Optimal Trajectories for Planar Rigid Bodies with Switching Costs

Yu-Han Lyu[1] and Devin Balkcom[2]

[1] Department of Computer Science, Dartmouth College, Hanover, NH 03755, USA
yuhanlyu@cs.dartmouth.edu
[2] Department of Computer Science, Dartmouth College, Hanover, NH 03755, USA
devin@cs.dartmouth.edu *

Abstract. The optimal trajectory with respect to some metric may require very many switches between controls, or even infinitely many, a phenomenon called *chattering*; this can be problematic for existing motion planning algorithms that plan using a finite set of motion primitives. One remedy is to add some penalty for switching between controls. This paper explores the implications of this *switching cost* for optimal trajectories, using rigid bodies in the plane (which have been studied extensively in the cost-free-switch model) as an example system. Blatt's Indifference Principle (BIP) is used to derive necessary conditions on optimal trajectories; Lipschitzian optimization techniques together with an A* search yield an algorithm for finding trajectories that can arbitrarily approximate the optimal trajectories.

1 Introduction

Consider an example problem, inspired by a problem from Mason [14]: a mover wants to move a refrigerator from one location and orientation to another. The refrigerator is too heavy to move by lifting or pushing, but it can be lifted onto any of the four legs at the corners of the square base and rotated. If there are no obstacles, what is the fastest sequence of rotations (with time cost computed as the sum of the absolute values of the angles rotated through)?

For some configurations (moving the refrigerator in a straight line), there exists no optimal trajectory with a finite number of actions: for any trajectory with finitely many switches, there is a faster trajectory with more switches, a phenomenon called *chattering*. When chattering occurs, the refrigerator mover is required to run back and forth between legs of the refrigerator infinitely many times, rotating the refrigerator through an infinitely small angle.

The chattering phenomenon is a fundamental problem in robot motion planning. Sussmann showed that an extension of the well-known Dubins car [9] to include bounds on angular acceleration leads to chattering [26]. Desaulniers showed that chattering may occur if there are obstacles in the environment [8], even for systems that are well-behaved without obstacles.

* This work was supported in part by NSF grant IIS-0643476.

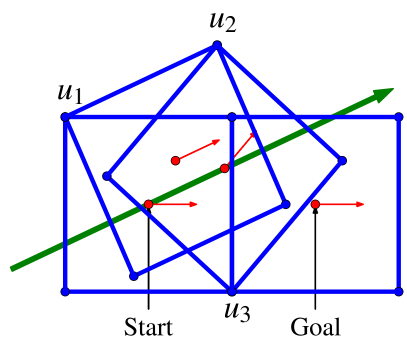

Fig. 1: An approximately optimal trajectory for a refrigerator robot starting at (-2, 0, 0), with unit cost for switching between any pair of controls. The green line is the control line.

A natural approach to avoiding trajectories that switch frequently between controls is to charge a fixed cost for switches. This fixed cost both avoids chattering, and penalizes otherwise un-modeled costs (such as the cost of wearing out a switching mechanism, or the time cost of running between legs in the refrigerator-mover's problem). We give an approximately optimal trajectory for refrigerator mover's problem with switching costs in Fig. 1; this trajectory was generated by the algorithm we will present in this paper. In the robotics community, the model of charging a fixed cost for discontinuous switches between controls has been used in practice [2, 25], but the implications of switching costs for optimal trajectories have perhaps not been thoroughly explored.

The main contributions of this paper are:

1. Necessary conditions for optimal trajectories for rigid bodies in the plane in the costly-switch model, i.e. Theorem 1 and 2.
2. A practical algorithm that finds approximately optimal trajectories for rigid bodies in the plane, with arbitrarily small additive error. This algorithm may also be trivially adapted to the zero-switch-cost model in the case that chattering does not occur.

Rigid bodies are building blocks for many models of robotic locomotion or manipulation systems, and rigid bodies provide a good example for study of optimal trajectories. We apply *Blatt's Indifference Principle* (BIP) to show existence of optimal trajectories, and to derive necessary conditions on these trajectories; for the simple case of rigid bodies in the plane, analytical integration of certain differential equations of BIP is possible. We then show that Lipschitzian optimization techniques can be applied to find trajectories between pairs of configurations, with arbitrarily small error in both final configuration and time cost.

We admit that this paper is quite technical, and builds on a body of previous work that is also quite technical. Nonetheless, we consider the costly-switch model to be fundamental, and Lipschitzian optimization techniques appear to provide a powerful approach to finding algorithms that provide guarantees of approximate optimality.

This paper extends work by Furtuna [10], which derived strong necessary conditions for optimal trajectories for rigid bodies in the plane with zero switching cost, but which did not provide algorithms to connect particular pairs of configurations, with bounds on error. This paper also extends [28], which does provide algorithms for the zero-switching-cost problem; however, these previous algorithms are inefficient for single-source, single-destination problems. Finally, this paper extends our work in [13], which derives some analytical solutions to simple versions of the costly-switch model using other techniques.

Related Work. For some models of mobile robots in the plane, optimal trajectories can be found analytically, including Dubins [9, 7], Reeds-Shepp [20, 26, 24] cars. We and many other researchers have tried to generalize techniques (typically based on Pontryagin's Maximum Principle [19]), aiming to gain a greater understanding of optimal motion for mobile robots [21, 5, 4, 6, 1, 22]. However, we are aware of little work in the robotics community providing strong results on optimal trajectories with a cost of switches; a notable exception is work by Stewart using a dynamic-programming approach to find optimal trajectories with a costly-switch model [25].

The problem of costly switches has been studied in the optimal control community with results dating back as far as the 1970s. One of the most powerful tools for solving optimal control problems, Pontryagin's Maximum Principle (PMP) [19], does not appear to be the right tool to characterize optimal trajectories in the costly-switch model due to the discontinuity with respect to time in the control and cost functions. In [3], Blatt proposed a model in which the control set contains certain primitives (a discrete set of actions), and there is some fixed cost associated with switching between controls. Blatt characterized a set of necessary conditions for optimal trajectories under this model; these necessary conditions are known as Blatt's Indifference Principle (BIP). Blatt showed that optimal trajectories always exist and the number of actions must be finite. Blatt's necessary conditions are similar to, but weaker than, those provided by PMP; using BIP to solve an optimal control problem is more challenging than using PMP in the cost-free-switch model. In Blatt's model, the control set is a discrete set, but other models have been proposed [16, 11, 15].

Although the costly-switch model was proposed in the '70s, no algorithms for finding optimal trajectories in costly-switch model were proposed until the '90s [27, 25]; several algorithms have been developed recently [12, 30]. These recent approaches are based on approximating the control function as a piecewise-constant functions, and applying global optimization techniques to find optimal solutions. These algorithms converge to optimal solutions as the number of iterations approaches infinity, but cannot guarantee a bound of error within finite time. In this paper, we provide a stronger result for a particular system; our algorithm guarantees a bound of error within finite time, for the restricted problem of finding optimal trajectories of rigid bodies in the plane.

In the costly-switch model, due to the similarity between BIP and PMP, by adapting Furtuna's analysis [10], we derive some general results that geometrically characterize optimal trajectories for rigid bodies in the plane with costly switches. Based on the necessary conditions for optimal trajectories, we also categorize optimal trajectories into several types that are similar to Furtuna's categorizations.

Although our conditions seem similar to Furtuna's conditions, our conditions are weaker due to the generality of the costly-switch model. We show that the problem of finding optimal trajectories has two important parts: one is to determine an optimal sequence of actions (the discrete structure of the trajectory), and the second is to determine an optimal characteristic value $H \in \mathcal{R}$ which in some sense parameterizes the shape of the trajectory.

Model and Notation. We use $q = (x, y, \theta) \in SE(2)$ to denote a configuration, and $u = (v_x, v_y, \omega) \in \mathcal{R}^3$ to denote a control: x and y velocities in a frame attached to the body (robot frame), and angular velocity. Let U be the control space containing a finite number of primitives: constant-control actions. For example, one action might be $(v_x, v_y, \omega) = (1, 0, 0)$, corresponding to driving in a straight line.

For a configuration q_0, if we apply a sequence of actions $\mathbf{u} \in U^n$ with a sequence of durations $\mathbf{t} \in \mathcal{R}_+^n$, then the result is a configuration $q' = q(q_0, \mathbf{u}, \mathbf{t}) \in SE(2)$, where q is a continuous function that integrates the control over time in the world frame and then adds q_0 to obtain the resulting configuration q'. Hence, a *trajectory* can be represented as a pair of sequences $(\mathbf{u}, \mathbf{t})$.

We model the cost of control switches as a function $C : U \times U \to \mathcal{R}_+$ that depends on the control applied before and the control applied after. Furthermore, we assume that for any three controls u_a, u_b, and u_c, the cost of switching satisfies the *triangle inequality*, $C(u_a, u_b) + C(u_b, +u_c) \geq C(u_a, u_c)$, to ensure that switching from u_a to u_c directly is always faster than switching to u_c through other intermediate controls. The cost of a trajectory is the summation of all durations and all switch costs of a trajectory.

Problem statement: given a start configuration q_s, a final configuration q_f, a finite control set U, and a cost function C, find a trajectory $(\mathbf{u}, \mathbf{t})$ with the minimum time cost, subject to $q(q_s, \mathbf{u}, \mathbf{t}) = q_f$.

2 Necessary Conditions of Optimal Trajectories

In this section, we will derive necessary conditions for optimal trajectories for rigid bodies in the plane in the costly-switch model. Based on these necessary conditions, we classify optimal trajectories into several classes.

2.1 Extensions of Previous Results

Due to the similarity between BIP and PMP, several results in [10] in the cost-free-switch model can be extended to the costly-switch model by similar mechanisms. Hence, we list these results here and omit their proofs.

Theorem 1. *Any optimal trajectory* $(\mathbf{u}^*, \mathbf{t}^*)$ *with* n *actions in the costly-switch model satisfies the following property: there exist four constants* $H > 0$*,* k_1*,* k_2*,*

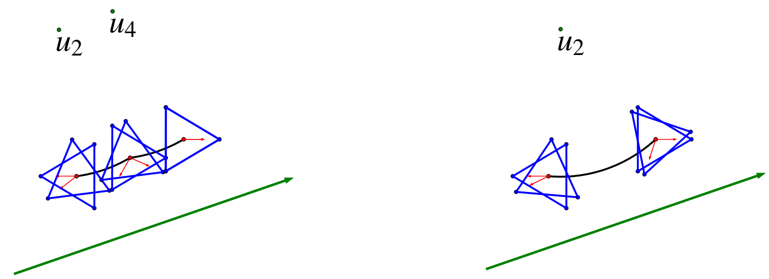

(a) Cost-free-switch model. (b) Costly-switch model with switch cost 1.

Fig. 2: Trajectories for an omni-directional vehicle starting at (-3, -1 π). For the cost-free-switch model, the optimal trajectory takes 5 actions. For the costly-switch model, the (approximately) optimal trajectory takes 3 actions. Green lines are control lines.

and k_3, such that for any control u_i^*, $1 \leq i \leq n$, with the instantaneous velocity (v_x, v_y, ω) in the world frame when u_i is applied at a configuration (x, y, θ), we have

$$k_1 v_x + k_2 v_y + \omega(k_1 y - k_2 x + k_3) = H, \textit{where } k_1^2 + k_2^2 \in \{0, 1\}. \tag{1}$$

A trajectory $(\mathbf{u}, \mathbf{t})$ is called *extremal*, if there exist four constants $H > 0$, k_1, k_2, and k_3, such that Equation 1 is satisfied. Equation 1 is virtually identical to the necessary condition derived using PMP for the cost-free-switch problem, except that there is no requirement that controls maximize the Hamiltonian H. Instead, H need only be constant throughout the trajectory.

An extremal trajectory with constants H, k_1, k_2, and k_3, is called a *control line trajectory*, if $k_1^2 + k_2^2 = 1$. An extremal trajectory with constants H, k_1, k_2, and k_3, is called a *whirl* trajectory, if $k_1^2 + k_2^2 = 0$.

Control Line Trajectories. There is a nice geometric interpretation for Theorem 1 when $k_1^2 + k_2^2 = 1$, related to the control line interpretation in [10]. For a control line trajectory $(\mathbf{u}, \mathbf{t})$, we define its corresponding *control line*, represented as (k_1, k_2, k_3) as a line in the plane with heading (k_1, k_2) and distance k_3 from the origin. Now, consider Equation 1. The term $k_1 v_x + k_2 v_y$ becomes the translational velocity along the vector (k_1, k_2) and the term $k_1 y - k_2 x + k_3$ becomes the *signed distance* from the reference point of the robot to the control line. By Corollary 1 in [10], when a rotation is applied, the signed distance from the rotation center to the control line is always H/ω. Similarly, when a translation is applied, the dot product between (k_1, k_2) and (v_x, v_y) must be H. See Fig. 2 for an (approximately) optimal trajectory for an omni-directional vehicle with control lines in the cost-free-switch model and in the costly-switch model. When the switch cost is introduced, optimal trajectories tend to use fewer number of switches.

Whirl Trajectories. For whirl trajectories, Equation 1 only implies that all angular velocities are equal. We can also extend the result in [10] to the costly-switch model. Due to space limitations, we do not include the result here.

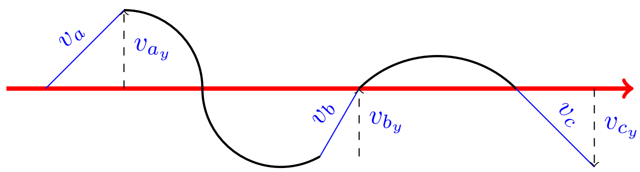

Fig. 3: Illustration of proof of theorem 2: a trajectory containing three actions of translations, v_a, v_b, and v_c. The sign of v_{ay} and v_{by} are the same.

2.2 Necessary Conditions for Control Line Trajectories

We can prove a further necessary condition for a control line trajectory to be optimal.

Theorem 2. *In the costly-switch model, any optimal control line trajectory has either zero translation actions, one translation action, or two non-parallel translation actions.*

Proof. Let $g = (\mathbf{u}, \mathbf{t})$ be a control line trajectory. Suppose that g is optimal but has two parallel translation actions. Let v_a and v_b be the velocity vectors in the world frame of two non-parallel translation actions of g. We can remove the action of v_b from g and increase the duration of v_a to $t_a + t_b$. The resulting trajectory still reaches the goal but has one fewer control and hence has smaller cost. This contradicts the optimality of g.

Suppose that g is optimal but has more than two non-parallel translation actions Let v_a, v_b, and v_c be the velocity vectors in the world frame of three translation actions of g. By Equation 1, we know that the projection of v_a, v_b, and v_c onto the control line must be H. Let v_{a_y}, v_{b_y}, and v_{c_y} be the projection of v_a, v_b, and v_c onto the norm of the control line. By the Pigeonhole Principle, we know that at least two of v_{a_y}, v_{b_y}, and v_{c_y} have the same sign.

Without loss of generality, assume that v_{a_y} and v_{b_y} have the same sign; let their durations be t_a and t_b respectively. See Fig. 3. If $v_{a_y} = v_{b_y}$, then the velocity vectors v_a and v_b are identical. This contradicts the assumption that v_a and v_b are non-parallel. If $v_{a_y} \neq v_{b_y}$, then without loss of generality, we assume $|v_{a_y}| > |v_{b_y}|$.

Since the projections of v_a and v_b onto the control line are the same, we can remove the actions of v_b from g and increase the duration of v_a to $t_a + \frac{t_b|v_{b_y}|}{|v_{a_y}|}$. Let u be the control corresponding to the translation vector v_b. Let u_p and u_q be the control before and after u in the trajectory. The new trajectory will decrease cost by $\frac{t_b|v_{b_y}|}{|v_{a_y}|} - t_b - C(u_p, u) - C(u, u_q) + C(u_p, u_q)$, which is strictly larger than zero. Hence, the resulting trajectory has smaller cost but still reaches the goal. This also contradicts the optimality of g.

We call a control line trajectory that has either zero translation actions, one translation action, or two non-parallel translation actions an *extremal control line trajectory.*

Singular, TGT, and Regular Trajectories. In [10], Furtuna classified trajectories with control lines into four classes: singular, TGT, generics, and regular. Here, we also classify extremal control line trajectories into four subtypes and we name trajectories by the names of their counterpart trajectories in the cost-free-switch model. An extremal control line trajectory is called *singular* if there exists a non-zero measure interval along the trajectory that multiple controls have the same Hamiltonian value within this interval.

As an extension of a result in [10], any singular trajectory in costly-switch model contains exactly one translation with velocity vector parallel to the control line, or contains a switch from one translation to another translation. Hence, by Equation 1, the Hamiltonian values either equal to the velocity of the only translation, or can be computed from the pair of consecutive translations. Since the control set U is given, the set of all possible Hamiltonian values for singular trajectories is finite.

An extremal control line trajectory is called *generic* if the trajectory is not singular. A generic trajectory is called TGT if both the first control and the last control are translations. For a TGT trajectory, when the initial configuration and goal configuration are given, we can obtain the Hamiltonian value analytically, using methods from [10]. A generic trajectory is called *regular* if it either starts or ends with a rotation. For regular trajectories, we do not have enough information to determine the Hamiltonian value, and hence finding optimal regular trajectories is the most challenging task.

2.3 Taxonomy of Optimal Trajectories

We summarize the taxonomy of optimal trajectories as Fig. 4. Since the Hamiltonian values for whirl, TGT, and singular trajectories can be determined, the problems of finding optimal trajectories in these three classes is equivalent to finding an optimal sequence of controls, a discrete search problem. For these three classes, we have designed three different A* search algorithms to find candidate optimal trajectories by searching over discrete trajectory structures; due to space limitations, we omit the details, and focus on the most challenging case, regular trajectories.

The problem of finding optimal regular trajectories has two ingredients: one is finding the Hamiltonian value H, which is a continuous variable, and another one is finding the sequence of controls, chosen from a finite set.

3 Optimal Regular Trajectories

A regular trajectory is a generic trajectory either starting or ending with a rotation.

3.1 Extensions of Previous Results

Due to the similarity between BIP and PMP, several results in [10] in the cost-free-switch model can be extended to the costly-switch model with a few simple modifications. We list these results here and omit their proof.

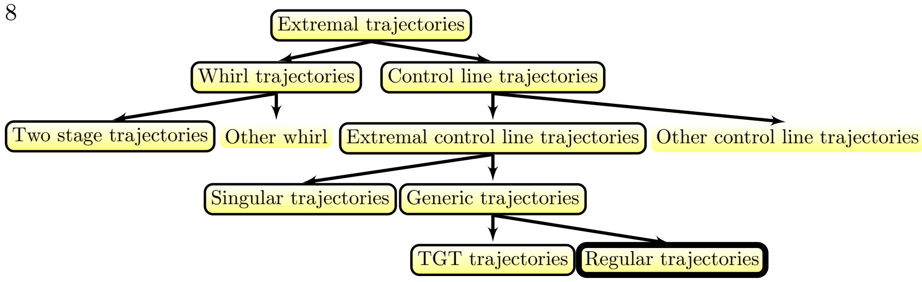

Fig. 4: Taxonomy of optimal trajectories. Each node corresponds to a type of optimal trajectories; each leaf node without border is not necessary for optimality. All leaf nodes with single border can be solved analytically. For the leaf node with double border, regular trajectories, we provide a search algorithm that can find a trajectory arbitrarily close to optimal trajectories.

For a fixed first control u_s, a fixed last control u_f, and a given Hamiltonian value, H, there exist at most two control lines; the mapping from the Hamiltonian values to the control lines can be represented by two continuous functions. Furthermore, for a control line $L = (k_1, k_2, k_3)$, $k_1^2 + k_2^2 = 1$, there exists a transformation T_L from the world frame to the *control line frame*. For a configuration q in the world frame, we use $q^L = T_L(q)$ to denote its representation in the control line frame whenever the control line L is clear from the text.

For a configuration q^L representing the configuration of the rigid body with respect to the control line, and a candidate control u, we would like to determine the duration that u may be applied before switching to some other control u', while satisfying Equation 1. It can be shown that there are at most two candidate durations such that at the time of switch, both controls u and u' have the same Hamiltonian value H. The mapping from configurations in the control line frame to durations (possibly ∞) can therefore be represented by two continuous functions.

In order to describe which duration function is being considered, we use a *duration selector*, which is a binary number associated with each pair of controls for a discrete trajectory structure. We call a pair of $(\mathbf{u}, \mathbf{s})$ as a *tentative structure* if $\mathbf{u}$ is a sequence of controls with length n and $\mathbf{s}$ is a sequence of duration selectors with length $n-1$, where s_i is the duration selector for control pair (u_i, u_{i+1}). For a given configuration q^L, the duration for each control in a tentative structure is fully determined except for the last one.

If we knew the tentative structure for a trajectory, together with the Hamiltonian value H (chosen from a continuous range), then the configuration(s) of the control line(s) consistent with the necessary conditions could be computed exactly, based on techniques in [10]. In fact, the control line can be found using only the Hamiltonian, the identity of the first and last controls, and the initial and the goal configurations; this allows computation of of durations of each control (except the last).

Given a configuration of the rigid body with respect to the control line, applying a particular control u will give a trajectory along which the Hamiltonian value computed by Equation 1 is constant. Along this trajectory, the body may or may not reach other configurations such that some other control can be applied giving this same Hamiltonian value. It can be shown that for a particular value H, there is some set of feasible control switches. Let $Q(H)$ be the set of all possible pairs of controls (u, u') such that there is a feasible switch, for a given H. Then,

Theorem 3. *There exists a finite set of critical values of $\mathcal{R}$ that partition the Hamiltonian values into a finite set of open intervals, such that for each interval D, if two Hamiltonian values H and H' are in D, then $Q(H) = Q(H')$. The set of critical values of the Hamiltonian values can be computed by analyzing the control set U.*

3.2 Reduction to a Lipschitzian Optimization Problem

It is easy to see that if there is a finite-time trajectory (found by any simple planner) between a pair of configurations, there exists a computable bound, B, for the number of actions in any optimal trajectories between those configurations in the costly-switch model.

Together with Theorem 3, the bound B can be used to show that there are finitely many discrete trajectory structures that must be considered for optimality. The basic approach enumerates all candidate starting and final controls (u_s, u_f) for a trajectory. Given each (u_s, u_f), we pose a Lipschitzian optimization problem to solve for H values with time and position error at most ϵ, for any desired $\epsilon > 0$. Then, we pick the best trajectory among all approximately optimal trajectories.

The problem of finding optimal regular trajectories with the first control u_s and the last control u_f has two parts: one is to determine a optimal tentative structure and another is to determine a Hamiltonian value H that approximately minimizes error and time. We first show how to find optimal trajectories for a fixed H value and then show how to determine H.

3.3 Finding Optimal Trajectories for a fixed *H*

Let D be an open interval of the partition of the Hamiltonian values containing H. Let $G_{u_s,u_f}(D)$ be the set of tentative structures in $G(D)$ with first control u_s and final control u_f. When u_s, u_f, and H are fixed, there are at most two control lines. For a fixed control line, for any tentative structure in $G_{u_s,u_f}(D)$, the duration of each control is fully determined except for the last control, and the duration of the last control will be determined by the goal configuration. Thus, finding optimal trajectories for the fixed u_s, u_f, and H is equivalent to finding the optimal tentative structures in $G_{u_s,u_f}(D)$. Although $G_{u_s,u_f}(D)$ is a finite set, the size of $G_{u_s,u_f}(D)$ may be large and hence we cannot enumerate all tentative structures in $G_{u_s,u_f}(D)$ to find optimal trajectories. Our approach is to use A* search guided by the distance function described below.

Distance Function. Let q_s^L and q_f^L be the initial and goal configurations, and $g = (\mathbf{u}, \mathbf{s})$ be a tentative structure. For any trajectory that reaches q_f^L with the last two controls of u_{n-1} and u_n, there are at most two possible configurations of switching control from u_{n-1} to u_n, and these two possible configurations only differ in x-coordinates in the control line frame. Let $q^L_{z_1}$ and $q^L_{z_2}$ be these two configurations. Let q^L_{n-1} be the configuration at which g will switch control from u_{n-1} to u_n. Note that q^L_{n-1} also only differ from $q^L_{z_1}$ and $q^L_{z_2}$ in x-coordinate. We define the *distance* between g and the goal, $d(L, g)$, to be the minimum difference from q^L_{n-1} to $q^L_{z_1}$ and $q^L_{z_2}$ in x-coordinate.

We still need to determine the duration of the last control in order to compute the cost. Let q^L_n be the final configuration of the trajectory. We require that $q^L_{n,y} = q^L_{f,y}$ and $q^L_{n,\theta} = q^L_{f,\theta}$ and the duration of the last control is determined by this restriction. We define the *cost* of g, $c(L, g)$ to be the sum of durations and switching costs.

Determining the Hamiltonian, H. In order to determine the best Hamiltonian value H, we first compute the partition of the Hamiltonian values according to Theorem 3. For each open interval, D, we determine the best Hamiltonian value $H \in D$. Then we pick the optimal Hamiltonian value among all best Hamiltonian values for each D.

Let D be an open interval of the partition of the Hamiltonian values. Since the first control, u_s, and the last control, u_f, are fixed, there are two mappings from H to the control lines; let $L(H)$ be one of the mappings from H to the control lines. The problem of finding the best Hamiltonian value $H \in D$ is as follows:

$$\begin{aligned} \min\ & c(L(H), g) \\ & d(L(H), g) = 0 \\ & g \in G_{u_s,u_f}(D), H \in D. \end{aligned} \tag{2}$$

We will use Lipschitzian optimization techniques to solve Problem 2. Here, we briefly introduce Lipschitzian optimization.

Lipschitzian Optimization. The goal of global optimization is to find optimal solutions of constrained optimization problem even for non-linear, non-continuous problems.

A function $f : \mathcal{R} \rightarrow \mathcal{R}$ is called *Lipschitz continuous* if there exists a constant $L \geq 0$, such that for all pairs x, y in the domain we have $|f(x) - f(y)| \leq L|x - y|$, where L is called the *Lipschitz constant.* Given a Lipschitz continuous function $f(x)$, the problem of finding the global minimum $\min_x f(x)$ is called a *Lipschitzian optimization* problem. For Lipschitzian optimization problems, there exists efficient algorithms to find globally (approximately) optimal solutions with arbitrarily small error in a finite time [17].

The Lipschitzian optimization algorithm we used for solving Problem 2 is Piyavskii's algorithm [18]. The idea of Piyavskii's algorithm is to iteratively subdivide a domain D into several intervals. For each interval, Piyavskii's algorithm determines the lower bound of the objective function based on Lipschitz constant, and decides whether to further subdivide this interval or disregard this interval based on the lower bound information. For any error bound $\epsilon > 0$, Piyavskii's algorithm guarantees to find a solution with additive an error at most ϵ within a finite number of iterations.

We will show that Problem 2 can be modeled as Lipschitzian optimization problem in the next section.

4 Lipschitz Continuity

Fix the first control to be u_s and the last control to be u_f. Let D be an open interval of the partition of the Hamiltonian values. Remember that the problem is $\min_{g \in G_{u_s,u_f}(D), H \in D} c(L(H), g)$, subject to $d(L(H), g) = 0$, where $G_{u_s,u_f}(D)$ is a finite set of tentative structures.

Since Lipschitz continuity is closed under the minimum operation, it suffices to prove that for any $g \in G_{u_s,u_f}(D)$, both the distance function $d(L(H), g)$ and cost function $c(L(H), g)$ are Lipschitz continuous with respect to $H \in D$ for any $g \in G_{u_s,u_f}(D)$.

4.1 Lipschitz Continuity of $d(L(H), g)$ and $c(L(H), g)$

Let $g = (\mathbf{u}, \mathbf{s}) \in G_{u_s,u_f}(D)$ be a tentative structure, where the length of $\mathbf{u}$ is n. We first consider the cost function $c(L(H), g)$, which depends on the durations of each control and switch cost. Since the number of controls is n, the switch cost will not change and hence we focus on durations. Let $t_i(H)$ be the duration for the i-th control u_i with respect to H. Since $c(L(H), g)$ is just a summation of all t_i, we only need to argue that each $t_i(H)$ is Lipschitz continuous.

Next, we consider the distance function $d(L(H), g)$. For control u_i and its corresponding sub-trajectory, we use d_i to denote the length of the sub-trajectory projection onto the control line. The distance function $d(L(H), g)$ can be rewritten as $|q^L_{s,x} + \sum_{i=1}^{n} d_i - q^L_{f,x}|$. It suffices to show that each $d_i(H)$ and the mapping T_L is Lipschitz continuous.

Durations $t_i(H)$ and projections $d_i(H)$, $1 < i < n$ are easier to analyze, since they depend on H directly. However, durations $t_1(H)$ and $t_n(H)$ depend on H, initial configuration q^L_s, and final configuration q^L_f in the control line frame. Hence, $t_1(H)$ and $t_n(H)$ depend on H not only directly but also indirectly through q^L_s and q^L_f. Similarly, $d_1(H)$ and $d_n(H)$ also depend on H directly and indirectly. The analysis of $t_1(H)$, $t_n(H)$, $d_1(H)$, and $d_n(H)$ should be separated from the analysis of $t_i(H)$ and $d_i(H)$, $1 < i < n$. Due to space limitations, we only show the analysis of $t_i(H)$ and $d_i(H)$, $1 < i < n$.

Analysis of $t_i(H)$ and $d_i(H)$, $1 < i < n$.

Theorem 4. *Let $D = (a, b)$ be an open interval of the partition of the Hamiltonian values. Let $g = (\mathbf{u}, \mathbf{s}) \in G_{u_s,u_f}(D)$ be a tentative structure with n actions. Let $t_i(H)$ be the duration for the u_i and $d_i(H)$ be the length of projection of the sub-trajectory corresponding to u_i onto the control line. For any δ,*

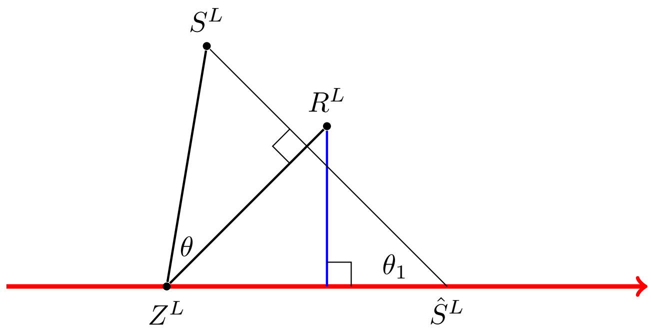

Fig. 5: Illustration of proof of Theorem 4: u_i is a translation.

$0 < \delta < (b-a)/2$, both functions $t_i(H)$ and $d_i(H)$ are Lipschitz continuous with respect to $H \in (a+\delta, b-\delta)$ for all $1 < i < n$.

Proof. The duration $t_i(H)$ and length $d_i(H)$ are fully determined by u_{i-1}, u_i, u_{i+1}, and H. Let q_i^L be the configuration in the control line frame at which the trajectory switches control from u_{i-1} to u_i. Let q_{i+1}^L be the configuration in the control line frame at which the trajectory switches control from u_i to u_{i+1}. Here, we use a result from [10] that there exists a point $p_i = p(u_{i-1}, u_i)$ rigidly attached to the robot, such that p_i will lie on the control line when the robot is at q_i^L. Similarly, when the robot is at q_{i+1}^L and switches from u_i to u_{i+1}, there exists a point $p_{i+1} = p(u_i, u_{i+1})$ attached to the robot such that p_{i+1} is on the control line.

We introduce some notation for the remainder of the proof. Let $Z^L = (Z_x^L, 0)$ be the location of p_i attached to the robot at q_i^L, which is on the control line. Let $S^L = (S_x^L, S_y^L)$ be the location of p_{i+1} attached to the robot at q_i^L. Let $\hat{S}^L = (\hat{S}_x^L, 0)$ be the location of p_{i+1} attached to the robot at q_{i+1}^L. By considering the position of S^L we can determine the t_i and d_i.

Depending on whether u_i is a translation or not, there are two cases:

u_i is a Translation. Let v_i be the velocity of u_i. By Theorem 1, the magnitude of the projection of the velocity onto the control line is H. Consequently, the magnitude of velocity in the y-coordinate in the control line frame is $v_y^L = \sqrt{v_i^2 - H^2}$. The duration of t_i can be computed as S_y^L / v_y^L. Consequently, the length of the projection of the trajectory onto the control line, $d_i(H)$, can be computed as $t_i H$. Hence, it suffices to prove t_i is Lipschitz continuous.

The control u_{i-1} must be a rotation, since if u_{i-1} is a translation, then u_i and u_{i-1} have the same Hamiltonian value along the sub-trajectory corresponding to u_i and the trajectory is a singular trajectory. Let $R^L = (R_x^L, R_y^L)$ be the location of the rotation center of control u_{i-1}. Let l_{SZ} be the distance between S^L and Z^L. Let θ be the angle rotating from vector $Z^L S^L$ to vector $Z^L R^L$ counterclockwise. Since the the mutual distance among S^L, R^L and Z^L is independent from H, l_{RZ} and θ are independent from H.

Let θ_1 be the angle between segment $S^L \hat{S}^L$ and the control line; the value of θ_1 is $\operatorname{acos}(H/v_i)$. Furthermore, it can be shown that the line $Z^L R^L$ is perpendicular to the line $S^L \hat{S}^L$ [10]. By geometric reasoning, S_y^L can be computed as

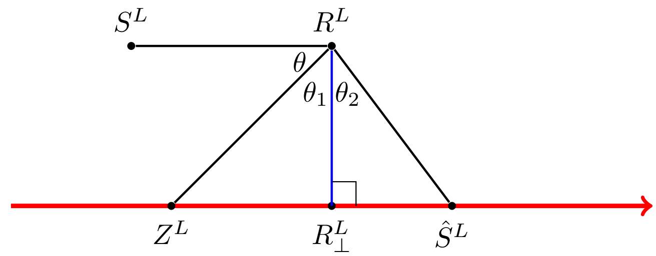

Fig. 6: Illustration of proof of Theorem 4: u_i is a rotation, $\omega_{i-1} > \omega_i$, and $\omega_{i+1} > \omega_i$.

$l_{SZ}\cos(\theta - \text{acos}(H/v_i)) = (l_{SZ}/v_i)(H\cos\theta + \sqrt{v_i^2 - H^2}\sin\theta)$. Hence,

$$t_i = \frac{S^L_y}{v^L_y} = \frac{l_{SZ}}{v_i}\left(\sin\theta + \frac{H\cos\theta}{\sqrt{v_i^2 - H^2}}\right).$$

We know a differentiable function is Lipschitz continuous if, and only if, this function has a bounded first derivative.

$$\frac{\partial t_i}{\partial H} = \left(\frac{v_i l_{SZ}\cos\theta}{(v_i^2 - H^2)^{1.5}}\right)$$

When $H \in (a+\delta, b-\delta)$ and $H < v_i$, the derivative of $t_i(H)$ and $d_i(H)$ are bounded.

u_i is a Rotation. Let $R^L = (R^L_x, R^L_y)$ be the location of the rotation center of control u_i and let $R^L_\perp = (R^L_x, 0)$ be the projection of R^L on the control line. We want to compute the angle, φ_0, between the control line to the vector R^LS^L, and the angle φ_1, between the control line to the vector $R^L\hat{S}^L$; these angles are measured in counterclockwise direction. The duration $t_i(H)$ can be computed as $(\varphi_1 - \varphi_0)/\omega_i$, where the subtraction wrapping around 2π and the result has the same sign as ω_i. Let r be the distance between the reference point of the robot and R^L when robot is at q^L_i. The projection of the trajectory on the control line, $d_i(H)$, can be computed as $r(\cos\varphi_1 - \cos\varphi_0)$. Thus, it suffices to show that φ_0 and φ_1 are Lipschitz continuous with respect to H.

Let l_{RZ} be the distance between R^L and Z^L and let l_{RS} be the distance between R^L and S^L. Let θ be the angle rotating from vector R^LZ^L to R^LS^L counterclockwise. Note that θ, l_{RZ}, and l_{RS} are independent from H.

Let θ_1 be the angle between the segment R^LZ^L and $R^LR^L_\perp$, which equals $\text{acos}(H/(l_{RZ}\omega_i))$. Let θ_2 be the angle between the segment $R^L\hat{S}^L$ and $R^LR^L_\perp$, which equals $\text{acos}(H/(l_{RS}\omega_i))$. Let ω_{i-1} and ω_{i+1} be the angular velocity of u_{i-1} and u_{i+1} respectively. Based on θ_1 and θ_2, we can compute φ_0 and φ_1 as follows:

		φ_0		φ_1
$\omega_i > 0$	$Z^L_x \leq R^L_x$	$3\pi/2 - \theta_1 + \theta$	$\hat{S}^L_x \geq R^L_x$	$3\pi/2 + \theta_2$
$\omega_i > 0$	$Z^L_x > R^L_x$	$3\pi/2 + \theta_1 + \theta$	$\hat{S}^L_x < R^L_x$	$3\pi/2 - \theta_2$
$\omega_i < 0$	$Z^L_x > R^L_x$	$\pi/2 - \theta_1 + \theta$	$\hat{S}^L_x < R^L_x$	$\pi/2 + \theta_2$
$\omega_i < 0$	$Z^L_x \leq R^L_x$	$\pi/2 + \theta_1 + \theta$	$\hat{S}^L_x \geq R^L_x$	$\pi/2 - \theta_2$

Thus, we have

$$\left|\frac{\partial \varphi_0}{\partial H}\right| \leq ((l_{RZ}\omega_i)^2 - H^2)^{-0.5} \text{ and } \left|\frac{\partial \varphi_1}{\partial H}\right| \leq ((l_{RS}\omega_i)^2 - H^2)^{-0.5}.$$

Consequently,

$$\left|\frac{\partial t_i}{\partial H}\right| \leq \frac{((l_{RZ}\omega_i)^2 - H^2)^{-0.5} + ((l_{RS}\omega_i)^2 - H^2)^{-0.5}}{|\omega_i|}.$$

$$\begin{aligned}\left|\frac{\partial d_i}{\partial H}\right| \leq & \frac{r}{l_{RZ}|\omega_i|}\left(|\sin\theta_1| + \left|\frac{H\cos\theta_1}{\sqrt{(l_{RZ}\omega_i)^2 - H^2}}\right|\right) \\ & + \frac{r}{l_{RS}|\omega_i|}\left(|\sin\theta_2| + \left|\frac{H\cos\theta_2}{\sqrt{(l_{RS}\omega_i)^2 - H^2}}\right|\right)\end{aligned}$$

When $H \in (a+\delta, b-\delta)$, H is smaller than $|l_{RZ}\omega_i|$ and $|l_{RS}\omega_i|$, and the derivatives of $t_i(H)$ and $d_i(H)$ are bounded.

5 Implemetation

We implemented the algorithm described in C++. Our testing environment is a desktop system with an Intel Xeon W3550 3.07 GHz CPU.

In the costly-switch model, we used three test cases. First, we used the bench mover's problem proposed in [13] as one test case. We compare our program's result with the result of analytical solver. Except for some cases in which the Hamiltonian value is close to the upper bound (for which numerical instability becomes a problem), our results coincide with the result of exact solver.

We used the refrigerator-movers problem as the second test case; one approximately optimal trajectory is shown in Fig. 1. Third, we used omni-directional vehicle as a test case; one approximately optimal trajectory is show in Fig. 2b.

In the cost-free-switch model, we compared our program with the exact solver proposed in [29], which determines optimal trajectories for the omni-directional vehicle analytically. Although our program is a general solver that can solve all problems of finding optimal trajectory for rigid bodies in the plane, our program is only about ten times slower; one approximately optimal trajectory is show in Fig. 2a.

6 Conclusion and Future Work

By adding a cost for switching between controls, we ensure existence of solutions for optimal control problems that do not involve chattering. By applying Blatt's Indifference Principle and Lipschitzian optimization approach, we can find approximately optimal trajectories and the error can be forced to be arbitrarily small.

The most exciting area of future work is to explore the application of BIP to systems other than rigid bodies in the plane. It is particularly interesting that optimal trajectories with costly switches exist even in the presence of obstacles.

There are at least two challenges in applying a BIP-based approach to finding optimal trajectories. The first challenge is that the potential number of optimal trajectory structures can be huge in the costly-switch model. In the costly-switch model, an algorithm might potentially need to explore a number of structures that is exponential in the number of controls in order to find solutions. For example, in order to find approximately optimal trajectories for omni-directional vehicle, whose control set contains fourteen controls, it takes about an hour to find an approximately optimal trajectory for an initial configuration and goal configuration.

The second challenge is numerical instability. When the Hamiltonian value of optimal trajectories is close to the boundary of the open interval in the partition of the Hamiltonian values, the Lipschitz constant for the duration function may be very large. Consequently, the numerical error in the computation also increases significantly and is inherently unstable. This is an issue for our solver in the costly-switch model and in the cost-free-switch model as well.

References

1. Pankaj K. Agarwal, Therese Biedl, Sylvain Lazard, Steve Robbins, Subhash Suri, and Sue Whitesides. Curvature-constrained shortest paths in a convex polygon. *SIAM Journal on Computing*, 31(6):1814–1851, June 2002.
2. Jérîme Barraquand and Jean-Claude Latombe. Robot motion planning: a distributed representation approach. *International Journal of Robotics Research*, 10(6):628–649, Dec 1991.
3. John M. Blatt. Optimal control with a cost of switching control. *Journal of the Australian Mathematical Society*, 19:316–332, 6 1976.
4. H. R. Chitsaz. *Geodesic problems for mobile robots*. PhD thesis, University of Illinois at Urbana-Champaign, 2008.
5. Hamid Reza Chitsaz, Steven M. LaValle, Devin J. Balkcom, and Matthew T. Mason. Minimum wheel-rotation paths for differential-drive mobile robots. *International Journal of Robotics Research*, 28(1):66–80, 2009.
6. M. Chyba and T. Haberkorn. Autonomous underwater vehicles: Singular extremals and chattering. In F. Ceragioli, A. Dontchev, H. Furuta, K. Marti, and L. Pandolfi, editors, *Systems, Control, Modeling and Optimization*, volume 202 of *IFIP International Federation for Information Processing*, pages 103–113. Springer US, 2006.
7. E. J. Cockayne and G. W. C. Hall. Plane motion of a particle subject to curvature constraints. *SIAM Journal on Control*, 13(1):197–220, 1975.
8. Guy Desaulniers. On shortest paths for a car-like robot maneuvering around obstacles. *Robotics and Autonomous Systems*, 17(3):139 – 148, 1996.
9. L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):pp. 497–516, 1957.
10. Andrei Furtuna. *Minimum Time Kinematic Trajectories for Self-Propelled Rigid Bodies in the Unobstructed Plane*. PhD thesis, Dartmouth College, June 2011.

11. K. Kibalczyc and S. Walczak. Necessary optimality conditions for a problem with costs of rapid variation of control. *Journal of the Australian Mathematical Society*, 26:45–55, 7 1984.
12. Ryan Loxton, Qun Lin, and Kok Lay Teo. Minimizing control variation in nonlinear optimal control. *Automatica*, 49(9):2652 – 2664, September 2013.
13. Yu-Han Lyu, Andrei Furtuna, Weifu Wang, and Devin Balkcom. The bench mover's problem: minimum-time trajectories, with cost for switching between controls. In *IEEE International Conference on Robotics and Automation*, 2014.
14. Matthew T. Mason. *Mechanics of Robotic Manipulation*. MIT Press, Cambridge, MA, USA, 2001.
15. Joanna Matula. On an extremum problem. *Journal of the Australian Mathematical Society*, 28:376–392, 1 1987.
16. E. S. Noussair. On the existence of piecewise continuous optimal controls. *Journal of the Australian Mathematical Society*, 20:31–37, 6 1977.
17. János D. Pintér. *Global Optimization in Action: Continuous and Lipschitz Optimization: Algorithms, Implementations and Applications (Nonconvex Optimization and Its Applications*. Springer, 2nd edition, December 2010.
18. S. A. Piyavskii. An algorithm for finding the absolute minimum of a function. *USSR Computational Mathematics and Mathematical Physics*, 12:13–24, 1967.
19. L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *Mathematical Theory of Optimal Processes*. John Wiley & Sons Inc, January 1962.
20. J. A. Reeds and L. A. Shepp. Optimal paths for a car that goes both forwards and backwards. *Pacific Journal of Mathematics*, 145(2):367–393, 1990.
21. David B. Reister and Francois G. Pin. Time-optimal trajectories for mobile robots with two independently driven wheels. *International Journal of Robotics Research*, 13(1):38–54, February 1994.
22. Marc Renaud and Jean-Yves Fourquet. Minimum time motion of a mobile robot with two independent, acceleration-driven wheels. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 2608–2613 vol.3, Apr 1997.
23. J.J. Slotine and S.S. Sastry. Tracking control of non-linear systems using sliding surfaces with application to robot manipulators. In *American Control Conference, 1983*, pages 132–135, June 1983.
24. P. Souères and J.-D. Boissonnat. Optimal trajectories for nonholonomic mobile robots. In J.-P. Laumond, editor, *Robot Motion Planning and Control*, pages 93–170. Springer, 1998.
25. David E. Stewart. A numerical algorithm for optimal control problems with switching costs. *Journal of the Australian Mathematical Society*, 34:212–228, 10 1992.
26. Héctor J. Sussmann and Guoqing Tang. Shortest paths for the Reeds-Shepp car: A worked out example of the use of geometric techniques in nonlinear optimal control. Technical report, Department of Mathematics, Rutgers University, 1991.
27. K.L. Teo and L.S. Jennings. Optimal control with a cost on changing control. *Journal of Optimization Theory and Applications*, 68(2):335–357, February 1991.
28. Weifu Wang and Devin Balkcom. Sampling extremal trajectories for planar rigid bodies. In Emilio Frazzoli, Tomas Lozano-Perez, Nicholas Roy, and Daniela Rus, editors, *Algorithmic Foundations of Robotics X*, volume 86 of *Springer Tracts in Advanced Robotics*, pages 331–347. Springer Berlin Heidelberg, 2013.
29. Weifu Wang and Devin J. Balkcom. Analytical time-optimal trajectories for an omni-directional vehicle. In *IEEE International Conference on Robotics and Automation*, pages 4519–4524, May 2012.
30. Changjun Yu, Kok Lay Teo, and Teng Tiow Tay. Optimal control with a cost of changing control. In *Australian Control Conference*, pages 20–25, Nov 2013.