

Generalizing Dubins curves: minimum-time sequences of body-fixed rotations and translations in the plane

Andrei A. Furtuna, Devin J. Balkcom

March 17, 2010

Abstract

This paper presents the minimum-time sequences of rotations and translations that connect two configurations of a rigid body in the plane. The configuration of the body is its position and orientation, given by (x, y, θ) coordinates, and the rotations and translations are velocities $(\dot{x}, \dot{y}, \dot{\theta})$ that are constant in the frame of the robot. There are no obstacles in the plane. We completely describes the structure of the fastest trajectories, and present a polynomial-time algorithm that, given a set of rotation and translation controls, enumerates a finite set of structures of optimal trajectories. These trajectories are a generalization of the well-known Dubins and Reeds-Shepp curves, which describe the shortest paths for steered cars in the plane.

1 Introduction

The problem of moving a rigid body between two configurations efficiently is of particular interest in robotics because the simplest model of a mobile robot or an object in a robotic manipulation task is often a rigid body. Path planning, controller design, and robot design may all benefit from precise knowledge of optimal trajectories for a set of permitted controls.

The minimum-time trajectories depend on the available controls, and on the goal configuration. For example, the fastest way to move a kinematic model of a car with bounded steering angle a short distance sideways might be to execute a particular parallel-parking type motion. To move a wheelchair to a distant location, the fastest trajectory might be to turn to face the goal, and drive to the goal.

The problem Dubins solved in 1957 was to find the shortest curve in the plane connecting two points with prescribed tangent vectors at those points, with the constraint that the average curvature over any section of the curve be less than some maximum constant $1/r$. Dubins showed that these shortest curves, called *R-geodesics*, are composed only of circles and straight lines, connected together in certain ways; figure 1 illustrates a few examples.

Dubins curves are also the fastest trajectories for a simple model of a steered car that moves with unit velocity. In the language of control theory, Dubins' problem can be restated as follows. Let the state of the car in \mathbf{SE}^2 be $q = (x, y, \theta)$. Given initial configuration q_0 and desired final configuration q_1 , find a vector-valued control function $u : [0, t_1] \rightarrow U_D$ such that t_1 is minimized and

$$q_1 = q_0 + \int_0^{t_1} R(\theta)u(t) dt \quad (1)$$

where

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (2)$$

and

$$U_D = \{1\} \times \{0\} \times [-1, 1]. \quad (3)$$

The components of the control vector, u_1 , u_2 , and u_3 correspond to the forward velocity of the car, the sideways velocity, and the angular velocity respectively.

For any rigid body in the plane,

$$\dot{q} = R(\theta)\dot{\hat{q}}, \quad (4)$$

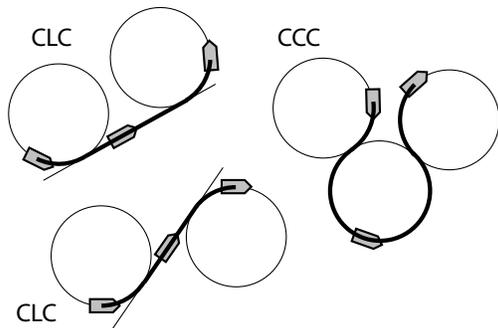


Figure 1: Example Dubins curves.

where $\dot{q} = (\dot{x}, \dot{y}, \dot{\theta})$ is the generalized velocity of the body in its own frame of reference. The generalization we consider is to let u to be \dot{q} . We consider the cases where U is either a finite set, or a closed polyhedron. Although some of the systems we consider are omni-directional, the velocity control constraints are non-holonomic, and cannot be expressed as constraints on the configuration of the system.

1.1 Applications and example systems

Dubins-like curves arise in many contexts. The car studied by Reeds and Shepp can go forwards as well as backwards. In our model, this system is obtained with

$$U_{RS} = \{-1, 1\} \times \{0\} \times [-1, 1]. \quad (5)$$

Dubins and Reeds-Shepp curves arise in many robotics problems. For example, the optimality of Reeds-Shepp curves is the motivation for Barraquand and Latombe's choice of discrete controls for their motion planner for a non-holonomic cart [5]. Recent work by Alterovitz *et al* [2] models the motion of a surgical needle through the body as a Dubins curve. LaValle's [16] work on rapidly-exploring random trees relies on a metric (or pseudometric) between configurations in the free space, and Dubins and Reeds-Shepp curves can be used to generate a metric for steered cars.

We hope that this paper provides a new look at Dubins and Reeds-Shepp curves, but there are also many other systems that can be modeled as rigid bodies in the plane with bounds on the generalized velocities.

Wheeled differential-drive vehicles have two powered wheels. Assume that there are bounds on the speeds of the wheels. The optimal trajectories are known for a simple kinematic rigid-body model [4], and only contain the discrete controls spin in place to the left or right, forwards, and reverse. In the original paper, the controls were chosen to be the speeds of the two driving wheels, which each are constrained to the set $[-1, 1]$. Since the forward differential kinematics equations that map between wheel speeds and generalized velocities of the vehicle in its own frame are linear in the wheel speeds, we find that the generalized velocities of the vehicle fall in a polygon.

Omnidirectional vehicles like Acroname's three-wheeled Palm-pilot robot kit or Segway's new four-omni-wheeled RMP robotic platform directly drive wheels that have rollers allowing sideways slippage. Although these vehicles are omnidirectional, some directions are faster than others, and driving these vehicles efficiently may require a more complex strategy than simply driving directly to the goal. The optimal trajectories are already known for a simple three-wheeled symmetric design [3], but not for four-wheeled or asymmetric designs. Bounding the wheel speeds again constrains the generalized velocity of the rigid body to fall in a polyhedron, since the forward kinematics are again linear in the wheel speeds.

Stable pushing is a manipulation strategy where a pushing device pushes a polygonal rigid body in the plane along one of the flat edges. For each edge, there is a polygon of rotation centers around which the pusher can be rotated without slip occurring between the pusher and the rigid body, as shown by Lynch and Mason [17].

Microrobots may potentially be modeled as rigid bodies, with a discrete set of control inputs attached to the robot. For example, McGray *et al.*'s microrobot [11] is essentially driven as a Dubins car that can only turn to the right.

The time-optimal trajectories for simple models of these systems are all special cases of the curves we will examine.

1.2 Main results

Our first main result is the following necessary condition on time-optimal trajectories, which will be proven in section 2.

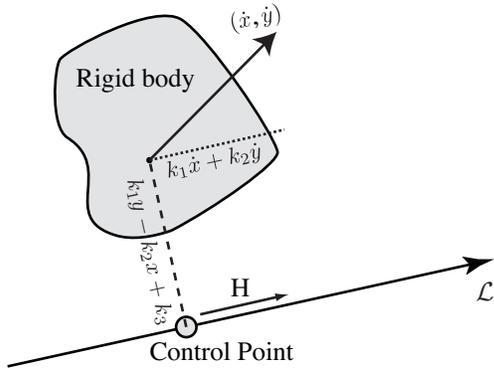


Figure 2: A rigid body relative to some control line. Trajectories that satisfy the maximum principle maximize the speed of the control point along the line. For the purpose of computing this speed, the control point is considered to be instantaneously attached to the rigid body. The constants k_1 , k_2 , and k_3 describe the position of the control line.

Theorem 1 Consider a rigid body in the unobstructed plane, with configuration $q = (x, y, \theta)$, controls u in some set of constant vectors, and system equations

$$\dot{q} = R(\theta)u, \quad (6)$$

For any time optimal trajectory of the system, there exist constants k_1 , k_2 , and k_3 , not all 0, such that at every time, the controls maximize the Hamiltonian equation:

$$H = k_1\dot{x} + k_2\dot{y} + \dot{\theta}(k_1y - k_2x + k_3). \quad (7)$$

Furthermore, the Hamiltonian is constant and positive over the trajectory.

Given a start and goal configuration, if we knew the constants k_1 , k_2 , and k_3 , we would know almost everything about the trajectory connecting the two configurations. Using the initial configuration $q_0 = (x_0, y_0, \theta_0)$ and the constants, we could write out the equation for the Hamiltonian, and from the available controls, choose $(\dot{x}, \dot{y}, \dot{\theta})$ to maximize the Hamiltonian. The body would follow this trajectory until some other controls became maximizing, or until the goal was reached.

The difficulty is that given a start and goal configuration, it is often difficult to determine the values of these constants. Therefore, we attack an easier problem – we

vary the constants, and study the possible structures of the optimal trajectories.

We will see that in most cases, there is a geometric interpretation of the Hamiltonian and the constants, shown in figure 2. There is a line in the plane, the *control line*, such that the rotations and translations maximize the speed of points where the control line intersects the extended rigid body, in the direction of the line. (By extended rigid body, we mean points that need not be actually part of the body, but are carried along with the body, as though there were an invisible glass plate attached to the body.) The location of the line depends on the start and goal configurations, but typically is not the line directly connecting the x and y locations of the start and goal. The control line essentially describes the trade-off between rotation and translation: bodies that are far from the control line will tend to ‘choose’ rotation to maximize the speed of the point, while bodies that are close to the line will tend to choose translation.

We will show that there are four kinds of trajectory, each with their own properties:

1. **Generic** trajectories are trajectories for which the value of the constants and the initial configuration completely identify the trajectory.
2. **Tacking** trajectories. If there are multiple translation controls at the same heading of the body, trajectories may freely switch between translation controls at that heading, due to the commutativity of translation.
3. **Tangent** trajectories. In these trajectories, all translation segments are parallel to the control line, or there exists a pair of sequential rotation centers that describe a line perpendicular to the control line.
4. **Constant-angular-velocity** trajectories. If $k_1 = k_2 = 0$, then theorem 1 immediately implies that $\dot{\theta}$ must be either maximized or minimized over the trajectory, and constant.

The Dubins CLC trajectories shown in figure 1 are examples of tangent trajectories, and the CCC trajectory is a generic trajectory. Tacking trajectories do not occur for the Dubins problem, and the only constant-angular-velocity trajectory is an arc of a circle with no control switches.

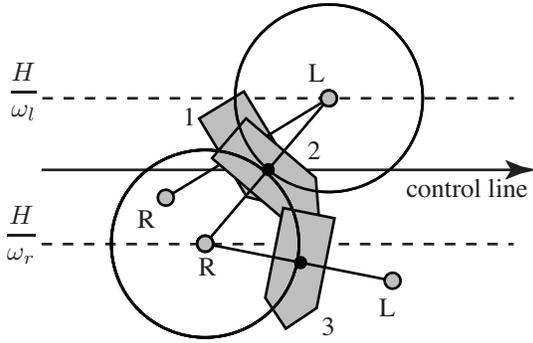


Figure 3: A Dubins CC trajectory and associated control line. The control switches from the left rotation center when the right rotation center reaches the line at signed distance H/ω_r from the control line.

Figure 3 gives an example of a trajectory that maximizes the Hamiltonian at each time, for some choice of constants. It turns out that the value of the Hamiltonian, which depends on the value of the constants, the maximizing control, and the state (x, y, θ) is the same no matter which reference point is picked. Any control that is not a pure translation has a rotation center, a point that does not move during the rigid-body transform. Therefore, choosing the rotation center as the reference point for computing the Hamiltonian is often very convenient, since (\dot{x}, \dot{y}) of this reference point is zero. In this case, it turns out that the value of the Hamiltonian is simply the signed distance of the rotation center from the control line, multiplied by the angular velocity of the body.

We can use this observation to describe the trajectory in figure 3 fairly simply. A Dubins car has two rotation controls, and two rotation centers; one to the right, and one to the left. The value of the Hamiltonian is some constant, H . At a distance of H/ω_l from the control line, the left rotation center may be ‘active’, and at a distance H/ω_r the right rotation center may be active. In configuration 1, the left rotation center is further from the control line, so the value of H computed if we apply the left rotation center is greater; this is the maximizing control, and the corresponding value of the Hamiltonian will remain constant over the trajectory. As the body spins about the left rotation center, the right rotation center moves. At any instant, we can compute the Hamiltonian for the right

rotation center, and for a while it is less than the Hamiltonian over the trajectory, so the left rotation center remains active. Eventually, the right rotation center reaches a line at distance H/ω_r from the control line (configuration 2 in figure 3). At this time, both controls maximize the Hamiltonian, since both rotation centers are on their associated lines, but whichever control is applied, at the next instant, only the right rotation center will be maximizing. We therefore say that at this control switch that the control associated with the right center is *sustainable*. For the remaining section of the trajectory shown, the right rotation center remains maximizing.

We will present an algorithm that takes the set of controls as input, and outputs all possible path structures, described as sequences of controls. This reduces the problem of finding the optimal trajectory between a particular start and goal to the problem of considering each possible path structure, finding the best path of each structure, and then the best path among all structures. We show that the number of generic trajectory structures is polynomial in the number of control inputs.

Along the way, there are many details to consider, but the basic idea is that shown in figure 3: trajectories essentially roll on rotation centers, each of which activates on its own line a particular distance from the control line. If there are translations, there may also be switches to translations at particular angles. We can see that the value of the Hamiltonian in some way characterizes the trajectory – it gives the distances from the line at which rotation centers switch, and the angle at which switches to translations occur.

There are also some degenerate cases, for which switching configurations occur at more than a finite set of times during the trajectory. These cases turn out to occur at particular values of the Hamiltonian, and given a start and goal configuration, the location of the control line can be computed for each degenerate case. For non-degenerate (i.e., generic) cases, the location of the control line is harder to determine from the start and goal, but we can enumerate a finite number of structures that the trajectory may take.

Section 2 reviews details of the Maximum Principle, and proves theorem 1, which gives the Hamiltonian for time-optimal trajectories of a rigid body in the plane with rotation and translation controls attached to the frame of the body. Sections 3 through section 6 treat the case where

at least one of the constants k_1 and k_2 in the Hamiltonian equation is nonzero, and section 7 treats the case where both are zero. Section 3 analyzes switches between controls on extremal trajectories, shows that controls are piecewise constant for almost all values of the Hamiltonian, and shows how the values of the Hamiltonian for degenerate cases may be computed. The section then shows how to compute the distance and orientation of the body from the control line during switches, for a given value of the Hamiltonian. Section 4 shows that there is a certain continuity in trajectory structures between critical values of Hamiltonian. Section 5 then gives techniques for generating generic extremal trajectory structures with values of the Hamiltonian other than at critical values, and section 6 explores degenerate trajectory structures resulting from critical values.

1.3 Related work

In 1957, Dubins characterized the shortest paths between two points in an obstacle-free plane, with the constraint that the path be tangent to given vectors at the start and goal, and with the average curvature over any interval along the path bounded by a positive real number R^{-1} [12]. These *R-geodesics* are composed of sequences of up to three segments, each of which is a line or an arc of a circle of radius R , and there are further constraints on how these segments may be connected together. *R-geodesics* are of particular interest in robotics because they describe the shortest paths for a simple model of a car with bounded steering angle and velocity.

There have been several extensions to Dubins' work. In 1990, Reeds and Shepp characterized the trajectories for a car that was also allowed to reverse [19]. Sussman and Tang described a general methodology for solving problems of this type [26], and Souères, Boissonnat, and Laumond [23, 24] discovered the mapping from pairs of configurations to optimal trajectories. The approaches developed enabled the discovery of the time-optimal trajectories for two other simple models of vehicles: the differential drive [4], and a particular three-wheeled robot that can drive sideways as well as forwards [3]. The present authors derived the structure of optimal trajectories for a general model of three-wheeled omni-directional vehicles [13]. This paper attempts to generalize the results for these various systems, drawing on techniques developed

by Sussman, Tang, Souères, Boissonnat, Laumond, and others.

There are also results that extend or generalize Dubins results in directions other than that taken in this paper. Chitsaz [6] explored the optimal trajectories for diff-drives with a particular distance function that is not equivalent to time. The optimal paths have also been explored for some examples of vehicles without wheels. Coombs and Lewis [8] consider a simplified model of a hovercraft, and Chyba and Haberkorn [7] consider underwater vehicles. The Maximum Principle is the starting point for these papers, as it is for the current paper.

The problem of finding optimal trajectories has also been studied for more complex models of vehicles, for which the state includes the configuration and generalized velocities, with bounded-acceleration controls. The optimal-control problem for dynamic vehicles appears to be very difficult – the differential equations describing the trajectories do not have recognizable analytical solutions, and in some cases, the optimal trajectories involve chattering, an infinite number of control switches in a finite time [25]. Papers by Reister and Pin [20], and Renaud and Fourquet [21] present numerical and partial geometric results for steered cars, and Kalmár-Nagy *et al.* [15] present algorithms for numerically computing approximately optimal trajectories for a bounded-acceleration model of the symmetric omnidirectional robot.

This paper considers the problem of finding trajectories in the obstacle-free plane without visibility constraints. Desaulniers [10] showed that in the presence of obstacles, shortest paths may not exist between certain configurations of steered cars. Optimal distance metrics may still give some useful information about how obstacles may interfere with desired motions. Vendittelli *et al.* [27] developed an algorithm to obtain the shortest non-holonomic distance from a robot to any point on an obstacle. Agarwal *et al.* [1] derive an algorithm to find the shortest curvature-constrained paths in a convex polygon. Optimal paths between pairs of points in configuration space may not exist in the presence of visibility constraints. Salaris *et al.* [22] give the optimal control words for a unicycle with a limited FOV camera. Hayet *et al.* [14] give necessary and sufficient conditions for the existence of paths between pairs of configurations for holonomic and differential drive systems.

2 The Maximum Principle

Pontryagin's Maximum Principle [18] places necessary conditions on the structure of optimal trajectories. Applying the Maximum Principle requires two steps. First, there is an *adjoint* vector, λ , that represents a privileged direction in the space of body velocities at each point on the trajectory. The adjoint vector is computed using Pontryagin's adjoint equation, a differential equation involving the body controls, the system equations, and the objective function (in this case, time) to be minimized. Along any trajectory of the body, the adjoint vector is non-zero, and a continuous function of both time and configuration. The expression for the adjoint also includes constants of integration that depend on the start and goal configurations of the system.

Second, the Maximum Principle states that along any optimal trajectory, the controls must be selected to maximize the Hamiltonian, which is the dot product of the adjoint and the generalized velocity of the body – in this case, $(\dot{x}, \dot{y}, \dot{\theta})$.

The controls for steered cars, differential drives, and omnidirectional vehicles are different. However, we can show that the adjoint is the same for all these vehicles; this is the key step in the proof of theorem 1.

Proof of Theorem 1. From the Maximum Principle, the adjoint equation is

$$\dot{\lambda} = -\frac{\partial}{\partial q} \langle \lambda, \dot{q}(q, u) \rangle \quad (8)$$

$$= -\begin{pmatrix} 0 \\ 0 \\ \lambda^T \left(\frac{\partial}{\partial \theta} R \right) f(u) \end{pmatrix}. \quad (9)$$

By direct integration, $\lambda_1 = k_1$ and $\lambda_2 = k_2$. Substitute these values back into the definition for $\dot{\lambda}_3$:

$$\dot{\lambda}_3 = k_1(s\dot{x} + c\dot{y}) - k_2(c\dot{x} - s\dot{y}), \quad (10)$$

where c and s are shorthand for $\cos \theta$ and $\sin \theta$. From equation 4,

$$\dot{x} = c\dot{\hat{x}} - s\dot{\hat{y}} \quad (11)$$

$$\dot{y} = s\dot{\hat{x}} + c\dot{\hat{y}}. \quad (12)$$

Substitute into equation 10,

$$\dot{\lambda}_3 = k_1\dot{y} - k_2\dot{x}, \quad (13)$$

and integrate:

$$\lambda_3 = k_1y - k_2x + k_3. \quad (14)$$

Application of the Maximum Principle completes the proof: the Hamiltonian to be maximized along time-optimal trajectories is

$$H = k_1\dot{x} + k_2\dot{y} + \dot{\theta}(k_1y - k_2x + k_3). \quad (15)$$

The generalized velocity of the body $(\dot{x}, \dot{y}, \dot{\theta})$ is a function of the state q and the current control u . Define

$$H_i(x, y, \theta) = H(u_i, x, y, \theta) \quad (16)$$

to be the Hamiltonian function associated with a control u_i . We call any trajectory satisfying the Maximum Principle an *extremal* trajectory. Along an extremal trajectory, the value of the Hamiltonian is $H_* = H(u_*, x, y, \theta)$, where u_* is some control that maximizes the Hamiltonian.

At most points along the trajectory, we expect a single control to maximize the Hamiltonian, and at these points, the control is fully determined by the Maximum Principle. At other points, multiple controls may maximize the Hamiltonian; at such points the trajectory may *switch* between controls, and there is (at least) a pair of controls u_i and u_j such that $H_i = H_j$.

2.1 The control line

We will deal with the case where $k_1 = k_2 = 0$ in section 7. For now, assume that at least one of k_1 or k_2 is nonzero. Without loss of generality, we may choose a positive scaling for the constants so that $k_1^2 + k_2^2 = 1$.

Define the *control line* to be a line with heading (k_1, k_2) , and signed distance k_3 from the origin. The first part of the Hamiltonian

$$k_1\dot{x} + k_2\dot{y} \quad (17)$$

is the component of the translational velocity of the rigid body along the vector (k_1, k_2) , and the term $-k_2x + k_1y + k_3$ is the distance from the reference point of the rigid body to the control line.

We now have a geometric interpretation of the Hamiltonian. Define the 'control line frame' to be a frame attached to the control line with x axis aligned with the control line, fixed anywhere along the line (see fig. 2). Then

in the control line frame, y is the distance of the rigid body from the control line, and θ is the angle the body frame makes with the control line. \dot{x} is the component of the body's velocity along the control line. In these coordinates, the Hamiltonian becomes

$$H = \dot{x} + y\dot{\theta}. \quad (18)$$

This expression holds for the signed distance y of a particular reference point from the control line. The next lemma will show that the value of H actually turns out to be independent of the choice of the reference point on the body.

Lemma 1 *Given a rigid body trajectory that obeys the Maximum Principle, the same value of the Hamiltonian will be obtained for any point of reference in the body frame.*

Proof: If the instantaneous motion is a translation, the result is immediate. Let the instantaneous motion be a rotation of center O and angular velocity ω . Let y_O be the distance between O and the control line. Let O' be O 's projection onto the control line and P be an arbitrary point in the frame of the vehicle. From the perspective of P , the Hamiltonian is

$$H_P = \dot{x}_P + y_P\omega = \|OP\|\omega \cos \angle POO' + y_P\omega \quad (19)$$

$$H_P = (\|OP\| \cos \angle POO' + y_P)\omega = y_O\omega = H_O \quad (20)$$

Therefore calculating the Hamiltonian at any point in the frame of the body is the same as calculating it at the center of rotation. ■

The fact that the Hamiltonian is independent of the choice of reference point suggests choosing a reference point where the expression for the Hamiltonian is simplified. Particularly, we can interpret the control law from the perspective of a point in the body frame that happens to be on the control line (see fig. 2). For such a point, $y = 0$ and the control law only requires to pick the control that maximizes this point's velocity \dot{x} along the control line.

Another interesting result is obtained by choosing the reference point, during a rotation control, to be the rotation center, where $\dot{x} = 0$. Therefore, given a value for H ,

we can compute the distance of the rotation center from the control line. A similar result can be obtained describing the angle a translation makes with the control line in terms of the value of H .

Corollary 1 *If the control corresponding to rotation center O and angular velocity ω is active at time t on an extremal trajectory of Hamiltonian value H , then at this time the signed distance from O to the control line is $y_O = \frac{H}{\omega}$.*

Corollary 2 *If a translation control of velocity v and forming angle α with the horizontal axis is active at time t on an extremal trajectory of Hamiltonian value H , then at this time $\cos(\alpha + \theta) = \frac{H}{v}$, where θ is the orientation of the body frame with respect to the control line.*

Figure 4 gives an example of how these two corollaries may be applied. The vehicle has four controls: rotation at angular velocity ω_l or ω_r about the center of the body, forwards translation, and reverse. In the figure, the body starts slightly below the distance H/ω_l from the control line, and forward translation maximizes the quantity $\dot{x} + y\omega$ over all controls. The body drives until the rotation center hits the line H/ω_r , spins to the right until reaching a critical angle, reverses until the line at H/ω_l , spins to the left, and then repeats the process.

Polar coordinates (v, α, ω) will be useful for the controls, since rotating the rigid body rotates the 'forward' direction, and all of the controls. That is, $\alpha = \text{atan2}(\hat{y}, \hat{x})$, the angle that the control makes with the horizontal axis in the frame of the body, $v = \sqrt{\hat{x}^2 + \hat{y}^2}$, the translational velocity of the body, and ω is the rotational velocity. In polar coordinates, the Hamiltonian becomes:

$$H(u, y, \theta) = v \cos(\theta + \alpha) + \omega y. \quad (21)$$

In these coordinates, the value of the Hamiltonian function corresponding to a control u_i is

$$H_i(y, \theta) = H(u_i, y, \theta), \quad (22)$$

along an extremal trajectory. The value of the Hamiltonian is $H_* = H(u_*, y, \theta)$, where u_* is some control that maximizes the Hamiltonian.

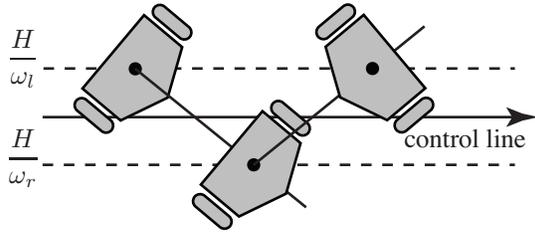


Figure 4: A differential-drive ‘forwards-right-reverse’ trajectory and associated control line.

3 Control switches on extremal trajectories

We have a necessary condition on the trajectory – the control must be chosen to maximize the speed of a point on the control line, attached to the body, along the control line. We don’t know where the line is for a given start and goal, but if we consider all possible placements of the rigid body relative to the line, we may enumerate the types of trajectories that occur: the extremal trajectories.

At some configurations along an extremal trajectory, multiple controls may maximize the Hamiltonian. These are the configurations where the control may switch. For example, assume that over some extremal trajectory, u_1 has been the maximizing control, but as the body follows that control, the configuration changes until u_2 also maximizes the Hamiltonian, and then at the next instant, only control u_2 maximizes the Hamiltonian. As long as the system is “well behaved”, it is possible to start from arbitrary configurations and generate the possible shapes of extremal trajectories. We will delineate and treat separately the special cases. The main result of this section is that “most” extremal trajectories have piecewise continuous controls and deterministic switches. The exceptions will be shown to only occur at *critical values* of the Hamiltonian, which we can calculate.

3.1 Chattering

Unfortunately, depending on the control inputs available, the class of problems we study do not always have solutions, or may have solutions that involve an infinite number of switches between controls.

Consider a rigid body in the plane that can translate in

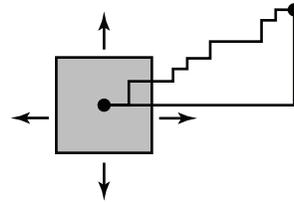


Figure 5: Two optimal trajectories for a translational platform.

any of the four directions aligned with the axes, north, south, east, and west, with speed one, as shown in figure 5. The fastest trajectory to move to a point to the east is unique and simple to describe: drive east, with no control switches. How about a trajectory to drive to the point $(2, 1)$? The minimum time required is 3 (the manhattan distance to the point from the origin), and one optimal trajectory is to drive east for time two, then north for time one. However, any other trajectory to the goal that uses only the controls ‘east’ and ‘north’ is also optimal. There could be an arbitrary number of switches between the two controls. For this system, optimal trajectories clearly exist; but it is easiest to describe the trajectories with a minimum number of switches, with the understanding that permutations of translation sections of the trajectory will also be optimal.

There are worse situations that can arise. Consider a refrigerator whose projection onto the plane is a square, with supporting legs at the corners of the square. Assume we can rotate the refrigerator about any of the legs with angular velocity 1 or -1 . What is the fastest way to move the refrigerator along the positive x axis? It turns out that the solution may be to ‘chatter’, or switch infinitely quickly, between two rotation controls, approximating a straight-line motion. In some sense, this trajectory does not exist, because it approximates a control that was not in the set we were originally given.

Chattering occurs when an arbitrarily short application of control u_1 causes control u_2 to be maximizing, and an arbitrarily short application of control u_2 causes u_1 to be maximizing. Such trajectories are both hard to implement on a control system, and hard to analyze.

We will show first that cases where controls *chatter* can be clearly identified and have a simple structure. Chattering can occur on control-line trajectories in two cases: if there are switches between translation controls, or if there

are translations parallel to the control line. We will analyze these cases in greater detail later in the paper, but first we will show that these are the only cases in which trajectories do not have piecewise-continuous controls. We will also see that chattering trajectories occur at particular discrete values of the Hamiltonian.

Theorem 2 *Let k_1 , k_2 , and k_3 be constants corresponding to an extremal trajectory, with $k_1^2 + k_2^2 \neq 0$. Let H be the value of the Hamiltonian over the trajectory, and let U be a discrete set of controls used to generate the trajectory. If the trajectory satisfies the following two conditions*

- (i) *There are no translation-translation switches;*
- (ii) *There is no time interval during which the trajectory is a translation parallel to the control line*

then the controls are piecewise constant over the trajectory, with a minimum positive time between control switches that may be computed as a function of the Hamiltonian and the controls.

Proof: The basic idea of the proof is to consider all possible pairs of controls (e.g. u_1u_3 or u_7u_2). For each pair of controls (a potential switch), and given H , we will see that there is a discrete set of (y, θ) values at which this switch can occur. This is true for switches that involve at least one rotation (condition (i)).

Then consider each possible pair of switches. (For example, u_1u_3 followed by u_3u_6 .) For each switch in a pair, the (y, θ) values are distinct, or they are not. If they are distinct, then the time between switches may be computed based on the control that occurs between the switches. If they are not, then this section is a translation parallel to the control line (condition (ii)).

We need to prove that for each pair of controls (or switch) and a given value of H , that there is a discrete set of (y, θ) values at which this switch can occur. Let the pair of controls be labelled as u_1, u_2 .

If either u_1 or u_2 is a pure translation ($\omega = 0$), the angle θ is constant over that section of the trajectory, and can be computed from equation 21.

If neither control is a pure translation, then the angle at which the control switches may be computed using the

Hamiltonian equations:

$$H = v_1 \cos(\theta + \alpha_1) + y\omega_1 \quad (23)$$

$$H = v_2 \cos(\theta + \alpha_2) + y\omega_2. \quad (24)$$

By multiplying the top equation by ω_2 , the bottom by ω_1 , and subtracting, we may eliminate y :

$$\omega_2 H - \omega_1 H = \omega_2 v_1 \cos(\theta + \alpha_1) - \omega_1 v_2 \cos(\theta + \alpha_2). \quad (25)$$

Using the cosine addition formula and rearranging, this equation can be written in the form

$$a \cos \theta + b \sin \theta = c, \quad (26)$$

where

$$a = \omega_2 v_1 \cos \alpha_1 - \omega_1 v_2 \cos \alpha_2 \quad (27)$$

$$b = -\omega_2 v_1 \sin \alpha_1 + \omega_1 v_2 \sin \alpha_2 \quad (28)$$

$$c = \omega_2 H - \omega_1 H. \quad (29)$$

Equation 26 is of standard form. Craig [9] gives the solution as

$$\theta = 2 \operatorname{atan} \left(\frac{b \pm \sqrt{a^2 + b^2 - c^2}}{a + c} \right). \quad (30)$$

(If $a + c = 0$, then $\theta = \pi$.) Notice that there are at most two solutions for θ . For each solution for θ , we can plug the values of θ , H , and the rotation control into equation 21 to compute the unique value of y . ■

We divide chattering into three cases: chattering between translation controls, chattering between pairs of rotation controls, and chattering involving three or more rotations. Lemmas 2, 3, and 5 show that the exact value of the Hamiltonian can be computed given a pair or triplet of controls in each of these cases, if chattering is possible between these controls. Therefore, to enumerate all of these cases, it is sufficient to consider each possible triplet and pair of controls.

Lemma 2 (Chattering between translations) *If an extremal trajectory contains a switch between two distinct pure translation controls u_i and u_j , then the value of the Hamiltonian is uniquely determined by those two controls.*

Proof: At the switch

$$H = v_i \cos(\theta + \alpha_i) \quad (31)$$

$$H = v_j \cos(\theta + \alpha_j) \quad (32)$$

We have two equations in two unknowns. First, we may solve for the angle θ that the body must make relative to the control line at the switch.

$$v_i \cos(\theta + \alpha_i) = v_j \cos(\theta + \alpha_j) \quad (33)$$

Use of the cosine addition identity and some algebraic manipulation leads to:

$$\tan \theta = \frac{v_j \cos \alpha_j - v_i \cos \alpha_i}{v_j \sin \alpha_j - v_i \sin \alpha_i} \quad (34)$$

We choose the solution for θ satisfying the requirement that $H > 0$, and compute the unique value for H from $H = v_i \cos(\theta + \alpha_i)$. ■

Lemma 3 (Chattering between a pair of rotations) *A chattering extremal trajectory parallel to the control line that contains no more than two rotation controls has a unique value of the Hamiltonian that can be computed from the pair of controls, and the rotation centers corresponding to these controls fall on a line perpendicular to the control line. Furthermore, the angular velocities at the rotation centers must have opposite signs, and the rotation centers must fall on opposite sides of the control line.*

Proof: Let the controls be u_1 and u_2 . The reference point of the body is arbitrary, and each rotation has a corresponding rotation center. Choose the reference point of the body to be the rotation center corresponding to u_2 . Over any non-empty interval, both y and θ are constant, so the measure of the each sets for which u_1 or u_2 is active is greater than zero. Notice that y is unaffected by control u_2 , so the y motion of the reference point depends only on the control u_1 . If the rotation centers do not fall on a line perpendicular to the control line, y is non-constant over the interval. We may compute H using the following relations, where y_1 and y_2 are the y coordinates of the rotation centers:

$$y_1 \omega_1 = H \quad (35)$$

$$y_2 \omega_2 = H \quad (36)$$

$$|y_2 - y_1| = d_{12} \quad (37)$$

and d_{12} is the (constant) distance between rotation centers. Since the resulting motion must be a translation, the rotation centers must have opposite signs and lie on opposite sides of the control line. ■

We will next consider the case where three or more controls are involved in chattering. To do so, we will need the following lemma.

Lemma 4 *Let u_1, \dots, u_k , $k \geq 3$, be distinct controls on a line in control space $(\hat{x}, \hat{y}, \hat{\theta})$, not all translations. Then, in the frame of the body, all the corresponding rotation centers are on a line. There can be at most one translation and its direction is perpendicular to the line formed by the rotation centers.*

Proof: Since not all controls are translations, their line intersects the $\hat{\theta} = 0$ plane in at most one point. This is the only translation admissible in the set.

Consider the reference point in the frame of the body to be placed at one of the centers of rotation in the given set. Project its velocity under all given controls onto the (\hat{x}, \hat{y}) plane. All the projections must lie on a line L' that passes through the origin. Therefore, all the rotation centers must be located on a line that is perpendicular to L' and the translation, if it is given in the set, must be perpendicular to the line of rotation centers. ■

Lemma 5 (Chattering between three or more rotations) *A chattering extremal trajectory parallel to the control line that contains three or more controls has a unique value of the Hamiltonian that can be computed from a triplet of these controls.*

Proof: For each control i that maximizes the Hamiltonian over the chattering segment we can write an equation of the form

$$v_i \cos(\theta + \alpha_i) + y \omega_i = H \quad (38)$$

Note that this equation contains three unknowns: θ , y and H . By assembling a system of three such distinct equations, we should expect, in the general case, to obtain a set of measure zero of solutions for this system of equations. This should yield the critical H values that are sought. We will next proceed to characterize the vehicles for which the system is singular and to derive the critical values of H for such vehicles.

Re-write the equation 38 as

$$v_i \cos \alpha_i c - v_i \sin \alpha_i s + \omega_i y - H = 0 \quad (39)$$

where c and s are the cosine and sine of θ respectively. Note that this is equivalently

$$\dot{x}_i c - \dot{y}_i s + \dot{\theta}_i y - H = 0 \quad (40)$$

Assume we knew that controls i, j and k are used during a chattering segment. Then we can assemble the linear system

$$\begin{bmatrix} \dot{x}_i & -\dot{y}_i & \dot{\theta}_i & -1 \\ \dot{x}_j & -\dot{y}_j & \dot{\theta}_j & -1 \\ \dot{x}_k & -\dot{y}_k & \dot{\theta}_k & -1 \end{bmatrix} \begin{pmatrix} c \\ s \\ y \\ H \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad (41)$$

Consider the rank of the matrix on the left-hand side of this system. There are two cases: either it is of rank three, or it is of rank less than three.

Case 1: The matrix is of rank three. Assume that i, j , and k have been chosen so that their corresponding row vectors, when taken as a matrix, have rank three. By elementary row operations on the system, we eliminate variables H and y and are left with an equation of the form $ac + bs = 0$, where a and b are constants, not both zero. In conjunction with $c^2 + s^2 = 1$, we calculate two possible values of the (c, s) pair. Replacing into the original system, we obtain two critical values of H . Only these two values of the Hamiltonian allow controls i, j and k to be simultaneously maximizing.

Case 2: The matrix is of rank less than three. Then one row is a linear combination of the other two; i.e., there exist constants a and b , not both zero, such that

$$a \begin{pmatrix} \dot{x}_i \\ -\dot{y}_i \\ \dot{\theta}_i \\ -1 \end{pmatrix} + b \begin{pmatrix} \dot{x}_j \\ -\dot{y}_j \\ \dot{\theta}_j \\ -1 \end{pmatrix} = \begin{pmatrix} \dot{x}_k \\ -\dot{y}_k \\ \dot{\theta}_k \\ -1 \end{pmatrix} \quad (42)$$

Equivalently, denoting by \hat{q}'_i the first three components of each vector, we obtain the system

$$a\hat{q}'_i + b\hat{q}'_j = \hat{q}'_k \quad (43)$$

$$a + b = 1 \quad (44)$$

This indicates that the three vectors lie on a line in the control space. Therefore, the points in control space that correspond to controls i, j and k also lie on the reflection of this line over the \dot{y} axis. Thus, vehicles that can have three or more maximizing controls at arbitrary values of the Hamiltonian need to have all of these controls on a single line in control space. From lemma 4, rotation centers corresponding to these controls must also lie on a line. Let L be the line through these rotation centers.

We will calculate the critical values of H in this case by requiring that the chattering control be sustainable, i.e. given a control function $u(t)$ that takes values u_i, u_j and u_k within arbitrarily small neighborhoods of a given moment we will require the Lebesgue integral of this function on any such neighborhood to be zero in its y and θ components.

Let P be the intersection of L with the control line. As shown above in the proof of theorem 2, the chattering control will result in a motion that translates P along the control line.

We want to show that the angle between L and the control line must be equal to $\pi/2$ along the chattering extremal.

Assume first that the angle between L and the control line is less than or equal to $\pi/2$. Note that all the controls corresponding to rotation centers above the control line have positive angular velocities, and all those below negative, because of corollary 1. No matter which control we choose, the Lebesgue integral of the motion of the point P will have a negative y component. The only possible way to obtain the required value of zero for the y component of the motion is to have the line L perpendicular to the control line. This determines the θ component of the state, which allows us to uniquely determine the value of the Hamiltonian at which it is possible to sustainably apply a chattering control as described, in the manner described above for the two controls case in lemma 3.

The case for which the angle between L and the control line is greater than $\pi/2$ is analogous. ■

3.2 Discretizing polyhedral control spaces

The lemmas in the last subsection allow possible chattering cases to be identified by considering all pairs and triplets of controls. For each pair or triplet for which chattering may occur, there is a corresponding direction and

speed of translation. If this translation is already in the set of permitted controls, then any extremal trajectory with chattering between these controls is equivalent to an extremal trajectory without this chattering. If the translation is not in the set of permitted controls, and these translations cannot be otherwise ruled out of the set of optimal trajectories, then there may be optimal trajectories that require infinitely many control switches in a finite time – a situation not well-described by the kinematic model.

The lemmas also indicate how to reduce a continuous polyhedral space of generalized velocities that the body can follow to a discrete set of controls sufficient to model the time-optimal trajectories. The adjoint function is a privileged direction in velocity space, such that at almost every time, the control must be chosen to maximize dot product of the generalized velocity of the rigid body with the adjoint. If the space of generalized velocities is polyhedral, we expect for most possible direction of the adjoint vector that a single vertex is maximizing. Therefore the discretized set of controls should contain a control corresponding to each vertex.

However, there are also directions of adjoint vectors that maximize generalized velocities on a face or edge of the polyhedron. In this case, we must consider chattering between the vertices of the face or edge. The lemmas show that for each such chattering, there is a single translation control that may be added to the discretized set, removing the need for chattering.

As a concrete example, consider the optimal trajectories for a Dubins car. The trajectories may include the left and right turns, which are vertices of the control space, but also may include straight lines, which are not vertices. The fact that these straight lines appear in optimal trajectories may be deduced by chattering the controls that are vertices.

3.3 Sustainable extremal controls

We now turn to non-chattering trajectories. These trajectories have piecewise-constant controls, and this section will show that at every point it is possible to write equations to determine which control precedes the switch, and which follows the switch. For example, in figure 3, in configuration 2, rotations about either of points L or R gives the same value of the Hamiltonian; both controls are maximizing. However, regardless of which control is fol-

lowed, at the next instant only one control will continue to satisfy the Maximum Principle. Given a (y, θ) value at a point where multiple controls maximize the Hamiltonian, under what circumstances is it possible to determine the ‘next’ control?

We define a *sustainable extremal control* at time t and configuration q to be a control such that there exists a strictly positive constant δ such that applying the control on $[t, t + \delta)$ generates an extremal trajectory with the same constants. To analyze the circumstances under which extremal controls are sustainable, we will need to look at derivatives of the Hamiltonian functions for each control with respect to time. We will choose the convention that along a piecewise-constant-control trajectory, the value of control exactly at the time of a switch is the same as the control in the interval after the switch – *i.e.* controls are applied over half-open intervals that are closed on the left side.

If a control u_j is extremal over some interval (sustainable), it must be the case that this control continues to maximize the Hamiltonian over the entire interval. For every other control u_i , either the Hamiltonian must remain below the maximum value, or if it does attain a maximum, not exceed that maximum – that is, the derivative of the Hamiltonians for other controls must be non-positive.

Define

$$\frac{d}{dt}H_i^j(q) \quad (45)$$

to be the right derivative of the Hamiltonian for control i at the configuration q relative to the control line, assuming the control u_j is applied for some half-open interval containing the current time. We can compute this value as follows.

$$H_i^j(q) = v_i \cos(\theta + \alpha_i) + y\omega_i \quad (46)$$

$$\frac{d}{dt}H_i^j(q) = -\dot{\theta}v_i \sin(\theta + \alpha_i) + \dot{y}\omega_i \quad (47)$$

$$\frac{d}{dt}H_i^j(q) = -\omega_j v_i \sin(\theta + \alpha_i) + v_j \omega_i \sin(\theta + \alpha_j) \quad (48)$$

Define a *generic* point of a non-chattering extremal trajectory to be a point such that there is a *single* sustainable extremal control. A *generic trajectory* is a trajectory such that every point along the trajectory is generic.

More formally, on a generic trajectory, for every time t , there exists a unique control u_i such that

$$H^* = H_i \quad (49)$$

and further for every j such that $H_j = H_i$ with $j \neq i$,

$$\frac{d}{dt} H_i^j(q) < 0. \quad (50)$$

3.4 Rotation-rotation switches

Along generic trajectories, the control is completely determined by the configuration of the body relative to the control line. Our goal in section 5 will be to generate these types of trajectory. To do this, we want to be able to compute the configuration of the body at the next switch, if we know the configuration at the current switch. In fact, given the current switch (or any other configuration), we can compute the value of H , and then use this value to compute the configuration at the next switch. There are two basic kinds of switches that can occur on generic trajectories, rotation-rotation switches, and rotation-translation switches, and we will analyze them in this section and the next.

Analysis of rotation-rotation switches is simplified by considering the rotation center associated with each control. In the frame of the body, rotation center (\hat{x}_i, \hat{y}_i) corresponding to control u_i is

$$\hat{x}_i = -\frac{\dot{\hat{y}}_i}{\omega_i} \quad (51)$$

$$\hat{y}_i = \frac{\dot{\hat{x}}_i}{\omega_i} \quad (52)$$

Consider a switch from control u_i to control u_j , with $i \neq j$. Let (x_i, y_i) be the coordinates of rotation center i with respect to the control line. From corollary 1, the Hamiltonian corresponding to control i at this configuration is

$$H = y_i \omega_i. \quad (53)$$

Geometrically, as we can see in figure 3 in configuration 2, the switch corresponds to a situation where each rotation center lies on its own line at a distance H/ω from the control line. Since the distance between the rotation centers is fixed by their associated controls, the problem is to fit a line segment in such a way that the endpoints contact

two parallel lines. There are only two solutions to this geometric problem, each corresponding to a direction of the switch; this is in fact the main idea behind the proof of theorem 2. The main result in this section will show how to identify the correct solution for the configuration, given a value for H , and a directed switch between a pair of controls.

We will need the following small result.

Proposition 1 *At any switch between two rotation controls i and j , the locations of the rotation centers are distinct.*

Proof: Assume the two rotation centers were not distinct. Then from $H = y_i \omega_i$ we may compute ω_i , and these two controls are identical. ■

Therefore the displacement vector between the two rotation centers at a switch, $(x_j - x_i, y_j - y_i)$ has nonzero length. Let d_{ij} be the this distance and let γ_{ij} be the angle this vector makes with the control line. d_{ij} may be easily computed from equations 51 and 52, and γ_{ij} is given by

$$\gamma_{ij} = \theta + \text{atan2}(\hat{y}_j - \hat{y}_i, \hat{x}_j - \hat{x}_i). \quad (54)$$

Lemma 6 (Rotation-rotation switching configurations) *For a given value of the Hamiltonian, with $k_1^2 + k_2^2 \neq 0$, the angle θ of the body with respect to the control line at which the control switches from rotation control i to rotation control j is unique, as is the distance y from the control line.*

Proof: By definition,

$$y_j - y_i = d_{ij} \sin \gamma_{ij} \quad (55)$$

From equation 53,

$$y_j - y_i = H \left(\frac{1}{\omega_j} - \frac{1}{\omega_i} \right) \quad (56)$$

Combining,

$$\sin \gamma_{ij} = \frac{H}{d_{ij}} \left(\frac{1}{\omega_j} - \frac{1}{\omega_i} \right) \quad (57)$$

This gives two solutions for γ_{ij} . However, one of the candidates can be eliminated. Notice that to the right of the switch, the right derivative $\frac{d}{dt} \omega_i y_i$ exists and is negative; otherwise there exists a point after the switch such that

choosing u_i maximizes the Hamiltonian. Since after the switch the body is rotating around rotation center j with angular velocity ω_j ,

$$\frac{d}{dt}y_i = \omega_j d_{ij} \cos \gamma_{ij}. \quad (58)$$

Premultiply by ω_i , giving

$$\omega_i \frac{d}{dt}y_i = \omega_i \omega_j d_{ij} \cos \gamma_{ij}. \quad (59)$$

Therefore

$$\omega_i \omega_j \cos \gamma_{ij} \leq 0 \quad (60)$$

This gives a constraint on the sign of the cosine of γ_{ij} , which allows us to select the correct solution from equation 57, yielding a single value for θ , the angle at which the control switch occurs. The distance y from the control line may then be computed from equation 53. ■

3.5 Rotation-translation switches

We will use a similar technique to compute the configurations of the body relative to the control line at which a switch involving a rotation and a translation may occur.

Lemma 7 (Rotation-translation switching configurations)

For a given value of the Hamiltonian, with $k_1^2 + k_2^2 \neq 0$, the angle θ of the body with respect to the control line at which the control switches from control i to control j , where exactly one of the controls is a translation, is unique, as is the distance y from the control line.

Proof: The translation corresponds to control $(v, \alpha, 0)$.

$$H = \dot{x} + 0y = v \cos(\theta + \alpha) \quad (61)$$

Solving this equation yields two solutions for $(\theta + \alpha)$, and thus two solutions for θ . We now want to show that only one of these solutions is valid. Assume the first control is a rotation. Notice that the right derivative

$$\frac{d}{dt}H_i^j(q) = v\omega \sin(\theta + \alpha) \quad (62)$$

exists and is non-positive; otherwise choosing control u_i increases the value of H beyond the maximum for this trajectory. So,

$$v\omega \sin(\theta + \alpha) \leq 0. \quad (63)$$

This allows the choice of the correct (and unique) value for θ from the two solutions previously computed.

Once θ is known, the value for y may be computed from the expression for the Hamiltonian. If the switch is from a translation to rotation control, the analysis is similar, with a sign change on inequality 63. ■

4 Critical values of the Hamiltonian

In the previous section, we showed how switches follow each other on generic trajectories, and showed that non-generic trajectories only occur at particular critical values of the Hamiltonian. The main result of this section is that between these critical values of the Hamiltonian, small variations in the (y, θ) configuration of the body relative to the control line do not change the structure of trajectories, in the sense that the sequence of control switches does not change. This does not mean that every trajectory with the same value of the Hamiltonian has the same structure, since various controls may lead to the same value of the Hamiltonian at certain configurations. However, we will see that every trajectory with the value of the Hamiltonian in some range between critical values, going through a particular control switch $u_i u_j$, has the same structure. In the following section, we will use this fact to generate all trajectory structures, by first partitioning the Hamiltonian, then considering a finite number of possible switches within each interval of the Hamiltonian, and for each switch, constructing a trajectory structure.

Lemma 8 *For a switch from control u_i to u_j , where at least one of the controls is a rotation, there is a set of values of H that is the intersection of the strictly positive real numbers with some closed interval, such that for each value of H in this set, there is a single value of (y, θ) at which the switch may occur, ignoring other controls, and outside the interval, there are no values of (y, θ) for which the switch may occur.*

Proof: If the controls are both rotations, the result follows from equations 57 and 60, since the range of the sine function in equation 57 is $[-1, 1]$. If one of the controls is a translation, then equation 61 and 63 give similar results. ■

Theorem 3 *There is a partitioning of the values of the Hamiltonian into a finite set of open intervals, such that every trajectory with a Hamiltonian within a single interval, containing the same control switch, will contain the same sequence of control switches, following the switch that the trajectories have in common.*

Proof: There are some discrete special values of the Hamiltonian, and we will build a partitioning of the Hamiltonian by these values. First, if a trajectory contains any translation-translation switch, we can compute the unique value of the Hamiltonian. So for each pair of translation controls in the input set, compute the Hamiltonian, and add these values to the partitioning. Also, for each pair of controls in the input set, add the boundaries of the sets computed in lemma 8. Finally, for each triplet of controls in the input set, compute the values of the Hamiltonian for which all three controls may maximize the Hamiltonian, using the linear system 41 and the strategy employed in the proof of lemma 5.

Now we want to show that outside of these special values, continuous perturbation of the value of H a trajectory starting with some control switch u_1u_2 does not change the structure of the trajectory.

Consider an extremal with a value of H not in the partitioning. For each control switch $u_i u_j$, consider the corresponding set of H values computed in lemma 3. H is either on the interior of the closed interval for this control, or exterior to the closed interval, since H cannot be zero. No continuous perturbation of H that does not cross one of values in the partitioning will allow the ‘exterior’ controls to become active during the trajectory, so we may ignore these controls.

Starting from switch u_1u_2 , consider each possible switch u_2u_k , with $k \neq 2$. Compute the unique value of (y, θ) at which the control may switch as a function of H . Using control u_2 , we can compute the potential time to each next switch, $t_{1,2,k}$. Each such time is a continuous function of H , and the next switch to occur will be the one with the smallest time. If no two of these times are equal, then continuous perturbation of H will not change the chosen next control until two times become equal (and minimum). If two such times are equal, then at the instant of the switch, three controls maximize the Hamiltonian. In this case, the linear system 41 from the proof of lemma 5 can be used to compute the value of the Hamil-

tonian, and this value has already been added to the partitioning. ■

4.1 Switching space

Given the set of controls, we enumerate the trajectory classes by looking at all the possible configurations that the rigid body can have with respect to the control line. Only the body’s distance and orientation relative to the control line determine the maximizing controls. It makes sense, therefore, to project extremal trajectories onto their y and θ dimensions, which we will call the *switching space*. The topology of the switching space is cylindrical, as the θ coordinate wraps around after 360 degrees.

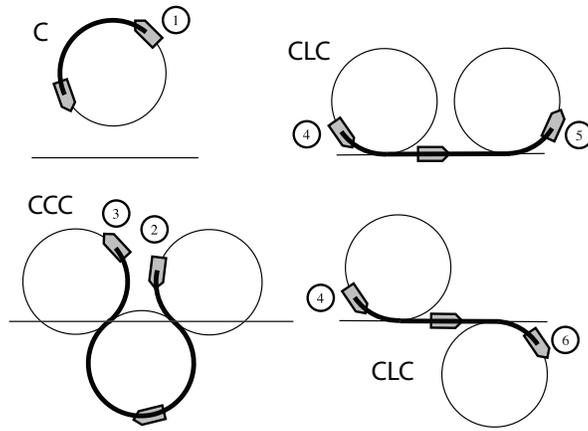
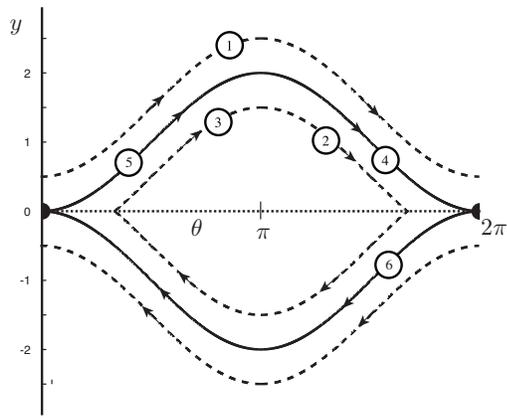
Each body configuration (y, θ) in switching space has associated with it a value H_i for each control, and one or more of these controls is maximizing. An extremal trajectory attains a set of (y, θ) values, and there is a corresponding curve in (y, θ) space. Generic trajectories have piecewise constant controls, so the curve corresponding to a generic trajectory is a union of smooth curves that intersect at configurations where the controls switch.

We have seen that there is a discrete set of H values such that between these H values, all extremal trajectories have the same structure, if they contain the same control switch. We will call trajectories that have a Hamiltonian from this discrete set *separating* trajectories.

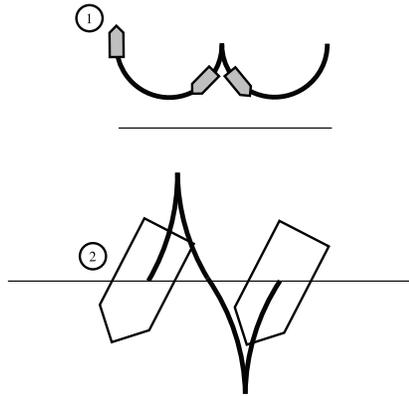
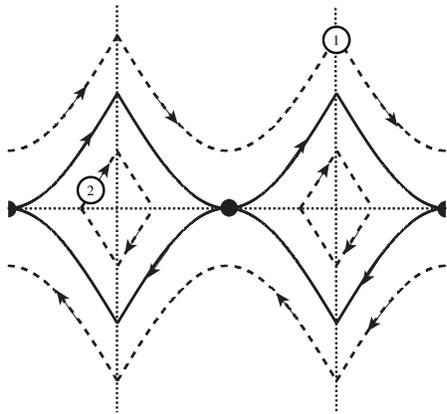
Points in the switching space may be classified into different types:

1. Switching configurations
2. Points on separating trajectories
3. Other points: points on generic trajectories that do not correspond to a switch

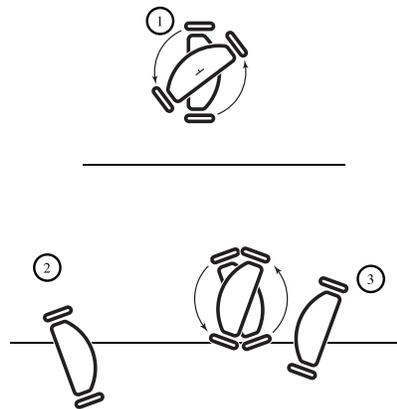
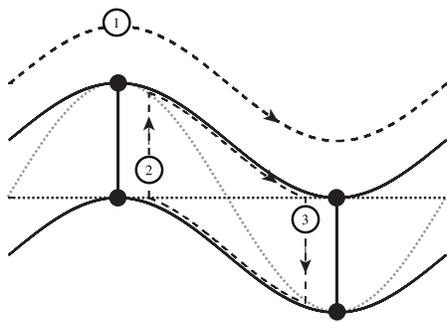
Figure 6a shows the switching space for a Dubins car. The trajectory labelled ‘1’ corresponds to a generic trajectory over which the car drives in a circle. The corresponding (y, θ) values are a sinusoid in the switching space. The trajectory from configurations ‘2’ to ‘3’ traces a circle until the car crosses the $y = 0$ axis, switches to a different circular arc, and switches back again to the first rotation at the axis. We say the $y = 0$ axis is a *switching curve* for the Dubins car. Finally, consider the two trajectories ‘4’ to ‘5’ and ‘5’ to ‘6’. These special trajectories contain translation parallel the control line, and



(a) The Dubins car.



(b) The Reeds-Shepp car.



(c) The differential-drive.

Figure 6: Switching spaces and example trajectories for standard robotic vehicles

their curves in switching space separate the C and CCC trajectories that we considered first.

Drawing all of the switching curves and separating trajectories in switching space gives a visualization of all possible extremal trajectories for the rigid body. Figures 6b and 6c give the switching curves and separating curves for Reeds-Shepp and differential-drive vehicles.

Given a new system, how can we compute a switching curve, which is the set of (y, θ) at which a switch may occur? Consider a pair of controls (u_1, u_2) . At any switch, the Hamiltonians of these controls must be equal; $H_i = H_j$.

$$v_1 \cos(\theta + \alpha_1) + y\omega_1 = v_2 \cos(\theta + \alpha_2) + y\omega_2 \quad (64)$$

After some algebraic manipulation, this gives the condition

$$(\hat{x}_1 - \hat{x}_2) \cos \theta - (\hat{y}_1 - \hat{y}_2) \sin \theta = y(\omega_1 - \omega_2). \quad (65)$$

Consider the vector that is the difference of the vectors (\hat{x}_1, \hat{y}_1) and (\hat{x}_2, \hat{y}_2) . Let Δv_{12} be the length of this vector, and let δ_{12} be the angle this vector forms with the horizontal axis. Then the above equation becomes

$$\Delta v_{12} \cos(\theta - \delta_{12}) = y(\omega_1 - \omega_2). \quad (66)$$

If the controls are distinct, at least one of Δv_{12} or $(\omega_1 - \omega_2)$ is nonzero. In the general case, this curve is a sinusoid in y, θ space. In degenerate cases, it is either the horizontal axis (as for the Dubins car) or a pair of vertical lines (as appear for the Reeds-Shepp car in figure 6b).

Since the switching curves are sinusoids or straight lines, the union of switching curves divides the space up into regions within which a single control maximizes the Hamiltonian. Notice that trajectories in this space that are not tangent to any switching curve can be perturbed without changing the structure (the order and type of controls) of the trajectory. Trajectories that are tangent to a switching curve therefore separate the space by type of trajectory.

How do we calculate the (y, θ) separating curves? It turns out that we have already seen them in a previous section; these curves are the pre-image of the discrete values of the Hamiltonian that form the partition in theorem 3.

Separating curves separate trajectories in (y, θ) space, and switching curves separate control regions. If we want

to draw a particular trajectory in switching space, there are two approaches. We have already seen the first, direct method – arcs of circles in the workspace of the body map to sinusoidal curves in (θ, y) switching space, and the sinusoidal curves are joined at points that cross the three switching curves.

The second way to draw trajectories in switching space relies on two observations: the Hamiltonian depends only on the value of y and θ , and the Hamiltonian is constant across any minimum-time trajectory. Therefore, dashed and solid curves in figure 6 can be thought of as being the level sets of the Hamiltonian function $H(\theta, y)$.

Each level curve contains a potentially infinite number of optimal trajectories, since it is possible to start and end anywhere on it, as long as travel proceeds continuously, with the possibility of periodicity as well. We can distinguish the *optimal classes* appearing on a level curve by first separating the unconnected regions, and then picking start and end regions on each.

5 Generic trajectories

The analysis of switching space provides a compact geometric way of describing all extremal trajectories for a given set of controls. In this section, we will describe an algorithmic approach to enumerate trajectory types.

5.1 Generating generic trajectories

We have said that generic trajectories are trajectories for which the configuration of the rigid body relative to the control line completely determines the trajectory. Given such an initial configuration, lemmas 6 and 7 can be used to generate the trajectory.

Given initial y and θ , compute the current maximizing control. Given the current control, y , and θ , compute a list of times to potentially switch to each of the other controls. From this list, choose the control for which the switching time is soonest. This is the next control. Compute the y and θ value at the switch by applying the original control for this amount of time, or alternately, using the value of H and the control switch. Repeat to find each new control.

If the current control is a rotation, the potential time to switch to another control is computed by computing the θ value at the switch, subtracting the current θ value from

the switching θ , and dividing by the current angular velocity ω . If the current control is a translation, the potential time to switch to another control is computed by computing the y value at the switch, subtracting the current y value from the switching y , and dividing by the current velocity \dot{y} .

The extremal trajectories generated are of infinite duration. Also, we expect these trajectories to be periodic, in the sense that they will return to the same control, and the same values of y and θ .

Lemma 9 (Periodicity of generic extremals) *Assume the rigid-body has n discrete controls. Any generic extremal trajectory becomes periodic after no more than $n(n - 1)$ switches.*

Proof: The trajectory starts at some initial configuration, and passes through a sequence of switches. Since there are n controls, there are $n(n - 1)$ unique pairwise combinations of controls. By the pigeonhole principle, the control pair from the first switch must repeat after no more than $n(n - 1)$ switches. Since the value of H is constant over any extremal trajectory, we can use lemmas 6 and 7 to compute the (y, θ) values at the first switch, and at the repetition of the first switch. ■

This algorithm is quite simple to implement, and we have done so in a few hundred lines of code of Javascript; the implementation is available from the authors' website.

5.2 Classifying generic trajectories

Each trajectory is described by a sequence of controls. The generic extremal trajectories generated by the algorithm in the previous section are of infinite duration, but are all periodic. We can describe the trajectory by writing out the sequence of controls corresponding to a single period. For example, $u_3u_1u_7u_1$ would describe the trajectory "... $u_1u_3u_1u_7u_1u_3u_1u_7u_1$...". We say that two infinite-duration extremals are of the same class if the control sequences that describe them are circular permutations of each other. (Thus, $u_1u_7u_1u_3$ would also be in the same class.)

We would also like to classify trajectories of finite duration. If such a trajectory contains at least one period (*i.e.*, repeats a switch), then the trajectory can be classified using the above system. Trajectories shorter than one period

may fall into multiple classes. An eventual goal is to use the classification to search first for optimal trajectories of each class between a start and goal, and then to find the fastest trajectories from among these. Short trajectories may show up several times in such a search, but do not otherwise cause a problem.

The generative algorithm suggests a method of describing all classes of generic trajectories. Choose many initial (y, θ) configurations, run the simulator for each, and report all unique results. The primary difficulty is how to sample the (y, θ) configurations so as to ensure that exactly one representative of each trajectory structure is found. Since every infinite-duration extremal trajectory passes through several switches, we can restrict the initial configurations to those that occur at a switch.

If we know the controls involved in the first switch, then we can find different (y, θ) values by varying the value of the Hamiltonian continuously. We have proven that varying the Hamiltonian will not change the structure of the trajectory, except at certain critical values. Therefore, for each (possibly open) interval in the partition of the Hamiltonian, choose one value of the Hamiltonian, building the discrete set \bar{H} . For each pair of controls in the input set, and value in \bar{H} , compute a (y, θ) value, and use this to generate one period of a trajectory.

As described, the algorithm will find many duplicate trajectory structures, since circular permutations of trajectory structures are considered to be in the same class. To avoid this, as each trajectory is generated, for each member of \bar{H} , mark the already-visited pairs of controls.

How many generic trajectory structures might there be? Let $n_{\bar{H}}$ be the cardinality of \bar{H} . We can see from the length of the visited list in the algorithm that the total bound on the number of trajectories is $n_{\bar{H}} \times n(n - 1)$, where n is the number of controls in the control set. Since the critical values of the Hamiltonian are obtained by considering all pairs and triplets of controls, $n_{\bar{H}}$ is upper bounded by $O(n^3)$. Therefore, $O(n^5)$ is an upper bound on the number of trajectory types. In practice, we expect the number of trajectory types to be much smaller, since most triplets of controls will not generate a critical value of the Hamiltonian, and since long trajectory structures will visit several locations in the visited list. For example, the Dubins car has three controls, and only three generic trajectory structures (left C, right C, and CCC). The Reeds-Sheep car has six controls, and only

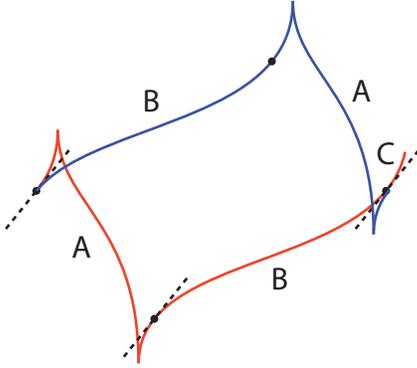


Figure 7: Construction showing that optimal trajectories for which the image of $\theta(t)$ is not S^1 , and for which $\dot{\theta}(0) \neq 0$, contain no more than one period.

four generic trajectory structures.

5.3 Limitation on the number of periods

We have seen that generic trajectories are periodic. Trajectories return to the same control switch repeatedly, and because H is constant over the trajectory, return in fact to the same (θ, y) values. The Maximum Principle only gives necessary conditions on trajectories – although the extremal trajectories that satisfy the Maximum Principle are periodic, we do not typically expect trajectories with many periods to be optimal. An example is the parallel-parking trajectory for the Reeds-Shepp car shown in figure 6b. One period may be optimal to move the car slightly sideways, but to move a long distance sideways, we expect the car to turn and drive rather than executing a sequence of parallel-parking moves.

In this section, we will show that under many circumstances, optimal trajectories contain no more than one period. We will need the following lemma.

Lemma 10 *Generic trajectory segments with constant $\theta(t)$ have constant controls.*

Proof: Over any optimal trajectory, H is constant and non-zero. The Hamiltonian is $H = \dot{x} + \dot{\theta}y$. If $\theta(t)$ is constant, $H = \dot{x}$. Therefore \dot{x} is a constant. Since $\dot{y} = \dot{x} \tan(\theta)$, \dot{y} is also constant. ■

Theorem 4 *Generic trajectories for which the image of $\theta(t)$ is not S^1 , and for which $\dot{\theta}(0) \neq 0$, contain no more than one period.*

Proof: Consider a candidate trajectory that contains more than a full period. We will prove that this trajectory is not optimal by constructing another trajectory from the start to the goal that takes an equal amount of time, but does not satisfy the Maximum Principle (and is therefore not optimal). Figure 7 illustrates the idea.

Any trajectory achieves both a minimum and a maximum value for θ . For now, assume that $\theta(0)$ is not either of the extreme values. By lemma 10, trajectories for which the image of $\theta(t)$ is a point contain zero periods.

Let T be the duration of the first period. Since $\theta(t)$ is continuous, it must achieve all the values between the minimum and maximum values of θ in $(0, T)$. Therefore there exists $t_1 \in (0, T)$ such that $\theta(t_1) = \theta(0)$. Also, $y(0) \neq y(t_1)$, since $t_1 < T$. Let A be the section of the trajectory on the interval $[0, t_1]$, let B be the section of the trajectory on $[t_1, T]$, and let C be the remainder of the trajectory. The controls at the start of A and the start of C are the same.

Now construct the trajectory BAC . This trajectory takes the same duration as ABC , is feasible, and reaches the goal. On this new trajectory, we have the same controls at the beginning of A (time $T - t_1$) and beginning of C (time T), but different y values. If we compute the Hamiltonians at these times,

$$H(T - t_1) = \dot{x}(T) + \dot{\theta}(T)y(T - t_1) \quad (67)$$

$$H(T) = \dot{x}(T) + \dot{\theta}(T)y(T), \quad (68)$$

we find that the Hamiltonian is not constant of the trajectory, since $y(T) \neq y(T - t_1)$ and $\dot{\theta}(t) \neq 0$. ■

Theorem 5 *Generic trajectories for which the image of θ is S^1 for rigid bodies with symmetric control bounds, so that the body can reverse along any trajectory at full speed, contain no more than one period.*

Proof: Consider a candidate trajectory that contains at least one period. We will prove that this trajectory is not optimal by constructing another trajectory from the start to the goal that takes an equal amount of time, but does not satisfy the Maximum Principle (and is therefore not optimal). Figure 8 illustrates the idea.

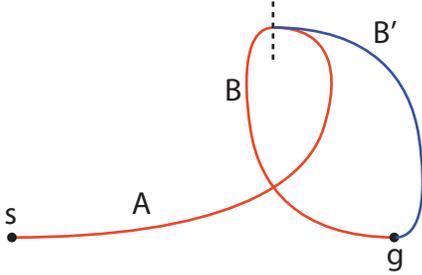


Figure 8: Construction showing that rigid bodies with symmetric control bounds have optimal trajectories containing no more than one period.

Choose $t_1 \in (0, T)$ such that $|\theta(t_1) - \theta(0)| = \pi$. Let A be the portion of the trajectory on $[0, t_1]$ and let B be the portion of the trajectory on $[t_1, T]$. The trajectory is of the form AB . Let B' be B reversed in time: *i.e.* the controls at the start of B' are the reverse of the controls at the end of B . Consider the trajectory AB' . This trajectory takes the vehicle from the start to the goal.

The constructed trajectory is not extremal. Consider the time t_1 . Notice that since θ is strictly monotonic over AB , the signs of both \dot{x} and $\dot{\theta}$ must change at $T - t_1$ on AB' . We can therefore compute two different values for the Hamiltonian at this time. ■

5.4 Restricting the location of the control line for given initial and final motions

We have seen that the problem of finding the optimal generic trajectory between two configurations is practically synonymous with find the values of the constant k_1 , k_2 , and k_3 . Arbitrary positive scalings of the constants do not affect extremal trajectories, so without loss of generality, we assume $k_1^2 + k_2^2 = 1$ (still assuming for the moment that at least one of k_1 or k_2 is nonzero), reducing the problem to search over two dimensions. In this section, we will show that knowing the structure of the trajectory gives a constraint that Hamiltonian be equal on the first and last segments, and essentially reduces this to a search in one dimension.

Since the space of controls is discrete, we can consider in turn all possible combinations of two controls, corresponding to the initial and final motions. For each one of these combinations, we will determine one parameter of

the control line for the corresponding optimal trajectory: either a point on the control line, or the control line's direction.

The method used to determine one parameter of the control line varies, depending on the kind of end segments we are considering. These are all possible cases:

1. *The initial and final motions are rotations of the same angular velocity.* The conservation of the Hamiltonian imposes that the two rotation centers be equidistant from the control line. Therefore the control line is parallel to the line determined by the two rotation centers.
2. *The initial and final motions are rotations of different angular velocities.* The Hamiltonian equations on both end segments are of the form

$$H = -y\omega \quad (69)$$

Since the angular velocities are not equal, let ω_1 be the smaller one. Let d be the distance between the two rotation centers and let P be the intersection point of the control line and the IC_1IC_2 line. Then the distance between P and IC_1 is $d \frac{\omega_1}{\omega_2 - \omega_1}$. This does not depend on H , so the location of P can be determined, which constrains the control line to pass through a point.

3. *The initial and final motions are a rotation and a translation.* This is very similar to the case above. The perpendicular from the rotation center to the translation's direction serves the role of the line uniting the two centers.
4. *The initial and final motions are translations of different speed vectors.* Let \vec{v}_1 and \vec{v}_2 be the two translation speed vectors. Let \vec{c} be a unit vector parallel to the control line. Then $H(t_1) = \vec{v}_1 \vec{c}$ and $H(t_2) = \vec{v}_2 \vec{c}$. Therefore $\vec{c}(\vec{v}_1 - \vec{v}_2) = 0$, so the control line is perpendicular to $\vec{v}_1 - \vec{v}_2$.
5. *The initial and final motions are translations of the same speed vector.* For such trajectories, it is always possible to construct an infinite number of time-equivalent trajectories by increasing the initial translation and decreasing the final translation by the

same amount; \dot{x} is the same for each translation segment. If the initial translation is of zero length, the start is at a switch and the control line passes through one of the initial locations of the switching points.

For any given start and goal, the above considerations transform the search for the location of the control line into a one-dimensional search. The search can be resolved through numerical methods. In each case, the second parameter of the control line is a real number p . For each p , there is a corresponding trajectory, beginning at the start configuration. Let $D(p)$ be the minimum distance, in the plane, between the goal and a configuration on the p -trajectory of the same θ as the goal. The problem can be re-stated as a requirement to minimize $D(p)$.

6 Tacking and tangent trajectories

The method to generate generic trajectories can also be applied starting in y, θ configurations that do correspond to critical values of H . However, in such cases the algorithm is no longer deterministic. Sooner or later, a configuration will be reached such that there exist multiple sustainable maximizing controls. Furthermore, there will exist a maximizing control, or a chattering of maximizing controls, such that the resulting motion is a translation that maintains this situation.

6.1 Tangent trajectories

Tangent trajectories contain switches at which there exist two or more sustainable controls, at most one of which is a translation.

There are thus two ways in which tangent trajectories are unconstrained by the Maximum Principle: first, in the choice of control at some switches; second, in the duration of certain controls that perpetuate critical y, θ configurations. However, the generating algorithm from the previous section can still determine some of the shape of tangent trajectories. Particularly, the *generic segments* between two critical points are fully determined, given an initial choice of a non-translating control. The following modifications to the generating algorithm make it output all generic segments for a given critical H :

1. At a critical state, ignore all controls that perpetuate this state.
2. Branch off a different computation for each one of the other controls that can be chosen.

Keeping a list of visited switches at generic points will allow the above computation to stop after $O(n^2)$ steps. Since all the translations parallel to the control line can be consolidated into a single segment to generate a time-equivalent trajectory, tangent trajectories have been shown to consist of a number of fully determined generic segments, appropriately connected, plus an arbitrary length translation, parallel to the control line, at one of the connections between generic segments.

6.2 Tacking trajectories

Tacking trajectories contain switches for which two or more sustainable controls are translations. Tacking trajectories are similar to the tangent trajectories, insofar as a chattering of the sustainable translations can generate a translation parallel to the control line. The determination of the generic segments can proceed in the manner of the previous subsection. The main difference is that the exit from the critical segment can also occur at a different y than the entrance. This problem can be dealt with by adding an extra branch to the computation: apply each of the translation controls alone.

It is possible to consolidate any number of translations at the same heading along a tacking trajectory into a pair of translations.

Theorem 6 *For every optimal trajectory, there exists an equivalent optimal trajectory that does not contain more than two translations at the same heading of the body.*

Proof: Let p be the net translation along a trajectory due to pure translations at some heading, and let T_p be the net duration of the pure translations at that heading. Let u_p be the translation control such that $T_p u_p = p$. If u_p is on the convex hull of the translation controls, then choose the two adjacent translation control vertices u_a and u_b ,

and apply the controls for times:

$$T_a = T_p \frac{\|u_p - u_b\|}{\|u_b - u_a\|} \quad (70)$$

$$T_b = T_p \frac{\|u_p - u_a\|}{\|u_b - u_a\|}. \quad (71)$$

Notice that if u_p is not on the convex hull of the translation controls, then the original trajectory was not optimal – scale u_p to the convex hull, and find controls u_a and u_b as above; the constructed trajectory is faster.

7 Constant-angular-velocity trajectories

So far, we have ignored the class of extremals corresponding to the case where $k_1 = k_2 = 0$. This section now deals with that case, and gives results that are essentially independent of the rest of the paper except theorem 1.

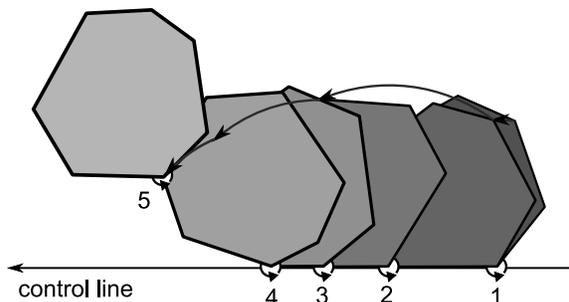
The Hamiltonian is simple:

$$H = k_3 \dot{\theta}. \quad (72)$$

Either k_3 is greater than zero, or less than zero. ($k_3 \neq 0$, since the Maximum principle restricts H from being identically zero.) If k_3 is positive, then any control with maximum ω satisfies the Maximum Principle. Otherwise, any control with minimum ω satisfies the principle.

In the simplest case, the controls for which the minimum and maximum values of ω are attained are unique. Then these trajectories are simple: constant controls, corresponding to pure rotations around a fixed rotation center. The more interesting case is when multiple controls maximize or minimize ω . The Maximum Principle does not directly give any information about when to switch between the controls, since the translation component is in the kernel of the Hamiltonian.

Under what circumstances might such a trajectory be optimal? The classic example is a Reeds and Shepp car that can reverse as well as go forwards. Consider the goal of spinning this car in place. A direct spin is not an available control, and a human driver would execute a three-point turn. The driver might move forwards around the left rotation center, with positive angular velocity, then backwards around the right rotation center, with positive



■ **Figure 9:** Example of a roll and catch trajectory. The polygonal control surface rolls along the control axis with constant angular velocity. When the last rotation center is put in place, the last motion is an off-axis rotation around this point (the “catch” stage). The trajectory of the last rotation center is shown, as well as the locations in the world frame of all the rotation centers used along the trajectory.

angular velocity, then forwards again around the left rotation center. Along this trajectory, angular velocity is positive and constant.

How long does the three-point turn take? It is simply the angle to be traversed divided by the angular velocity. Of course, the driver could also follow a four-, five-, or six-point turn, taking the same time, but following a very different trajectory. Therefore, we expect that there may be many optimal trajectories between configurations for which the amount of angle to turn through is the limiting factor, rather than the distance to be travelled. Rather than constructing all such trajectories, we will show that there is a canonical trajectory structure, which we call ‘roll-and-catch’ (see fig. 9 for an example) that we can use to always find one optimal trajectory. We also show that for this canonical trajectory structure, we can find the precise parameters of the optimal trajectory for every start and goal; this is something that appears to be much more difficult for typical generic trajectories with one of k_1 or k_2 nonzero.

We will refer to constant-angular-velocity trajectories by the short name ‘whirls’ in this section. A direct application of the Maximum Principle does not give much information about whirls past the fact that angular velocity is constant and maximized or minimized over the trajectory. However, we will show that we can apply the Maximum Principle to an alternate formulation restricting the

trajectories to a fixed structure, the roll-and-catch.

The problem of finding the optimal whirl trajectories can be restated equivalently in the following way. Consider a closed surface of rotation centers Z in the plane, containing at least two distinct points, and a vehicle that this surface is attached to. The vehicle can rotate at angular velocity 1 around any point in Z . (The clockwise case is symmetric.) Do time-optimal trajectories exist for such a vehicle, and if so, is there a method to construct an optimal trajectory for given start and end configurations, q_0 and q_f respectively?

7.1 Existence of optimal trajectories

Observe that if q_0 and q_f form an angle of $\alpha_{0f} \in [0, 2\pi)$, then any trajectory connecting q_0 to q_f will take time $2k\pi + \alpha_{0f}$, where k is a positive integer. If two trajectories from q_0 to q_f take different times, these times must differ by a multiple of 2π . Regarding the existence of optimal trajectories, it is therefore sufficient to establish controllability. We will then characterize a particular class of optimal trajectories that exist between any pair of configurations, and apply the non-autonomous version of Pontryagin's Maximum Principle to derive the shape of these kinds of trajectories.

We will show the vehicle is controllable as long as Z contains at least two distinct points, A and B . Consider A as the origin of the vehicle's frame and AB as the x axis and assume the length of the segment AB is 1. Rotating around A is therefore a spin in place. Translations in the y direction can be achieved by the following control sequence (see fig. 10):

1. Rotate around A for time ϵ .
2. Rotate around B for time $2\pi - 2\epsilon$.
3. Rotate around A for time ϵ .

The resulting translation is of length $2 \sin \epsilon$ in the y direction, therefore arbitrarily small. For $\epsilon = \frac{\pi}{2}$, the above control sequence results in a translation of length 2 in time 2π .

The vehicle can therefore be controlled in a manner similar to a unicycle, with turn-drive-turn trajectories that can reach any point in the plane. Furthermore, given a distance of d between the initial and target positions of point

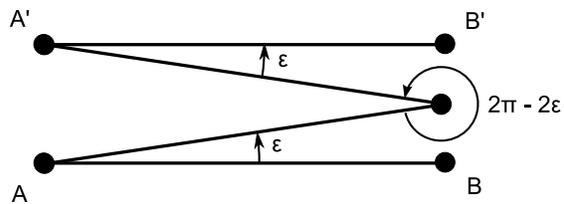


Figure 10: Achieving an infinitesimal translation by three positive rotations around points A and B . The first rotation is of angle ϵ around A , the second rotation of angle $2\pi - 2\epsilon$ around B , and the third rotation is of angle ϵ around A again.

A , turn-drive-turn control will reach the target in time less than 2π (initial turn) + $2\pi \lceil \frac{d}{2} \rceil$ ("straight" driving) + 2π (final turn), which is linear in the distance to be travelled. Therefore the vehicle is controllable and there exists an optimal trajectory between any pair of configurations.

Note that the vehicle is not small-time locally controllable, as even infinitesimal translations will need time at least 2π . However, if different targets are close together, the trajectories to them will also differ by very small amounts.

7.2 Convexification of the control surface

We will next show that we can, without loss of generality, replace Z by its convex hull. Given any point O on the convex hull of control surface Z , we will show that the vehicle can simulate any rotation around O by an equal time three-point turn. For any trajectory of a vehicle that is controlled by using the convex hull of Z , we can therefore generate an equal time trajectory for the original vehicle by replacing rotations around points outside Z by three-point turns around points in Z . The trajectory corresponding to the original non-convex surface will have no more than three times as many switches than the trajectory corresponding to the convex hull.

Given points A and B in Z , point O on the AB segment and a desired turning angle $\alpha \in [0, 2\pi)$, we build a trajectory of time α that results in a rotation around point O (see figure 11). Let A_0, B_0 be the initial positions in the world frame of points A and B ; let A_f, B_f be desired-final positions, after a rotation of angle α around O . Let O_0 be the initial position of O , and let b be the bisector of angle $\angle B_0 O_0 B_f$. Note that b is also the bisector of seg-

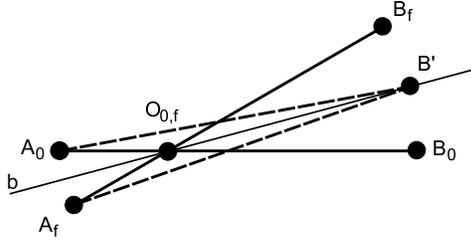


Figure 11: Achieving an arbitrary positive rotation around point O on the segment AB by a sequence of three positive rotations. The first rotation is around A and brings B on top of bisector b . The second rotation is around B and brings A into its intended final position. The third rotation is around A again and achieves the desired configuration.

ments A_0A_f and B_0B_f . Then the following three-point turn achieves the A_f, B_f configuration:

1. Rotate around A until B crosses b . Let the crossing point be called B' .
2. Since b is the bisector of segment A_0A_f , A_0 and A_f are equidistant from B' . Rotate around B until A is brought to A_f .
3. Rotate around A until B is brought to B_f .

It is easy to verify that this trajectory attains the desired configuration. Since B only crosses b once, the time taken by the three-point turn is less than 2π , and therefore the time of the three-point turn is α . So we can “simulate” rotations around an arbitrary point on the AB segment by using three-point turns around A and B . The vehicle is therefore capable of arbitrary rotations around any point inside the convex hull of Z , and given a vehicle we can, without loss of generality, replace Z by its convex hull (in fact, the corners of the convex hull are sufficient). Thus we will assume in the following that Z is convex.

7.3 A sufficient family of trajectories for optimality

In a direct application, the Maximum Principle does not place any constraints on whirl trajectories. We will identify a class of optimal trajectories that always exist and are composed of two stages, with the objective of applying the Maximum Principle to characterize the shape of

the first stage. Given a whirling vehicle with convex control surface Z , we will call an *xy stage trajectory for point A* ($A \in Z$) a trajectory between two configurations q_0, q_f that satisfies the following two conditions:

1. The first stage of the trajectory places A on its correct position in q_f in as short a time as possible. We will call this the *xy stage*, as only the x and y coordinates for A need to be attained.
2. The second stage of the trajectory is a rotation around A , until q_f is attained.

xy stage trajectories always exist between a given pair of configurations, as the vehicle is controllable. Given two such configurations, q_0 and q_f , consider an optimal trajectory and an *xy stage trajectory* between them, respectively. Let t_f be the time taken by the optimal trajectory, let t_1 and t_2 be the respective times taken by the two stages of the *xy stage trajectory*. Since the optimal trajectory does place A in its correct location, $t_1 \leq t_f$; therefore $t_f \leq t_1 + t_2 < t_f + 2\pi$. But the times of two trajectories between the same pair of configurations must differ by a multiple of 2π ; therefore $t_1 + t_2 = t_f$ and the *xy stage trajectory* is optimal. *xy stage trajectories* are therefore a class of optimal trajectories that always exist. In the following, we will confine our efforts to characterizing this class of optimal trajectories and finding a method to always construct one such trajectory.

7.4 The Maximum Principle

The configuration space for the *xy stage* is two-dimensional, containing only the x and y coordinates. This makes it possible to remove the θ coordinate from the state, and re-apply the Maximum Principle. However, removing θ from the state makes the configuration space velocity depend on time: $\dot{q} = f(q, u, t)$. So we need to apply the non-autonomous version of the Maximum Principle. The Maximum Principle for time-optimal trajectories for non-autonomous vehicles is very similar to the version used previously in this paper, with the exception that the function H is only required to be positive and not necessarily constant. Taking the final rotation center as a reference point in the frame of the body, we obtain the condition that, along the *xy stage*, the function $H(x, u, t)$

needs to be maximized by the chosen control at each point along the trajectory, where

$$H = \lambda_1 \dot{x} + \lambda_2 \dot{y} \quad (73)$$

and the (λ_1, λ_2) vector is non-null. Since $\frac{d\lambda_1}{dt} = \frac{\partial H}{\partial x} = 0$, $\lambda_1(t)$ is a constant function. Similarly, $\lambda_2(t)$ is a constant as well, therefore

$$H = k_1 \dot{x} + k_2 \dot{y} \quad (74)$$

with $k_1^2 + k_2^2 > 0$. Therefore, for each xy stage optimal trajectory, there exists a control direction in the plane, given by the vector (k_1, k_2) such that the optimal control is maximal when projected onto this direction.

7.5 Shape of the xy stage

Consider, along an xy stage optimal trajectory, the time when the control switches from u_i to u_{i+1} , corresponding to rotation centers R_i and R_{i+1} respectively. Consider the functions $H_i = H(x, u_i, t)$ and $H_{i+1} = H(x, u_{i+1}, t)$. Both these functions are continuous. Immediately before the switch, $H_i \geq H_{i+1}$ and immediately after the switch $H_i \leq H_{i+1}$. Therefore, H_i and H_{i+1} are equal at the time of the switch.

Furthermore, at the switching time, let P be the reference point and v_i and v_{i+1} its velocity vectors immediately before and after the switch respectively. Since the angular velocity is constant, the lengths of the vectors v_i and v_{i+1} are proportional to the lengths of the segments PR_i and PR_{i+1} respectively. The angle between the pre-switch and post-switch velocity vectors is furthermore equal to the angle $\angle R_i P R_{i+1}$, as the velocities are perpendicular to the radii. So the triangle formed by the two velocity vectors is proportional to the triangle $\Delta R_i P R_{i+1}$ and these two triangles form an angle of $\frac{\pi}{2}$.

Since $H_i = H_{i+1}$ at the time of the switch, and these two functions are the projections of v_i and v_{i+1} onto the control direction, the third side of the triangle formed by v_i and v_{i+1} is perpendicular onto the control direction. This side corresponds to $R_i R_{i+1}$ in the proportional and rotated by $\frac{\pi}{2}$ triangle; therefore $R_i R_{i+1}$ is parallel to the control direction. This holds for all switches along the xy stage. Therefore, in the world frame, all the rotation centers used during the xy phase are found on a line passing

through the first rotation center and parallel to the control direction.

Setting the world reference frame on this axis, we notice by a similar argument that, if R_j is placed higher, in respect to the control line, than R_k , then $H_j > H_k$. Since the first rotation center used is on the control line, assuming the control direction points right to left, all the other rotation centers must be above the control line in the initial state; this condition is then propagated along the trajectory. We have therefore proven the following:

Lemma 11 *For each xy stage optimal trajectory, there exists a control line such that the xy stage is a rolling of control surface Z in the positive direction along the control line.*

The above is a necessary but not a sufficient condition for xy stage optimality. Because the xy stage is shown to be a roll, we will alternatively call xy stage trajectories “roll and catch” trajectories.

7.6 The position of the control line for known initial and final controls

Number the corners of the convex hull consecutively clockwise as R_1, R_2, \dots, R_m . Assume we knew that the initial control is a rotation around R_1 and the final two controls are rotations around R_k and R_f respectively. The optimal control is therefore R_1, R_2, \dots, R_m repeated n times (where n is an unknown integer) and then $R_1, R_2, \dots, R_{k-1}, R_k, R_f$.

In the world frame, let R'_1 be the initial location of R_1 , let R'_f be the final location of R_f and let R'_k be the location, at the final switch, of R_k . We are given the position of R'_1 , and we know the position of R'_f as the final motion is a rotation around this point. Let d'_{1f} be the length of the segment $R'_1 R'_f$. In order to determine the structure of the trajectory (if it exists), it is sufficient to find the position of R'_k , which determines the control line $R'_1 R'_k$.

Let r_{ij} be the distance, in the body frame, between two arbitrary rotation centers R_i and R_j . Let $l_i = r_{i,i+1}$, i.e. the length of the i th side of the control surface. Let $p = \sum_{i=1}^m l_i$ be the perimeter of the control surface Z . Since the trajectory is a roll along the control line, the length of the segment $R'_1 R'_k$ is:

$$d'_{1k} = np + l_1 + l_2 + \dots + l_{k-1} \quad (75)$$

In the triangle $\Delta R'_1 R'_k R'_f$, the triangle inequality must hold:

$$d' - r_{kf} < np + l_1 + l_2 + \dots + l_{k-1} \leq d' + r_{kf} \quad (76)$$

(The left-hand side is a strict inequality because, as shown above, if two rotation centers are on the control line at the same time, the one that is used on the immediately preceding interval must have a smaller x coordinate.)

Since r_{kf} is a section through Z , $2r_{kf} \leq p$. Note that relation 76 has a span of $2r_{kf}$ between the leftmost and rightmost side, and the middle changes in increments of p . Therefore relation 76 has at most one solution for the unknown integer n , which we can obtain by subtracting and dividing appropriately and taking the floor function:

$$n = \lfloor \frac{d' + r_{kf} - (l_1 + l_2 + \dots + l_{k-1})}{p} \rfloor \quad (77)$$

Furthermore, n needs to satisfy the left-hand side of 76 above. By replacing this solution into equation 75 above, we determine the length of d'_{1k} . This fully determines the triangle $\Delta R'_1 R'_k R'_f$ and, by extension, the position of the control line and the structure of the trajectory. In order for the xy stage to be extremal, we only need to check observance of the Maximum Principle at its final point, i.e. calculate the configuration at the switch from R_k to R_f and verify that no point of Z is above the control line in this configuration.

Since there is at most one solution for the location of the control line, if such a solution exists then the corresponding xy stage is the fastest way to get R_f into its final position by using R_1 as the first rotation center and R_k as the last.

7.7 Constructing an xy stage trajectory for given initial and final configurations

Given a control surface and initial and final configurations q_0 and q_f respectively, we have proven above that there exists a “roll and catch” optimal trajectory between these configurations, for any choice of reference point on the control surface (the reference point being the location of the last rotation center that is used on the trajectory). We have also shown how to find this trajectory, if the initial

and the second to last controls were known; these two controls have to be rotations around corners of the control surface, which we have shown can be assumed without loss of generality to be a convex polygon. Therefore, the following simple algorithm is valid:

1. Enumerate all possible ordered pairs of corners of the control surface.
2. For each such pair, construct the “roll and catch” trajectory (there can exist at most one) that corresponds to the chosen initial and second to last controls.
3. Pick the fastest trajectory.

The algorithm evidently runs in time that is $O(m^2)$, where m is the number of corners of the polygonal control surface. For a given control surface, the running time is constant. However, if the control surface has a very high number of corners (e.g. it is bounded by a smooth convex curve), this algorithm is impractical.

8 Conclusion

Although many of the details of the proofs are technical, or required significant algebraic manipulation, most of the results simply verify things that we might suspect almost immediately from theorem 1. There are some constants of integration in the necessary condition for optimal trajectories; varying those constants changes the structure of trajectories. Geometrically, the constants describe a line in the plane. Optimal trajectories can be described by the fact that they in some sense maximize ‘effort’ along this line, where effort is defined as the speed of some point rigidly attached to and sliding along the line, pushed by the rigid body.

If there are discrete controls for the rigid body, then over most periods of time one of the controls is active, and the body translates or rotates. This happens until the orientation of the body or the distance of the body from the control line changes sufficiently that some new control maximizes effort.

For most values of the constants and initial configurations of the rigid body (or equivalently, for most initial configurations of the body relative to the control line), the extremal trajectory that is generated is well-behaved

– piecewise-constant in the controls, with switches between controls that occur at well-defined points. Varying the configuration of the body slightly relative to the line usually does not change the structure of trajectories. There are a finite number of trajectory structures, and we can enumerate them by considering certain specific initial configurations of the body relative to the control line.

However, there are also special cases, and these cases turn out to be an important class of solutions to many minimum-time problems. For example, to move a differential-drive or a steered car to a distant configuration, the fastest trajectory is to turn to face the goal, drive to the goal, and turn to the goal angle. This trajectory is a ‘tangent’ trajectory, and at any point while the body translates along the line, all three of the controls ‘forwards’, ‘left’, and ‘right’ maximize the Hamiltonian equation. Thus, the time of the translation segment cannot be determined from the value of the constants – trajectories with a long translation period in the middle satisfy the necessary condition for optimality, as do trajectories with a short translation.

Other special cases occur due to the commutativity of pure translation controls, or from chattering due to three or more controls. Fortunately, in each of these cases, the location of the control line (and the value of H) can be computed directly from the Hamiltonian equation. Unfortunately, this is not always enough to specify the shape of the trajectory.

There is a final special case, in which the control-line geometry cannot be applied in the same way – constant-angular-rotation trajectories. We are familiar with a particular example of this, three-point parking of a steered car. The goal is to turn the vehicle through some angle, and the duration of the trajectory depends on that angle and the angular velocity. As long as the amount of translational driving done while the body spins is kept below some threshold, translation is ‘free’ – so there may be many very different trajectories to the goal, all of which are optimal. However, we have shown how to derive at least one optimal trajectory between each start and goal configuration.

Although we have described the structure that optimal trajectories must follow, we have left open the problem of computing the particular trajectory that connects a start and goal pair. For symmetric systems with only a few controls (including the car-like, differential-drive,

and omnidirectional systems studied in previous work), it is straightforward to derive the optimal trajectory of each structure between a pair of configurations, and then choose the fastest of those trajectories. For more complicated systems, it may be necessary to rely on numerical optimization. Nonetheless, we expect that knowing the types of trajectories that may arise will be useful in the design of planners and efficient robots.

References

- [1] Pankaj K. Agarwal, Therese Biedl, Sylvain Lazard, Steve Robbins, Subhash Suri, and Sue Whitesides. Curvature-constrained shortest paths in a convex polygon. In *SIAM J. Comput.*, pages 392–401. ACM Press, 1998.
- [2] Ron Alterovitz, Michael Branicky, and Ken Goldberg. Motion planning under uncertainty for image-guided medical needle steering. *International Journal of Robotics Research*, 27(11):1361–1374, 2008.
- [3] Devin J. Balkcom, Paritosh A. Kavatkar, and Matthew T. Mason. Time-optimal trajectories for an omni-directional vehicle. *International Journal of Robotics Research*, 25(10):985–999, 2006.
- [4] Devin J. Balkcom and Matthew T. Mason. Time optimal trajectories for differential drive vehicles. *International Journal of Robotics Research*, 21(3):199–217, 2002.
- [5] Jérôme Barraquand and Jean-Claude Latombe. Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles. *Algorithmica*, 10:121–155, 1993.
- [6] Hamidreza Chitsaz, Steven M. LaValle, Devin J. Balkcom, and Matthew T. Mason. Minimum wheel-rotation paths for differential-drive mobile robots. In *IEEE International Conference on Robotics and Automation*, 2006. To appear.
- [7] M. Chyba and T. Haberkorn. Designing efficient trajectories for underwater vehicles using geometric control theory. In *24rd International Conference on Offshore Mechanics and Arctic Engineering*, Halkidiki, Greece, 2005.

- [8] A. Theo Coombs and Andrew D. Lewis. Optimal control for a simplified hovercraft model. Preprint.
- [9] John J. Craig. *Introduction to Robotics: Mechanics and Control*. Addison-Wesley, second edition, 1989.
- [10] Guy Desaulniers. On shortest paths for a car-like robot maneuvering around obstacles. *Robotics and Autonomous Systems*, 17:139–148, 1996.
- [11] B.R. Donald, C.G. Levey, C.D. McGray, I. Paprotny, and D. Rus. An untethered, electrostatic, globally controllable mems micro-robot. *Microelectromechanical Systems, Journal of*, 15(1):1–15, Feb. 2006.
- [12] 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*, 79:497–516, 1957.
- [13] Andrei A. Furtuna, Devin J. Balkcom, Hamidreza Chitsaz, and Paritosh Kavathekar. Generalizing the Dubins and Reeds-Shepp cars: Fastest paths for bounded-velocity mobile robots. In *IEEE International Conference on Robotics and Automation*, pages 2533–2539, 2008.
- [14] Jean-Bernard Hayet, Claudia Esteves, and Rafael Murrieta-Cid. A motion planner for maintaining landmark visibility with a differential drive robot. In *Workshop on the Algorithmic Foundations of Robotics*, December 2008.
- [15] Tamás Kalmár-Nagy, Raffaello D’Andrea, and Pritam Ganguly. Near-optimal dynamic trajectory generation and control of an omnidirectional vehicle. *Robotics and Autonomous Systems*, 46:47–64, 2004.
- [16] Steven M. LaValle. *Planning Algorithms*. Cambridge Press, 2006.
- [17] Kevin M. Lynch and Matthew T. Mason. Stable pushing: Mechanics, controllability, and planning. *International Journal of Robotics Research*, 15(6):533–556, December 1996.
- [18] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes*. John Wiley, 1962.
- [19] 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.
- [20] 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.
- [21] Marc Renaud and Jean-Yves Fourquet. Minimum time motion of a mobile robot with two independent acceleration-driven wheels. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 2608–2613, 1997.
- [22] Paolo Salaris, Felipe A. W. Belo, Daniele Fontanelli, Luca Greco, and Antonio Bicchi. Optimal paths in a constrained image plane for purely image-based parking. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1673–1680, 2008.
- [23] 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.
- [24] Philippe Souères and Jean-Paul Laumond. Shortest paths synthesis for a car-like robot. *IEEE Transactions on Automatic Control*, 41(5):672–688, May 1996.
- [25] Héctor Sussmann. The Markov-Dubins problem with angular acceleration control. In *IEEE International Conference on Decision and Control*, pages 2639–2643, 1997.
- [26] Héctor 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. SYCON 91-10, Department of Mathematics, Rutgers University, New Brunswick, NJ 08903, 1991.
- [27] M. Vendittelli, J.P. Laumond, and C. Nissoux. Obstacle distance for car-like robots. *IEEE Transactions on Robotics and Automation*, 15(4):678–691, 1999.