
5.1 The Simplified Voronoi Diagram 129
5.1.1 A Generalized Distance Function 133
5.1.2 Completeness for Motion Planning 135
5.2 Complexity bounds 146

6 Lower Bounds for Motion Planning 148
6.1 Lower Bounds for the Shortest Path Problem 149
6.1.1 Free Path Encoding for Shortest Paths 149
6.1.2 The Environment 150

6.1.3 Lower Bounds for 3-d Shortest Path 156
6.1.4 The Virtual Source Approximation 159
6.1.5 Environment Size 161
6.2 Lower Bounds for Dynamic Motion Planning 164
6.3 Motion Planning with Uncertainty 168
6.3.1 Blot Motion 169
6.3.2 The Environment 169
6.3.3 Legal Move Filter 170
6.3.4 Tape and State Logic Structures 173
6.3.5 Initialization and Termination 177

7 Conclusions 180
7.1 Algebraic Decision Methods 180
7.2 Robot Motion Planning Algorithms 182
7.3 Lower Bounds 185

References 187

Index 196

Chapter 1

Introduction

The objective of robot motion planning is to automatically guide a robot around in an environment filled with obstacles, and to choose forces to be applied when assembling objects. Economic considerations often mandate the use of a cost function on the robot's motion, such as time duration or the amount of energy required. When planning a *trajectory* for the robot, which gives position as a function of time, planning must take into account physical limitations of the robot. For example, the motors that drive the robots's joints will have limited torque. The sampling rate of the digital controllers for the motors usually limits the maximum speed of a trajectory that the robot can track successfully. So there are many possible variations on the motion planning problem, where some or all of these extra constraints are taken into account. This thesis studies the complexity of several fundamental classes of problem, and introduces some novel geometric techniques for solving them.

The most difficult planning problems arise in assembly, where clearances between parts are typically smaller than the robot's positioning or sensing accuracy. The robot cannot consistently execute motions below this accuracy, and when it attempts to do so, the results of its actions are no longer deterministic. Such a motion can lead to any of several qualitatively different situations. There is no obvious alternative open to the planner except the enumeration of all the possibilities, leading to a combinatorial explosion with the number of plan steps.