SLICE-BASED PATH PLANNING

by

Michael Charles McHenry

A Dissertation Presented to the
FACULTY OF THE GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA
In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(Computer Science)

Dec 1998

Copyright 1998 Michael Charles McHenry

Acknowledgements

I thank Professor George Bekey for being my advisor and supporter for many years. His energy and infectious enthusiasm never ceases to amaze me. I also thank the other members of my committee; Professors Michael Arbib, Wei-Min Shen, and Bartlett Mel.

During the years I worked in industry, I was lucky to work with a group of people whose intellectual breadth and curiosity really motivated me to pursue a Ph.D. In particular, my mentor during those years, Dr. Thomas John, showed me the joy of intellectual pursuit and helped imbue me with math envy.

During my many years of graduate school, I was lucky enough to car pool with a series of people who made Los Angeles freeway traffic always bearable and often times great fun. Joe Devlin and the rest of the Goodman Gang, including Joe Clougherty, Michele De Sando, and Marc Dillard, all gave me a taste of academic interests far removed from typical computer science. Mariana Jackson breathed life into my otherwise sedate existence.

The endless passion and energy of my current car pool mate and lab mate James Montgomery helped get me through many hours of commute to say nothing of years of graduate school. Monty instilled in me a greater appreciation of nature and nutrition. He showed me the many joys of gardening for which I will forever be grateful. And perhaps most significantly, he showed me that kindness and enthusiasm can serve as a magnet which attracts fun and interesting people to accomplish Herculean tasks.

When I arrived in the Robotics Lab, Andy Fagg taught me how to read scientific papers critically and helped set me on the right track. In fact, Andy pretty much built the track I rode on! Andy is among the most selfless and generous people I have ever met. And his seeming immunity to cynicism helped make the rough spots bearable.

Many other fellow lab minions helped make the late nights enjoyable. Jason Chu on occasion let me think I was wise. Parag Batavia showed me what my undergraduate education should have been like. Scott Cozy was a master craftsman and an inspiration. Steve Goldberg patiently held my hand through my limited travels into Gates-land. Stergios Roumeliotis gave me both a burning desire to visit Greece and a greater appreciation for the United States of America.

Gaurav Sukhatme was an immovable rock of wisdom and sanity in a sea of craziness. I hope Gaurav has a long tenure as a Professor as he is the best teacher I have had the pleasure of studying with.

Julieta DeLa Paz keeps the department running. Her can-do attitude and problem solving nature helped insulate me from the cold machinery of the USC bureaucracy.

Kusum Shori was a source of assistance a thousand different ways and a thousand different times. It is hard to describe the positive effect of her kind words and constant smile.

Of all the fellow students and academics I met while in graduate school, none came close to demonstrating the pure joy of learning and intellectual curiosity that my housemate Richard Worthington did. He always made me feel lucky to be in graduate school. I also thank Richard for so generously sharing his home with me for so many years. I must also express my gratitude to my other two housemates; Ginger and Sage. How anybody ever makes it through graduate school without the unbounded affection of a dog I cannot fathom.

When I finished my undergraduate education I was made uneasy by the realization that for the first time in my life there was no fixed extent to my current life stage. If I did nothing I might be living in the same place and be doing approximately the same thing until I retired. Now I find I am once again entering the nebulousness of the real world. Yet the earlier unease has been replaced by hope, excitement, and anticipation. This change is entirely due to my loving wife Denise, to whom I give my thanks for her patience and for sharing our mutual future.

Finally, my thanks go to my family. My sisters Teresa and Pat were constant and loving supporters and friends through all the years. My parents Robert and Rose nurtured and supported me for all my life. Their faith and love made it possible. Thank you.

Contents

List Of Tables

List Of Figures

Abstract

This study addresses the problem of incremental path planning. A novel algorithm called Slice-Based Path Planning is described. This algorithm is based on constraining the search for a collision free path to lie within a slice, a two-dimensional non-planar manifold of the robot's full configuration space. Constraining the search space in this way allows the algorithm to scale to problems which would otherwise be infeasible. In addition, the smaller search space allows the application of a global (with respect to a slice) path planning method such as dynamic programming. The slice is adapted to a particular set of workspace obstacles and goal positions using an iterative optimization procedure. The slice-based approach offers two significant advantages. First, an existing slice can incrementally adapt to changes in the locations or number of obstacles. Secondly, the global intra-slice path planner can find a satisfactory free path quickly which is then continually improved upon as the slice optimization proceeds. The implementation uses a grid-based representation of a slice which makes the algorithm well suited to a massively parallel implementation. The algorithm has been evaluated using three sets of problems. The first set measures the basic algorithm's performance in environments with static obstacles. The second set measures the advantage of incrementally updating an existing slice when new obstacles are added to the workspace. The third set evaluates the algorithm's ability to continually adapt slices in environments with moving obstacles. The planner has also been demonstrated performing a simple reaching and grasping task using a Belgrade/USC hand mounted on a PUMA 560 manipulator.

Chapter 1

Introduction

This research addresses the problem of incremental path planning and describes a novel algorithm called Slice-based Path Planning (SBPP). The SBPP algorithm can be applied to two situations not commonly addressed by path planning algorithms. The first occurs when the available planning time is finite but unknown. The SBPP algorithm is able to produce a satisfactory plan quickly and then gradually improve that plan for the remaining available time. The second situation occurs when the obstacles are frequently or continually moving in the workspace. In this situation two approaches are possible. One approach is to develop a planner which can incorporate predictions of obstacle motion directly in the produced plan. The alternative approach embodied in the SBPP algorithm is to continually and incrementally update the path plan based on the most recently sensed positions of the obstacles. In this way SBPP tracks ongoing changes in the workspace rather than relying on potentially inaccurate predictions of obstacle motion. The prediction-based approach and the incremental approaches are not mutually exclusive and future work will consider incorporating both.

The two situations described above occur frequently in real-world applications. Consider for example, two loosely coordinated robots which share a common tool. While one robot is using the tool the other robot can continually refine its plan until the other robot completes its task and releases the tool. This is an example of the "finite but unknown planning time" situation.

As another example, consider a mobile robot planning a complicated path across a large and densely occupied factory floor. Suppose further that the path planning computation requires five minutes to produce a satisfactory plan but that new obstacles are periodically added or existing obstacles are occasionally removed. If the path planning algorithm is required to replan from scratch each time the obstacles change, then the planner will frequently be frustrated in its effort to produce a plan which accurately reflects the current workspace. The incremental SBPP algorithm is able to make efficient use of the result of its earlier computation and thus avoids having to replan from scratch each time portions of the workspace change.

Algorithms such as the SBPP algorithm with the ability to improve the quality of results gradually as computation time increases have been referred to as "anytime algorithms"[15]. Such algorithms offer another potential advantage in situations with limited computation resources. Higher level control software can more intelligently allocate scarce resources to those anytime components which are most critical or likely to have the greatest utility.

The SBPP algorithm offers two other key advantages. First is its scalability. Many existing path planning algorithms cannot feasibly be applied to systems with more than three or four degrees of freedom. We will show an application of SBPP to a six degree of freedom PUMA manipulator. The second key advantage is the algorithm's ability to produce paths which reflect more complicated path quality metrics than simply being collision-free.

The basic heuristic embodied in the algorithm is to decompose the path planning problem into two parts. The first uses information about the current obstacles, goal specifications, current configuration, and kinematic constraints, such as joint limits, to determine a two-dimensional manifold within the entire configuration space. This two-dimensional non-planar surface is called a slice. The second part produces a path plan contained entirely within the slice. This intra-slice path planning problem is the comparatively simple problem of planning the path of a point in two dimensions. The slice computation and the intra-slice planning can be performed concurrently.

The slice computation is implemented as an iterative numerical optimization based on a criterion containing four terms. The first term reflects a measure of the distance between the robot and the surrounding obstacles (summed over all the configurations contained within the slice). The second term is based on the distance between any of the joint variables and their physical limits, again for all the configurations contained within the slice. The third and fourth terms involve measures of the continuity and smoothness of the slice. This helps to produce an intra-slice path which is itself continuous and smooth.

The overall algorithm has two major limitations. The first is that the particular form of the slice is currently restricted to being a vector valued functions of x and y. Many potential paths (such as those including a pure vertical motion) can not be represented using manifolds of this form. In addition, the slice is computed using an iterative optimization technique. Any such optimization technique is prone to problems of local minima[6].

The intra-slice path planning problem is small enough that nearly any path planning algorithm can be applied. In fact the computation required for intra-slice path planning problem is negligible in comparison to slice optimization. For this reason, the current implementation of the overall Slice-based Path Planner uses a non-incremental dynamic programming approach to intra-slice planning. [1]

In some sense the algorithm compiles information about the workspace obstacles in the formation of the slice. This compilation represents an investment of processing. The benefits achieved by that compilation will vary depending on particular aspects of the robot and of the distribution of workspace obstacles. For example, consider a change in the environment which quickly cuts off what had been the best intra-slice path. If the slice contains multiple satisfactory paths then a new satisfactory path can be found with a single iteration of the intra-slice component. If the slice doesn't contain any other satisfactory paths the path quality will decrease

[1] With the addition of a path distance component to the path quality metric, an incremental dynamic programming approach is possible. However this approach was found experimentally to be less efficient than the non-incremental version. This issue is discussed in greater detail in Chapter 3.

sharply. Subsequent improvement will require adaption of the slice which requires significantly more computation. Furthermore, the slice optimization process finds only local minima.

The rest of this chapter is organized as follows. The next section provides a more detailed problem statement. Section 1.2 covers the motivations for incremental motion planning in more detail than above. Section 1.3 outlines the algorithm and provides a brief summary of the results. The final section summarizes the specific contributions made by this thesis.

1.1 Problem Statement

It is important to distinguish between "path planning" and "motion planning." The term "path planning" is used to describe the problem of producing a collision free path with no specification of velocity. In other words, a path plan describes the path traced out during a motion, independent of time. "Motion planning" on the other hand, describes the problem of producing a collision free motion which includes a description of the time at each point of the motion. More formally, a path plan is a curve in configuration space, while a motion plan is a curve in the Cartesian product of the configuration space and time. [2] Any motion planning system can be used as a path planning system by disregarding the temporal specification.

The basic non-incremental path planning problem is shown in Problem 1.

[2]A feasible motion plan must be monotonic in time.

Problem 1 (Non-Incremental Path Planning Problem)

Given:	*Obstacle Geometries*
	Obstacle Positions
	Robot Geometry and Kinematics – determines configuration space $\mathcal{C}$
	Start and Goal Configurations c_s, c_g
Output:	*Collision-free Path* $p(s)$ *Connecting Start and Goal* [a]
Computation Metric:	*Computation Time*

[a] $p(s) : s \mapsto \mathcal{C}\ \forall s \in (0, 1),\ p(s)$ is not a collision causing configuration, $p(0) = c_s,\ p(1) = c_g$

The collision-free path must also obey the specified kinematic constraints. In the case of serial link manipulators these constraints will normally consist of physical joint limits.

Computation Metrics

The computation metric describes how the path planning algorithm itself will be evaluated. The most common measure for a non-incremental path planner is the time elapsed between when the algorithm begins computing and the time when a collision free path is produced. As will be discussed in Chapter 2, many algorithms have an initial preprocessing step during which a compact representation of freespace is computed. For these kinds of algorithms it is reasonable to evaluate the algorithm based on two times: the pre-processing time and the time to produce a path between a start and goal configuration.

The above metric is empirical and thus makes it difficult to extrapolate performance to different problem instances. Another common metric which is better able to relate novel problem instances to algorithm performance is the worst case time complexity. Typically this is expressed as a function of parameters describing the complexity of the obstacle and robot geometry such as the number of polygons and the number of total vertices.

Both metrics are ill-suited to evaluating incremental algorithms such as SBPP for several reasons. First, since the computed plan is continually changing it does not complete at a fixed moment of time. Secondly, since the algorithm is incremental, its performance will depend on the initial conditions of the algorithm's internal data structures. In the SBPP algorithm this internal data structure consists primarily of a representation of the current slice. Another difficulty is that the worst-case time complexity metric is normally based on a polygonal description of the robot's and the obstacle's geometry, whereas SBPP uses a discretized representation. There is no concise measure of a discretized representation which can easily be related to expected performance of an incremental path planning algorithm. The final reason that neither the worst-case bound nor the empirical complete time is a satisfactory metric is that we are also interested in evaluating the produced plans according to more complicated path metrics. The metric we will use includes evaluations of clearance between the robots and obstacles, path distance, and how close joint angles are to their physical limits. The particular metrics used in the current implementation are described in Section 3.3.

The metric we will use to evaluate SBPP is path quality expressed as a function of computation time. The metric is fundamentally empirical, hence it is difficult to separate the influence of specific aspects of a problem instance from the general capabilities of the algorithm. We address this issue by collecting information from many instances of the same general class of planning problems and providing a statistical measure of path quality expressed as a function of time. These profiles which map from a computation time to a probability distribution of the quality of the result are called "performance distribution profiles". This performance metric is very similar to that described by Zilberstein & Russel [56].

Since a computational metric based on completion time is ill-suited to incremental path planning, we will address a variant of Problem 1. This variant, Problem 2, differs from Problem 1 in two ways: the additional input of a path metric which maps an entire path to a scalar value; and the computation metric now describes how the path quality evolves during the computation. t_c denotes computation time, i.e. elapsed time since the computation began.

Problem 2 (Incremental Path Planning Problem)

Given:	*Obstacle Geometries* *Obstacle Positions* *Start and Goal Configurations* *Robot Geometry and Kinematics* *Initial Slice* *Path Metric* $J_p : p \mapsto \mathbb{R}$
Output:	$p(t_c)$ *Collision-free Path Connecting Start and Goal Configurations Computed in Time* t_c
Computation Metric:	$J_p(t_c)$ *Path Quality as a Function of Computation Time*

Problem 3 removes the restriction that obstacles be fixed. This will demonstrate how the incremental nature of SBPP allows the computed path to track changes in the obstacle locations and goal specification.

The focus of the described research is on the incremental aspects of the planner. The planner is not in any way modeling the motion of the obstacles. At each iteration the planner produces a path which is based on the obstacles' most recently sensed positions and assumes that the obstacles are now fixed. It is as if the planner were continually producing a path based on a robot which could move infinitely fast. Clearly if the robot moves very slowly in comparison to the motion of obstacles then the plan produced at the time the motion begins execution will not succeed. Future work will explore two strategies for addressing this problem. Under the first strategy, the planner would continue to replan during execution of the movement. Under the second strategy, the planner would incorporate a prediction of the obstacles future movement and knowledge of the robot's maximum speed directly into the produced motion plan.

Problem 3 (Incremental Path Planning Problem with Changing Obstacles)

Given:	*Obstacle Geometries*
	Obstacle Positions and Orientation as Functions of Time
	Start Configuration
	Goal Configuration as Function of Time
	Robot Geometry and Kinematics
	Initial Slice
	Path Metric $J_p : p \mapsto \mathbb{R}$
Output:	$p(t_c)$ *Collision-free Path Connecting Start and Goal Configurations*
Computation Metric:	$J_p(t_c)$ *Path Quality as a Function of Computation Time*

1.2 Motivation

Despite the considerable research applied to the problem of path planning, many difficulties persist. These difficulties include problems of scale, representation, speed, and path quality. In this section we will briefly examine each of these difficulties, critique some existing approaches to them, and indicate how the SBPP approach addresses them. A more comprehensive treatment of past work appears in Chapter 2.

Scalability

Path planning is generally accepted to have a worst case time complexity which is exponential in the number of degrees of freedom of the robot [27]. Thus many existing algorithms are restricted to systems with few degrees of freedom (four or less). This has led to the development of many algorithms which are incomplete, i.e.

they are not guaranteed to find a collision free path even if one exists. The potential field method is one such incomplete method. It essentially performs a local search in the full configuration space and thus avoids having to form an internal representation of the entire configuration space.

The SBPP algorithm described in this document is also an example of an incomplete algorithm. The algorithm is incomplete because there is no guarantee that the computed slice contains a collision free path. However, the algorithm uses a global intra-slice path planning method, a locally minimal slice may still lead to satisfactory path plans. This is most likely when there are many possible solutions.

Representation

The worst case time complexity of many path planning algorithms is expressed as a function of the geometric complexity of the obstacles in the environment and of the robot itself. Typically this is measured in terms of numbers of polygons and the total number of vertices. Yet most real-world sensor data requires sophisticated processing to produce an accurate polygonal description of the environment. The Slice-based Path Planner avoids some of these complications by using a discretized representation. Section 4.3.1 describes the process by which the current implementation constructs the discretized representation (a 3-D obstacle bitmap) using a single overhead camera.

Path Quality Criterion

Many existing path planning algorithms are unable to reflect any measure of path quality beyond simple collision avoidance. The SBPP algorithm is based in part on variational methods. This enables it to produce paths which reflect a wider range of path quality metrics. These metrics include distance to obstacles, distance to physical joint limits, rate of change, and smoothness. [3]

[3] The rate of change and smoothness criteria are evaluated in spatial terms rather than temporal. This is discussed in more detail in Chapter 3.

Speed

Developing planners for robots is a challenging task for many reasons: sensor data is usually noisy and incomplete; control is imperfect; and the environment is constantly changing. Given the difficulty of constructing a perfect model of the environment most practical planning problems require frequent (if not constant) replanning. Yet oftentimes the planner must share very limited computational resources with other tasks such as sensor processing, communication, and actuator controls. The complexity of the task has led some researchers to eschew planning altogether and adopt pure reactive control policies. Many others have adopted hierarchical strategies which place a high level planner on top of a lower level reactive component [53].

The SBPP approach addresses the aforementioned difficulties in several ways. First, its incremental nature lowers the computational cost of replanning. Real-world environments (described by a set of obstacles and their locations) do not change discontinuously. At any point in time the set of obstacles present and their locations are very similar to the situation a few moments earlier. Yet most existing path planning algorithms (and the metrics with which those algorithms are evaluated) are designed as if motion planning problems were presented like a sequence of unrelated snapshots. Since path planning can be computationally expensive and require a substantial amount of time, it is useful to adopt path planning algorithms which can make efficient use of the continuous nature (in time) of the environment. In addition, the SBPP algorithm can be interrupted at any time and a path read-off. Thus planning can be performed whenever spare computer cycles are available without committing future cycles. Occasionally a small change in one of the environmental obstacles will necessitate a new path which lies far from the original path. Under such circumstances the portions of the slice which are not affected by the obstacle change will still contain high quality configurations. Thus it is these situations for which the slice-based path planning approach is best suited.

Secondly, the slice defines a two dimensional subspace of the entire configuration space. Subsequent search for a path is restricted to lie within this space. Plans can be generated much more quickly in this two dimensional space than in the entire configuration space.

Parallel programming is another strategy for improving the speed of path planning algorithms. The decomposition of the path planning problem employed by the SBPP algorithm (into the slice computation process and the intra-slice planning process) lends itself naturally to a coarsely parallel implementation. In addition, the discretized representation of the slice enables a massively parallel implementation of the slice computation process. The current implementation is strictly serial and a parallel implementation is left as future work.

As an anytime algorithm the SBPP can also be incorporated into more intelligent robot control architectures which can dynamically adapt the computational resources devoted to individual tasks such as sensor processing and planning [55].

1.3 Algorithm and Results

This section briefly describes the basic strategy embodied in the slice-based path planner. The key idea is to compute a good subset of the entire configuration space in which to look for collision free paths. The path plan which is ultimately produced will lie entirely within this subspace. An important aspect of a "good" subset is that it contain at least one collision free path.

The slice-based path planning algorithm can be viewed as a hybrid algorithm which combines local and global search in a unique way. The slice optimization process is a local search in the space of all slices. The intra-slice path planning algorithm (dynamic programming) is a global search method within the slice. Like most local search algorithms the slice optimization process can become stuck in local minima; but the global nature of the intra-slice search mitigates the problems of local minima slices.

As background for the discussion that follows we define two terms: configuration space and manifold. Configuration space is the set of all possible positions and orientations of the robot [33]. Each element of this set corresponds to a particular location and orientation of every part of the robot. This space can be parameterized in different ways. For example, the configuration space of a planar two rotational link manipulator can be described by $\theta_1 \times \theta_2$ or equally well by the position of the end effector, $X \times Y$.[4] Which parameterization is better will depend on the actual computation being performed. A N-dimensional manifold is a subspace which *locally* looks like R^N[51]. Thus a 2-dimensional manifold can be viewed (locally) as a sheet contained in some higher dimensional space.

Definition of a Slice

One early technique for addressing the path planning of systems with many degrees of freedom was to hold some of those degrees of freedom fixed during the motion. This technique has been called "slicing" [27]. In configuration space it corresponds to restricting search to a subspace of configuration space. In the simple parameterization of configuration space determined by each joint angle, this subspace corresponds to an axis-aligned hyperplane. The novel aspect of the SBPP algorithm is that a slice need not be planar, much less axes-aligned. Instead, the SBPP algorithm adapts the slice to reflect the obstacles in the environment. The algorithm assumes a particular kind of manifold. In particular, it uses a slice in the form of a vector-valued function of the horizontal position, (x, y). This (x, y) position would typically describe the center of mass for a mobile robot or the end-effector position for a serial link manipulator. An N degree of freedom robot will require N-2 slice variables to specify the robot's configuration (since each point in the slice determines a unique x and y position). The approach developed here assumes a parameterization of configuration space which includes the horizontal workspace coordinates (x, y) as two configuration variables. The slice's value $\vec{f} = (f_1(x, y), \ldots, f_{N-2}(x, y))$ at each position determines the configuration for that particular (x, y) position. However, depending on the parameterization of configuration space it isn't always the case that $(x, y, f(x, y)$ determines a single configuration. For example, a planar two link

[4]These parameterizations aren't actually exactly equivalent since (in the absence of joint limits) there are two sets of joint angles which produce a single end effector position.

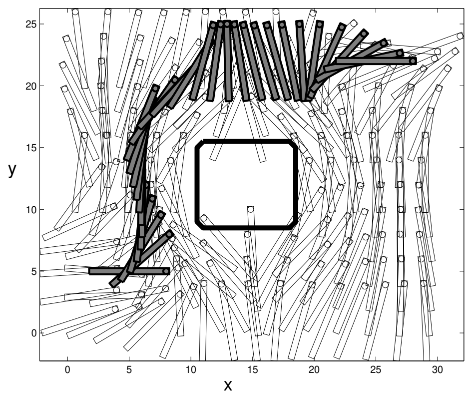

(a) The rectangle in the middle is an obstacle. Each rectangular outline shows the slice value (θ) for a particular (x, y) location.

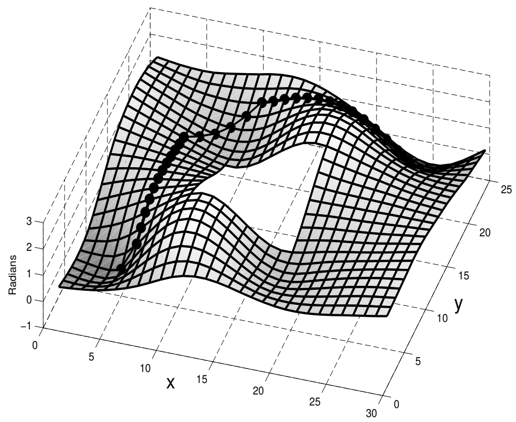

(b) Surface plot of θ. The thicker dark line shows an intra-slice path.

Figure 1.1: Two illustrations of the same slice for a robot which can translate and rotate in the plane. .

robot has in general two inverse kinematic solutions (elbow left and elbow right). In these cases the current implementation arbitrarily (but consistently) chooses a single solution.

Figure 1.1 illustrates a slice for a simple rectangular robot which can translate and rotate in the plane. In this case the slice is simply $\theta(x, y)$. This slice represents a commitment as to the robot's orientation, θ at each discretized (x, y) location.

A slice in the configuration space of a rigid body which can translate and rotate in three-space would be described as follows:

$$
\begin{aligned}
z &= f_1(x, y) \\
\gamma\ (roll) &= f_2(x, y) \\
\beta\ (pitch) &= f_3(x, y) \\
\alpha\ (yaw) &= f_4(x, y)
\end{aligned}
$$

The motivation for choosing this particular class of subspaces is two-fold. First, for many classes of robots and many classes of path planning problems the spatial dimensions (x, y, z) are in some sense "more crucial" to path planning than the orientation dimensions (γ, β, α). [5] Secondly, as will be shown later, slices parameterized by the spatial dimensions are relatively easy to compute and represent in a massively parallel fashion.

Slice-Based Path Planning Algorithm

The remaining portion of this section is devoted to a brief sketch of the slice-based path planning algorithm. Figure 1.2 illustrates the basic algorithm flowchart. The workspace is represented as a bitmap describing the obstacles $Obs_{\mathcal{W}}$. If a rectangle centered at (x_i, y_j) contains any portion of a workspace obstacle the corresponding bit in $Obs_{\mathcal{W}}$ is set equal to 1.

As part of the slice computation it is necessary to obtain a measure of the distance between the robot in a particular configuration and the surrounding obstacles. This is achieved as follows: First a potential field is computed based on $Obs_{\mathcal{W}}$. This potential field is used only to direct the slice computation and is *not* used for navigation. At each discretized (x, y) position, the corresponding configuration is read off from the current slice. The metric of how close the robot (in that particular configuration) is to an obstacle is simply the summed potential field across a set of control points located at pre-determined positions on the robot. Thus all other things being equal, a configuration which positions the robot such that the control points are at locations of low potential (further from obstacles) will be preferred.

The actual slice computation is defined as an optimization problem where the optimization criterion incorporates four components/terms, including the potential term just outlined. The second and third terms describe the magnitudes of the first and second derivatives respectively $(|(\frac{\partial f_i}{\partial x}, \frac{\partial f_i}{\partial y})|$ and $|(\frac{\partial^2 f_i}{\partial x^2}, \frac{\partial^2 f_i}{\partial y^2})|)$. These terms

[5] The volume the robot sweeps out when the x or y variables are moved from their minimum to their maximum is in general larger than the volume swept out by changes in orientation. This indicates that obstacles have a greater impact on configuration quality along the spatial axes. Obviously the degree to which this is true depends on the robot's geometry and kinematics.

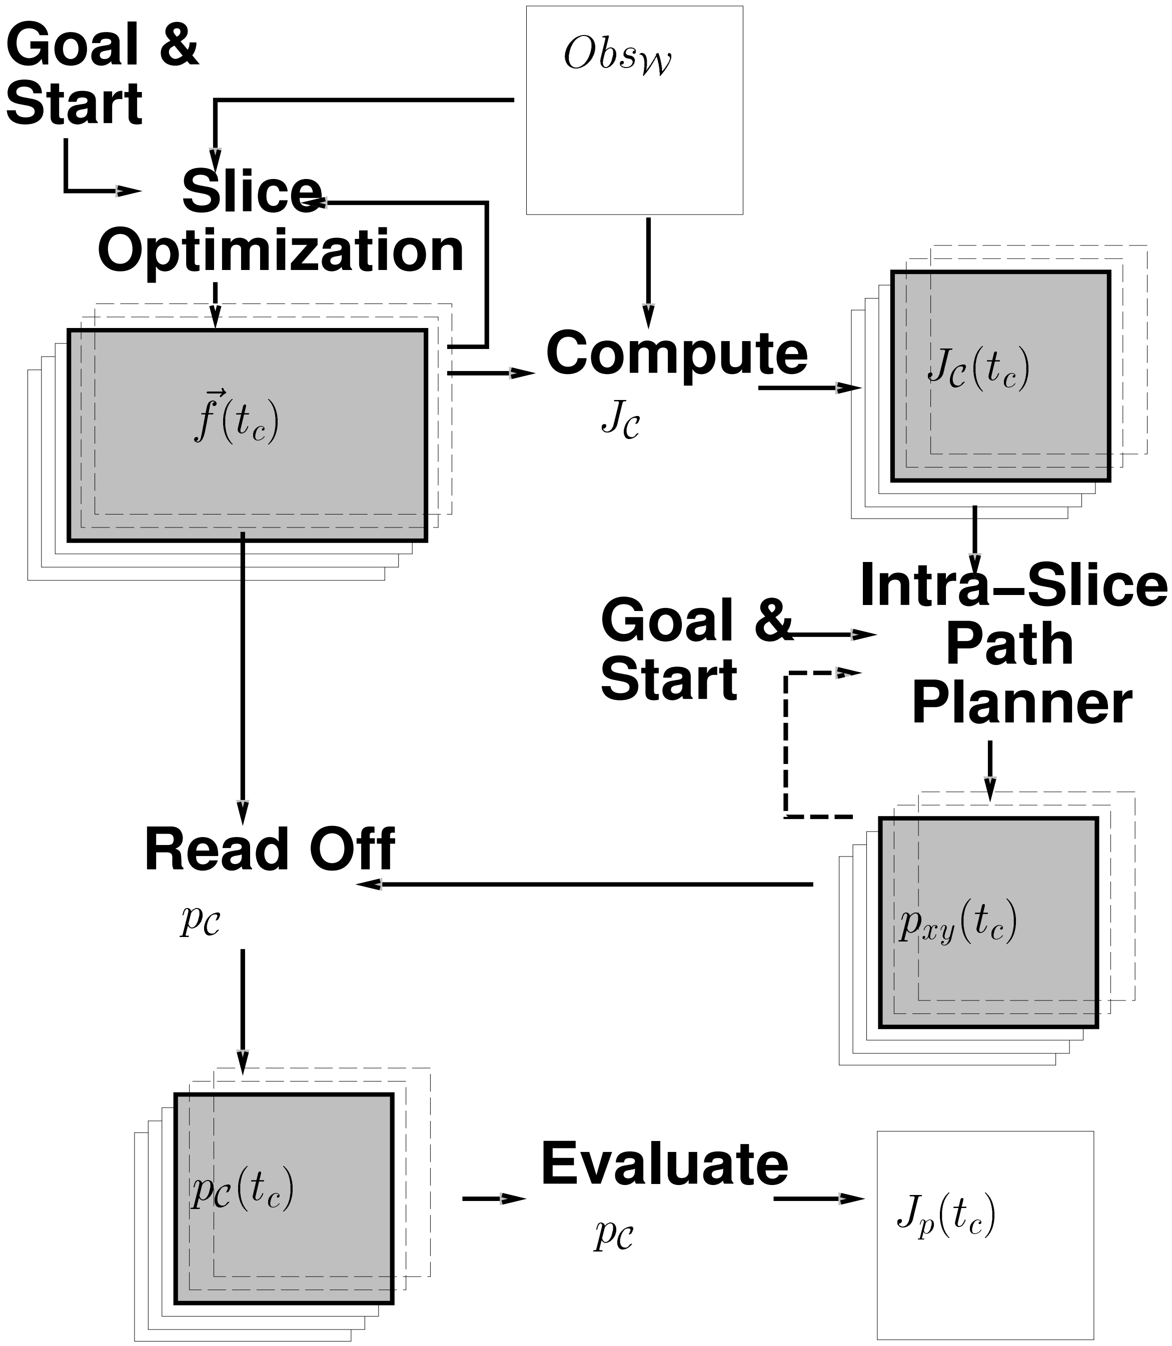

Figure 1.2: SBPP flowchart. The main input is the obstacle bitmap $Obs_{\mathcal{W}}$. At each point in the computation (represented as computation time t_c), the existing slice is modified $\vec{f}(t_c)$. $J_{\mathcal{C}}(t_c)$ is the measure of configuration quality which is applied to each element of the slice. The intra-slice path planner produces an intra-slice path$p_{xy}(t_c)$. The path containing all the configuration information $p_{\mathcal{C}}(t_c)$ is read from the slice. The final step is to evaluate the quality of the path itself $J_p(t_c)$.

help ensure that the slice is both continuous and smooth, which helps produce paths requiring limited torques and accelerations. [6] The final term reflects any joint limits.

Computing a slice is one component of the slice-based algorithm. The other key component is the computation of a collision free intra-slice path. In most situations it is not possible to produce a slice which contains only collision free configurations. However, as long as there exist connected regions of collision free configurations containing the start and goal configurations, the SBPP algorithm will still be able to produce a collision free path. Thus, finding an optimal slice is not necessary. But the slice computation needs to be fast, incremental and produce slices which satisfy the constraints.

The intra-slice path planning component requires determining which configurations contained in the slice are collision free and which are not. In other words, the intersection between the slice and the configuration space obstacles must be computed. The most direct strategy would be to use detailed information of robot and obstacle geometry to perform direct collision detection. This information could then be represented in a discretized fashion by a bitmap $(Obs_f(x, y)$. If the robot in configuration $(f_i, \ldots, f_{N-2}, x, y)$ collides with any obstacle (along any portion of the robot), then $Obs_f(x, y)$ would be set to 1. The strategy currently implemented is considerably simpler. Obs_f is replaced by a two dimensional grid of scalar values, J_C. $J_C(x, y)$ is a more complicated measure of configuration quality than the binary value representing the presence of a collision. This measure is derived directly from the slice optimization criterion. This greatly simplifies the intra-slice path planning but prevents the planner from being able to determine exactly if the best intra-slice path is collision free or not.

The two dimensional intra-slice path planning problem is small enough that global techniques such as dynamic programming can be applied. Dynamic programming has the advantage of being able to use the measure of configuration quality (J_C)

[6] The current implementation does not produce paths for which the x and y components are smooth. Thus the overall path is itself not smooth. Thus executation of the produces paths could require large torques. This problem could be addressed by using a more sophisticated intra-slice planner.

to produce a path which is optimal (according to the path criterion described later and with respect only to intra-slice paths). The final step is to translate the intra-slice path to a path in the entire configuration space by reading off the remaining configuration variables from the slice.

1.4 Contributions

The primary contribution of this work is the development of a path planning system with the following features:

1. Small changes in obstacles positions and goal configurations are incrementally incorporated into existing plans without requiring replanning from scratch.

2. The path planner produces satisfactory plans quickly and gradually improves those plans with additional computation time. [7]

The above features result from two novel aspects of the slice-based path planning approach described. The first novel aspect is the extension of existing notions of configuration space slice to the non-planar case. The second novel aspect is the decomposition of the path planning problem into a slice-optimization component and an intraslice path planning component.

A secondary contribution is the recognition of the importance and some initial examples of performance metrics which are applicable to incremental path planners. More specifically, the graphs of path quality as a function of computation time which appear in Chapter 4 illustrate concrete measures of the relationship between solution quality and computation time which should be applicable to all incremental path planners.

[7] What constitutes a satisfactory plan will be discussed in Chapter 3.

1.5 Thesis Overview

The rest of this document is structured as follows. The next chapter reviews some of the key existing literature. Many classical (non-grid based approaches) will be briefly mentioned but the primary focus will be on the most closely related methods such as potential fields. Chapter 3 describes the SBPP algorithm in detail. The following chapter describes the results of applying the algorithm to a simple planar robot and to a PUMA 560 manipulator. The final chapter summarizes the thesis contributions and outlines avenues for future development.

Chapter 2

Literature Review

In this chapter we review previous work in path planning and in other topics related to the SBPP approach. The first section outlines the field of anytime algorithms. The next section describes classical graph-based path planning approaches such as road-maps and cell decomposition. While these techniques are comparatively well developed and have resulted in many analytic results concerning complexity and completeness, the focus of this work is on grid-based techniques and thus the coverage of graph-based methods is comparatively shallow. The third section describes potential field methods and the related harmonic function based methods. The final section covers a range of work including variational methods, dynamic programming, and neural networks.

2.1 Anytime Algorithms

Anytime algorithms are algorithms which can exchange execution time for quality of results [15]. They address the problem of time-dependent planning, where the available planning time varies. The application area of anytime algorithms is broad and includes diagnosis, belief networks, databases, web searching and others [20]. Surprisingly, few applications of the basic idea to path and motion planning exist. Zilberstein studied the problem of interleaved sensing and acting for a simple simulation involving a point robot moving amongst a set of obstacles without dead ends [55]. The path quality metric in this case was the ratio of computed path length to the optimal path length. His work also includes a theory for the combination of interdependent anytime components in an optimal fashion [56]. Such work could eventually be applied to the SBPP algorithm and used to dynamically alter the balance of cpu cycles allocated to slice optimization and to intra-slice planning.

2.2 Graph-based Methods

Path and motion planning problems can be naturally expressed as a search problem in configuration space (path planning) or in the Cartesian product of configuration space and time (motion planning). There are two easily identifiable strategies for making search in such large spaces feasible. One is to form a compact representation of the collision free regions of configuration space. The second is to guide the search so that only a small part of the overall space needs to be evaluated. The graph-based methods described in the next section typically emphasize the former, while the potential field methods described in the subsequent section emphasize the latter.

Many of the most well developed path planning methods are based on forming a compact representation of the free regions of the workspace or of the collision-free portions of configuration space. This compact representation typically takes the form of a graph. The two primary approaches are roadmap methods and cell decomposition methods. The primary difference between the two is what the nodes of the graphs represent. In roadmap methods each node identifies an individual point in the configuration space, while in cell decomposition methods each node represents an entire area in configuration space.

2.2.1 Roadmap Methods

Roadmap algorithms involve graphs where each node represents a significant configuration in $\mathcal{C}_{free}$. Each graph link represents a one-dimensional curve (often a straight line, arc or small degree polynomial) lying entirely in freespace. Roadmap algorithms frequently assume a polygonal description of the obstacles.

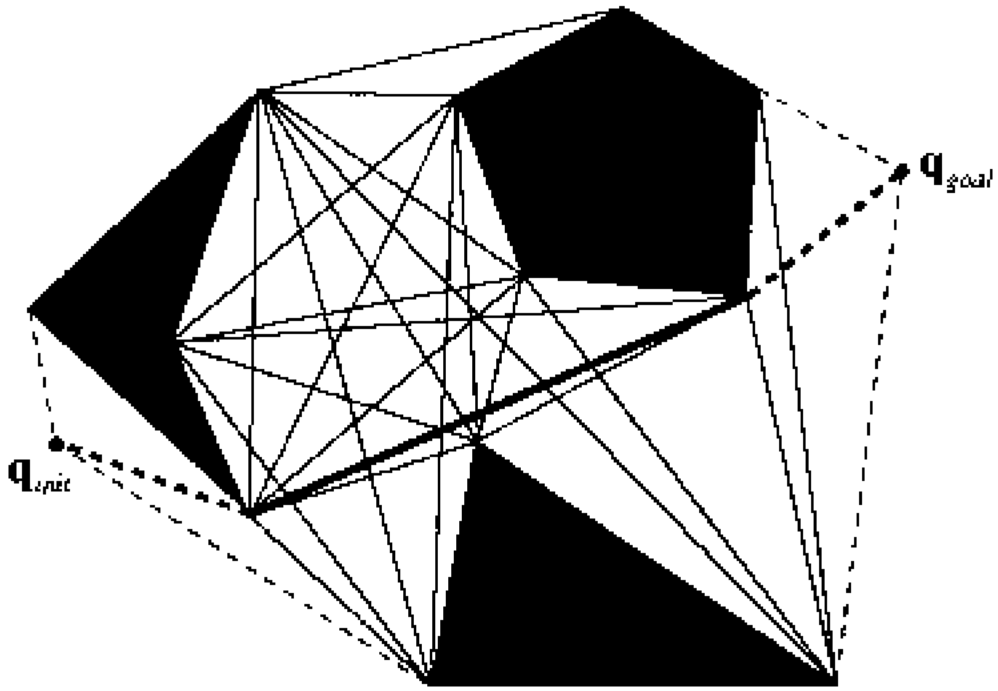

Figure 2.1: Example of visibility diagram (from Latombe 1991)

Visibility Graph

One of the earliest roadmap algorithms was the visibility graph [36]. This algorithm is most naturally applied to two-dimensional configuration spaces with polygonal obstacles. The graph nodes represent the obstacle vertices and the starting and goal configuration. Graph links are constructed between nodes for which a straight line connecting them does not pass through an obstacle. Figure 2.1 illustrates the visibility graph for a simple two-dimensional case.

In the two dimensional case the visibility graph is often used to solve minimum distance problems. It can be applied to three dimensional problems (but in this case the minimum distance path does not in general lie within the visibility graph). But with the exception of a circular mobile robot which can always move in any direction (ie has holonomic motion constraints), few real-world path planning problems have been solved using the visibility graph method.

Voronoi Diagram

Many roadmap methods are based on the Voronoi diagram. In this method the graph nodes represent points which are equidistant from three obstacles. The links between nodes represent curves for which each point is equidistant from two obstacles. Voronoi diagrams are one member of a class of methods referred to as retraction methods. As a mathematical entity the Voronoi diagram dates back to the 1600's [10]. Rowat was the first to apply the technique to the path planning of a mobile robot [43]. Originally the Voronoi diagram was defined only on point obstacles. Gowda described an incremental method for updating a Voronoi diagram in this case [19]. Ó'Dúnlaing and Yap [37] described the "generalized Voronoi diagram" which extended the point obstacle Voronoi diagram to the case of polygonal obstacles. Figure 2.2 illustrates the Voronoi diagram for a simple two-dimensional case.

Several extensions to higher dimensional spaces have been developed. Choset [10] describes the Hierarchical Generalized Voronoi Graph (HGVG) which can be applied to higher dimensional work spaces. The HGVG can be constructed in an incremental fashion but the algorithm is incremental in a different sense than has been discussed here. Choset addresses the sensor based planning problem and the HGVG is a technique for building an environmental map based on line of site sensor information collected during robot exploration. It does not address the problem of obstacles being added or replaced to the environment nor to continuous motions of obstacles. In addition, the algorithm is not applicable to the planning of robots with rotational degrees of freedom.

Brooks introduced a heuristic-based method called the freeway method [7]. This algorithm addresses the problem of the motion of a polygonal robot which can translate and rotate in a plane populated with polygonal obstacles. The basic concept is to annotate a roadmap with information describing the collision free orientations. By making a few assumptions it can be shown that it is only necessary to record this orientation information at a finite number of points within the roadmap. Those

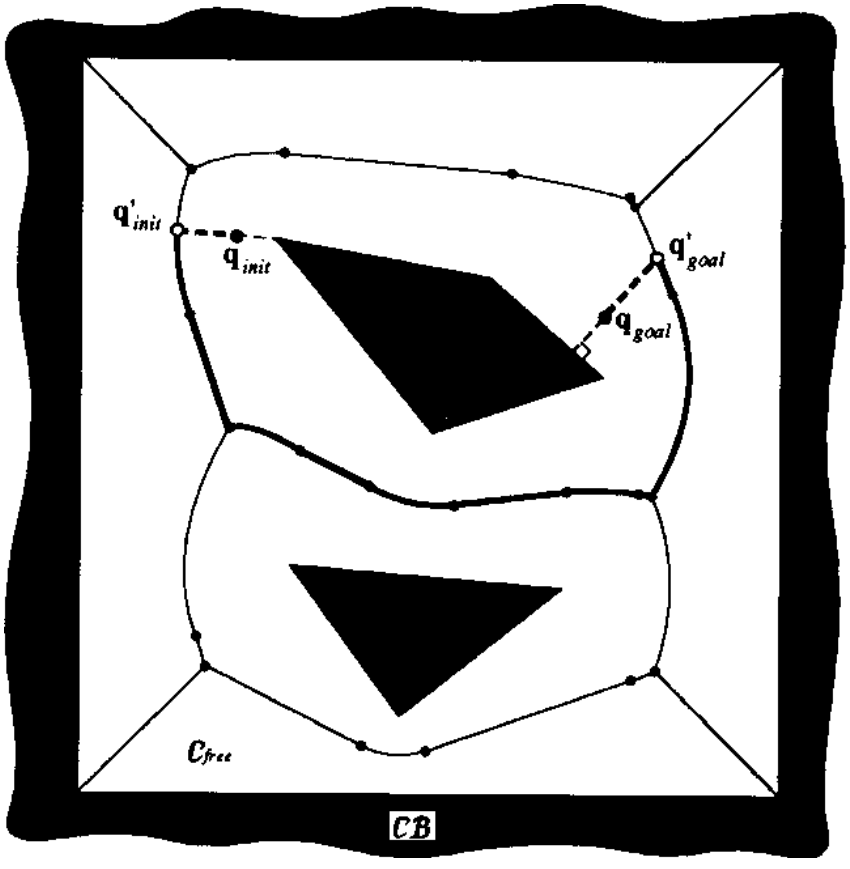

Figure 2.2: Example of Voronoi diagram (from Latombe 1991)

assumptions are in fact overly restrictive which causes the overall algorithm to be incomplete.

The Silhouette algorithm developed by Canny has many positive aspects [8]. It is complete and is not restricted to systems with few degrees of freedom. It is the only known algorithm whose time complexity is only singly exponential in the number of configuration variables. In addition, it assumes only that the obstacles are described as a semi-algebraic set (as opposed to the much narrower class of generalized polygons)[1]. But the algorithm is extremely complicated and no effective general implementation exists [28]. And although the algorithm is not restricted

[1]A set is semi-algebraic if its boundaries can be described by multivariate polynomials with rational coefficients

to polygonal objects, the processing necessary to compute the boundaries of the configuration space obstacles may itself be prohibitive for many kinds of robots.

Randomized Roadmaps

Several researchers have demonstrated path planners capable of addressing very large degree of freedom systems using randomly generated roadmaps. Kavraki with others developed the Probabilistic Roadmap Planner while Overmars demonstrated a similar technique [23, 38, 39]. Configurations are chosen randomly and checked for collisions. If there are no collisions the configuration is added as a node to the roadmap and an attempt is made to find a collision free path to the nearby existing nodes of the roadmap. When a collision free path is found a corresponding link is added to the roadmap. Various improvements have focused on applying heuristics to sample the configuration space more finely in some areas than in others[39]. Many other successful algorithms have been built using methods based on random search, see [24] for an overview.

These algorithms offer a different trade-off between computation and path quality: The longer the time allowed for building up the roadmap the greater the likelihood that the system will be be able to connect a particular start and goal pair. Thus the roadmap construction aspect is in some sense an anytime algorithm (where the computation performance metric is the probability of success). However a particular solution/path is not gradually improved. Also the algorithm can not incrementally update an existing roadmap when the workspace obstacles change slightly. Thus, unless other mechanisms are added the algorithms are unsuitable for online applications.

2.2.2 Cell Decomposition

Cell decomposition methods build a graph whose nodes represent entire regions of freespace and whose inter-node links represent adjacent regions of free space. The shape of these regions are chosen so that paths can be easily generated between

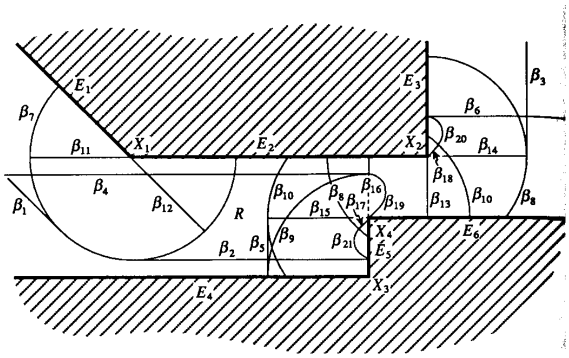

Figure 2.3: This figures illustrates the complexity of exact cell decomposition even for the relatively simple case if a rod which can translate and rotate in the plane. The lines represent potential boundaries of the exact cell decomposition regions method of Schwartz and Sharir (from Latombe 1991).

arbitrary configurations within two adjacent regions. The class of cell decomposition methods can be further divided into two categories. The exact decomposition methods represent freespace precisely. The inexact methods, on the other hand, use (typically rectangular) regions to approximate freespace.

Many exact methods have been developed for low dimensional workspaces and with polygonal representations of the robot and obstacles. One commonly addressed problem is the so called "ladder problem", that of moving a rod which can translate and rotate through a two dimensional workspace populated by polygonal obstacles. Each cell represents an area where the constraints on motion are similar. Even in this simple case the graph can become complicated very quickly. This is illustrated in Figure 2.3 which shows candidate boundaries for an exact cell decomposition in the case of a rod.

Schwartz and Sharir described a general algorithm (applicable to systems of any dimension) which requires only that the robot and obstacles be described by semi-algebraic sets [44]. But again the algorithm is extremely complicated and no general implementation exists [28]. Barbehenn used this algorithm as the basis for developing the only truly incremental path planning system this author is aware of

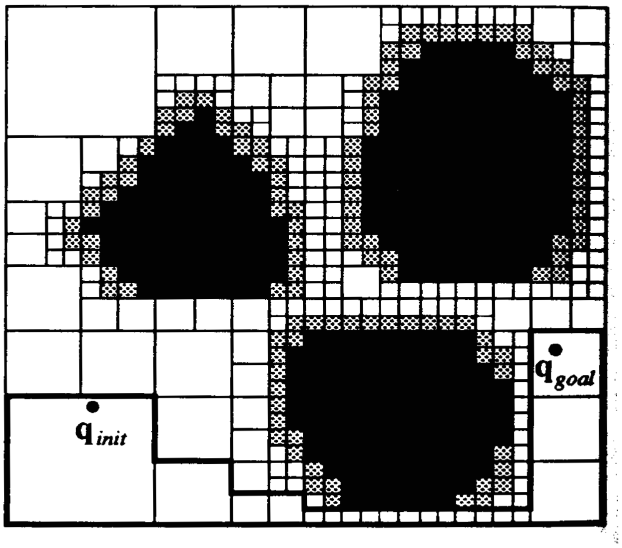

Figure 2.4: Example Inexact Cell Decomposion (from Latombe91). The black cells lie entirely within an obstacle while the gray cells are those that contain portions of an obstacle.

[1]. Barbehenn's algorithm is however restricted to the case of a rod moving in the plane.

The inexact cell decomposition methods introduced by Lozano-Pérez have the advantage of using a single simple shape (normally rectangular) for all cells [32]. Hence they are much easier to implement. The cells are used to form a conservative estimate of free space. Figure 2.4 illustrates a typical inexact cell decomposition.

Obviously the performance of the inexact methods will depend in part on the granularity of the discretization. Those methods which are guaranteed to find a collision free path if one exists assuming a fine enough granularity are said to be "resolution complete". By varying the granularity the obstacles can be matched quite precisely. Thus the fact that the inexact methods are only resolution complete is not often of practical significance.

Unlike the exact cell decomposition methods, the inexact methods have been applied to serial manipulators of moderate size. All these methods avoid constructing an explicit description of the configuration space obstacles. Instead they divide each joint's possible range into a finite set of intervals. Consider the inexact cell centered at $(q_1, \ldots, q_n)$ and described by $(q_1-\delta_1, q_1+\delta_1)\times\ldots\times(q_n-\delta_n, q_n+\delta_n)$. The workspace volume swept out by the robot when q varies across the total cell is computed. If this volume intersects an obstacle, then that cell is marked as a configuration space obstacle. Determining the swept volume is a complicated problem itself. Lozano-Pérez suggested approximating that volume by sweeping along the individual axes one at a time [34]. Since this approximation is not conservative it is also necessary to grow the robot geometry appropriately.

The key difficulty which must be addressed is the size of the constructed graph which grows exponentially with the configuration space dimension. This can be partially addressed by a hierarchical representation based on the observation that if the first link intersects an obstacle then the other joint angles are irrelevant [18, 34].

A common strategy is to avoid expanding cells unnecessarily. Rather than check every cell for collisions, only the cells that the search method traverses are checked. Many different search methods and heuristics have been proposed. The heuristics are in general limited and care must be taken to ensure that the search method does not require representing a very large graph. One common technique is to construct a two level algorithm where the higher level algorithm generates intermediate subgoals [16, 22].

Some motion planners share aspects of both exact and inexact cell decomposition methods. Lozano-Pérez described an algorithm for the planning of a planar rod which generates a polygonal approximation to the area swept out as the rod's orientation varies over some finite range [33].

2.3 Potential Field Methods

This section covers two related methods which do not attempt to construct a graph which represents freespace. Instead they construct a potential function for which the gradient provides useful information as to the best direction (in configuration space) to proceed.

2.3.1 Potential Field Methods

This section introduces a class of methods first introduced by Khatib, which do not attempt to construct a graph which represents freespace [25]. Instead it treats the robot as being subject to two forces: an attractive force pulling the robot to a goal and a repulsive force pushing the robot away from obstacles. The robot then moves in the direction of the net force. The method was originally proposed as an online navigation method but it has frequently been applied to offline planning as well. The method has the significant advantage of being very fast but is incomplete. It is incomplete because of the presence of local minima, where the attractive force and the repulsive force are equal. Many different potential functions have been designed to lower the number and depth of the local minima [52][26].

A potential function which is free of local minima is called a navigation function. The classical approach uses potentials which are functions of the distance to the nearest obstacle. Navigation functions based on this kind of potential function are only possible for very restrictive cases [42, 41]. In the following section several alternative constructions of navigation functions will be mentioned.

Although it is in general not possible to eliminate local minima, many hybrid strategies have been developed which extend the classical potential field method in order to escape local minima. Some of these algorithms have successfully addressed problems with more than 10 degrees of freedom.

Barraquand and Latombe described two planners which adopt interesting approaches [4, 27]. Each uses numerically computed workspace potentials which are free of local minima. A set of control points is attached to the robot and each has an associated potential field. The configuration space potential is defined in terms of the workspace potentials evaluated at the locations of each control point. Although each of the workspace potentials is free of local minima, competition between the various control points can lead to local minima in the configuration space potential. The two planners described below offer slightly different strategies for escaping these local minima.

The first planner applies what is essentially a best first search strategy in configuration space. The set of configurations along the path is stored and when a local minimum is found, the locally minimal configuration is marked, removed from the path list and the search begins again. This algorithm effectively "fills up" local minima. In general the number of configuration elements in the basin of attraction of a local minimum is exponential in the number of configuration space dimensions. Nonetheless the algorithm is able to produce path plans for systems with as many as four degrees of freedom.

The second planner introduced the application of random search to escape from local minima. The basic version performs a random walk whenever stuck in a local minimum. From the endpoint of that random walk a best first/gradient descent policy guides search to another local minimum or to the goal. If a local minimum with lower potential is found, the path is extended to the new local minimum and the process continues. If after some number of attempts a new lower local minimum is not found, a point along the existing path generated by an earlier random walk is chosen. The path is truncated at that earlier point and a new random walk begins. This algorithm has been successfully applied to very large systems having as many as 31 degrees of freedom [27].

An important advantage of both the above planners is the absence of an extensive preprocessing stage. By avoiding the construction of compact representations of freespace the planners are much more efficient and can easily be applied to situations where the environment is changing frequently. The disadvantage is that there remain many simple examples of workspaces for which random search has a very difficult time escaping local minima [9][54].

2.3.2 Grid-Based Potential Functions

As mentioned above it is in general not possible to compute a potential function which is free of local minima. However, it is possible to numerically compute potential functions which are free of local minima. These navigations functions use numeric procedures to propagate information from one grid element to the next and thus are in some sense global methods. Examples of this approach include wave-propagation methods [4, 30, 35] and the harmonic function based methods discussed below [11]. Since the potential at the grid points must be explicitly represented these approaches are limited to spaces of relatively low dimension ($<= 4$).

Distance Functions

Perhaps the simplest navigation function is the distance function which describes the distance from each grid element to the goal. Barraquand describes a simple algorithm for computing this function [4]. Lemmon and Glasious describe neural net approaches to computing a distance function [30, 17]. Muraca applied a distance function to a planar two-link manipulator [35].

Harmonic Functions

Connolly and others have made extensive use of a different type of navigation function based on harmonic functions [49, 13, 11]. Harmonic functions are the solutions to Laplace's equations and can be computed via a simple iterative procedure. The obstacle and goal locations are mapped to boundary conditions for a simple partial differential equation (Laplace's equation). The iterative calculation is simple enough that a massively parallel implementation is feasible [46]. Analog implementations have also been described [47, 46, 49]. Connolly has also investigated the relationship between biological motor control and harmonic function based motion planning [12, 14, 11].

While harmonic functions are easily computed and are free of local minima they are not without problems. One significant problem is that the gradients can be extremely small [49]. These small gradients combined with limited numerical precision in digital implementations or with ambient noise in analog implementations can cause the algorithm to fail to find an existing path. This problem is less severe if the obstacles are mapped to Neumann boundary conditions rather than Dirichlet boundary conditions.

2.4 Variational Methods

Many studies have approach the problem of path planning from the perspective of optimization rather than search. This approach has the advantage of being able to incorporate more diverse notions of path quality than those based solely on distance and obstacle avoidance. Indeed many of these studies address the more general and challenging problem of optimal control. Indeed the numerical optimization methods on which many of these algorithms are based are often inherently incremental in nature. Most of these approaches can still be viewed as search but not as search in configuration space. Rather these variational approaches involve local search in the infinitely dimensioned space of curves connecting the start to the goal. And like most local search methods there is the problem of local minima.

Lee [29] described a technique which uses neural networks to implement a variational approach. A simulated annealing process is used in an attempt to address the problem of local minima. The obstacles are represented directly in the weights of the neural network. Quinlan described a similar system which combines elements of potential field methods and a variational approach [40]. Neither approach was applied to a serial link robot.

Barraquand et al describe a variational approach based on dynamic programming [2]. The algorithm randomly chooses a lower dimensional neighborhood around an existing path within which dynamic programming is applied. This has the advantage of being able to use more complicated path quality metrics but is still susceptible to problems of local minima.

2.5 Summary

Anytime algorithms allow an evaluation of solution quality as a function of computation time. This approach has thus far found little application in path planning, perhaps because so few path planning algorithms are capable of making any sort of trade-off between path quality and computation time. The work described herein represents an attempt both to develop a scalable path planning algorithm which can make such a trade-off and also to quantitatively evaluate that trade-off for the described algorithm. The particular path planning approach described uses aspects of many existing path planning algorithms. The control point/workspace potential strategy for constructing a configuration space potential is adopted from potential field methods. The slice optimization framework combines elements of variational methods and harmonic function-based methods. The notion of restricting search to a smaller subset of configuration space is virtually ubiquitous with the field of path planning, but is most directly connected to the pre-existing notion of planar slices. The non-planar slice developed in this work represents a much broader class of subsets and hence can contain collision-free paths which are not possible in planar subsets.

Chapter 3

The Slice-based Path Planning Algorithm

In this chapter we describe the Slice-based Path Planning (SBPP) algorithm in detail. The first section describes the overall algorithm and its application. The algorithm has two key components: the slice optimization and intra-slice path planning which are described in the second and third sections respectively.

3.1 Algorithm Overview

The first particular problem addressed by this research is the incremental path planning problem (Problem 2 of the Introduction chapter but repeated here). This problem focuses on the gradual improvement of the produced path plan and assumes static obstacles.

Note that the actual measures of configuration quality J_C, slice quality J_f, and path quality J_p discussed below are expressed as cost criteria. Thus there is an inverse relation between quality and the criterion evaluation. For example a high quality path has a low J_p.

Problem 2 (Incremental Path Planning Problem)

Given:	*Obstacle Geometries*
	Obstacle Positions
	Start and Goal Configurations
	Robot Geometry and Kinematics
	Initial Slice
	Path Metric $J_p : p \mapsto \mathbb{R}$
Output:	$p(t_c)$ *Collision-free Path Connecting Start and Goal Configurations*
Computation Metric:	$J_p(t_c)$ *Path Quality as a Function of Computation Time*

Figure 3.1 illustrates the overall architecture of the algorithm.

The obstacle bitmap Obs_W is constructed on the basis of the input information describing the obstacle geometries and position. This bitmap is of two (in the case of a planar workspace) or three dimensions. The robot geometry is represented as a set of control points distributed across the surface of each link of the robot.

Since the produced path is constrained to lie entirely within the slice, the slice must pass through the start and goal configurations (c_s and c_g). Thus these configurations determine boundary conditions for the slice at the corresponding grid elements $((x_s, y_s, \vec{f}(x_s, y_s)) = c_s$ and $(x_g, y_g, \vec{f}(x_g, y_g)) = c_g$. There can in fact be multiple goal configurations in which case the intra-slice path planning component will choose the best one according to the path criterion described below.

The initial slice, f_0, is somewhat arbitrary since the slice optimization process will continually adapt the slice based on the current obstacle information. However, a poor choice for the initial slice will result in an initial path which is poor and will require many iterations before a local minimum is reached. In the experiments described in this document the initial slice was computed by allowing the slice optimization to reach a local minimum using an empty obstacle bitmap (except for

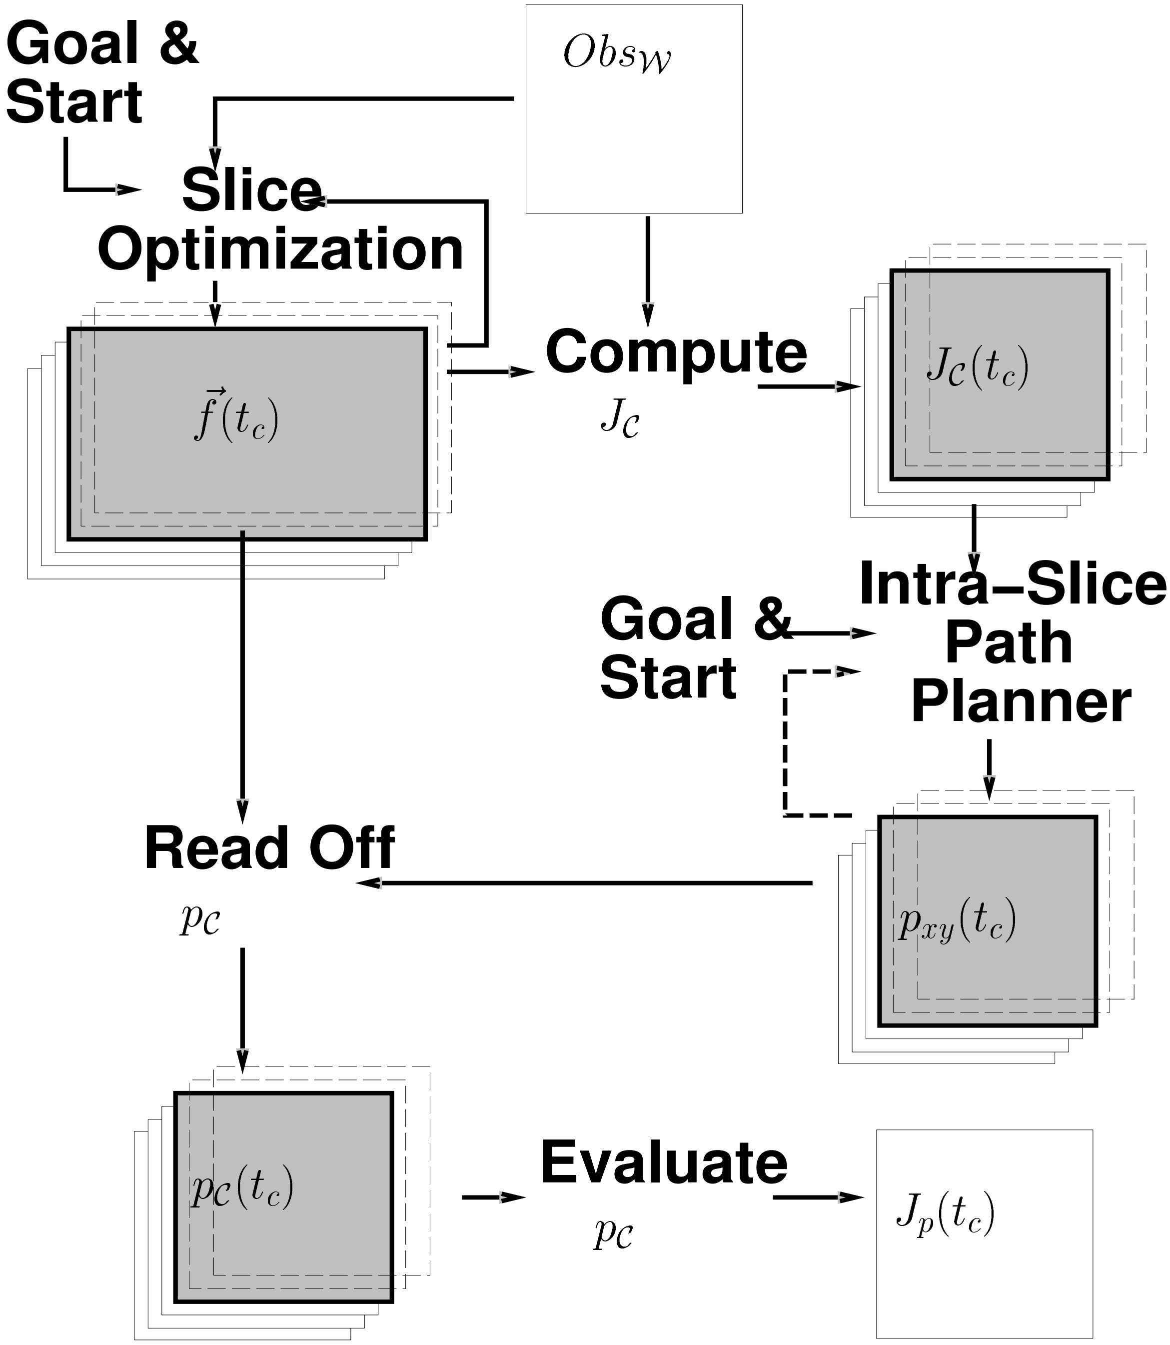

Figure 3.1: Overall SBPP architecture. The main input is the obstacle bitmap $Obs_{\mathcal{W}}$. At each point in the computation (represented as computation time t_c), the existing slice is modified $\vec{f}(t_c)$. $J_{\mathcal{C}}(t_c)$ is the measure of configuration quality which is applied to each element of the slice. The intra-slice path planner produces an intra-slice path$p_{xy}(t_c)$. The path containing all the configuration information $p_{\mathcal{C}}(t_c)$ is read from the slice. The final step is to evaluate the quality of the path itself $J_p(t_c)$.

permanent obstacles such as the border around the planar workspace and a desktop in the PUMA).

It is not in general possible to compute a slice which contains only collision free configurations. This is true for several reasons. First, there may be (x, y) locations for which no configuration is collision free. Secondly, the slice might be caught in a local minimum. For instance, imagine two spherical obstacles placed one on top of the other. A slice could get caught between them. Thirdly, the optimization criterion terms containing measures of the gradient and Laplacian magnitudes can overcome the potential term. Figure 1.1 illustrates just this situation. Observe the robot centered at $(15, 7)$ which intersects the lower portion of the rectangular obstacle. The potential term would normally force the corresponding slice element to a collision free orientation. However, in the example the robot's orientation at the slice elements on the left and right have been driven in opposite directions. Thus the gradient term counteracts the potential term producing a configuration which is not collision free.

Some measure of quality of each configuration contained within the slice is needed in order to perform the intra-slice path planning. One possibility is a simple binary measure of whether the configuration is collision free or not. This would correspond to finding the intersection between the slice and the workspace obstacles. The other possibility, and the one adopted in the current implementation, is to use a more complicated metric which includes measures of: how close the obstacles are to the robot, how close the joints are to their physical joint limits, and how quickly the configuration variables are changing. This choice allows the intra-slice path planning component to produce paths which reflect all those criteria. In the current implementation this path criterion is based on the same configuration quality measure $J_C(x, y, \vec{f}(x, y))$ used in the slice optimization.

The actual intra-slice planning could be solved by many different planning algorithms. The algorithm applied in the current implementation is dynamic programming using value iteration [48]. This allows us to produce paths which are optimal according to the optimization criterion defined on paths. This path criterion is based directly on the slice optimization criterion, hence the utility of $J_C(x, y)$. The intra-slice path planning component is described in more detail in Section 3.3.

The full path can be computed easily from the intra-slice path simply by reading off the remaining configuration variables from the slice. The overall measure of algorithm performance adopted is the path quality (as defined by the same path optimization criterion used to compute the intra-slice path) as a function of time. This metric makes clear the trade-off between computation time and path quality. The following chapter will quantify this trade-off in several different ways.

The total computation proceeds as follows: First the start and goal configurations and the bitmap Obs_W are fed to the SBPP planner. The planner generates an initial slice and ensures that it passes through the start and goal configurations. Then the two main processes, the slice optimization and the intra-slice planner, begin executing concurrently. At each iteration of the computation, denoted by t_c, the current slice $f(t_c)$ and Obs_W are used to produce $J_C(t_c)$. The gradient of J_C with respect to all the slice variables is used to optimize the slice. $J_C(t_c)$ is itself used by the intra-slice path planner to produce a new path $p_{xy}(t_c)$. The final step uses $p_{xy}(t_c)$ together with slice $f(t_c)$ to produce a plan with all the configuration space dimensions $p(t_c)$. In order to evaluate the SBPP algorithm, the evaluation of path quality J_p is recorded at each iteration and used to produce a graph of path quality as a function of computation time.

The second problem, incremental path planning with changing obstacles, is addressed in an almost identical fashion. The key difference is that Obs_W will vary over time and must be reread at each iteration. Similarly, the potential field must also be updated at each time Obs_W changes. Since the obstacles will be constantly moving, the measure of path quality will not approach an asymptote.

3.2 Slice Optimization

A slice describes a choice of the remaining degrees of freedom when the robot's (x, y) position is determined. Thus the slice of an N degree of freedom system will represent an $(N-2)$ valued function. A slice will be denoted by $\mathbf{f}(x,y) = (f_1(x,y),\ldots,f_{N-2}(x,y))$ but at times we will simply use f. The tuple $(x,y,\mathbf{f}(x,y))$ consisting of the (x,y) position and the $N-2$ slice variables completely determines the robot's configuration.

Since subsequent path planning is restricted to lie within the computed slice, it is clear that the computed slice has a large effect on the quality of the full configuration path ultimately produced. Ideally the slice would always contain the optimal path (according to the chosen criterion). Since this is not in general possible the SBPP algorithm uses a heuristic based optimization criterion as the basis for computing a slice. The main heuristic is simply to prefer slices which contain few collision-causing configurations. This is intended to increase the likelihood that at least one collision free path exists within the slice.

Mathematically the slice optimization criterion is described by Equation 3.1. It has the form of an integral of a more basic criterion (denoted J_C and defined in Equation 3.2) which describes the quality of a particular configuration. This integral is taken over the reachable portions of the horizontal reachable workspace $\mathcal{W}$.

$$J_f(\vec{f}(x,y)) = \iint_W J_C(\vec{f}(x,y),x,y)\; dx\; dy \tag{3.1}$$

In the existing implementation the slice optimization process does not use any information regarding the current plan and weights the importance of all its constituent configurations equally. Future work will explore the advantages of placing greater importance on the configurations which are part of the current plan. Alternatively the slice criterion could place greater importance on those configurations which are more likely to be included in the plan such as the configurations in the vicinity of the start and goal configurations.

The start and goal configurations determine boundary conditions on the slice optimization. The slice must contain those two points. It is possible however to describe only portions of the goal configuration. For example, if the purpose of the planned motion is to grasp a vertical cylinder the goal can be specified as a desired height and roll. The slice optimization procedure will then select the yaw and pitch.

At any point in time the computed slice is unlikely to be a *global* minimum for two reasons. First, the numerical procedure (gradient descent) which implements the slice optimization can require a large number of iterations before it approaches a minimum. And if the obstacles are constantly moving the slice computation will always be "following behind" the minimal slice. And as mentioned above, the space of slices contains many local minima thus making finding a global minimum unlikely. Fortunately, the overall motion planning method does not require the slice to be optimal. In addition, the optimization criterion itself is heuristic thus making the notion of optimality somewhat artificial.

Configuration Criterion

The configuration criterion J_C contains four terms which measure the following quantities: the distance between the links of the robot and any obstacles; the magnitude of the first and second derivatives of the slice at that configuration; the distance of each joint variable to its joint limit. The w terms are weights which reflect the importance of each individual term. $w_{k,j}$ denotes the weight (a scalar value) applied to the kinematic term of joint j. $w_{d,i}$ and $w_{s,i}$ denote the weights of the derivative and smoothness terms respectively applied to each slice component. Each term is described in more detail below.

$$
\begin{aligned}
J_{\mathcal{C}}(\mathbf{f}(x,y),x,y) =& w_p \sum_b \sum_{cp_b} Pot(T_b(\mathbf{f}(x,y),x,y)\mathbf{x}_{cp_b}) \\
&+ \sum_{j=1}^{N} w_{k,j} \frac{(\frac{\theta_{max}-\theta_{min}}{2} - \theta_d(\mathbf{f}(x,y),x,y))^3}{\theta_d(\mathbf{f}(x,y),x,y)} \\
&+ \sum_{i=1}^{N-2} w_{d,i} \left| \left(\frac{\partial f_i}{\partial x}, \frac{\partial f_i}{\partial y} \right) \right|^2 \\
&+ \sum_{i=1}^{N-2} w_{s,i} \left| \left(\frac{\partial^2 f_i}{\partial x^2}, \frac{\partial^2 f_i}{\partial y^2} \right) \right|^2
\end{aligned}
\tag{3.2}
$$

Potential Term

Obviously a configuration which results in a collision will not be part of a collision free path. So one term of the configuration quality metric is based on the distance between the robot in that configuration and the surrounding obstacles. Rather than perform explicit collision detection using detailed geometric knowledge of the robot and the obstacles, we adopt a simpler strategy based on potential fields. This term is based on the sum of the potential field at a set of control points distributed across the robot's surface. For the planar rod these points are evenly distributed across the length of the rod. The index of the first sum, b, identifies a particular link or other rigid body making up the robot. The inner sum is taken over the number of control points distributed on rigid body b. The tuple $(\mathbf{f}(x,y),x,y)$ determines a particular configuration. Thus $T_b(\mathbf{f}(x,y),x,y)$ is the homogeneous transform describing rigid body b with respect to the world frame. Since $\mathbf{x}_{cp_b}$ describes the location of a control point with respect to a frame attached to rigid body b, the product of $T_b(\mathbf{f}(x,y),x,y)$ and $\mathbf{x}_{cp_b}$ describes the position in the workspace of control point cp_b when in the configuration $(\mathbf{f}(x,y),x,y)$. The potential function itself is described by Equation 3.3 and illustrated in Figure 3.2. This function was chosen so as to increase sharply near obstacles and to smoothly decay to zero at a distance d_z from an obstacle. $d(\mathbf{x})$ is simply the Euclidean distance from $\mathbf{x}$ to the closest obstacle. $d(\mathbf{x})$ is zero inside the obstacles.

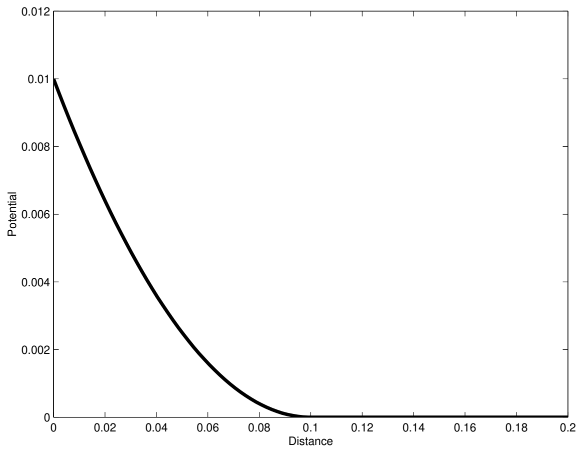

Figure 3.2: Potential as a function of the distance from an obstacle located at 0. $d_z = 0.3$

$$Pot(\mathbf{x}) = max(0, d_z - abs(d(\mathbf{x})))^2 \tag{3.3}$$

Joint Limits

The second term of J_C (shown below in Equation 3.4) serves to enforce the kinematic constraints imposed by any joint limits. It consists of the sum taken over m joints of a simple function of the distance between the joint value and the closest physical joint limit. This distance θ_d is itself a function of the joint values when in configuration $(\mathbf{f}(\mathbf{x}, \mathbf{y}), \mathbf{x}, \mathbf{y})$ as shown in Equation 3.5. These joint values are computed using the inverse kinematics of the robot. Figure 3.3 shows the joint limit term as θ varies between θ_{min} and θ_{max}. The exponent in Equation 3.4 was chosen simply to provide a larger flat region in the center of the joint range. The function goes to ∞ as θ approaches either joint limit (θ_d approaches 0). This ensures that no joint is driven out of range [1].

[1]However care must be taken to ensure that during the numerical optimization procedure this assumption is not violated.

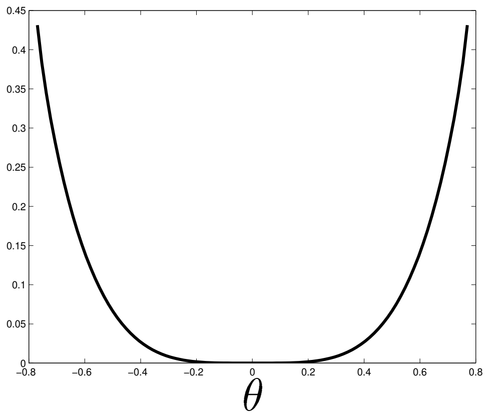

Figure 3.3: Graph of kinematic constraint term of configuration quality. $\theta_{min} = -\pi/4,\ \theta_{max} = \pi/4$

$$\sum_{j=1}^{N} w_{k,j} \frac{(\frac{\theta_{max}-\theta_{min}}{2} - \theta_d)^3}{\theta_d} \tag{3.4}$$

$$\theta_d(\mathbf{f}(\mathbf{x},\mathbf{y}),\mathbf{x},\mathbf{y}) = \mathbf{min}(\theta_{\mathbf{max}} - \theta(\mathbf{f}(\mathbf{x},\mathbf{y}),\mathbf{x},\mathbf{y}), \theta(\mathbf{f}(\mathbf{x},\mathbf{y}),\mathbf{x},\mathbf{y}) - \theta_{\mathbf{min}}) \tag{3.5}$$

Continuity and Smoothness

The third term describes the magnitude of each slice component's gradient. The fourth term describes the magnitude of the Laplacian. Each of these is calculated using a finite difference approximation. The elements which correspond to (x, y) locations occupied by obstacles are not included in the derivative approximations.

The third term describes the magnitude of the gradient of each slice component, f_i. This term helps ensure that each slice component is continuous and is not steeply sloped. This is particularly important since the algorithm uses a discretized representation of a slice. Consider for example a slice component describing vertical height, if the slope of a slice component were permitted to change rapidly, it would be possible for the slice to place the robot above an obstacle at one element and below the same obstacle at an adjacent element. The intra-slice path planner would not have sufficient information to determine that the robot cannot move from one configuration to the other without causing a collision.

The fourth term describes the smoothness of each slice component. This term helps to produce paths whose execution requires reasonable accelerations. If the slice is smooth (has continuous second derivatives $\frac{\partial^2 f_i}{\partial x^2}$ and $\frac{\partial^2 f_i}{\partial y^2}$ for all slice components i) then any intra-slice path p_{xy} which is temporally smooth (has continuous second derivatives $\frac{\partial^2 p_x}{\partial t^2}$ and $\frac{\partial^2 p_y}{\partial t^2}$) will produce an overall path which is also smooth. In fact, the current implementation of the intra-slice path planner uses dynamic programming and does *not* produce paths which are smooth with respect to x and y. The robot on which the algorithm has been tested moves slowly enough that this has not posed a problem. Future work will investigate different intra-slice path planning techniques which can produce smooth paths.

The slice optimization is currently implemented using a pre-conditioned quasi-Newton conjugate gradient technique[45] using the NAG™ Foundation Toolbox for Matlab™. While this technique was more efficient than a steepest descent strategy, further improvement is likely possible through the application of finite element or multi-grid approaches.

Two points about the efficiency of the slice optimization step should be noted. First, the slice is always a two-dimensional structure. Thus the basic computation is only loosely dependent on the configuration space dimension. In other words, the number of operations required for a single optimization iteration does not increase exponentially with the number of configuration space dimensions. But of course, finding a locally minimal slice in a larger configuration space could require significantly more time [2]. Secondly, the system can start computing a slice even before any goal positions are specified. In other words, the system can be constantly monitoring the current obstacle locations and updating the slice. When a goal position is

[2]As mentioned earlier, in most situations the path planning will succeed even with slices which are far from optimal.

specified, the system need only "stretch" the original slice so that it passes through the proper configuration at the goal location.

3.3 Path Planning within a Slice

The above discussion describes how the slice is computed. In this section, the second step of finding a path from the start to a goal position will be described. The path planning component of the algorithm is also formulated as an optimization problem. The criterion, shown in Equation 3.6 is simply the integral of the configuration quality metric along the path. The path is parameterized by s, the horizontal distance traveled (the distance of the path when projected onto the xy plane). The constant c_{dist} implements a preference for shorter paths. A larger value of c_{dist} will produce shorter paths which in general bring the robot closer to obstacles, joint angles closer to their limits, and allow more rapid and less smooth changes in slice variables.

$$J_p = \int \left[J_C(\vec{f}(x(s), y(s))) + c_{dist} \right] \, ds \tag{3.6}$$

At times during the slice optimization some configurations within the slice will get worse even though the overall slice is improving. In other words, it is possible for $J_C(x, y)$ at *some* elements to increase during the optimization process. This is because the configuration qualities of adjacent elements are interrelated through the smoothness and derivative terms. When the path planning component produces a path which contains configurations whose cost is increasing it is possible for the overall J_p metric to increase.

The result of the intra-slice path planning component is a path in just two dimensions p_{xy}. This path is approximated on the grid as a sequence of grid elements $(x_1, y_1), \ldots, (x_f, y_f)$. The final step of the SBPP algorithm is to read off the full path p_c. The path is also approximated on the x, y grid.

$$p_c = [(x_1, y_1, \mathbf{f}(x_1, y_1)), \ldots, (x_f, y_f, \mathbf{f}(x_f, y_f))] \tag{3.7}$$

Chapter 4

Results

In this chapter we describe the results of applying the SBPP to two robots: a rod which can translate and rotate in the plane and a PUMA 560 six degree-of-freedom serial manipulator. For each robot we describe the results of three experiments designed to illustrate the incremental nature of the SBPP algorithm and to quantify the relationship between computation time and path quality.

The first experiment is simply the problem described in Chapter 1 as the Incremental Path Planning Problem. The basic sequence of operations is shown below in the chart entitled Experiment 1. The results of this experiment are presented as plots of average path quality versus computation time (number of iterations). In addition to providing an illustration of the incremental nature of the algorithm, these kinds of performance profiles could also be used as the basis for a stopping criterion or as a basis for deciding how to allocate limited computational resources.

The second experiment measures the benefit of adapting a pre-existing slice when a set of obstacles is added to the workspace (see Experiment 2). An initial set of obstacles Obs_0 is created and slice optimization proceeds (the first **while** block). When the slice is no longer changing, a new set of obstacles Obs_{new} is added to the workspace. The SBPP algorithm commences and at each iteration the path quality $J_{p,Inc}$ is recorded. The second phase of the experiment measures the algorithm performance using $Obs_1 = Obs_0 \cup Obs_{new}$ and after the algorithm has been re-initialized to use the initial slice f_0.

Experiment 1 Static Obstacle Experiment. The slice optimization is performed in the while block.

$f_0 \Leftarrow$Initial Slice
Generate Obs_W

{Optimize Slice}
while $|\frac{\partial J_s}{\partial f}| > 0$ **do**
  $f_{t_c+1} \Leftarrow f_{t_c} - \frac{\partial J_s}{\partial f}$
  Compute Intra-slice Path p
  Compute J_{p,t_c+1}
end while

Plot J_p

Comparison of the $J_{p,Inc}$ and $J_{p,Non-Inc}$ demonstrates the performance improvement when incrementally adapting an existing slice as compared to beginning the computation from scratch. The results of the experiment are presented in two forms. The first form follows the style of Experiment 1 and consists of plots of average path quality versus computation time for both the incremental and non-incremental case. The second form shows the average improvement of the incremental case over the non-incremental case.

The last experiment measures the algorithm's ability to track continual movement of the obstacles. The basic procedure is described in Experiment 3. In this experiment the obstacles move in a random walk at each increment of time i. Each time the obstacles move the SBPP algorithm is allowed to proceed for a few iterations j. This allows us to artificially adapt the relative rate of obstacle motion with respect to the processing speed of the SBPP algorithm. The ratio of SBPP iterations to obstacle movements is described by r.

At each point in time the path quality produced by just r iterations of the SBPP algorithm is compared to the quality which would have been obtained had the slice optimization been allowed to continue indefinitely. In other words, we measure the difference between the path quality after a few iterations of the SBPP algorithm compared to the path quality which would have resulted from freezing the obstacle

Experiment 2 Incremental Change to Obstacles

$f_0 \Leftarrow$Initial Slice
Generate Obs_0
Compute $\frac{\partial J_s}{\partial f}$

{Optimize Slice}
while $|\frac{\partial J_s}{\partial f}| > 0$ **do**
  $f_{i+1} \Leftarrow f_i - \frac{\partial J_s}{\partial f}$
  Compute Intra-slice Path p
  Compute $J_{p,i+1}$
end while
$Obs_1 = Obs_0 \bigcup Obs_{new}$

{Incremental Case}
$f_j \Leftarrow f_i$
Optimize Slice {While recording $J_{p,Inc}$}

{Non-Incremental Case}
$f_j \Leftarrow f_0$
Optimize Slice {While recording $J_{p,Non-Inc}$}

Plot $J_{p,Inc}$ and $J_{p,Non-Inc}$

locations and allowing the slice computation to proceed until no further improvement was possible. Experiments like this can help to determine the computational resources necessary to achieve a certain level of performance given the normal velocity with which the workspace obstacles move. This kind of measure could also serve as the basis for dynamically controlling the balance between computational resources devoted to planning and to other processing requirements.

4.1 Planar Rod

Figure 4.1 illustrates the planar workspace and robot used in the experiments described in this section. The robot is a simple rod described as a set of 7 control points distributed evenly along its length. The robots length is 7 units. The workspace is represented as a two dimensional grid of size 30 by 30. In each case the robot starts in the upper right corner at position (25,25) and with zero degrees of rotation. The goal configuration is in the opposite corner, again with zero orientation. The initial slice f_0 used for computing the non-incremental test is shown pictorially via the robot outlines distributed across the workspace. This initial slice was generated by letting the slice optimization process run until steady state using a environment which contains no obstacles.

Figure 4.1 shows the path produced using the initial slice f_0 before any slice optimization has been performed. As can be seen there is no collision free path contained within the initial slice. Figure 4.2(a) shows this same problem after the SBPP algorithm has been applied. The right portion of Figure 4.2 shows the result of adding new obstacles to the existing set as described in Experiment 2.

The obstacles were generated as follows: 50 circle-shaped obstacles were placed randomly in the workspace. The obstacles don't necessarily appear as circles in the figures because of discretization. The radius of each circle is chosen randomly (using a uniform distribution) as either 1, 2 or 3. The circle is placed randomly within the workspace (again with a uniform distribution). The outline of the workspace is marked as obstacles. Any obstacles within a 10×10 box around the start and goal

Experiment 3 Continuous Change to Obstacles

$f'_0 \Leftarrow$Initial Slice
Generate Obs_0
Compute $\frac{\partial J_s}{\partial f'}$

{Optimize Slice}
while $\frac{\partial J_s}{\partial f'} > 0$ **do**
  $f'_{j+1} \Leftarrow f'_j - \frac{\partial J_s}{\partial f'}$
  Compute Intra-slice Path p
  Compute $J_{p,j+1}$
end while
$f_0 \Leftarrow f'_{j+1}$

r = Optimization Iterations / Obstacle Motion
for i = 1 to 100 **do**
  $Obs_{i+1} \Leftarrow RandomWalk(Obs_i)$

  {Incremental Case}
  for k = 1 to r **do**
    $f_{i+1+(k-1)/r} \Leftarrow f_{i+1+(k-1)/r} - \frac{\partial J_s}{\partial f}$
    Compute Intra-slice Path p
    Compute $J_{p,i+1+(j-1)/r}$
  end for
  Record Best Path $J_p(i)$

  {Non-Incremental Case}
  Optimize Slice {Complete Optimization}
  Record Best Path $J_{p,Opt}(i)$

end for
Plot J_p and $J_{p,Opt}$

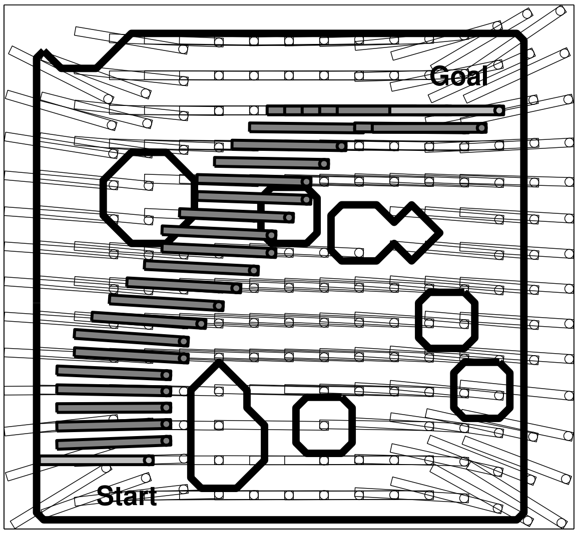

Figure 4.1: Pictorial representation of the workspace, goal and start configurations, and initial slice. The goal and start configurations are shown in light gray. The remaining configurations contained in the path are shown in dark gray. Note that the initial slice which is shown has not yet been adapted to the existing obstacles. Thus it is not surprising that this initial slice does not contain a collision free path.

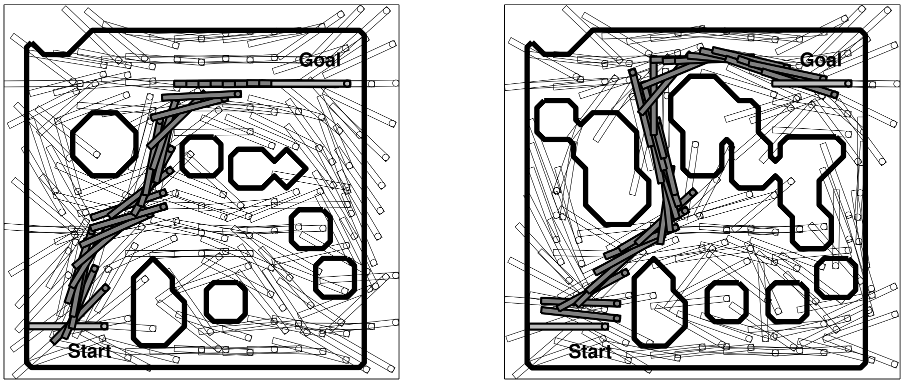

(a) Initial Obstacle Set

(b) Augmented Obstacle Set

Figure 4.2: The left shows the original obstacle set and the right the augmented set created by adding more obstacles.

w_s	0.5
w_d	10.0
w_k	0
w_p	0.1

Table 4.1: Cost function weights used to describe J_C.

positions are removed to ensure that the start and goal configurations are collision free.

Table 4.1 shows the parameters used to the describe the optimization criterion J_C (Equation 3.2) in all the planar robot experiments described below. The weights were manually selected so as to produce paths which reflected a reasonable balance between smoothness and obstacle clearance. w_k being set to zero is equivalent to saying that there are no joint limits and that the rod is can rotate freely through all orientations.

4.1.1 Static Obstacles

The planner was tested on 100 randomly generated obstacle sets. Figure 4.1.1 illustrates the results. Figure 4.3(a) shows a plot of average path quality as a function of time t_c. Since the path quality varies significantly between different randomly generated obstacle scenes, we also compute an adjusted metric. This adjusted metric represents the future improvement possible and is described by Equation 4.1.

$$J'_p = J_p - min(J_p); \tag{4.1}$$

Figure 4.3(b) plots this adjusted metric as a function of time. Since the minimal value of J_p usually occurs as time $t_c \to \infty$, the adjusted metric tends towards 0. Since the slice optimization is performed independently of the intra-slice path planning, it is possible that improving the slice can actually lead to a lower quality path. In practice we have found this occurs infrequently and with only minor negative changes in path quality. Negative changes could easily be avoided simply by always storing the best path computed thus far. It would also be necessary to reset this stored path whenever the obstacles moved. The vertical bars of Figure 4.3(b) indicate one standard deviation. This performance profile can be used as a basis for predicting the improvement of path quality with each additional iteration.

The greatest improvements in path quality occur during the first few iterations of the SBPP algorithm. In fact, for this particular problem, more than 50 percent of the improvement comes in the first 2 iterations.

Comparison with Randomized Path Planner of Barraquand and Latombe

This section compares the performance of the SBPP algorithm with the Randomized Path Planning (RPP) algorithm of Barraquand and Latombe when applied to the same example obstacle environment [3]. Figure 4.4 shows a path generated by each planning algorithm. The two paths differ in several respects. The RPP plan passes around the two adjacent rectangles in the lower left, while the SBPP algorithm navigates through the narrower region between the the same rectangles and the ellipse. In addition, the RPP produced path contains a nearly pure rotation in the upper area between the parallelogram and the triangle. While the SBPP algorithm cannot represent a pure rotation, in this example it isn't necessary since the robot enters the upper center area in the a different orientation.

Figure 4.5 helps illustrate why the SBPP algorithm was unable to find same path as the RPP approach. In this case the obstacles have been altered so that the passage between the ellipse and the rectangles is completely cut off. We see that the best path that the SBPP produces in this case is *not* collision free. This is an example of the slice optimization becoming stuck in the a local minima.

Figure 4.5(b) shows the configuration quality J_C. Notice that the narrow region between the left edge and the leftmost rectangle contains a small band of low quality configurations corresponding to collision causing configurations. The area around the rectangles in which the RPP algorithm produced a looping path around the obstacles constitutes a long narrow channel. This channel is not wide enough for the planar robot to fully rotate within. Thus the orientation at one end of the channel constrains

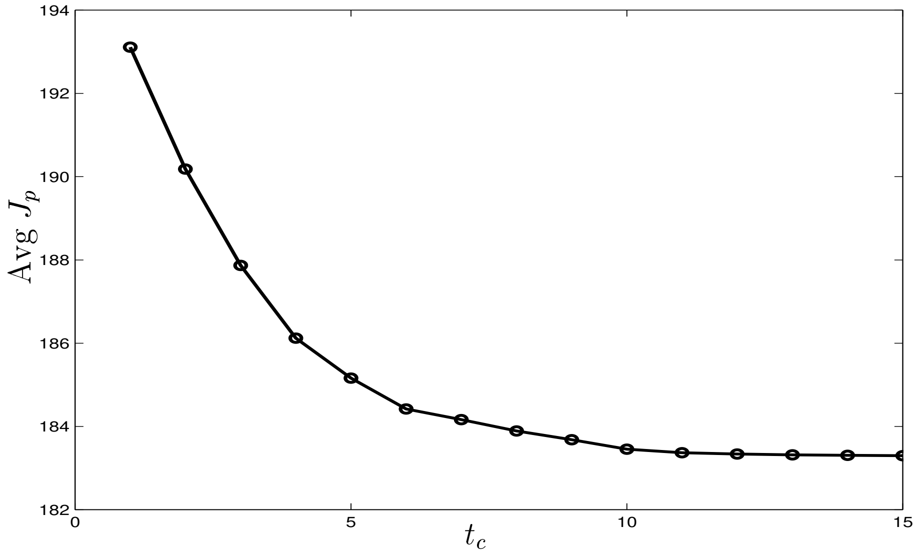

(a) Avg Path Cost

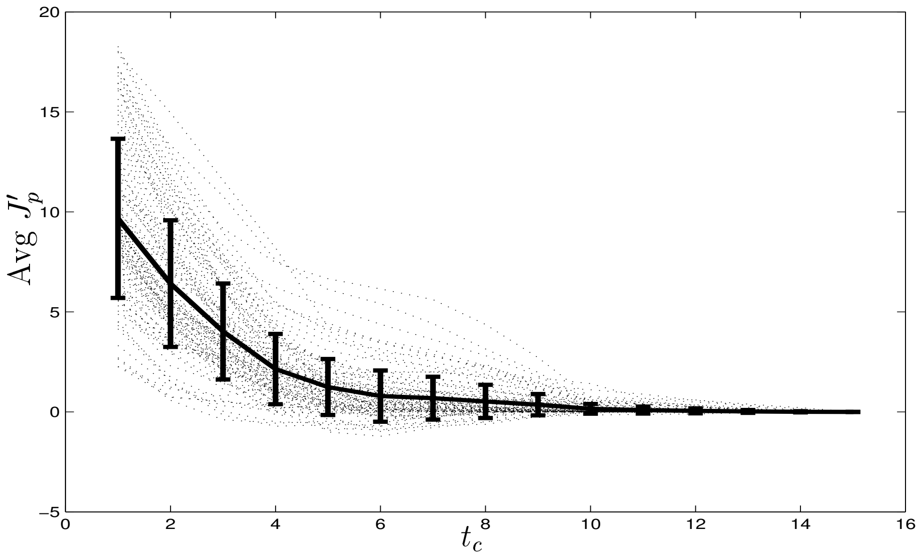

(b) Avg of $J'_p = J_p - min(J_p)$

Figure 4.3: Results of Experiment 1. t_c is measured as the number of completed iterations.

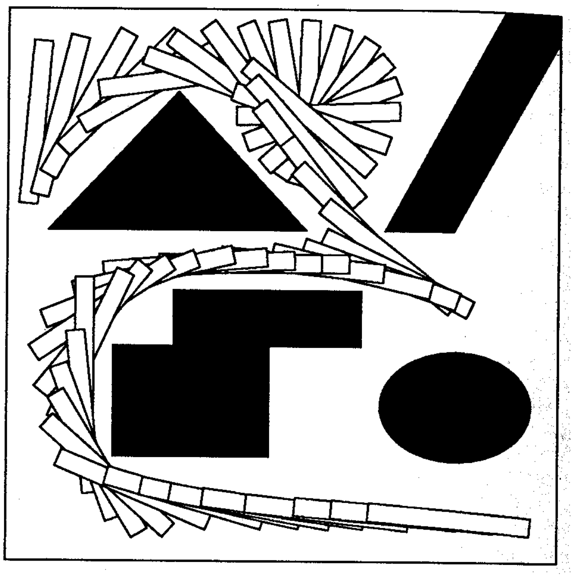

(a) Randomized Path Planner produced path

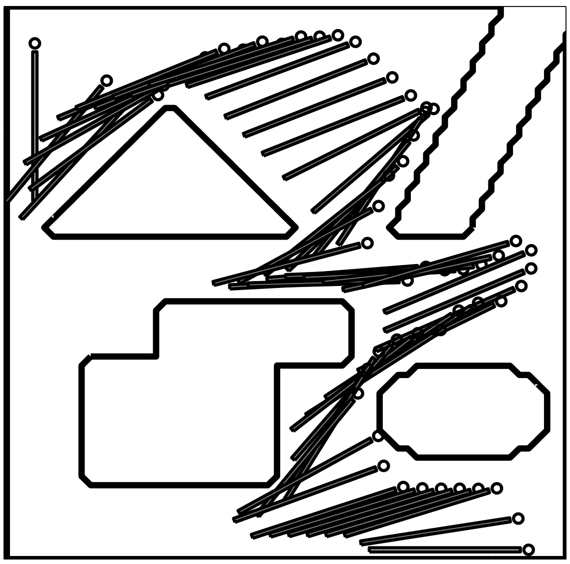

(b) Slice-based Path Planner produced path

Figure 4.4: Comparison of paths produced by Randomized Path Planning algorithm and Slice-based Path Planning algorithms in similar environments.

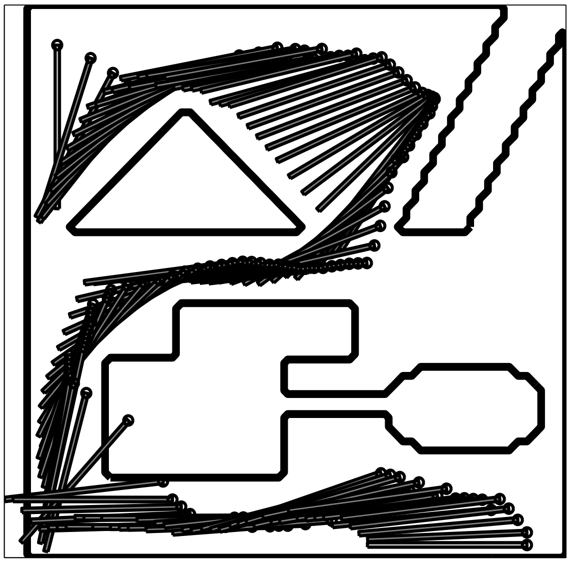

(a) SBPP produced path when previous path is blocked so as to lead SBPP algorithm to produce path around rectangles.

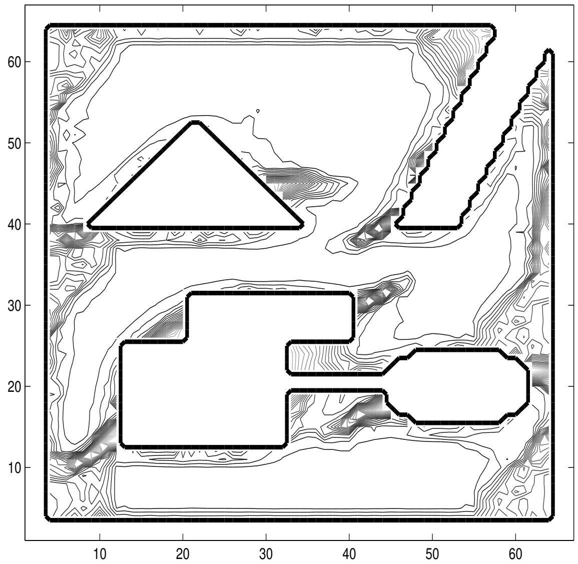

(b) Contour plot of J_c.

Figure 4.5: Illustration of SBPP produced path when "natural" path is blocked. The blocked area serves to force the path to the left region. The poor performance around the rectangles in the lower left help illustrate the difficulty the SBPP algorithm has in handling long narrow channels.

the orientation at the other end. If the slice does not meet this constraint then the slice (since it is continuous) must produce a rotation within the channel. This corresponds to the collision causing configurations shown in Figure 4.4. In order to eliminate these collisions the slice optimization would have to alter the configuration at each (x, y) location between the existing collision area and one end of the channel so as to cause a 180 degree rotation at each configuration ($f'(x, y) = f(x, y) + 180$). This is possible in some situations but in the environment shown, it is unlikely that it can be done *without* lowering the overall slice quality during the rotation.

In order to provide a measure of the depth of the local minima we performed the following experiment. The RPP produced path was used to manually create an initial slice containing configurations similar to those constituting the RPP produced path. The SBPP algorithm was then run until a new minimal slice was obtained. The resulting path is shown in Figure 4.6. In this case the path and slice quality (J_s and J_p) were 20046 and 44906 respectively. While in the case shown in Figure 4.4(b) the corresponding metrics were 16908 and 41626. So in fact the according to the chosen path metric the SBPP produced path is better than the RPP path. This evaluation is based entirely on the particular form of the path quality metric chosen. The currently chosen path quality metric described by Equation 3.6 is a simple path integral of J_C. This metric instills a strong bias towards shorter paths.

The SBPP implementation is significantly less efficient than the RPP implementation described in Barraquand and Latombe [3]. Whereas the RPP approach required approximately 3 seconds to produce the shown path, the SBPP algorithm required on the order of 5 minutes to first produce a collision free path. Some portion of this difference is due to implementation itself. The SBPP algorithm is implemented in an interpreted language (Matlab$^{\mathrm{TM}}$). In addition, RPP algorithm uses a pyramid of grids of progressively finer resolutions. The current implementation of slice optimization must calculate the partial derivate of slice quality with respect to each slice configuration component $\frac{\partial J_s}{\partial f_i}(x, y)$ across the entire two dimensional slice surface. The RPP algorithm on the other hand uses a random search method to approximate the gradient of the configuration space potential.

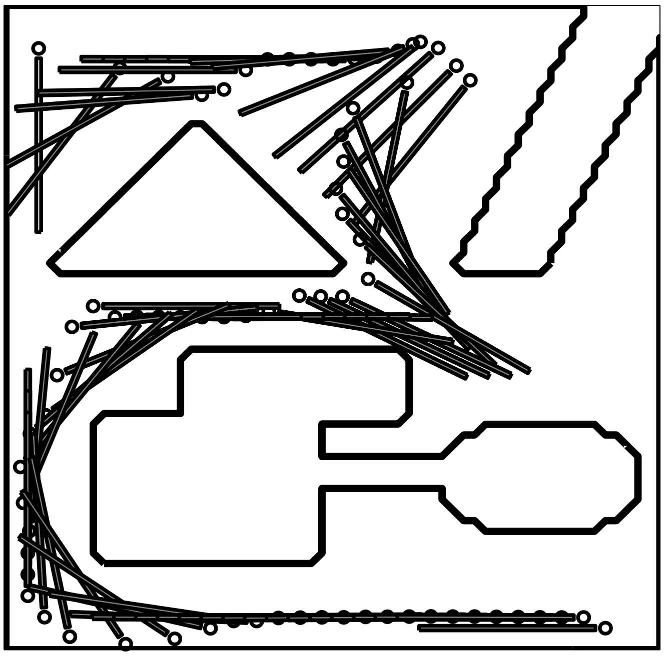

Figure 4.6: SBPP produced plan using manually selected initial slice.

Addressing these implementation differences could bring the SBPP algorithm's performance much closer but fundamentally the slice optimization procedure will possibly always be more computationally intensive than the RPP algorithm. But the key advantage of the SBPP algorithm (that it can incrementally adapt to changing obstacle positions) remains. If the computation cycles spent before an external go signal are regarded as being "free", then a hybrid strategy could apply the slice optimization procedure of the SBPP algorithm (before the go signal) in order to continually update the slice representing a good subset of the entire configuration space. Once the external go signal is received, the RPP algorithm could be applied to the two dimensional intra-slice path planning problem. The RPP algorithm applied to this two dimensional problem will be faster than when applied to the full configuration space. If the produced slice does not contain a collision free path than the RPP algorithm could then be applied to the more difficult full configuration space problem.

4.1.2 Adding New Obstacles

This section demonstrates the results of Experiment 2. This experiment has two parts. Each part involves a test using a large set of obstacles Obs_1. In one part we measure path quality as a function of time when using an f_0 (the slice which was calculated assuming there were no obstacles) as an initial slice. In the other part, an intermediate slice is computed using just a subset of Obs_1 (denoted Obs_0 in the description of Experiment 1). This intermediate slice is used as the initial slice of the second part of the experiment.

In this experiment, 30 initial objects are generated in the same manner as in Experiment 1. After the slice optimization has reached a local minima, 10 additional objects are added using the same random distribution. This new larger set of obstacles is used to test the SBPP with two different initial conditions (slices). The first initial condition (the incremental case) uses the asymptotic slice produced using the original smaller set of obstacles. The second initial condition uses the normal initial slice f_0 produced using no obstacles. The experimental results are shown in Figure 4.7.

Not surprisingly, the incremental case clearly performs better than the non-incremental case. The slice represents a compilation of the available obstacle information. In those regions of the workspace where no new obstacles are added, the slice contains configurations which are still of high quality (as measured by J_C). The availability of higher quality configurations within the slice allows the intra-slice planner to find paths which are in general of higher quality.

4.1.3 Tracking Obstacles

This section illustrates the ability of the planner to track continuous changes in the workspace. In these experiments the ratio of slice optimization increments to obstacle motion (r) was set to 3. Thus each time the obstacles moved the slice optimization procedure was run for three iterations. As before we plot the difference between the quality of the produced plan with that of the plan produced if the

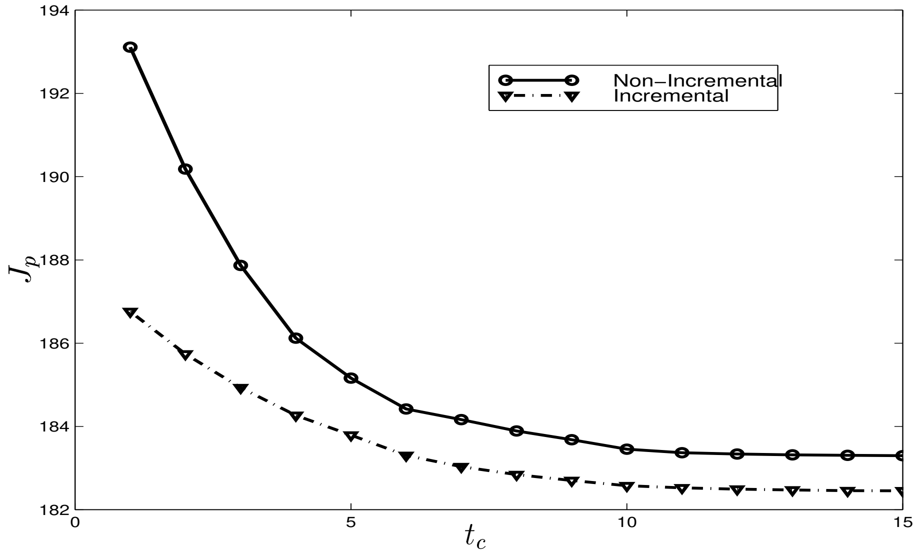

(a) Avg Path Cost

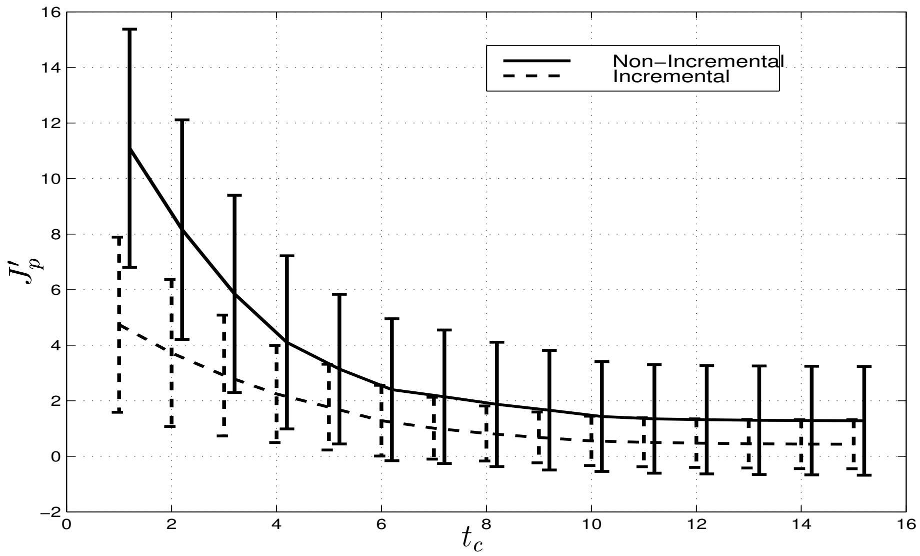

(b) Avg of $J'_p = J_p - \hat{J}_p$

Figure 4.7: Plots Comparing Performance of Incremental and Non-Incremental Cases

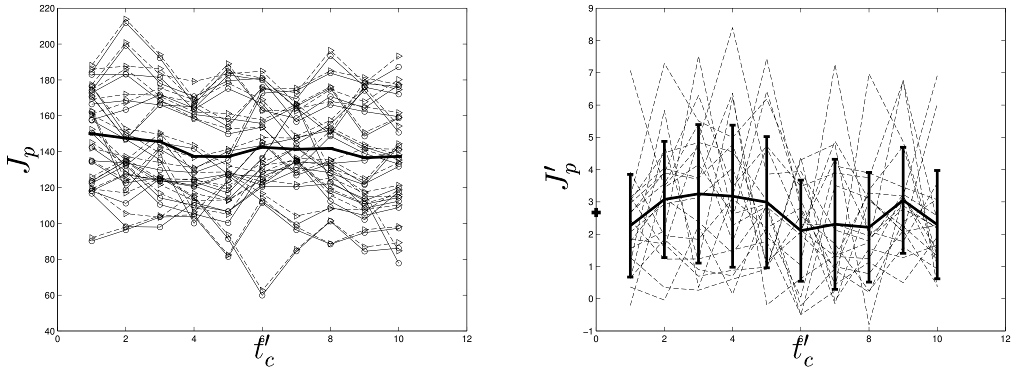

(a) J_p for each obstacle set. The heavy solid line shows $Avg(J_p)$

(b) J'_p for each obstacle set. The + on the vertical axis shows the average taken across each obstacle set and each iteration.

Figure 4.8: Results of Experiment 3

obstacles were frozen and the algorithm was allowed to run until a local minimum was reached.

The experiment results are plotted in Figure 4.8. The left shows plots of J_p over 10 iterations for 20 different obstacle sets. The right shows the difference between the actual path cost (after having iterated just r times) and the asymptotic path cost ($J_p(\infty)$). Thus we see that with just 3 iterations the computed path cost is on average approximately 2.7 units higher than the path quality which would have been obtained if the obstacles were frozen and the SBPP algorithm allowed to run for an extended period of time.

4.2 PUMA Manipulator

This section describes the results of applying the algorithm to a PUMA 560 serial link manipulator (Figure 4.9). This robot is significantly more challenging than the rod for several reasons. First the kinematics of a serial link manipulator are much more complicated than for a simple planar rigid body such as the rod. In addition, the PUMA's workspace is three-dimensional requiring significantly more

Figure 4.9: A PUMA 560. This robot is used for the experiments described in this section and the reaching and grasping demonstration described in the following section.

memory resources than the planar case. Finally, the PUMA's joints all have physical limitations which need to be incorporated into the path planning.

The workspace was defined to be ($(-1.0, 1.3)\times(-0.2, 1.4)\times(-0.8, 1.0)$ in meters). This workspace was described by a three dimensional bitmap of dimensions 30 by 24 by 24. The reachable workspace is actually much smaller so that a slice requires only 339 elements rather than $30 \times 24 = 720$ elements.

In all the experiments described, the goal configuration described as (x, y, height, roll, pitch, yaw) was $(x_g = 0.51, y_g = 0.77, z_g = -0.47, \gamma_g = 0, \beta_g = \pi/8, \alpha_g = atan2(y_g, x_g))$. The start configuration was $(x_s = -0.52, y_s = 0.77, z_s = -0.47, \gamma_g = 0, \beta_s = \pi/8, \alpha_s = atan2(y_s, x_s))$. The origin of the world frame lies within the shoulder approximately 2/3 of a meter above the table surface. The initial slice as well as the start and goal configurations are illustrated in Figure 4.10. The apparent discontinuity in the mesh plot of the yaw in (b) is caused by the yaw wrapping past π. The SBPP algorithm performs all angular arithmetic modulo 2π so the yaw slice component is still continuous.

The goal configuration specifies the position and orientation of a coordinate frame attached to the end effector. In this case the end effector is a Belgrade/USC hand. The transform which describes the relation between the hand frame and the robot's tool frame is shown in Equation 4.2. The hand's geometry was described using 7 evenly spaced control points ((0.12 0 0) ... (-0.3466 0 0) described in the hand's coordinate frame). The robot's kinematics and link lengths were adopted from the data files of the RCCL robot control library [31].

$$\begin{pmatrix} 0 & -1 & 0 & 0.03 \\ 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 0.3466 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{4.2}$$

The obstacles were generated as follows: 40 spherical obstacles were placed randomly across the workspace and with random radii between 0.025 and 0.15 meters (both with uniform distributions). The optimization criterion weights are described in Table 4.2. These weights were chosen manually. The differing w_k values reflect individual joint's range of possible motion. The third value of w_d corresponds to the z slice component and its higher value was necessary to keep the slice relatively flat. This avoids problems caused by the relatively coarse discretization.

4.2.1 Static Obstacles

The planner was tested on 10 randomly generated obstacle sets. Figure 4.11 illustrates the results. The left hand portion shows a plot of average path quality as a function of time t_c. The right hand plot shows the same adjusted path quality

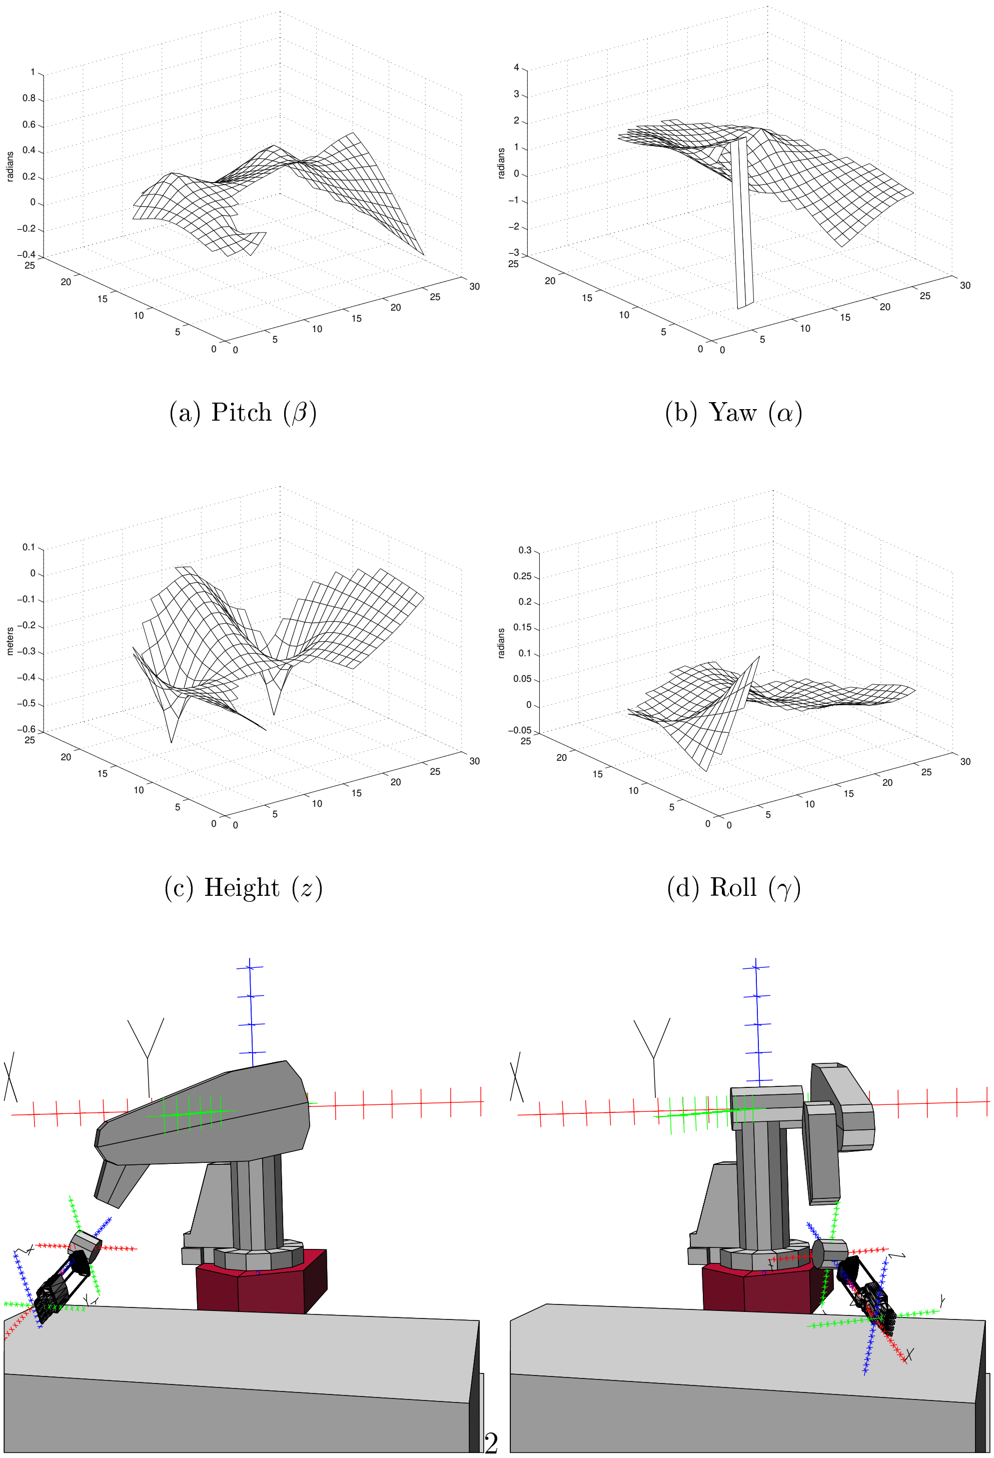

(a) Pitch (β)

(b) Yaw (α)

(c) Height (z)

(d) Roll (γ)

(e) Start Configuration

(f) Goal Configuration

Figure 4.10: (a-d) Mesh plots of each component of the initial slice. (e) and (f) show the start and goal configurations respectively.

w_s	(1 0.1000 10 1)
w_d	(1 1 50 1)
w_k	(1 1 10 10 10 1)
w_p	20

Table 4.2: Cost function weights used to describe J_C.

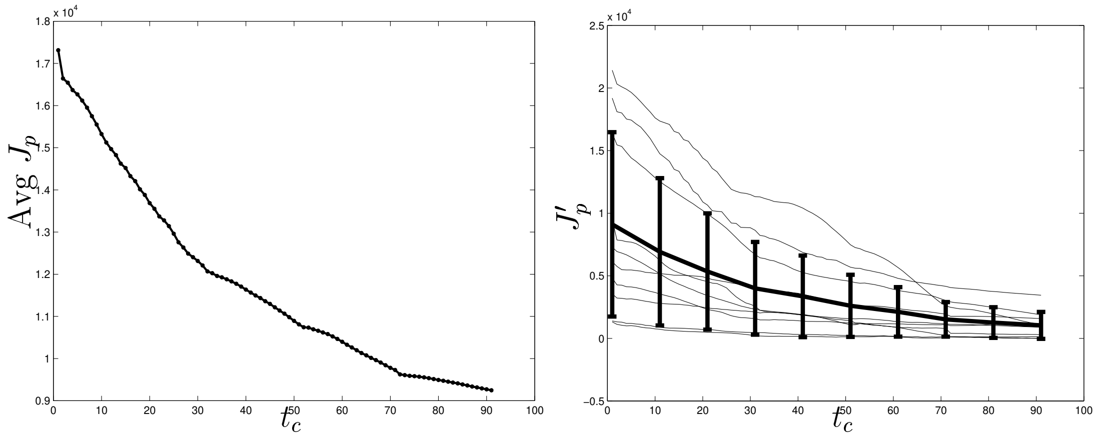

(a) Avg Path Cost

(b) Avg of $J'_p = J_p - min(J_p)$

Figure 4.11: Results of Experiment 1 applied to PUMA robot.

metric used in the previous section and described by Equation 4.1. The greatest improvement in path quality occurs during the first few iterations of the SBPP algorithm. On average 50 percent of the improvement comes after approximately 30 iterations.

4.2.2 Adding New Obstacles

This experiment was conducted in the same manner as the planar case. The slice optimization was allowed to reach a local minimum using 40 randomly distributed obstacles, then 20 additional obstacles were added. The performance was compared between the incremental (based on the slice obtained from the 40 obstacle environment) and the non-incremental case (based on the initial slice obtained from a zero obstacle environment). The results are shown in Figure 4.12. As before the incremental algorithm is superior.

4.3 Reaching and Grasping Demonstration

This section describes a demonstration of the SBPP algorithm applied to a simple reaching and grasping task. This demonstration uses the robot hand shown in Figure 4.13 which was developed at the University of Southern California and the University of Novi-Sad at Belgrade. The Belgrade/USC robot hand has five digits but requires only four degrees of freedom. A rocker arm mechanism couples two pairs of fingers mechanically [5, 21]. The hand is controlled by a Motorola 68332 micro-controller running a simple proportional-derivative (PD) control strategy.

The basic demonstration task is as follows:

1. Grab a visual image from an overhead mounted camera
2. Identify objects on a table from a known set
3. Translate object locations from visual frame to world frame
4. Find a path from the current robot configuration to a user specified target object which avoids collisions with any of the objects identified in step 2
5. Execute the planned path and grasp the target object

One of the objects is a large drawer and if the user selects the drawer as the grasped object the planned motion will include grasping the handle and opening the drawer a predetermined distance. Once the drawer is open the process can repeat and identify objects contained within the drawer. The object(s) in the drawer can then be chosen as the target object and removed from the drawer. The hand is nearly as large as the drawer and this motion requires precise planning to avoid collision with the sides of the drawer.

Each step of the demonstration task is described in more detail below.

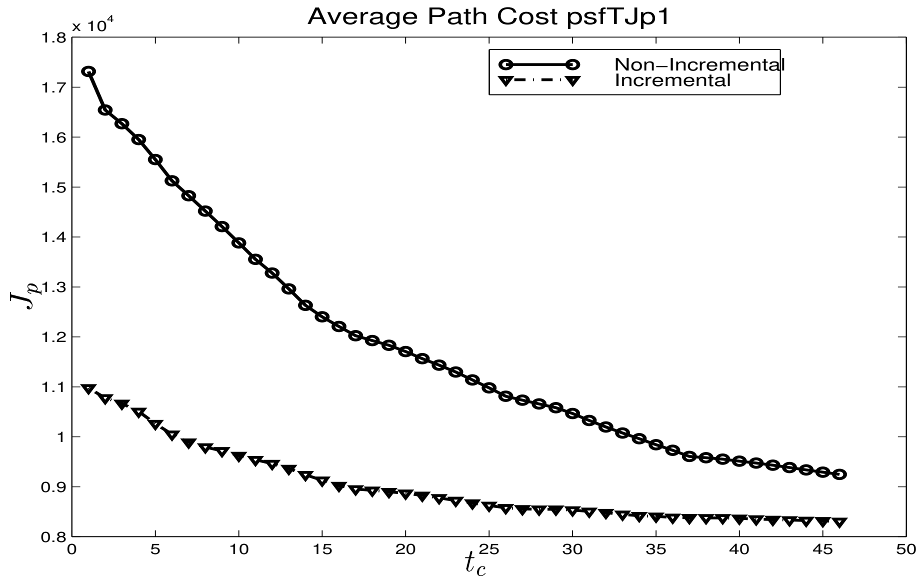

(a) Avg Path Cost

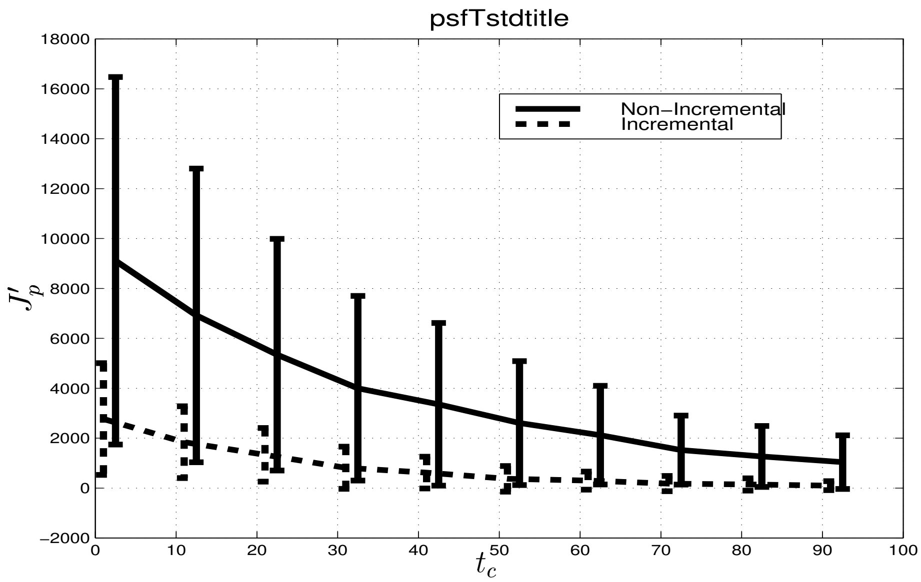

(b) Avg of $J'_p = J_p - \hat{J}_p$

Figure 4.12: Plots Comparing Performance of Incremental and Non-Incremental Cases

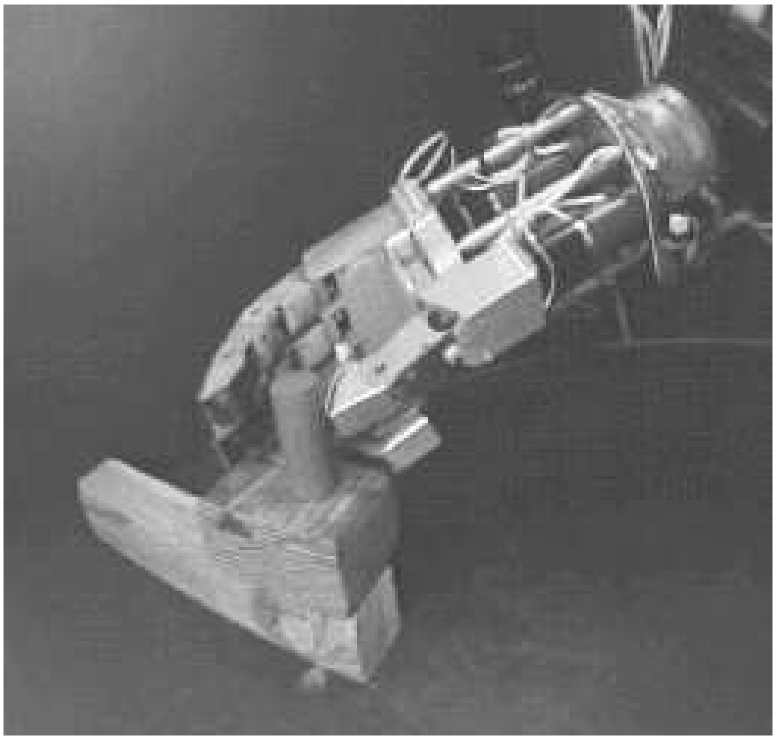

Figure 4.13: The four degree of freedom Belgrade/USC robot hand.

Figure 4.14: Example image from overhead camera. The small marks indicated the computed center and major axis of each identified object.

4.3.1 Grab an Image and Identify Objects

A CCD camera is mounted on the ceiling looking down on the table on which the obstacles are located. Figure 4.14 shows an example image containing several objects. The gray-scale image is processed, thresholded to form a binary image, and a simple region growing algorithm is applied to produce a set of connected regions. The size and second moment of each region are compared to a table of known objects and each is classified according to the nearest neighbor method. The set of known objects consists of: three wooden blocks, four cylinders and the drawer.

Clearly the image recognition process described above is somewhat rudimentary. The algorithm makes use of the fact that the objects are always located on a plane which is located at a fixed distance from the camera and that the set of objects to be identified is quite small. Nonetheless the algorithm is able to identify the objects and their locations accurately enough to allow the planned motion to successfully grasp the objects in most cases.

Translate Object Locations to World Frame

In general a single image is not enough to locate an object in three dimensional space. But since the position of the table top on which the objects sit is fixed with respect to the camera, the position of the objects described in the world frame are easily determined. We used the camera calibration technique of Tsai[50] and the public domain implemention available via the web[1] .

Goal Specification

The transformation from the target object id to a specific goal configuration is done differently depending on the target object. For cylindrical objects, the algorithm automatically chooses the yaw and pitch angle as part of the slice optimization process. Thus if object i is selected as the target object, the goal is specified as $(x_g = x_i, y_g = y_i, z_g = 0.9 * ht_i + z_{table}, \gamma_g = 0, \beta_g = \varnothing, \alpha_g = \varnothing)$ where $\varnothing$ denotes a variable which will be calculated as part of the slice optimization. (x_i, y_i) describe the object's horizontal location. The z value is chosen to be 9/10 of the way up the cylinder of height ht_i. This value allows the hand to avoid the table top and provides relatively stable grasps. z_{table} describes the z position of the top of the table. When the cylinder is specified as the goal the initial yaw value is normally set to $atan2(y_i, x_i)$.

For the blocks the goal is specified as $(x_i, y_i, 0.9 * ht_i + z_{table}, \pi/2, \pi/8, \theta_i + \pi/2)$, where θ_i describes the orientation of major axis and angular addition is modulo 2π. This causes the hand to grasp the object from the top perpendicular to the major axis. If the drawer is the target, the location of the drawer handle is computed (based on the known geometry) and a configuration is chosen which causes the robot hand's fingers to be placed between the handle and the drawer (with the back of the hand towards the drawer).

The system locates the drawer using a simple rectangular mark placed on the top of the drawer. The system assumes the handle is on the side closest to the

[1] http://www.cs.cmu.edu/afs/cs.cmu.edu/user/rgw/www/TsaiCode.html

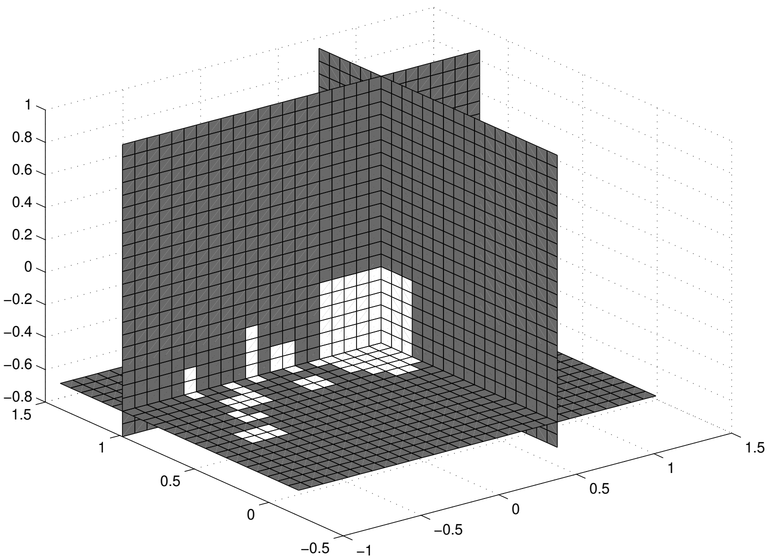

Figure 4.15: Three dimensional bitmap derived from camera image shown in Figure 4.14. The areas containing obstacles are denoted with white.

robot base. Since the system is unable to determine if the drawer is actually open or not, the current implementation assumes that the drawer is initially closed and updates the status only after being commanded to open it. If the drawer is open and the target object is contained within the boundaries of the open drawer, the z component of the goal configuration is adjusted upwards to compensate for the distance between the bottom of the drawer and the table top. In addition, if the target object is inside the drawer the initial yaw (at the slice element corresponding to the (x_g, y_g) is chosen to be parallel with the face of the drawer. This was necessary to avoid arriving at a locally minimal slice which caused the goal configuration to orient the hand perpendicular to and *through* the front face of the drawer.

Path Planning

Figure 4.15 shows an illustration of the Obstacle bitmap constructed from the image shown in Figure 4.14. The parameters describing the optimization criteria were the same as shown in Table 4.2.

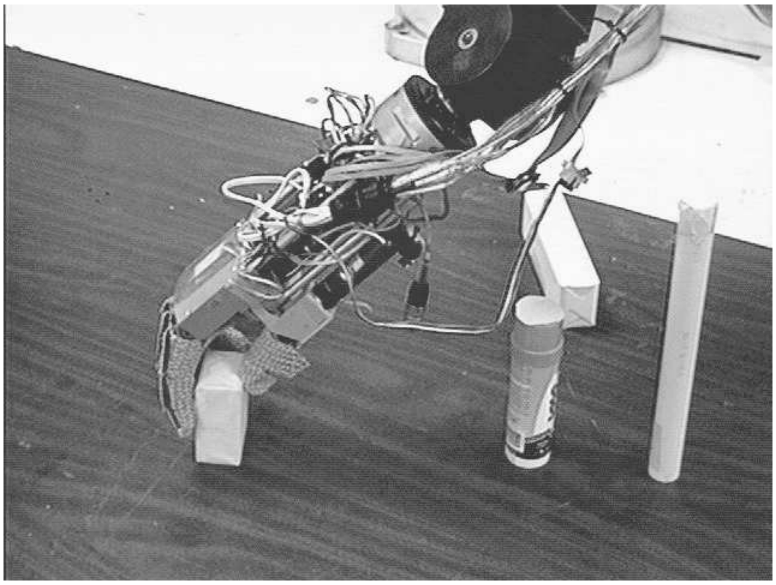

Figure 4.16: Robot hand picking up a wooden block.

The figures below show the robot executing plans with several different targets. The first, Figure 4.16 shows a block being picked up. Figure 4.17 shows the robot grasping the drawer handle prior to opening the drawer. Figure 4.18 shows the robot during grasping of an object in the open drawer.

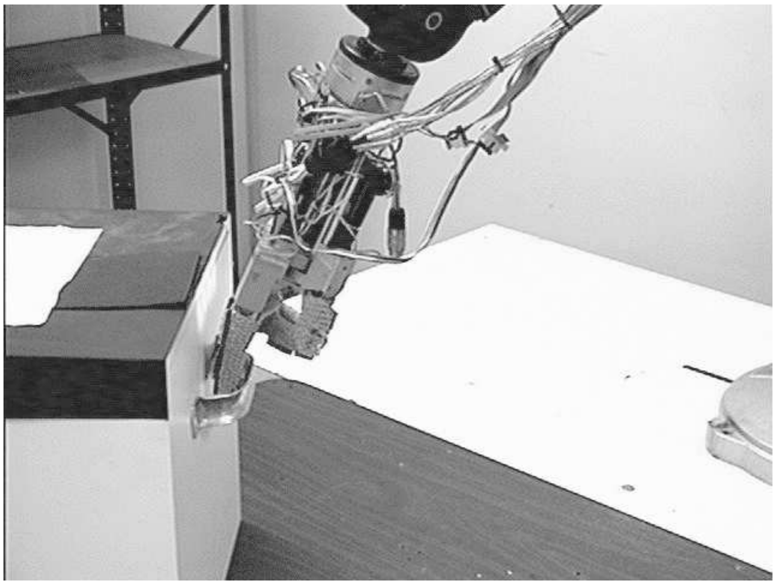

Figure 4.17: Robot hand grasping drawer handle prior to opening.

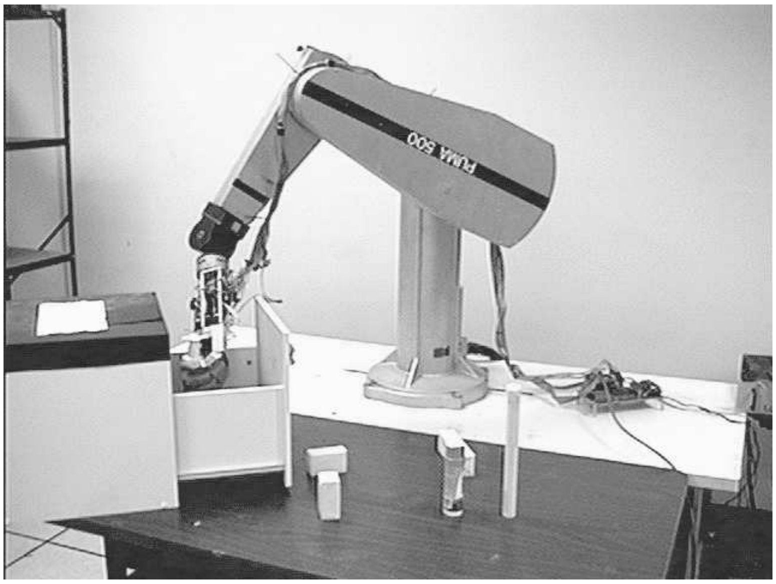

Figure 4.18: Robot hand grasping cylinder contained in drawer.

Chapter 5

Conclusions and Future Work

The concept of a non-planar subset of configuration space called a slice and a process by which it can be incrementally updated to reflect workspace obstacles has been shown. Furthermore the slice concept has been incorporated into the Slice-based Path Planning (SBPP) algorithm which has the following novel features:

Flexibility The path planner can reflect path quality constraints beyond simple collision avoidance. The algorithm was demonstrated using a cost criterion which included terms based on workspace obstacles as well as kinematic constraints, total path distance, first and second spatial derivatives.

Anytime Planning The algorithm can produce a satisfactory path plan quickly and can then incrementally improve that plan given additional computation time. Empirical measurements of the trade-off between computation time and path quality were demonstrated.

Scalability The SBPP algorithm can be applied to systems with many degrees of freedom and in three dimensional workspaces.

In addition, we have implemented the algorithm in a working path planner and applied the planner to a six degree of freedom serial link manipulator. This planner has been demonstrated in simulation on problems with both fixed and dynamic obstacles. The planner was also used to perform desktop reach and grasping motions using a PUMA 560 manipulator mounted with the Belgrade/USC hand.

Future Work

The SBPP algorithm is a step towards capable real-time path planning but many refinements are possible.

One of the original motivations for developing a motion planner that was based in part on optimization was the ability to integrate grasp selection and path planning. The implemented planner is capable of choosing from multiple goals. Thus, if a grasp planner could enumerate a set of possible grasps and the corresponding position and orientation of the robot's wrist, that information could be incorporated into both the slice and path optimization criterion thereby integrating grasp selection directly with path planning.

The current implementation could also be improved by adding the ability to track the ongoing execution of the motion. In addition, the current vision software is extremely inefficient. An implementation of the vision component which could track a moving object in real-time would better illustrate the online capability of the SBPP algorithm.

The current implementation is incremental but is far from being real-time. This is due in part to the fact that the system is implemented using uncompiled Matlab$^{\text{TM}}$. But besides reimplementing the algorithm in a faster programming language many other performance improvements are possible. The most computationally intensive portion of the algorithm is the slice optimization and this could be substantially improved in several ways. The current slice representation is based on a rectangular grid. Adopting a finite-element method would allow more efficient computation and potentially improve the quality of the produced paths. In addition, the overall algorithm and the slice optimization process are both well suited to parallel implemention. It would also be beneficial to control the allocation of processor cycles more carefully. Currently the ratio of time spent optimizing the slice to the time spent planning an intra-slice path is constant. Zilberstein has shown how such ratios of computational resources allocated to anytime algorithms can be optimized automatically [56] .

Another potential improvement would replace the existing intra-slice path planner with a planner capable of producing smooth paths. More significantly, nothing about the slice concept or the SBPP algorithm requires a two dimensional slice. Indeed, a three dimensional slice would allow the representation and execution of paths which contain pure vertical motions.

There is also substantial work to be done with regards to the SBPP algorithm. One important question is how the algorithm can be extended to handle motion planning problems as well as path planning problems. One possible approach is to add a temporal aspect to the intraslice path planning portion of the algorithm. Another approach would add time explicitly as another slice component. This would of course require changes to the intraslice path planning in order to ensure that the time specification was monotonic along the path. A related extension would incorporate a model of moving obstacles motions and the direct reflection of those predictions on the produced plans. This might also eventually lead to a method for coordinating the plans of multiple robots operating in a shared workspace.

Perhaps the most ambitious improvement in the algorithm would be to incorporate constraints in the feasible motions. Currently the intra-slice planner and the slice optimization process are based on the assumption that the robot can always move in any direction. For robots with non-holonomic constraints possible motions are restricted. Perhaps in a similar fashion the constraints describing the physical system's dynamics could be incorporated in a planner which embodies the notion of a slice in state-space rather than in configuration space.

Bibliography

[1] Michael Barbehenn and Seth Hutchinson. Toward an exact incremental geometric robot motion planner. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS '95)*, pages 39–44, 1995.

[2] Jérôme Barraquand and Pierre Ferbach. Path planning through variational dynamic programming. Technical Report 33, Paris Research Laboratory, Digital Equipment Corp., Sept 1993.

[3] Jérôme Barraquand, Bruno Langlois, and Jean-Claude Latombe. Numerical potential field techniques for robot path planning. Technical Report STAN-CS-89-1285, Stanford University, Oct 1989.

[4] Jérôme Barraquand and J.C. Latombe. Robot motion planning - a distributed representation approach. *International Journal of Robotics Research*, 10(6):628–649, Dec 1991.

[5] G. A. Bekey, R. Tomovic, and I. Zeljkovic. Control architecture for the Belgrade/USC hand. In *Dextrous Robot Hands*, pages 136–149. Springer-Verlag, 1990.

[6] George A. Bekey and S.F. Masri. Random search techniques for optimization of nonlinear systems with many parameters. *Mathematics and Computers in Simulation*, 25(3):210–213, 1983.

[7] Rodney A. Brooks. Solving the find-path problem by good representations of free space. *IEEE Transactions on Systems, Man, and Cybernetics*, 13(3):190–197, 1983.

[8] John F. Canny. *The Complexity of Robot Motion Planning*. MIT Press, Cambridge, MA, 1988.

[9] Daniel J. Challou and Maria Gini. Parallel robot motion planning. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 24–51, Atlanta, GA, 1993.

[10] Howie Choset and Joel Burdick. Sensor based motion planning: The hierarchical generalized Voronoi graph. In *Workshop on Algorithmic Foundations of Robotics*, pages 47–62, 1996.

[11] Christopher I. Connolly. *Harmonic Functions as a Basis for Motor Control and Planning*. PhD thesis, Univ. Mass. Amherst, 1994.

[12] Christopher I. Connolly and J. Brian Burns. The basal ganglia and the planning of actions. Technical Report 92-18, University of of Massachusetts, 1992.

[13] Christopher I. Connolly and Roderic A. Grupen. On the applications of harmonic functions to robotics. *Journal of Robotic Systems*, 10(7):931–946, 1993.

[14] C.I. Connolly and J.B. Burns. A model for the functioning of the striatum. *Biological Cybernetics*, 68(6):535–544, Apr 1993.

[15] Thomas L. Dean and Mark Boddy. An analysis of time-dependent planning. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 49–54, St. Paul, MN., 1988. Morgan Kaufmann.

[16] Bernard Faverjon and Pierre Tournassoud. A practical approach to motion-planning for manipulators with many degrees of freedom. Technical Report 951, Institut National de Recherche en Informatique et en Automatique, 1988.

[17] R. Glasius, A. Komoda, and C.C.A.M. Gielen. Neural network dynamics for trajectory formation and obstacle avoidance. *Neural Networks*, 8(1):125–133, 1995.

[18] L. Gouzénes. Strategies for solving collision-free trajectories problems for mobile and manipulator robots. *International Journal of Robotics Research*, 3(4):51–65, 1984.

[19] I.G. Gowda, D.G. Kirkpatrick, D. T. Lee, and A. Naamad. Dynamic Voronoi diagram. *IEEE Transactions of Information Theory*, (29):724–731, 1983.

[20] Joshua Grass. Reasoning about computation resource allocation. *Crossroads*, Sept 1996. Electronic publication: http://www.acm.org/crossroads/xrds3-1/racra.html.

[21] Thea Iberall, G.S. Sukhatme, D. Beattie, and G. A. Bekey. Control philosophy and simulation of a robotic hand as a model for prosthetic hands. In *Proc. IEEE/RSJ International Conference on Robots and Systems (IROS)*, pages 824–831, Yokohama, 1993.

[22] Pekka Isto. A two-level search algorithm for motion planning. In *Proc. IEEE International Conference on Robotics and Automation*, pages 2025–2031, 1997.

[23] Lydia Kavraki and Jean-Claude Latombe. Randomized preprocessing of configuration space for fast path planning. In *Proceedings of the IEEE Intl. Conf. on Robotics and Automation*, pages 2138–2139, 1994.

[24] Lydia Kavraki, Petr Švestka, Jean-Claude Latombe, and Mark Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.

[25] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. Rob. Res*, 5(1):90–99, 1986.

[26] Pradeep Khosla and Richard Volpe. Superquadric artificial potentials for obstacle avoidance and approach. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 1778–1784, Philadelphia, Pennsylvania, April 1988.

[27] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishing, Boston, USA, 1991.

[28] Jean-Claude Latombe. Class Notes - Stanford University CS326A, 1998.

[29] Sukhan Lee and Jun Park. Neural computation for collision-free path planning. *Journal of Intelligent Manufacturing*, 2:315–326, 1991.

[30] M.D. Lemmon. 2-DOF robot path planning using cooperative neural fields. *Neural Computation*, 3:350–362, 1991.

[31] J. Lloyd and V. Hayward. Real-time trajectory generation in multi-RCCL. *journal of robotic systems*, 10(3):369–390, 1993.

[32] Tomás Lozano-Pérez. Automatic planning of manipulator transfer movements. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(10):681–698, 1981.

[33] Tomás Lozano-Pérez. Spatial planning: a configuration space approach. *IEEE Transactions on Computers*, 32:108–120, 1983.

[34] Tomás Lozano-Pérez. A simple motion-planning algorithm for general robot manipulators. *IEEE Transactions on Robotics and Automation*, 3(3):224–238, 1987.

[35] P. Muraca, G. Raiconi, and T. Varone. Cooperative neural field for the path planning of a robot arm. *Journal of Intelligent and Robotic Systems*, 15:11–18, 1996.

[36] Nils J. Nilsson. A mobile automaton: An application of artificial intelligence techniques. In *Proceedings of the 1st International Joint Conference on Artificial Intelligence*, pages 509–520, 1969.

[37] Colm Ó'Dúnlaing and C.K. Yap. A retraction method for planning the motion of a disc. *Journal of Algorithms*, 6:104–111, 1982.

[38] Mark Overmars. A random approach to path planning. Technical Report "RUU-CS-92-32", Ultrecht University, Oct 1992.

[39] Mark Overmars and Petr Švestka. A probabilistic learning approach to motion planning. Technical Report RUU-CS-94-03, Ultrecht University, Oct 1994.

[40] S. Quinlan and O. Khatib. Elastic bands: Connecting path planning and robot control. In *Proceedings IEEE International Conference on Robotics and Automation*, pages 802–807, 1993.

[41] Elon Rimon and D. E. Koditschek. Robot navigation functions on manifolds with boundary. *Adv Appl. Math*, 5(1):412–442, 1990.

[42] Elon Rimon and D.E. Koditschek. The construction of analytic diffeomorphisms for exact robot naviation on star worlds. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 21–26, 1989.

[43] P.F. Rowat. *Representing the Spatial Experience and Solving Spatial Problems in a Simulated Robot Environment.* PhD thesis, University of British Columbia, 1979.

[44] J.T. Schwartz and M. Sharir. On the 'piano mover' problem: II. general techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics*, 4:298–351, 1983.

[45] J. R. Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. http://www.cs.cmu.edu/ quake-papers/painless-conjugate-gradient.ps, 1994.

[46] M. R. Stan, W. P. Burleson, C. I. Connolly, and R. A. Grupen. Analog VLSI for robot path planning. *Journal of VLSI Signal Processing*, 8(1):61–73, 1994.

[47] Mircea R. Stan and Waynge P. Burleson. Analog VLSI for robot path-planning. In *Asilomar Conference on Computers, Signals and Systems*, pages 915–919, 1992.

[48] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction.* MIT Press, 1998.

[49] L. Tarassenko and A. Blake. Analogue computation of collision-free paths. In *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, pages 540–545, 1991.

[50] Roger Y. Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics and Automation*, RA-3(4):323–344, Aug. 1987.

[51] Izu Vaisman. *A first course in differential geometry.* Marcel Dekker, Inc, 1984.

[52] R. Volpe and P. Khosla. Manipulator control with superquadric artificial potential functions: Theory and experiments. *IEEE Transactions on Systems, Man, and Cybernetics*, 20:1423–1436, 1990.

[53] X. Xiaodong and G. A. Bekey. SROMA: An adaptive scheduler for robotic assembly systems. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1282–1287, 1988.

[54] X. Zhu and K. Gupta. On local minima and random search in robot motion planning. Technical report, Simon Fraser Univ, Burnaby, BC, Canada, 1993.

[55] Shlomo Zilberstein. Resource-bounded sensing and planning in autonomous systems. *Autonomous Robots*, 3:31–48, 1996.

[56] Shlomo Zilberstein and S.J. Russel. Anytime sensing, planning and action: A practical model for robot control. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 1402–1407, Chambery, France, 1993.